# Using Vector Symbolic Architectures for Distributed Action Representations in a Spiking Model of the Basal Ganglia

**Madeleine Bartlett[1] (madeleine.bartlett@uwaterloo.ca),**
**P. Michael Furlong[2] (michael.furlong@uwaterloo.ca),**
**Jeff Orchard[1] (jorchard@uwaterloo.ca),**
**and Terrence C. Stewart[3] (terrence.stewart@nrc-cnrc.gc.ca)**
[1] Cheriton School of Computer Science, University of Waterloo, Waterloo, ON N2L 3G1, Canada
[2] Systems Design Engineering, University of Waterloo, Waterloo, ON N2L 3G1, Canada
[3] National Research Council of Canada, University of Waterloo Collaboration Centre, Waterloo, ON N2L 3G1, Canada

## Abstract

Existing models of the basal ganglia assume the existence of separate channels of neuron populations for representing each available action. This type of localist mapping limits models to small, discrete action spaces, since additional actions require additional channels, costing neural resources and imposing new connective tracts. In contrast, evidence suggests that the basal ganglia plays a role in the selection of both discrete action units, and continuously-valued action kinematics. In this work, we model the basal ganglia with distributed action representations, using high-dimensional vectors. This method lends itself to representing both discrete and continuous action spaces. Vectors that represent actions are weighted by a scalar value (their salience to the current task), and bundled together to form a single input vector. This paper provides an overview of the encoding method and network structure, as well as a demonstration of the model solving an action selection task using spiking neurons.

**Keywords: Action Selection, Basal Ganglia, Computational Modelling, Vector Symbolic Architecture, Semantic Pointer**

## Introduction

The basal ganglia has been established to play a role in action selection (Mink, 1996; Redgrave et al., 1999; Hikosaka et al., 2000; Klaus et al., 2019). Numerous computational models have been developed to explain this function (Gurney et al., 2001; Humphries et al., 2006; Bogacz & Larsen, 2011), often adopting localist representations of action spaces. These representations have been valuable in exploring the basal ganglia's role in selecting discrete units or sequences of actions (Graybiel, 1998; Jin & Costa, 2015; Gurney et al., 2001; Humphries et al., 2006). However, recent evidence has suggested that the basal ganglia is also involved in the selection of continuously valued action kinematics such as speed and vigor (Barbera et al., 2016; Lederman et al., 2021; Turner & Desmurget, 2010). Current methods for representing the action space adopted by existing models are not well-suited to modelling this function. In this paper, we shift the representation in long-standing basal ganglia models from localist to distributed by adopting a Vector Symbolic Architecture (VSA) approach which can be used to represent continuous action spaces in the same neural substrate that could process discrete action spaces. Thus this work is a first-step towards a single model of the basal ganglia capable of selecting from both discrete and continuous action spaces.

In general, models of action selection treat the basal ganglia as a signal selector. In many existing models (e.g. Gurney et al., 2001; Stewart et al., 2010; Chersi et al., 2013; Humphries et al., 2006), the basal ganglia receives as input a set of saliency signals, each specifying the salience of an available action. The basal ganglia's task is to accentuate the largest salience while reducing the rest. Since the output from the basal ganglia is proportional to the represented saliences, this 'pushing apart' of the saliences ensures an exaggerated down-stream disinhibition (and thereby selection) of the desired action, and the continued suppression of unwanted actions. However, such models are lacking biological fidelity because they rely on localist representations of actions, which require increasing (and potentially implausible) numbers of neurons to represent a complete action space (Stewart & Eliasmith, 2012). This assumption restricts these models from generalizing to continuous, open-ended action spaces.

In this work, we adapt a model of the basal ganglia to use a novel method for representing action spaces, while preserving its functional role. Namely, we use distributed vector representations of actions that can be applied to continuous or large dimensional action spaces, without the constraints of one-hot representations. Our goal is to provide a proof-of-concept for the use of this representation in the context of a model of the basal ganglia, and to provide a demonstration of basal ganglia-like behaviour following only minor adaptations to the Stewart et al. (2010) model. In the following sections we first describe Semantic Pointers (SPs) and then present the adapted model of the basal ganglia. We present a preliminary set of simulation results, including a demonstration of the network implemented in spiking neurons, and discuss the results and future directions.

