A Computational Model of Action Selection and Action Specification in the Basal Ganglia

Priyadarshini Saha (p24saha@uwaterloo.ca)[1], Madeleine Bartlett (maddy.bartlett@nrc-cnrc.gc.ca)[2] & Jeff Orchard (jorchard@uwaterloo.ca)[1]

[1]David Cheriton School of Computer Science, University of Waterloo, Waterloo, ON N2L 3G1, Canada

[2]National Research Council of Canada, University of Waterloo, Waterloo, ON N2L 3G1, Canada

Abstract

It has been proposed that the basal ganglia is involved in selecting which action is performed, while the motor cortex specifies how the selected action is carried out. However, recent electrophysiological evidence from a reach-to-pull task challenges this view by showing that the motor cortex and the dorsal striatum in the basal ganglia jointly specify continuous action movement parameters, such as reach angle and pull force. This finding implicates the basal ganglia in both action selection and action specification, creating a gap for a mechanistic model that unifies both functions in a single, grounded framework. We address this gap with a biologically plausible model of the basal ganglia that leverages dynamic neural fields applied to high-dimensional representations of action-salience distributions. The dynamic neural field converts high-entropy representations into low-entropy representations without altering the underlying action-salience distribution's peak—a property previously shown to support accurate readout of continuous action parameters (e.g., movement speed or force). Our model is an extension that applies the same dynamic neural field transformation across multiple actions so that we can simultaneously (i) resolve which action is chosen (e.g., lever pull vs. button press) and (ii) specify the continuous action parameter value with which the chosen action is to be executed (e.g., force or speed). In our simulations patterned on the reach-to-pull paradigm, the basal ganglia model performs both discrete action selection and action specification of a continuous action parameter in a single operation, offering a concise mechanism for how basal ganglia-cortex circuits decide what to do and how to do it.

Keywords: basal ganglia; action selection; action specification; dynamic neural fields; discrete action; continuous action parameter; salience; action-salience distributions; action-place cells; action-salience bundles; entropy-reduction transformation

Introduction

Motor behaviour requires solving two coupled problems: selecting an action from discrete alternatives, and determining how the chosen action is to be executed. Traditionally, the basal ganglia (BG) has been framed as a discrete action selector that chooses which action to initiate from a finite set of possibilities (hereafter, "action selection"), for example choosing between pulling a lever or pushing a button, (Gurney et al., 2001a; Humphries et al., 2006; T. Prescott et al., 2000; Redgrave et al., 1999; Stewart et al., 2012). At the same time, the motor cortex is proposed to specify how the chosen action is executed (hereafter, "action specification") by setting values for its continuous action parameter(s), such as pull force or speed (Lemon, 2008; Picard & Strick, 2001; Scott, 2004; Shadmehr & Krakauer, 2008; Shenoy et al., 2013; Teka et al., 2017). Despite this classical division of labour, how discrete action selection is integrated with continuous action specification—and the extent to which BG circuitry contributes to both computations—remains an open question.

The BG comprise a set of interconnected subcortical nuclei in which widespread cortical signals converge in the dorsal striatum (dSTR), and are transformed through pallidal and subthalamic circuitry to influence behaviour via the globus pallidus internus (GPi) output to the thalamus, which in turn modulates action execution through thalamo-cortical loops (Rocha et al., 2023). Most computational models of action selection in the BG adopt a localist, channelized representation, in which each candidate action occupies a separate, non-overlapping neural channel (Berns & Sejnowski, 1998; Cohen & Frank, 2009; Gurney et al., 2001a, 2001b; Lindahl et al., 2013; O'Reilly & Frank, 2006; Ponzi, 2008; Soni & Frank, 2025; Stewart et al., 2010; Suryanarayana et al., 2019; Ursino et al., 2020; Wiecki & Frank, 2013). In these models, the activity of each channel is driven by a scalar salience signal that indicates the perceived favourability of that action for a given task or goal. Selection is implemented via lateral inhibition in an off-center/on-surround competition that reinforces the most salient (center) channel, while broadly suppressing competing (surround) channels (Gurney et al., 2001a; Humphries et al., 2006; T. J. Prescott et al., 2024; Stewart et al., 2010; Suryanarayana et al., 2019). An action is selected when inhibitory GPi activity for that channel is suppressed—typically when direct (D1; "Go") pathway drive outweighs indirect (D2; "NoGo") pathway suppression—thereby disinhibiting thalamic targets and facilitating action execution. Together these localist, channelized architectures clarify how the BG acts as an action selector over a repertoire of discrete actions.

A growing body of work has challenged the view of the basal ganglia as a purely discrete action selector and implicating it in action specification (Dhawale et al., 2021; Dudman & Krakauer, 2016; Fobbs et al., 2020; Park et al., 2025; Turner & Desmurget, 2010). Most recently, Park et al. (2025) showed that dorsal striatum (dSTR) and primary motor cortex (M1) jointly encode continuous action kinematics that determine *how* an action is executed. The researchers designed a task to systematically vary continuous action parameters (target location and pull force) while keeping the gross action ("reach-to-pull") constant. This design allowed them to distinguish whether the striatum selects only between actions or

In G. Boleda, I. Dautriche, M. Maldonado, C. Saldana and T. Torrent (Eds.), *Proceedings of the 48th Annual Conference of the Cognitive Science Society*

also specifies their continuous action kinematic details. Simultaneous electrophysiological recordings from dSTR and M1 revealed that the dSTR population activity could classify task conditions (location × force) and decode continuous action kinematics with accuracy comparable to M1. This finding supports the view that BG circuitry contributes to action specification alongside its established role in action selection.

