Towards Continuous Action Representations in Models of the Basal Ganglia

Madeleine Bartlett[1], P. Michael Furlong[2], Terrence C. Stewart[3] Jeff Orchard[1],

[1]Cheriton School of Computer Science [2]Systems Design Engineering, [3]National Research Council of Canada, University of Waterloo

Introduction

Motivation:

- Accumulating evidence suggests that the Basal Ganglia plays a role in the selection of continuous action features (e.g. vigor, speed) [Dudman and Krakauer, 2016]
- Modelling this ability could lend insights into Basal Ganglia function and dysfunction, such as the development of bradykinesia in Parkinson's Disease
- In recent work we successfully implemented a model capable of handling continuous action spaces [Bartlett et al., 2024]
- This model uses **Spatial Semantic Pointers** (SSPs; [Komer et al., 2019]) to represent distributions of salience over action spaces
- However, the function implemented for transforming broad salience distributions to narrower distributions/specific actions was sub-optimal

Aim: To implement alternative models of Basal Ganglia function, capable of operating in continuous action spaces.
Result: Demonstrated several networks capable of taking broad salience distributions over continuous actions spaces and sharpening them.

What Are SSPs?

- Neurally instantiated, symbol-like representations of concepts that can be transformed in many ways to yield further representations that function to support cognitive processes [Eliasmith, 2013]
- High dimensional vectors
- Can be represented in the activities of populations of neurons - different neurons will be sensitive to different parts of the vector-space (Fig 1)
- Established Vector Symbolic Algebras (VSAs) exist that allow us to perform operations on vectors, combining them together to form more complex structures

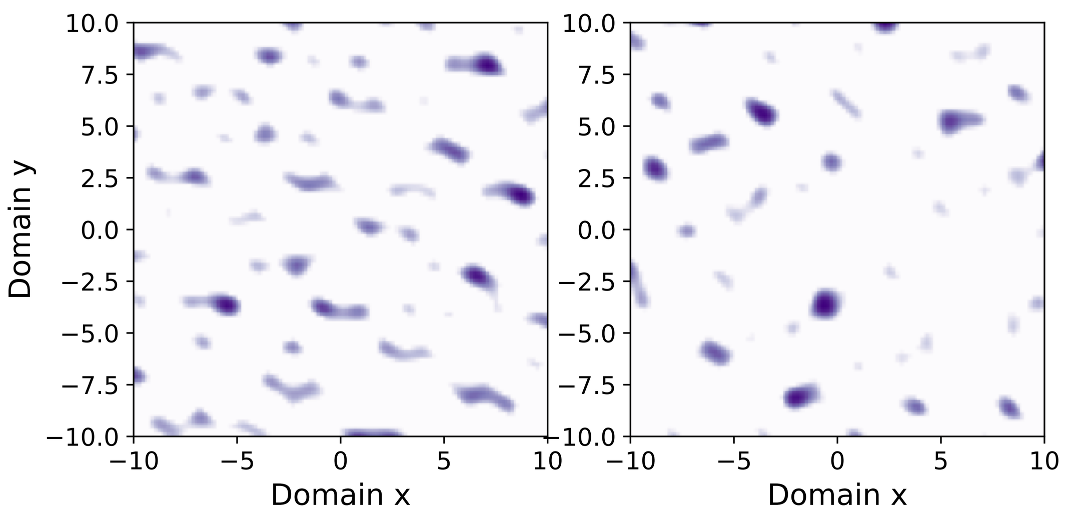

Figure 1: Receptive field of two neurons taken from a population of neurons encoding SSPs. The SSPs are encoding a 2D space. Darker patches show regions of the 2D space in which the neuron will fire.