## Methods

### Semantic Pointers

The semantic pointer hypothesis (Eliasmith, 2013, p78) states that higher-level cognitive functions depend on *semantic pointers* - high dimensional neural representations that carry



In L. K. Samuelson, S. L. Frank, M. Toneva, A. Mackey, & E. Hazeltine (Eds.), *Proceedings of the 46th Annual Conference of the Cognitive Science Society.* 

some semantic information, and that can be composed into higher-level representations that can be manipulated to perform cognition. SPs are provided by a VSA – a family of operations on high-dimensional vectors, and a mapping from symbols to the high-dimensional representation (Blouw et al., 2016; Voelker et al., 2021). These operations include a binding method for combining vectors to form new vectors that represent the conjunction of the bound vectors. Here, we use Plate's HRR Plate (1992, 1995), in which binding is defined as circular convolution $\circledast$. Binding can be used to construct role-filler pair vectors, or other structured representations. For example, we can encode position $k$ (like an index to a list) by creating an index vector $A$ and binding it to itself $k$ times:

$$A^k = A \circledast A \circledast ... \circledast A \tag{1}$$

In traditional VSAs, $k$ could be any integer value. Circular convolution, however, has a natural extension to real-valued $k$, which we call Spatial Semantic Pointers (SSPs) Komer et al. (2019). As a result, SSPs can be used to encode continuous-valued variables. A second operation, bundling, achieved via vector addition, is used to combine vectors into sets. It creates a new vector which maintains some similarity with its constituent parts. These vectors, or SPs can, in turn, be represented in a distributed manner by the collective activity of populations of neurons, where the input current to each neuron is a linear function of the value to be encoded (Stewart, 2012; Eliasmith & Anderson, 2003).

An important characteristic of these high-dimensional representations is their *pseudo-orthogonality* – the similarity, as measured by the vector dot product, between two different randomly generated vectors is small and, on average, equal to zero. Pseudo-orthogonality lets us represent and manipulate multiple SPs in one vector. We use this feature to store the saliences of different actions in one neural population, which permits us to consider a variable number of actions and their corresponding saliences without changing the dimensionality of the vector itself. Hence, a neural network could be used to represent an indeterminate number of salience values, without any change to the network. Thus, we offer a biologically plausible model of the basal ganglia that can represent a flexible number of action saliences in a distributed manner within a substrate of (potentially spiking) neurons, without imposing unbounded demands on neural resources.

The purpose of this exploratory work was to establish a viable method for encoding actions for the basal ganglia, that could be extended to encoding continuous-valued actions in future work. The input to our network consists of a $d$-dimensional vector that encodes all of the input actions and their saliences. We first create a $d$-dimensional SP, $A_{a_i}$, for each action, $a_i$, and scale it by its salience value, $s_i$. Thus the direction of these vectors corresponds to the action, and the magnitude to the salience. The scaled action SPs, $s_i A_{a_i}$, are then bundled together using vector addition in order to produce a single $d$ dimensional vector $B$,

$$B = s_1 A_{a_1} + s_2 A_{a_2} + ... + s_n A_{a_n} \ .$$

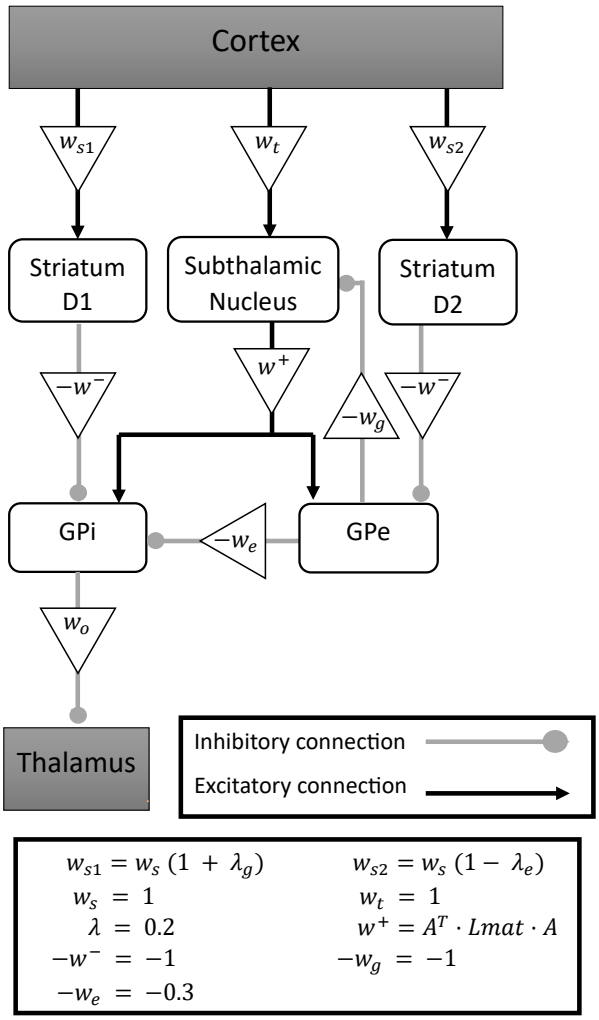


Figure 1: Schematic of the basal ganglia model

The benefit of this encoding is that the network does not need to be restructured to accommodate more actions. The size of the input bundle is determined only by the dimensionality of the SP representation, which is independent of the number of actions being encoded. Although, it should be noted that the total number of actions that can be represented is not independent of the dimensionality of the vector (see Plate, 2003, Appendix B). Additionally, this representation can be readily extended to encoding continuous-value variables