The aforementioned localist models are adequate when the BG is treated purely as an action selector, but they struggle when the circuit must also specify a value for a continuous action parameter. Representing continuous action parameters in a localist framework necessitates binning these variables into discrete categories, sacrificing precision. Achieving either a wider range or finer precision for continuous action parameters demands proportionally more bins, linearly increasing the number of representational units (e.g. neurons or ensembles), and leading to poor scalability. Furthermore, these localist encoding methods are brittle because each bin is an isolated representational "slot": a particular continuous action value is carried primarily by the activity of its dedicated unit, with little overlap from neighbouring bins. Thus, if that unit is corrupted due to damage or noise the corresponding value is effectively lost with no mechanism for smooth degradation to a nearby approximation. Although intermediate values can be approximated by external interpolation or weighted averaging across neighbouring bins, such methods are ad-hoc, noise-sensitive, and require fine-tuning. These limitations compound when representing both discrete actions and their continuous action parameters: encoding multiple continuous action parameter values for each discrete action (e.g., button press at 3N, 3.5N, etc.) results in a combinatorial explosion of units that scales poorly as precision requirements increase. Therefore, localist encoding schemes are ill-suited for smooth control and generalization over continuous action parameters.

The model presented by Bartlett et al. (2025) addresses limitations of representing continuous action parameters in a localist framework by using a distributed encoding approach. Using Vector Symbolic Algebra (VSA) and Spatial Semantic Pointers (SSPs), the model encodes action-salience distributions (i.e. probability densities) over a continuous action parameter space as high-dimensional vectors (called action-salience bundles) represented by spiking neuron populations. The authors posit the BG as a "clean-up" system: striatal dynamic neural fields (DNFs; Schöner et al., 2015) sharpen the encoded action-salience distribution by reducing entropy around its peak without shifting it, thereby narrowing the range of values available to downstream selection mechanisms for action specification. This captures the key features of Park et al.'s reach-to-pull paradigm: a single gross action (reach-to-pull) is represented by its own neural channel with force as the continuous action parameter provided as an action-salience bundle. However, Bartlett et al.'s model is limited to a single action channel, and therefore does not address how the BG can jointly decide which action to perform and how to execute it when multiple discrete alternatives are available. We address this shortcoming in our paper by extending their framework to multiple discrete action channels, each with its own continuous action parameter (action-salience) distribution (e.g., force for lever pull and button press), enabling joint action selection and action specification.

Collectively, prior work has not yet demonstrated a single BG mechanism that selects an action while simultaneously specifying its execution value along a continuously valued action parameter. A related proposal, the affordance competition hypothesis by Cisek (2006), proposes that frontoparietal cortex encodes multiple candidate actions as distributed population activity over a continuous action space (e.g., reach direction), where kinematic action specification is intrinsic to the population code and action selection emerges from biased competition via mutual inhibition and top-down influences (e.g., goals, context, value). Joint action selection and specification occur when one representation crosses a dynamic decision threshold and suppresses its competitors. However, Cisek (2006) places the computation in the cortex, leaving the BG's role in this integrated process open to question.

In this paper, we present a BG model that jointly performs action selection and action specification in a single operation. Building on Bartlett et al. (2025), we retain VSA/SSP distributed encoding and striatal DNF dynamics, but extend the framework from a single fixed action to multiple discrete action channels, each coupled to its own continuous action parameter distribution (e.g., force). Each channel receives a scalar salience and a high-dimensional bundle encoding an action–salience distribution over its continuous action parameter domain. Joint action selection and action specification is achieved by applying the striatal DNF's entropy-reduction dynamics to the concatenated multi-channel action-salience bundle representation, so the same computation identifies the winning action while sharpening the corresponding continuous action parameter estimate into a low-entropy output for action specification. Although developed independently, our integrated mechanism closely mirrors Cisek (2006)'s hypothesis, while situating the competition within the BG.

Methods

In the following sections we summarize the methods of Bartlett et al. (2025) for representing and manipulating continuous action parameters in a biologically plausible form, then present our extended, multi-channel BG architecture.

Spatial Semantic Pointers

Spatial Semantic Pointers (SSPs) are high-dimensional vector representations of continuous quantities (often spatial coordinates) that use the Holographic Reduced Representation (HRR) (Plate, 1995) VSA. An SSP maps a point $x \in \mathbb{R}^m$ to a D-dimensional vector $\phi(x) \in \mathbb{R}^D$ such that similarity in the vector space reflects proximity in the underlying continuous space: nearby values yield similar vectors, distant values yield dissimilar vectors (Eliasmith, 2013; Eliasmith & Anderson, 2003). SSPs are typically constructed via a Fourier-phase

(HRR-compatible) embedding; we refer readers to (Eliasmith, 2013; Eliasmith & Anderson, 2003) for the formal construction and properties.

SSPs provide a smooth, similarity-preserving code for continuously valued action parameters, and because they support dimensionality-preserving operations such as bundling, they enable compact representations of quasi-probability (or action-salience) distributions over continuous action parameters within a fixed-dimensional vector space (for detailed discussions see Dumont (2025), Eliasmith (2013), Furlong and Eliasmith (2023), and Furlong et al. (2024)).

Bundling

In VSA, bundling (superposition) is a method for representing a set of vectors in a single vector of the same dimensionality. Given N D-dimensional vectors $\mathbf{A_1}, \ldots, \mathbf{A_N}$, their bundle is $\mathbf{B} = \mathbf{A_1} + \mathbf{A_2} + \cdots + \mathbf{A_N} = \sum_{i=1}^{N} \mathbf{A}_i \in \mathbb{R}^D$. To encode a distribution of salience over a continuous action parameter (e.g., force), we start with a continuous domain, $\mathcal{A} = [a_{\min}, a_{\max}]$ (for our simulations, a scalar interval), and choose N evenly spaced samples $a_j \in \mathcal{A}$ for $j \in \{1, \ldots, N\}$. Let $\phi(a_j)$ denote the SSP encoding of a_j, and let a_j have an associated salience $s_j \geq 0$. The resulting action-salience bundle is,

$$\mathbf{B} = s_1\phi(a_1) + \cdots + s_N\phi(a_N) = \sum_{j=1}^{N} s_j\phi(a_j) \in \mathbb{R}^D \ .$$