VSAs

- Plate's Holographic Reduced Representations [Plate, 1995]
- Similarity = cosine similarity (dot product). Concepts that are semantically similar will be closer together in the vector space (Fig 2)
- Binding = circular convolution ($\circledast$; element-wise multiplication in Fourier domain). Combine 2 vectors into a new, dissimilar vector
- Bundling = vector addition. Combine vectors into sets, maintaining similarity
- Unbinding = binding with the inverse. For bound vector $A \circledast B$, unbinding A gives B (Fig 3)
- We can use this VSA to **encode both discrete and continuous variables**

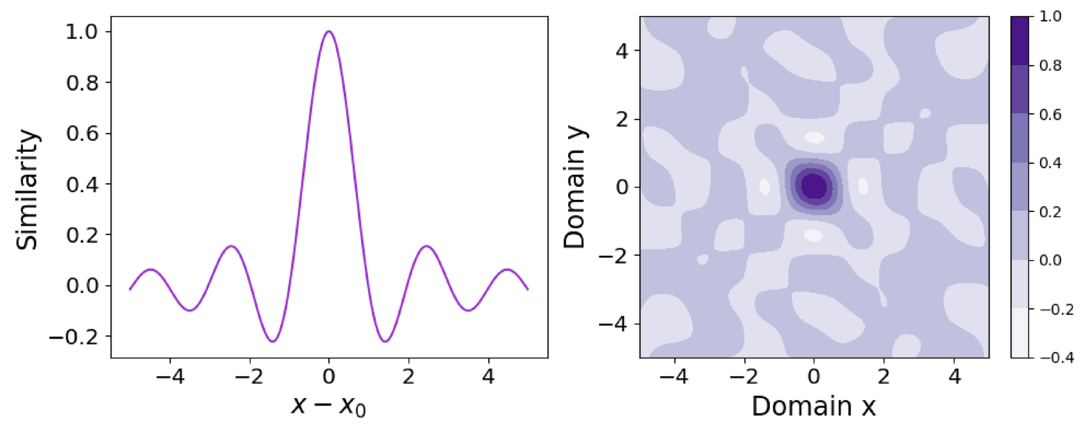

Figure 2: Similarity between an SSP encoding the origin (0) and (0,0) and SSPs encoding the whole domain space. Left: 1-D domain space, Right: 2-D domain space

$$\mathrm{MAP} = (\mathrm{STRAWBERRY} \circledast xy_{\mathrm{STRAWBERRY}}) + (\mathrm{APPLE} \circledast xy_{\mathrm{APPLE}})$$
$$xy_{\mathrm{STRAWBERRY}} = \mathrm{MAP} \circledast \mathrm{STRAWBERRY}^{-1}$$

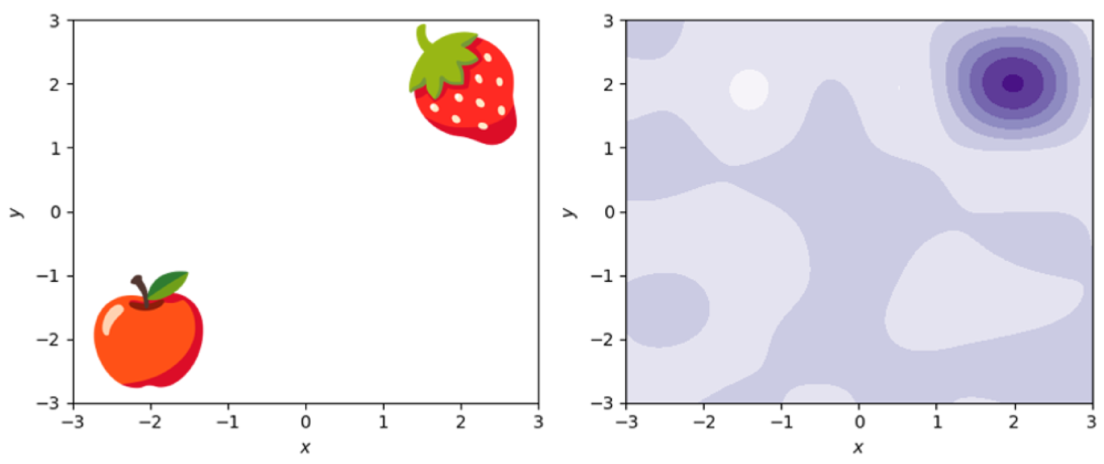

Figure 3: Objects Strawberry and Apple are bound with their locations on the x,y plane (left). These two bound vectors are bundled together. The bundle is queried (via unbinding) to find the location of the Strawberry (right)

Encoding Continuous Variables - Fractional Binding

We use SSPs to represent the action space because they:

- offer us a means of representing information that can be neurally instantiated
- can be used to represent both discrete and continuous variables in the same neural substrate