## Model of basal ganglia

The model developed for this exploration was an adaptation of that presented by Stewart et al. (2010), which was a spiking-neuron version of Gurney et al.'s (2001) non-spiking basal ganglia model. The model presented by Stewart et al. (2010) used groups of Leaky Integrate-and-Fire (LIF) neurons to approximate the non-spiking Rectified Linear Unit (ReLU) neurons of the original. The input for both of these

two original networks is a vector containing the scalar action saliences. The salience of each action was represented by a separate action channel within the network. In the current network, however, the collection of salience values is represented by a weighted sum of SP vectors, stored in the bundle $B$. Thus, instead of separate localist channels of neurons for each action, there is a single continuous representation.

Apart from some modifications to accommodate the new vector representation, the rest of the network remains mostly the same as in Stewart et al. (2010) and Gurney et al. (2001) (see Figure 1). The striatum receives input from the cortex consisting of the SP bundle of salience-weighted action vectors. These inputs are also subject to dopaminergic modulation via the parameters $\lambda_g$ and $\lambda_e$. The striatum consists of two populations corresponding to the D1 and D2 neurons. The D1 population sends inhibitory projections to the Globus Pallidus Internus (GPi) thus forming the direct pathway. The D2 population, as part of the indirect pathway, projects to the Globus Pallidus Externus (GPe). The GPe in turn inhibits the GPi and also sends inhibitory signals to the Subthalamic Nucleus (STN). The STN makes up the hyperdirect pathway, receiving inputs directly from cortex and sending excitatory signals out to the GPi and GPe. As shown in Figure 1, the connection weights are the same as those used in Stewart et al. (2010) and Gurney et al. (2001) except for the connection weights for signals coming out of the STN ($w^+$) and GPi ($w_o$). Additionally, the ReLU computation, present at each stage in the original model, was removed from the GPi in the new network because it interfered with the network's output, preventing it from successfully selecting the appropriate action. Here we provide a description of the new connection weights developed for connecting the STN to the GPe and GPi.

The STN-GPe loop of the basal ganglia plays a valuable role in balancing the direct and indirect pathways by supporting voluntary movements whilst suppressing unwanted actions (Mink, 1996; Nambu et al., 2002). In the model developed by Stewart et al. (2010) (hereafter, the 2010 model), the role of the STN-pallidial connections was to exaggerate the differences between the highest-value salience and its competitors. The new connection weights were developed to perform a similar function but acting on the bundled representation rather than individual salience channels. To sharpen the relative saliences of the SPs contained in the bundle, we constructed a matrix that has the effect of applying a Laplacian operator(Chung, 1997, §1.2) to the elements of the bundle. This exploits the pseudo-orthogonality of the SP representations to operate entirely in the high-dimensional vector space. Laplacian operators, when used in image processing, accentuate local maxima or minima in an image by computing the difference between the central element of an image kernel and its neighbours in a connected neighbourhood. In our approach, the elements of our bundle are represented by vectors on the unit hypersphere, which we can consider, abstractly, as all being neighbours. Thus, the weights from the STN to GPe and GPi are defined

$$w^+ = A^T \mathcal{L}\, A, \tag{2}$$

where $A$ holds the action SP vectors in its rows,

$$A = \begin{bmatrix} —A_{a_1}— \\ \vdots \\ —A_{a_n}— \end{bmatrix}$$

and

$$\mathcal{L} = 0.02 \begin{bmatrix} n-1 & -1 & \cdots & -1 \\ -1 & n-1 & \cdots & -1 \\ \vdots & \vdots & \ddots & \vdots \\ -1 & -1 & \cdots & n-1 \end{bmatrix}.$$

The above values are for a symmetrized Laplacian for actions that form a fully-connected graph of dimensionality $n$. In our experiments, we use $n = 20$ actions. We acknowledge that this is equivalent to assuming that the STN uses a localist representation of the action space. In future work we intend to find alternative solutions in order to move away from this assumption. However, we argue that relying on this assumption is acceptable for the current purposes of providing a proof-of-concept for the method of representing actions using SPs. Particularly because, in our model, we are not arguing against the existence of discrete action spaces but instead for the existence of both discrete and continuous spaces existing in the same neural substrate. The Laplacian matrix was produced using SciPy's (Virtanen et al., 2020) `laplacian` function.