To make the action-salience profile-to-bundle mapping explicit, we precompute an SSP basis matrix by stacking the sampled continuous action SSPs row-wise,

$$E \equiv \Phi \equiv \begin{bmatrix} \phi(a_1)^\top \\ \vdots \\ \phi(a_N)^\top \end{bmatrix} \in \mathbb{R}^{N \times D}, \quad \Phi_j := \phi(a_j),$$

and define the sampled action-salience profile as $\mathbf{s} = [s_1, \cdots, s_N]^\top \in \mathbb{R}^N_{\geq 0}$. Then bundling can be written compactly as $\mathbf{B} = E^\top \mathbf{s} \in \mathbb{R}^D$, so $E^\top$ acts as an encoder that maps an N-dimensional action-salience profile into a D-dimensional action-salience bundle.

The salience s_j for a_j can be decoded by dot-product: $\hat{s}_j = \mathbf{B} \cdot \phi(a_j)$, where $\hat{s}_j$ is the estimated salience of a_j. Equivalently, the full decoded action-salience profile is obtained as $\hat{\mathbf{s}} = E\mathbf{B} \in \mathbb{R}^N$ where $\hat{\mathbf{s}} = [\hat{s}_1, \ldots, \hat{s}_N]^\top$. Thus $E^\top$ and E form an explicit encoder-decoder pair between an action-salience profile and its bundle. Additionally, because SSP similarity varies smoothly with distance in $\mathcal{A}$, the same decoding mechanism supports interpolation at intermediate continuous action parameter values that are not explicitly encoded or sampled. For $a_j < a^* < a_{j+1}$, we estimate the salience of a^* as: $\hat{s}^* = \mathbf{B} \cdot \phi(a^*)$. Furthermore, since bundling is dimensionality-preserving (we use $D = 512$), large or dense continuous action parameter domains can be represented without increasing the vector representational size. Together, these properties make bundles compact representation of continuous action-salience distributions that can be transformed by downstream DNF dynamics for action specification.

Dynamic Neural Fields

Dynamic neural fields (DNFs) are recurrent neural networks defined over a continuous feature space (e.g., spatial location or force magnitude) (Schöner et al., 2015). Following Bartlett et al. (2025), we represent a continuous action parameter domain $\mathcal{A}$ by discretizing the DNF's action axis with N action-place (AP) cells—analogous to hippocampal place cells, but tuned to action parameter values rather than spatial locations. Each AP cell j is tuned to sampled continuous action $a_j \in \mathcal{A}$, so that its activity encodes the salience s_j of the action-salience distribution at a_j.

Given this representation, the DNF's recurrent dynamics implement an entropy-reduction (sharpening) transformation. External inputs set an initial high-entropy action-salience profile over the N AP cells. Then, the DNF's local excitation and global inhibition interactions (Schöner et al., 2015) sharpen this profile by concentrating activity around the most salient location while suppressing competitors by driving activity towards the baseline. The dynamics converge to a single, sharp peak centered at the most salient continuous action parameter value (i.e., the operation does not translate the peak), along the continuous action parameter axis, and thereby yielding a low-entropy action-salience profile suitable for robust downstream action specification.

Model Architecture

Our BG network (Figure 1) follows the canonical localist, channelized architecture (Frank, 2005; Gurney et al., 2001a, 2001b; Stewart et al., 2010). We omit lateral inhibition motifs; instead, competition is resolved by striatal DNF dynamics together with downstream GPi gating (Bartlett et al., 2025). Following the Bartlett et al. (2025) model, direct (D1) and indirect (D2) pathways use the same fixed gain, abstracting away pathway-specific asymmetries present in vivo (see Sec. Online Resources for model implementation).

Connections from cortex provide M high-entropy action-salience bundles $\{\mathbf{B}_m\}_{m=1}^{M}$, with $\mathbf{B}_m \in \mathbb{R}^D$ for $m \in \{1, \cdots, M\}$, as input to each action channel m. Each $\mathbf{B}_m$ encodes an action-salience distribution over a shared continuous action parameter domain $\mathcal{A}$, sampled at N evenly spaced points; for us, all channels encode the same domain space at the same sampling resolution.

Following Bartlett et al. (2025), for each channel m, we decode action-salience bundle $\mathbf{B}_m$ using decoder $E^\top$ to obtain an N-dimensional action-salience profile $\mathbf{s}_m \in \mathbb{R}^N$. We concatenate the M action-salience profiles into a single NM-dimensional high-entropy action-salience profile vector as $\mathbf{S} = [\mathbf{s}_1^\top | \cdots | \mathbf{s}_M^\top]^\top \in \mathbb{R}^{NM}$ which drives the striatal DNF dynamics. The DNFs are tiled by AP-cells arranged as M contiguous blocks of N AP-cells each. Within block m, AP-cell j is tuned to $a_j^{(m)}$ and its activity encodes salience $s_j^{(m)}$ for $j \in \{1, \ldots, N\}$. Together, these steps implement our key extensions: (i) an explicit multi-action channel design with M parallel discrete action channels, each encoding a continuous action parameter via an action-salience bundle; and (ii) con-