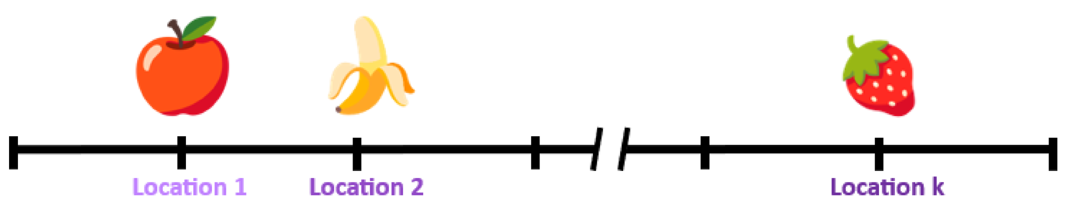

$\vec{V}$ = Location

Location 1 = $\vec{V}$

Location 2 = $\vec{V} \circledast \vec{V} = \mathcal{F}^{-1}\{\mathcal{F}\{V \times V\}\} = \mathcal{F}^{-1}\left\{\mathcal{F}\{\vec{V}\}^2\right\}$

Location k = $\vec{V} \circledast \vec{V} \circledast \ldots \circledast \vec{V} = \mathcal{F}^{-1}\left\{\mathcal{F}\{\vec{V}\}^k\right\}$

$k \in \mathbb{R}$

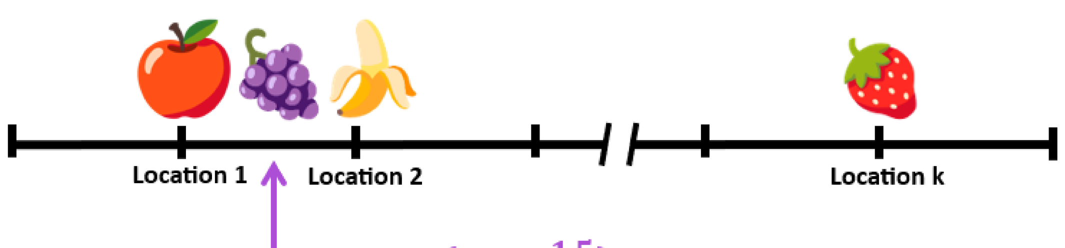

Location 1.5 = $\mathcal{F}^{-1}\left\{\mathcal{F}\{\vec{V}\}^{1.5}\right\}$

Experiments

We are using SSPs to represent actions in models of the Basal Ganglia.

- Traditional models select from discrete action spaces by choosing the action with the highest **salience**
- For continuous actions, we would either do the same, or produce a sharpened distribution that's easier to select from (Fig 4)

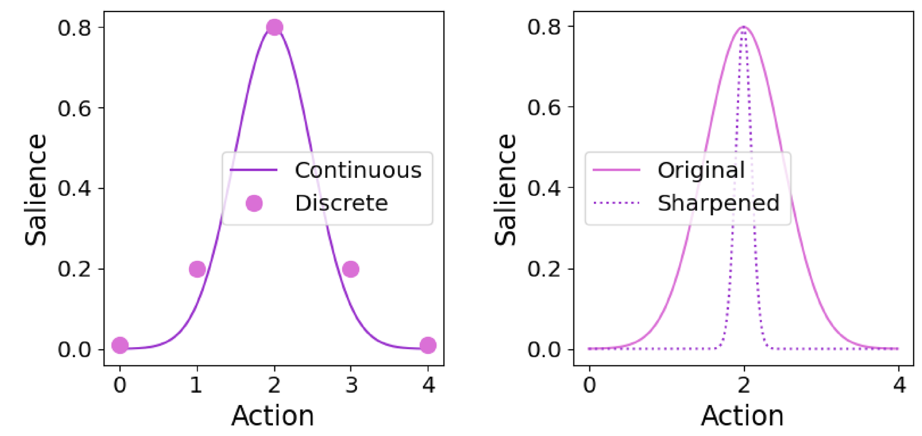

Figure 4: Left: example distributions of saliences over a discrete and a continuous action space. Right: an illustration of a sharpened distribution

So far we've explored using;

1. a Hopfield Network,
2. an Independent Accumulator, and
3. a Winner-Take-All Network