Consider what happens when we multiply $w^+$ by the saliency bundle, $B$, forming the product $w^+B$. When we apply the first matrix, $A$, the output is (approximately) a vector of the saliences, $\mathbf{s} = [s_1, \ldots, s_n]^T$. Then when we multiply by the Laplacian matrix, we get $\mathbf{s}' = \mathcal{L}\,\mathbf{s}$, which is a vector whose $i^{\text{th}}$ element is the application of the Laplacian operator centred on the $i^{\text{th}}$ saliency. This result is the sharpened saliency values, $s'_i = 0.02\left((n-1)s_i - \sum_{j\neq i} s_j\right) + \eta$, where $\eta$ is noise due to cross-talk between the SP representations. The noise $\eta$ can be controlled by increasing the dimensionality of the SP representation. When we then multiply $A^T\mathbf{s}'$, we produce a new vector that is equivalent to $\sum_i s'_i A_{a_i}$. This is a new bundle where the SPs that represent the actions are scaled by the sharpened saliences, but the arithmetic is performed entirely in the high-dimensional vector space.

The result is effectively a clean-up memory with a lateral inhibition between the stored actions, similar to the lateral inhibition implemented by Levy & Gayler (2009). Levy & Gayler (2009) avoid the need to extract the coefficients ($k$ there, $s$ here) by adopting the Multiply, Add, Permute (MAP) VSA (Gayler, 2004) and exploiting the permutation and self-cancellation properties. MAP, as well as many other VSAs, however, does not allow for fractional binding which is necessary for producing vector embeddings of continuous-valued variables. Thus, whilst this method is flawed in relying on a localist approach to perform lateral inhibition, it is ideally suited to representing continuous action spaces. Future work

will focus on developing models of basal ganglia function that do not rely on such localist methods.

## Simulations

The aim of these simulation experiments was to explore the network analytically. The operations being performed by the connection weights between nuclei are all linear transforms and can therefore be understood without simulations. However, the network inherited nonlinearities from its predecessors in the form of ReLU functions within the nuclei, applied to the represented values/saliences. In the original model, these nonlinearities serve to ensure that the output of the channel with the highest salience $= 0$, and $> 0$ for all other channels, thereby achieving the goal of having a model whose outputs match the pattern in the GPi where neurons are inhibited to the point of being 'turned off', allowing for disinhibition of the thalamus and other downstream nuclei. Given that the model is now working in a high-dimensional vector space, it is difficult to interpret the function of these nonlinearities, though it is unlikely that they will be able to ensure this same pattern of activity. We therefore first explored the utility of the ReLU functions by systematically removing the ReLU functions from individual and groups of nuclei, and assessing the effect this had on network performance.

Performance was assessed using two novel performance metrics. The goal of the network is not only to pick out the action with the greatest salience as the chosen action, but to increase the difference between the greatest salience and all other saliences. One way to measure this is to measure the difference between the maximum and second maximum saliences coming in, and then compare this with the difference between those same saliences as they come out of the network. We calculate this *Margin* metric as

$$Margin_{\text{in}} = \max(s_{\text{in}}) - 2^{nd}\max(s_{\text{in}})$$
$$Margin_{\text{out}} = \max(s_{\text{out}}) - 2^{nd}\max(s_{\text{out}})$$

where $s_{\text{in}}$ are the saliences going into the network, and $s_{\text{out}}$ are the corresponding salience values output by the network. The network's behaviour should result in $Margin_{\text{out}}$ being greater than $Margin_{\text{in}}$. Thus our final performance metric is

$$Margin_{\text{out}} - Margin_{\text{in}} \ .$$

Higher values of this metric would indicate good performance in that the network was successfully increasing the difference between the highest and second highest salience values.

A second metric measured the change in the margin between the largest and the smallest saliences, using the same equation as above, replacing $2^{nd}\max(s)$ with $\min(s)$. This analysis did not reveal any insights beyond what can be gleaned from Table 1. Whilst we do not report those results here, they are available in our online resources.

Obtaining the saliences coming out of the networks requires different approaches for each network. For the network developed by Stewart et al. (2010) (2010 network), the GPi contains one channel per action, and the output of each channel is the salience of the corresponding action. In contrast, for the new network, the output has to be decoded from $B_{\text{GPi}}$, the bundle of weighted action vectors encoded by the GPi. It is therefore first necessary to decode the scales from this bundle, which can be done by calculating the dot product between the bundle $B_{\text{GPi}}$ and the action vectors,

$$s = A B_{\text{GPi}} \ ,$$

yielding $s$, the vector of weights for each of the actions. Hence, $w_o$ in Figure 1 is equal to $-A$.