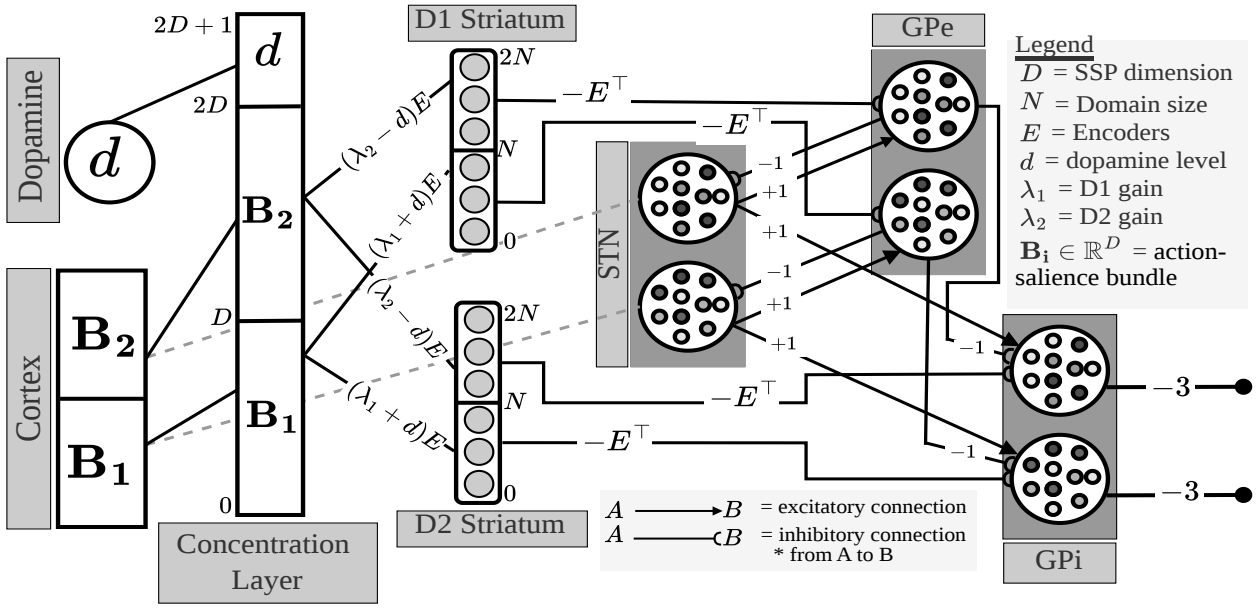

Figure 1: Basal ganglia model architecture for joint action selection and action specification.

catenation of action-salience profiles so that, operationally, the DNFs receive a single multi-modal NM-dimensional input spanning M action channels.

From $\mathbf{S}$, the DNF dynamics converge to a single localized peak. The block containing this peak identifies the winning channel, and the within-block peak index j^* specifies the continuous action value by mapping j^* back to the continuous action value $a_{j^*} \in \mathcal{A}$. Thus, action selection and action specification are resolved by a single entropy-reduction operation.

After DNF convergence, connection weights $E^\top$ (see Sec. Bundling) re-encode the NM-dimensional striatal activities into M D-dimensional low-entropy action-salience bundles, which then propagate through standard BG pathways. These low-entropy action-salience bundles then interact with their original high-entropy cortical action-salience bundle counterparts via the subthalamic nucleus (STN), providing a substrate for explore-exploit tradeoffs. The BG output is a concatenation of M low-entropy action-salience bundles (equivalently a DM-dimensional vector). We decode this output into an low-entropy NM-dimensional action-salience profile and apply an argmax readout to identify (i) the winning channel (action selection), and (ii) within-channel peak index j^*. In vivo, we propose that an analogous computation for action specification readout may be realized by the GPi-thalamus-cortex loop over the low-entropy action-salience bundle of the winning channel.

Constructing Basal Ganglia Input

To support simultaneous action selection and action specification, it is necessary to construct a BG input that encodes both the discrete channel salience across actions and within-channel continuous action-salience profile in such a form that DNF's dynamics would be sufficient to resolve competition. We therefore define these quantities separately and then combine them into a single fixed-dimensional representation.

For each discrete action channel $m \in \{1, \ldots, M\}$, we distinguish two quantities that must be jointly represented at the cortical input to the BG. First, each channel has a scalar channel salience $c_m \in [0, 1]$, which denotes the perceived favourability of discrete action m in relation to all other discrete alternatives. Second, each channel carries a continuous action-salience profile $\mathbf{s}_m \in \mathbb{R}^N_{\geq 0}$ over the shared continuous action parameter domain $\{a_j\}_{j=1}^N \subset \mathcal{A}$. This action-salience profile is represented as $\mathbf{s}_m = [s_1(a_1), \ldots, s_N(a_N)]^\top$. Here, the action-salience value $s_m(a_j)$ specifies how strongly execution value a_j is preferred for the continuous action parameter, conditional on selecting discrete action m.

A central challenge is to encode both the across-channel salience (c_m) and the within-channel continuous action parameter salience ($\mathbf{s}_m$) within a fixed-dimensional representation. Since our model resolves action selection and action specification in a single BG computation, both signals must be encoded simultaneously. We therefore employ two preprocessing steps to $\mathbf{s}_m$. This input construction is an engineering choice to reconcile established action selection and action specification modelling conventions within a single BG computation; we do not claim that cortex performs these exact preprocessing operations, only that multiple cortical processes could generate equivalent bundle inputs, and we adopt this procedure as a concrete instantiation.

First, we perform unit-height normalization that equalizes

the peak magnitude across all discrete action channels:

$$\tilde{\mathbf{s}}_m = \frac{\mathbf{s}_m}{\max_j s_m(a_j)} \in [0,1]^N,$$

This normalization procedure removes arbitrary scale differences in the action-salience profiles by rescaling each $\mathbf{s}_m$ so that its maximum value is 1, forcing all peaks to have equal height while preserving the distribution's shape and peak location. This step is crucial because, for action specification, the selected continuous execution value is determined by the peak location (given by index j^*), not the unnormalized peak height (given by $s_m(a_{j^*})$).