Results

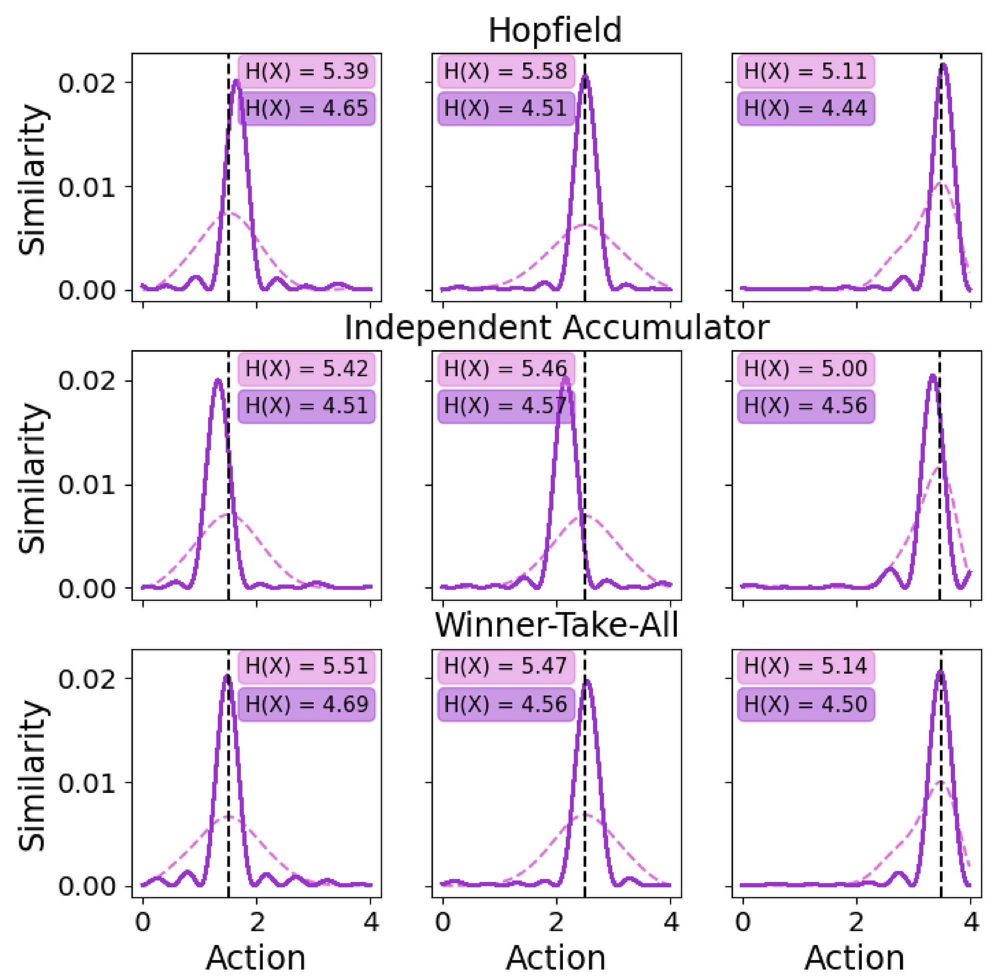

Figure 5: Plots of normalized similarity for the input (dashed) and output (solid) salience distributions from each network for 3 experiments. Similarity is a proxy for salience. Labels show the entropy ($H(X)$) of the Input (light pink) and Output (darker purpler) salience distributions

Conclusions

- The models we present were all effective at transforming wide distributions of salience over continuous action spaces to sharper distributions
- This work illustrates SSPs as a viable option for modelling selection from continuous action spaces in the Basal Ganglia

Acknowledgements

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

References

Madeleine Bartlett, Michael Furlong, Terrence C Stewart, and Jeff Orchard. Using vector symbolic architectures for distributed action representations in a spiking model of the basal ganglia. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 46, 2024.

Joshua T Dudman and John W Krakauer. The basal ganglia: from motor commands to the control of vigor. *Current opinion in neurobiology*, 37:158–166, 2016.

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

Brent Komer, Terrence C Stewart, Aaron Voelker, and Chris Eliasmith. A neural representation of continuous space using fractional binding. In *CogSci*, pages 2038–2043, 2019.

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