### Utility of ReLU

This set of experiments aimed to explore the utility of the ReLU activation functions given that the new network is operating on a very different input compared to its predecessors (Stewart et al., 2010; Gurney et al., 2001). The ReLU activation functions were systematically removed and reintroduced into the network's subpopulations to explore their function and usefulness. We also performed the same simulations on the 2010 network and compared the results.

Each network configuration was tested 10 times with 10 random seeds, and the results were averaged. For these initial explorations we implemented both networks as 'mathematical models'; signals are represented and transformed perfectly, rather than through a neural approximation. The mathematical models were tasked with choosing between three actions whose saliences were chosen from $\{0.1, 0.2, 0.3\}$. The action with the highest salience changed every second, and the network was run for 3 seconds. The network was assessed at the end of each 1 second trial, taking the output produced in the timestep 10ms before the end of the trial. The results of this exploration can be seen in Table 1.

Comparing performance across the two networks we see that the new network was not able to achieve the same level of performance as the 2010 model. This is arguably unsurprising since the network is no longer acting directly on the salience values as it was in the original case, but instead the network is operating on a vector bundle. A further difference is that the Laplacian matrix we adopt differs from the lateral inhibition weights used in the 2010 model. However, it should be noted that the new network still successfully identifies the action with the highest salience, and has the desired effect of exaggerating the difference between the saliences with the highest value at input and the other saliences. The goal here is to find the simplest network configuration that maximises the distance between the largest salience and the other, competing saliences. Based on the results shown in Table 1, the network that meets both of these criteria is one without any ReLU functions at all (see shaded cells).

### Spiking Neurons

The next step was to implement the new network in spiking neurons. Given enough neurons it is possible to create a neural network that can approximate any function. To establish how many neurons were required for the spiking network

Table 1: Difference between the change in the largest and the second largest saliences. Testing effect of removing ReLU from different parts of the network. Values shown are the mean margin differences calculated across the 10 random seeds. Standard Deviation for the experiments using the 2010 network was always $< 0.001$.

| Populations with ReLU | mean($Margin_{out} - Margin_{in}$) (s.d.) | | | | | |
|---|---|---|---|---|---|---|
| | New Network | | | 2010 Network | | |
| | $t < 1$s | 1s $< t <$ 2s | 2s $< t <$ 3s | $t < 1$s | 1s $< t <$ 2s | 2s $< t <$ 3s |
| All (Original) | 0.083 (0.032) | 0.067 (0.009) | 0.083 (0.033) | 0.188 | 0.188 | 0.188 |
| All but StrD1 | 0.083 (0.032) | 0.067 (0.009) | 0.083 (0.033) | -0.100 | -0.100 | -0.100 |
| All but StrD2 | 0.072 (0.027) | 0.059 (0.009) | 0.072 (0.027) | 0.188 | 0.188 | 0.188 |
| All but STN | 0.083 (0.032) | 0.067 (0.009) | 0.083 (0.033) | 0.042 | 0.042 | 0.042 |
| All but GPe | 0.109 (0.045) | 0.086 (0.013) | 0.108 (0.046) | 0.188 | 0.188 | 0.188 |
| All but GPi | N/A | N/A | N/A | 0.188 | 0.188 | 0.188 |
| StrD1 | 0.109 (0.045) | 0.086 (0.013) | 0.108 (0.046) | 0.188 | 0.188 | 0.188 |
| StrD2 | 0.109 (0.045) | 0.086 (0.013) | 0.108 (0.046) | 0.188 | 0.188 | 0.188 |
| STN | 0.109 (0.045) | 0.086 (0.013) | 0.108 (0.046) | 0.188 | 0.188 | 0.188 |
| GPe | 0.072 (0.027) | 0.059 (0.009) | 0.072 (0.027) | 0.188 | 0.188 | 0.188 |
| GPi | N/A | N/A | N/A | 0.188 | 0.188 | 0.188 |
| None | 0.109 (0.045) | 0.086 (0.013) | 0.108 (0.046) | 0.188 | 0.188 | 0.188 |

to closely approximate the performance of the mathematical model, the network with no ReLU functions was converted to a spiking network using Leaky Integrate-and-Fire (LIF) neurons and tested with 10, 20, 50, 100, or 200 neurons per dimension of the SP bundle ($d = 512$), in each subpopulation of the network. This resulted in a total of 25,600, 51,200, 128,000, 256,000, or 512,000 neurons respectively.

Table 2 presents the performance metrics for these experiments, calculated for both the maximum vs. minimum, and the maximum vs. second maximum salience. Unsurprisingly, performance improves with greater numbers of neurons, with 200 neurons per SP dimension giving the best performance and most closely approximating the performance of the mathematical model. The following exploration is of this network, employing a total of 512,000 spiking LIF neurons.