Second, we apply a channel-salience scaling step, where we encode the across-channel salience by scaling the normalized action-salience distribution $\tilde{\mathbf{s}}_m$ by c_m for discrete action m:

$$\mathbf{q}_m = c_m \tilde{\mathbf{s}}_m \in [0,1]^N.$$

The scaling of $\tilde{\mathbf{s}}_m$ by c_m sets the peak height of $\mathbf{q}_m$ to reflect how salient the discrete action m is relative to other action channels. Simultaneously, this scaling procedure preserves the shape and peak location of the within-channel action-salience distribution $\tilde{\mathbf{s}}_m$ required to recover the preferred continuous execution value for the continuous action parameter. Thus, channel identity is encoded before concatenation: c_m specifies action salience, whereas $\tilde{\mathbf{s}}_m$ specifies the continuous action parameter profile.

Finally, we encode $\mathbf{q}_m$ into a D-dimensional cortical action-salience bundle using the encoder $E \in \mathbb{R}^{N\times D}$ (see Sec. Bundling): $\mathbf{B}_m = E^\top \mathbf{q}_m \in \mathbb{R}^D$, yeilding M high-entropy cortical input action-salience bundles $\{\mathbf{B}_m\}_{m=1}^M$ that are delivered to the BG.

Experimental Design

We adapt the reach-to-pull paradigm of Park et al. (2025) to a modeling setting that explicitly separates discrete action selection from continuous action specification. Instead of two reach targets and two loads, our model has two discrete action channels—lever-pull and button-press—and a shared continuous action parameter: force, $f \in [0, 15]$g (load, in grams), sampled at 0.01g resolution ($N = 1500$).

To mirror the two-load structure, each channel receives a bimodal force–salience profile with peaks at 3g and 12g. On each trial, the channel salience c_m scales the action-salience profile amplitude for each discrete action channel, thereby setting its relative salience against its competitors. Within each channel, the relative peak heights determine whether 3g or 12g is preferred. The effective channel saliencies (lever-pull, button-press) are: Trial 1 $(0.93, 0.61)$, Trial 2 $(0.95, 0.63)$, Trial 3 $(0.62, 0.98)$, and Trial 4 $(0.65, 0.93)$ (see Figure 2). Each scaled bimodal force–salience profile is then encoded as a 512-dimensional force-salience input bundle.

We run a 2×2 design (action × load) across four trials: lever-pull at 3g and 12g (trials 1–2), and button-press at 3g and 12g (trials 3–4). Input bundles drive a BG circuit with a striatal DNF of 1500 units per channel (3000 total), presented for 1.5 seconds, with fixed dopamine level of 0.2. This enables a direct test of whether BG dynamics jointly resolve action selection and action specification.

Results

As shown in Figure 2, across all four trials, the BG network's striatal DNF transformed the concatenated high-entropy action-salience distribution (dot-dashed green line) into a single localized peak at the intended channel (solid purple line). In Trials 1 and 2, the peak appears in the lever-pull (LP) block at the instructed load (3g, then 12g); in Trials 3-4, the peak emerges in button-press (BP) at 3g and 12g (respectively), while LP is suppressed. The decoded action-salience distribution from the BG output (dotted blue line) mirrors this sharpening: only the winning channel retains a structured low-entropy action-salience distribution, and its within-channel peak index j^* maps directly to the specified continuous action value $a_{j^*} \in \mathcal{A}$.

Discussion

Our work develops a biologically grounded BG circuit that couples discrete action selection with continuous action specification within a single computation. Each candidate action is represented as a channel carrying an SSP-encoded action-salience distribution over a continuous action parameter (e.g., force). Striatal DNFs compete over the concatenated action–salience profiles to perform an entropy-reduction transformation, concentrating activity around the most salient region. In a complete coritcal-BG-thalamo-cortical implementation, this competition would be expected to supress GPi activity in the winning channel, yielding disinhibition of the corresponding thalamic channel (Rocha et al., 2023); within that disinhibited channel, action specification would be read out by mapping the within-channel peak location back to the continuous action parameter domain (Bartlett et al., 2025).

Although our theoretical inputs are idealized bimodal action–salience profiles, realized action-salience bundles exhibit small encoding and decoding distortions. Despite this representational noise, the model converges to the correct action and continuous action value across all trials, indicating robustness that is consistent with the stability properties of DNF recurrent dynamics (Schöner et al., 2015).

In this architecture, action selection and action specification occur in lockstep by construction: the same competition that identifies the winning channel also determines the within-channel peak location, which the thalamo-cortical pathway can then read out as the specified continuous action value. This coupling is enforced by our model connectivity pattern and representational scheme rather than from a biological derivation. Moreover, our construction of $\mathbf{q}_m$ (see Sec. Constructing BG Input) is an engineering choice rather than a mechanistic claim: $\mathbf{q}_m$ need not be constrained to $[0,1]^N$, and equivalent joint encoding of c_m and $\mathbf{s}_m$ could be achieved with alternative scalings, ranges, or transformations. How the cortex might implement any such mapping or encoding remains an open question. Thus, while we show that a single

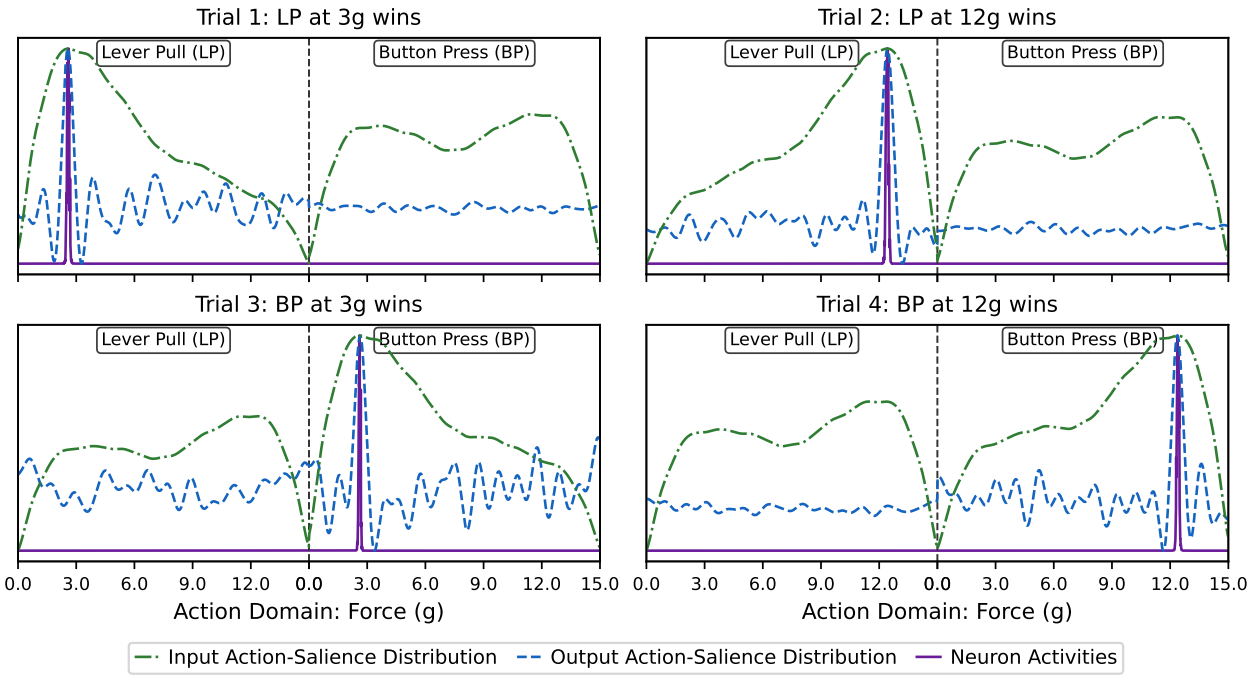

Figure 2: Basal ganglia action–force representations across four trials over time. Each panel shows the force domain (LP left, BP right) with input action–salience (green, dot-dashed line), decoded BG output (blue, dashed line), and DNF AP-cell activity (purple, solid line). Y-axis labels are omitted to emphasize peak location and distribution shape along the action domain, rather than absolute magnitude.

striatal substrate can support both computations given suitable representations and connectivity, we do not yet establish that biological circuits must realize this unification under realistic noise, delays, and recurrent cortico–BG–thalamic dynamics.

Several limitations of our current model point to direct extensions. First, the model specifies only a single continuous action attribute per discrete action, whereas natural movements require multi-attribute action specification (e.g., a combination of force, direction, and speed). While SSPs naturally generalize to higher-dimensional continuous action parameters, the DNF would need to be adopted—either via a factorized design (separate, smaller field for each attribute) or a joint field spanning all possible attribute combinations with appropriately defined interaction kernels. Second, our inputs were held static for 1.5 seconds, which bypasses time-varying decision formation and evidence accumulation. Incorporating an STN-mediated "brake" that delays commitment under uncertainty would allow inputs to evolve and accumulate evidence until a decision threshold is reached, better mimicking real-world decision making processes (Ballanger et al., 2009; Frank et al., 2007). Finally, we abstract away recurrent cortico–thalamic loops and plasticity; future work should include an explicit thalamo-cortical readout, short-loop pallidal feedback, and experience-dependent shaping of DNF kernels.

Two open questions remain. The first regards symmetry breaking: if two channels encode identical action–salience distributions, does the circuit resolve ties stochastically, via history dependence, through delayed commitment, or fail to resolve competition? While noise in the DNF may eventually break symmetry, animals likely resolve such cases more slowly through evidence accumulation and top-down biases (Cisek & Kalaska, 2010; Ratcliff & McKoon, 2008). Testing this would require constructing matched alternatives and then introducing controlled perturbations to quantify their effects on choice and latency. The second concerns dopaminergic control: we held dopamine constant to isolate DNF-driven dynamics, but tonic and phasic dopamine are expected to modulate pathway gains, exploration–exploitation, and learning (Dreyer et al., 2010; Gilbertson & Steele, 2020; Humphries et al., 2012; Keeler et al., 2014). A natural extension is to couple the circuit to reinforcement learning (e.g., actor–critic or temporal difference learning) with phasic dopamine as a teaching signal and test how systematic shifts in tonic dopamine regulate the stability and speed of joint action selection and specification. In this setting, reinforcement signals would shape both the selected action and the continuous action parameter values required to execute that action.

In conclusion, we present a BG circuit that jointly performs discrete action selection and continuous action specification in a single operation. This extends localist, channelized BG models from selecting "which action" to also determining "how much", and yields testable hypotheses for how the cortico–BG–thalamo–cortical loop can couple choice and vigour. The framework further motivates principled extensions to multidimensional specification, time-varying decision dynamics, and dopamine-dependent learning.

Acknowledgements

This work was funded by National Research Council of Canada's Artificial Intelligence for Design program via grant AI4D-151-1.

Source Data Sex Reporting

Park et al. (2025), whose reach-to-pull paradigm motivates this work, reported using both male and female mice, but did not specify the exact sex breakdown. We therefore use their findings as biological motivation for a general corticostriatal computation and do not make sex-specific claims.

Online Resources

The code implementation of the proposed basal ganglia model, along with the data and scripts used to generate results presented in this paper, is publicly available on Github: **https://github.com/2002sahapriya/Basal-Ganglia**. This repository includes all necessary resources to reproduce the experiments and facilitate further exploration of the proposed approach.