**Spiking Behaviour** The network was tasked with selecting between 3 actions with saliences $\in [0.1, 0.2, 0.3]$. The action with the highest salience changed after 1 second of simulation time. The GPi sends tonic inhibition to the thalamus and therefore must itself be inhibited in order for action selection to occur (Alexander & Crutcher, 1990; DeLong, 1990). This pattern, where subsets of neurons are being 'turned off' when different actions are selected, is evident in the spiking data shown in Figure 2. This pattern was also seen in Stewart et al. (2010). However, because we adopt a distributed representation, the present model does not run into the problem of having localist 'grandmother cells' (Gross, 2002; Quiroga et al., 2005) that previous models have. In fact, prior work has shown that, when using this style of representation, it is possible to remove the neurons which change the most for different values (i.e. the neurons that 'turn off' in Figure 2) and the output vector can still be accurately recovered due to the distributed representation (Stewart et al., 2011).

A central advantage of the vector representation adopted here is that the model does not require additional resources (i.e. neurons or signal channels) to represent more actions. The current version of the model does require that the $A$ matrix be large enough to accommodate the action space. However, future work will move away from this localist approach. To illustrate this we tested the model with a 6-action choice task. As shown in Figure 3, the network was successfully able to solve this task. Importantly, no changes were made to our basal ganglia model to implement this task of selecting from 6 rather than 3 actions - the only change was the number of actions and saliences used to generate the input vector.

## Discussion

In this paper we present a novel approach to modelling the basal ganglia. The proposed approach builds on work by Gurney et al. (2001) and Stewart et al. (2010), sharing the same network structure. However, we adopt a novel distributed representation of the action space that can extend the model to representing both discrete and continuous action spaces, without changing the neural circuitry. Using hyperdimensional computing methods we encode each action as a high-dimensional vector called a Semantic Pointer (SP). These vectors are then scaled by their associated salience before being bundled together into a single bundle vector, which is provided as input to the basal ganglia network. In order to perform the function of the basal ganglia in a winner-takes-all fashion, we applied a Laplacian operator to the bundle, with the effect of accentuating the largest salience and separating it from its neighbour saliences.

The network was successfully implemented in spiking LIF neurons and shown to be capable of selecting the highest salience action in both a 3- and a 6-action choice task. No

Table 2: Network Performance. Testing effect of different numbers of neurons in a spiking network.

| N Neurons | mean($Margin_{out} - Margin_{in}$) (s.d.) | | | | | |
|---|---|---|---|---|---|---|
| | Max vs. Min | | | Max vs. 2nd Max | | |
| | $t < 1$s | 1s $< t <$ 2s | 2s $< t <$ 3s | $t < 1$s | 1s $< t <$ 2s | 2s $< t <$ 3s |
| 10 | 0.091 (0.118) | 0.076 (0.087) | 0.062 (0.114) | 0.052 (0.090) | 0.049 (0.084) | 0.028 (0.074) |
| 20 | 0.107 (0.066)) | 0.073 (0.057) | 0.094 (0.063) | 0.087 (0.053) | 0.027 (0.053) | 0.078 (0.054) |
| 50 | 0.139 (0.027) | 0.130 (0.036) | 0.121 (0.033) | 0.074 (0.025) | 0.073 (0.022) | 0.064 (0.025) |
| 100 | 0.139 (0.025) | 0.142 (0.031) | 0.147 (0.035) | 0.082 (0.033) | 0.068 (0.015) | 0.086 (0.032) |
| 200 | 0.151 (0.028) | 0.153 (0.033) | 0.156 (0.038) | 0.088 (0.030) | 0.074 (0.015) | 0.089 (0.030) |

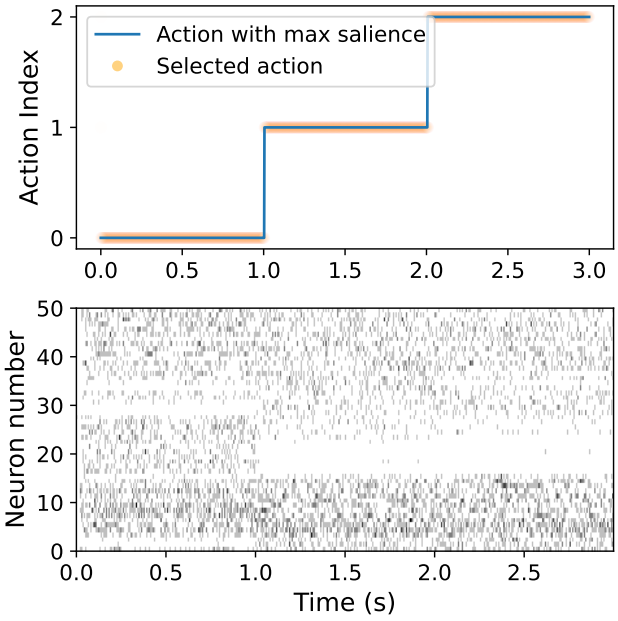


Figure 2: Top: Plot of the index of the action with the highest salience (blue line), and the action selected by the network (orange dots). Bottom: Raster plot of spiking neuron activity in the GPi during a 3-action choice task.

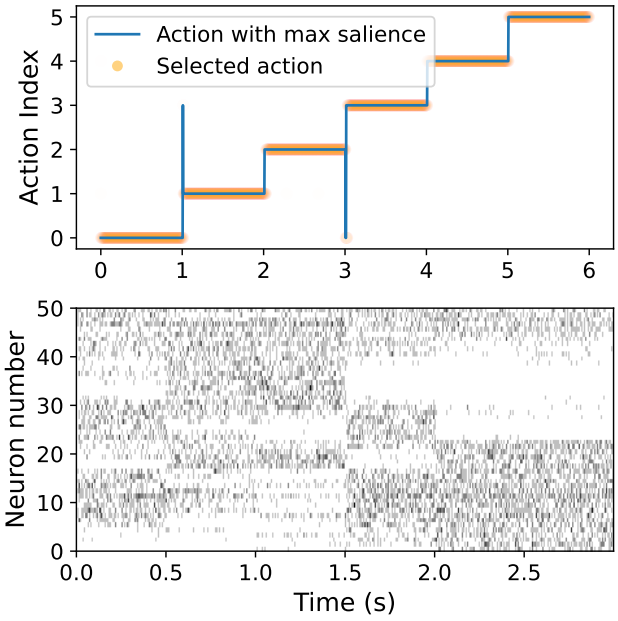


Figure 3: Top: Plot of the index of the action with the highest salience (blue line), and the action selected by the network (orange dots). Bottom: Raster plot of spiking neuron activity in the GPi during a 6-action choice task.

change to the network was needed to adjust the number of actions, due to the distributed nature of the representation. Interestingly, it was observed that the spiking behaviour within the GPi appears to replicate the pattern seen in biology, wherein subpopulations of neurons are inhibited to the point of being 'turned off', allowing the thalamus to excite the cortex, and for an action to be performed.

These findings are promising in demonstrating a model of the basal ganglia that uses VSA-represented actions and saliences. We demonstrate that the model is able to function as an action selector with only a few modifications to accommodate the new representation. Additionally, the model was successfully implemented in spiking LIF neurons, and showed spiking behaviour in the GPi reflective of patterns seen in biological systems.

For future work, we plan to leverage more of the advantages of VSAs. VSAs can be applied to encoding continuous variables through the use of fractional binding (Plate, 1992, 1995; Voelker et al., 2021). Revisiting equation 1; when encoding discrete SPs, $k$ is any integer value. However, with fractional binding, $k$ can be any real-valued number; for instance you could have $A^{1.5}$. Thus there is a trivial extension from discrete representations to continuous representations, meaning we should be able to construct one basal ganglia model for both discrete and continuous action spaces. In future work we intend to explore how this method of representing continuous variables might my applied to action selection in continuous action spaces. We will also continue to improve the fidelity of the model, introducing more biological detail.

## Acknowledgments

This project was supported by collaborative research funding from the National Research Council of Canada's Artificial Intelligence for Design program (AI4D-151-1). PMF was supported by AFOSR grant FA9550-17-1-0026.

## Online Resources

Experiment and analysis scripts can be found in the github repository: `https://github.com/maddybartlett/VSABasalGanglia`.

## References

Alexander, G. E., & Crutcher, M. D. (1990). Functional architecture of basal ganglia circuits: neural substrates of parallel processing. *Trends in neurosciences*, *13*(7), 266–271.

Barbera, G., Liang, B., Zhang, L., Gerfen, C. R., Culurciello, E., Chen, R., ... Lin, D.-T. (2016). Spatially compact neural clusters in the dorsal striatum encode locomotion relevant information. *Neuron*, *92*(1), 202–213.

Blouw, P., Solodkin, E., Thagard, P., & Eliasmith, C. (2016). Concepts as semantic pointers: A framework and computational model. *Cognitive science*, *40*(5), 1128–1162.

Bogacz, R., & Larsen, T. (2011). Integration of reinforcement learning and optimal decision-making theories of the basal ganglia. *Neural computation*, *23*(4), 817–851.