References

Ballanger, B., van Eimeren, T., Moro, E., Lozano, A. M., Hamani, C., Boulinguez, P., Pellecchia, G., Houle, S., Poon, Y. Y., Lang, A. E., & Strafella, A. P. (2009). Release your horses: Deep Brain Stimulation of the subthalamic nucleus improves motor functions at the expense of response inhibition: An H2150 PET study [Abstract from the 15th OHBM Annual Meeting]. *NeuroImage*, *47*(Suppl 1), S115. https://doi.org/10.1016/S1053-8119(09)71055-9

Bartlett, M., Furlong, P. M., Stewart, T. C., & Orchard, J. (2025). A computational model of action specification in the basal ganglia [Manuscript submitted for publication]. *bioRxiv*. https://doi.org/10.1101/2025.08.12.669938

Berns, G. S., & Sejnowski, T. J. (1998). A computational model of how the basal ganglia produce sequences. *Journal of Cognitive Neuroscience*, *10*(1), 108–121. https://doi.org/10.1162/089892998563824

Cisek, P. (2006). Integrated neural processes for defining potential actions and deciding between them: A computational model. *Journal of Neuroscience*, *26*(38), 9761–9770. https://doi.org/10.1523/JNEUROSCI.5605-05.2006

Cisek, P., & Kalaska, J. F. (2010). Neural mechanisms for interacting with a world full of action choices. *Annual Review of Neuroscience*, *33*, 269–298. https://doi.org/10.1146/annurev.neuro.051508.135409

Cohen, M. X., & Frank, M. J. (2009). Neurocomputational models of basal ganglia function in learning, memory and choice. *Behavioural Brain Research*, *199*(1), 141–156. https://doi.org/10.1016/j.bbr.2008.09.029

Dhawale, A. K., Wolff, S. B. E., Ko, R., & Ölveczky, B. P. (2021). The basal ganglia control the detailed kinematics of learned motor skills [Epub 2021 Jul 15]. *Nature Neuroscience*, *24*(9), 1256–1269. https://doi.org/10.1038/s41593-021-00889-3

Dreyer, J. K., Herrik, K. F., Berg, R. W., & Hounsgaard, J. D. (2010). Influence of phasic and tonic dopamine release on receptor activation. *Journal of Neuroscience*, *30*(42), 14273–14283. https://doi.org/10.1523/JNEUROSCI.1894-10.2010

Dudman, J. T., & Krakauer, J. W. (2016). The basal ganglia: from motor commands to the control of vigor [Epub 2016 Mar 21]. *Current Opinion in Neurobiology*, *37*, 158–166. https://doi.org/10.1016/j.conb.2016.02.005

Dumont, N. S.-Y. (2025). *Symbols, dynamics, and maps: A neurosymbolic approach to spatial cognition* [PhD thesis]. University of Waterloo.

Eliasmith, C. (2013, June). *How to build a brain: A neural architecture for biological cognition*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199794546.001.0001

Eliasmith, C., & Anderson, C. H. (2003). *Neural engineering: Computation, representation, and dynamics in neurobiological systems*. MIT Press.

Fobbs, W. C., Bariselli, S., Licholai, J. A., Miyazaki, N. L., Matikainen-Ankney, B. A., Creed, M. C., & Kravitz, A. V. (2020). Continuous representations of speed by striatal medium spiny neurons. *The Journal of Neuroscience*, *40*(8), 1679–1688. https://doi.org/10.1523/JNEUROSCI.1407-19.2020

Frank, M. J. (2005). Dynamic dopamine modulation in the basal ganglia: A neurocomputational account of cognitive deficits in medicated and non-medicated Parkinsonism. *Journal of Cognitive Neuroscience*, *17*(1), 51–72. https://doi.org/10.1162/0898929052880093

Frank, M. J., Samanta, J., Moustafa, A. A., & Sherman, S. J. (2007). Hold your horses: Impulsivity, deep brain stimulation, and medication in Parkinsonism. *Science*, *318*(5854), 1309–1312. https://doi.org/10.1126/science.1146157

Furlong, P. M., & Eliasmith, C. (2023). Bridging cognitive architectures and generative models with vector symbolic algebras. *Proceedings of the 2023 AAAI Fall Symposium*. https://ojs.aaai.org/index.php/AAAI-SS/article/view/27686

Furlong, P. M., Simone, K., Dumont, N., Bartlett, M., Stewart, T. C., Orchard, J., & Eliasmith, C. (2024). Biologically-plausible Markov Chain Monte Carlo sampling from Vector Symbolic Algebra-encoded Distributions. *ICANN*. https://compneuro.uwaterloo.ca/files/furlong2024.VSA.sampling.ICANN.pdf

Gilbertson, T., & Steele, D. (2020). Tonic dopamine, uncertainty and basal ganglia action selection. *bioRxiv*. https://doi.org/10.1101/2020.11.10.376608

Gurney, K., Prescott, T. J., & Redgrave, P. (2001a). A computational model of action selection in the basal ganglia. i. a new functional anatomy. *Biological Cybernetics*, *84*(6), 401–410. https://doi.org/10.1007/PL00007984

Gurney, K., Prescott, T. J., & Redgrave, P. (2001b). A computational model of action selection in the basal ganglia. ii.

analysis and simulation of behaviour. *Biological Cybernetics*, *84*(6), 411–423. https://doi.org/10.1007/PL00007985

Humphries, M. D., Khamassi, M., & Gurney, K. (2012). Dopaminergic control of the exploration-exploitation tradeoff via the basal ganglia. *Frontiers in Neuroscience*, *6*, 9. https://doi.org/10.3389/fnins.2012.00009

Humphries, M. D., Stewart, R. D., & Gurney, K. N. (2006). A physiologically plausible model of action selection and oscillatory activity in the basal ganglia. *Journal of Neuroscience*, *26*(50), 12921–12942. https://doi.org/10.1523/JNEUROSCI.3486-06.2006

Keeler, J., Pretsell, D., & Robbins, T. (2014). Functional implications of dopamine D1 vs. D2 receptors: A 'prepare and select' model of the striatal direct vs. indirect pathways [The Ventral Tegmentum and Dopamine: A New Wave of Diversity]. *Neuroscience*, *282*, 156–175. https://doi.org/https://doi.org/10.1016/j.neuroscience.2014.07.021

Lemon, R. N. (2008). Descending pathways in motor control. *Annual Review of Neuroscience*, *31*, 195–218. https://doi.org/10.1146/annurev.neuro.31.060407.125547

Lindahl, M., Kamali Sarvestani, I., Ekeberg, Ö., & Kotaleski, J. H. (2013). Signal enhancement in the output stage of the basal ganglia by synaptic short-term plasticity in the direct, indirect, and hyperdirect pathways. *Frontiers in Computational Neuroscience*, *7*, 76. https://doi.org/10.3389/fncom.2013.00076

O'Reilly, R. C., & Frank, M. J. (2006). Making working memory work: A computational model of learning in the prefrontal cortex and basal ganglia. *Neural Computation*, *18*(2), 283–328. https://doi.org/10.1162/089976606775093909

Park, J., Polidoro, P., Fortunato, C., Arnold, J., Mensh, B., Gallego, J. A., & Dudman, J. T. (2025). Conjoint specification of action by neocortex and striatum. *Neuron*, *113*, 620–636. https://doi.org/10.1016/j.neuron.2024.12.024

Picard, N., & Strick, P. L. (2001). Imaging the premotor areas. *Current Opinion in Neurobiology*, *11*(6), 663–672. https://doi.org/https://doi.org/10.1016/S0959-4388(01)00266-5

Plate, T. (1995). Holographic reduced representations. *IEEE Transactions on Neural Networks*, *6*(3), 623–641. https://doi.org/10.1109/72.377968

Ponzi, A. (2008). Dynamical model of salience gated working memory, action selection and reinforcement based on basal ganglia and dopamine feedback. *Neural Networks*, *21*(2–3), 322–330. https://doi.org/10.1016/j.neunet.2007.12.017

Prescott, T., Redgrave, P., & Gurney, K. (2000). Layered control architectures in robots and vertebrates. *Adaptive Behavior*, *7*. https://doi.org/10.1177/105971239900700105

Prescott, T. J., Montes González, F. M., Gurney, K., Humphries, M. D., & Redgrave, P. (2024). Simulated dopamine modulation of a neurorobotic model of the basal ganglia. *Biomimetics*, *9*(2), 139. https://doi.org/10.3390/biomimetics9020139

Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computation*, *20*(4), 873–922. https://doi.org/10.1162/neco.2008.12-06-420

Redgrave, P., Prescott, T. J., & Gurney, K. (1999). The basal ganglia: A vertebrate solution to the selection problem? *Neuroscience*, *89*(4), 1009–1023. https://doi.org/10.1016/S0306-4522(98)00319-4

Rocha, G. S., Freire, M. A., Britto, A. M., Paiva, K. M., Oliveira, R. F., Fonseca, I. A., Araújo, D. P., Oliveira, L. C., Guzen, F. P., Morais, P. L., et al. (2023). Basal ganglia for beginners: The basic concepts you need to know and their role in movement control. *Frontiers in Systems Neuroscience*, *17*, 1242929. https://doi.org/10.3389/fnsys.2023.1242929

Schöner, G., Spencer, J., & Research Group, D. (2015, November). *Dynamic thinking: A primer on dynamic field theory*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199300563.001.0001

Scott, S. H. (2004). Optimal feedback control and the neural basis of volitional motor control. *Nature Reviews Neuroscience*, *5*(7), 532–545. https://doi.org/10.1038/nrn1427

Shadmehr, R., & Krakauer, J. W. (2008). A computational neuroanatomy for motor control. *Experimental Brain Research*, *185*(3), 359–381. https://doi.org/10.1007/s00221-008-1280-5

Shenoy, K. V., Sahani, M., & Churchland, M. M. (2013). Cortical control of arm movements: A dynamical systems perspective. *Annual Review of Neuroscience*, *36*, 337–359. https://doi.org/10.1146/annurev-neuro-062111-150509

Soni, A., & Frank, M. J. (2025). Adaptive chunking improves effective working memory capacity in a prefrontal cortex and basal ganglia circuit. *eLife*, *13*, RP97894. https://doi.org/10.7554/eLife.97894

Stewart, T. C., Bekolay, T., & Eliasmith, C. (2012). Learning to select actions with spiking neurons in the basal ganglia. *Frontiers in Neuroscience*, *6*, 2. https://doi.org/10.3389/fnins.2012.00002

Stewart, T. C., Choo, X., & Eliasmith, C. (2010). Dynamic behaviour of a spiking model of action selection in the basal ganglia. *Proceedings of the 10th International Conference on Cognitive Modeling (ICCM)*, 235–240. https://compneuro.uwaterloo.ca/files/publications/stewart.2010.pdf

Suryanarayana, S. M., Kotaleski, J. H., Grillner, S., & Gurney, K. N. (2019). Roles for globus pallidus externa revealed in a computational model of action selection in the basal ganglia. *Neural Networks*, *109*, 113–136. https://doi.org/10.1016/j.neunet.2018.09.014

Teka, W. W., Hamade, K. C., Barnett, W. H., Kim, T., Markin, S. N., Rybak, I. A., & Molkov, Y. I. (2017). From the motor cortex to the movement and back again. *PLOS ONE*, *12*(6), e0179288. https://doi.org/10.1371/journal.pone.0179288

Turner, R. S., & Desmurget, M. (2010). Basal ganglia contributions to motor control: A vigorous tutor [Epub 2010 Sep 17]. *Current Opinion in Neurobiology*, *20*(6), 704–716. https://doi.org/10.1016/j.conb.2010.08.022

Ursino, M., Véronneau-Veilleux, F., & Nekka, F. (2020). A non-linear deterministic model of action selection in the basal ganglia to simulate motor fluctuations in Parkinson's disease. *Chaos*, *30*, 083139. https://doi.org/10.1063/5.0010458

Wiecki, T. V., & Frank, M. J. (2013). A computational model of inhibitory control in frontal cortex and basal ganglia. *Psychological Review*, *120*(2), 329–355. https://doi.org/10.1037/a0031542