Chersi, F., Mirolli, M., Pezzulo, G., & Baldassarre, G. (2013). A spiking neuron model of the cortico-basal ganglia circuits for goal-directed and habitual action learning. *Neural Networks*, *41*, 212–224.

Chung, F. R. (1997). *Spectral graph theory* (Vol. 92). American Mathematical Soc.

DeLong, M. R. (1990). Primate models of movement disorders of basal ganglia origin. *Trends in neurosciences*, *13*(7), 281–285.

Eliasmith, C. (2013). *How to build a brain: A neural architecture for biological cognition*. Oxford University Press.

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

Gayler, R. W. (2004). Vector symbolic architectures answer Jackendoff's challenges for cognitive neuroscience. *arXiv preprint cs/0412059*.

Graybiel, A. M. (1998). The basal ganglia and chunking of action repertoires. *Neurobiology of learning and memory*, *70*(1-2), 119–136.

Gross, C. G. (2002). Genealogy of the "grandmother cell". *The Neuroscientist*, *8*(5), 512–518.

Gurney, K., Prescott, T. J., & Redgrave, P. (2001). A computational model of action selection in the basal ganglia. ii. analysis and simulation of behaviour. *Biological cybernetics*, *84*, 411–423.

Hikosaka, O., Takikawa, Y., & Kawagoe, R. (2000). Role of the basal ganglia in the control of purposive saccadic eye movements. *Physiological reviews*, *80*(3), 953–978.

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.

Jin, X., & Costa, R. M. (2015). Shaping action sequences in basal ganglia circuits. *Current opinion in neurobiology*, *33*, 188–196.

Klaus, A., Alves da Silva, J., & Costa, R. M. (2019). What, if, and when to move: basal ganglia circuits and self-paced action initiation. *Annual review of neuroscience*, *42*, 459–483.

Komer, B., Stewart, T. C., Voelker, A., & Eliasmith, C. (2019). A neural representation of continuous space using fractional binding. In *Cogsci* (pp. 2038–2043).

Lederman, J., Lardeux, S., & Nicola, S. M. (2021). Vigor encoding in the ventral pallidum. *Eneuro*, *8*(4).

Levy, S. D., & Gayler, R. W. (2009). Lateral inhibition in a fully distributed connectionist architecture. In *Proceedings of the ninth international conference on cognitive modeling*.

Mink, J. W. (1996). The basal ganglia: focused selection and inhibition of competing motor programs. *Progress in neurobiology*, *50*(4), 381–425.

Nambu, A., Tokuno, H., & Takada, M. (2002). Functional significance of the cortico–subthalamo–pallidal 'hyperdirect'pathway. *Neuroscience research*, *43*(2), 111–117.

Plate, T. A. (1992). Holographic recurrent networks. *Advances in neural information processing systems*, *5*.

Plate, T. A. (1995). Holographic reduced representations. *IEEE Transactions on Neural networks*, *6*(3), 623–641.

Plate, T. A. (2003). *Holographic reduced representation: Distributed representation for cognitive structures* (Vol. 150). CSLI Publications Stanford.

Quiroga, R. Q., Reddy, L., Kreiman, G., Koch, C., & Fried, I. (2005). Invariant visual representation by single neurons in the human brain. *Nature*, *435*(7045), 1102–1107.

Redgrave, P., Prescott, T. J., & Gurney, K. (1999). The basal ganglia: a vertebrate solution to the selection problem? *Neuroscience*, *89*(4), 1009–1023.

Stewart, T. C. (2012). A technical overview of the neural engineering framework. *University of Waterloo*, *110*.

Stewart, T. C., Bekolay, T., & Eliasmith, C. (2011). Neural representations of compositional structures: Representing and manipulating vector spaces with spiking neurons. *Connection Science*, *23*(2), 145–153.

Stewart, T. C., Choo, X., Eliasmith, C., et al. (2010). Dynamic behaviour of a spiking model of action selection in the basal ganglia. In *Proceedings of the 10th international conference on cognitive modeling* (pp. 235–40).

Stewart, T. C., & Eliasmith, C. (2012). Compositionality and biologically plausible models.

Turner, R. S., & Desmurget, M. (2010). Basal ganglia contributions to motor control: a vigorous tutor. *Current opinion in neurobiology*, *20*(6), 704–716.

Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., ... SciPy 1.0 Contributors (2020). SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, *17*, 261–272. doi: 10.1038/s41592-019-0686-2

Voelker, A. R., Blouw, P., Choo, X., Dumont, N. S.-Y., Stewart, T. C., & Eliasmith, C. (2021). Simulating and predicting dynamical systems with spatial semantic pointers. *Neural Computation*, *33*(8), 2033–2067.