Learning in large-scale models of biological cognition (FA9550-17-1-0026)

PI: Chris Eliasmith (University of Waterloo)

Senior Personnel: Brent Komer, Peter Duggins, Jan Gosmann, Thomas Lu, Nicole Dumont, Andreas Stoeckel

AFOSR Program Review:
Computational Cognitive and Machine Intelligence
(Oct 2020, Washington, DC)

Learning in Large-scale Models of Biological Cognition

Chris Eliasmith, University of Waterloo

Objective: Adaptive neural cognition

- Develop and test functional models of hippocampus for working and long-term memory consolidation
- Develop and test functional models of basal ganglia for hierarchical reinforcement learning (HRL)
- Integrate these models within a large-scale neural model of the brain (Spaun)

Approach: Learning in large spiking nets

- NEF for neural dynamics and non-linear compute
- SPA for high-D cognitive repns and action selection
- Various NEF/SPA methods for online learning in spiking networks: hPES, voja, etc.

DoD Benefits:

- Human-like cognitive behaviour for more natural artificial collaborators
- Continuous adaptive control systems

Progress:

- Extensions to world's largest functional brain model (Spaun)
- Initial HRL and hippocampal stand-alone systems

Progress (cont):

- Large-scale spiking model of working, long-term and episodic memory
- New neural representation of metric spaces (Spatial Semantic Pointers), state-of-the-art for policy learning
- State-of-the-art RNN (Legendre Memory Unit) for temporal memory, from hippocampal model

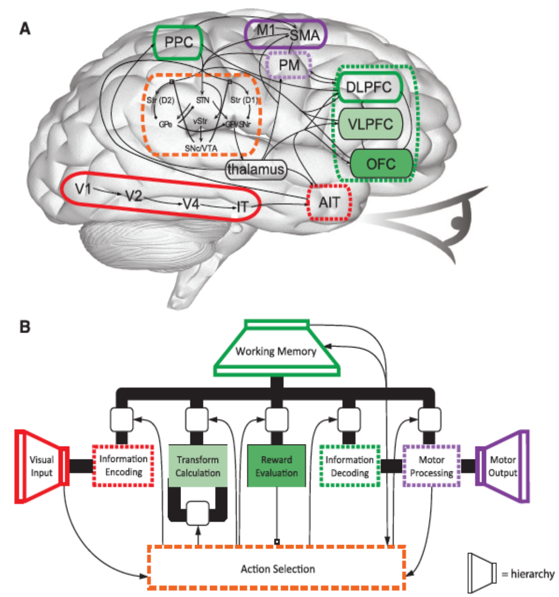

Spaun depicted A) anatomically and B) functionally.

List of Project Goals

1. Extend models of learning and representation in the hippocampus, known to be central to episodic and working memory function
2. Extend models of learning in basal ganglia and related forebrain structures, known to be involved in model-based and hierarchical reinforcement learning (HRL)
3. Integrate these approaches in a large-scale model of the brain, and examine low-level biological interventions on human cognition
4. Construct biologically plausible models of the consolidation of information to long-term memory and its use
5. Model how the brain learns complex cognitive tasks using reinforcement learning and related techniques

Progress Towards Goals

1. Unifying model of working and long-term memory (Gosmann)
2. Legendre Memory Units (LMUs) for ML and neuro temporal learning (Voelker, Dejong, Chilkuri)
3. SSPs and grid cells for accurate continuous representation (Komer, Dumont)

hippocampus

4. Successor repn using SSPs for A2C RL (Dumont)
5. Complex single neuron models in Spaun 2.0 (Duggins)
6. Adaptive filtering using LMU repns (Stoeckel, Nat)
7. Spatial Semantic Pointers (SPPs) (Komer, Voelker, Stewart) [last report]
8. Delay Networks (DNs) (Voelker) [last report]
9. …+8

CUE (Context unified encoding)

- "A unified spiking neuron model of short-term and long-term memory" (Psych Rev, 2020)
- Self-contained (intrinsic control)
- Free recall, serial recall
- Rapid episodic learning
- Hebb effect
- Spiking neurons

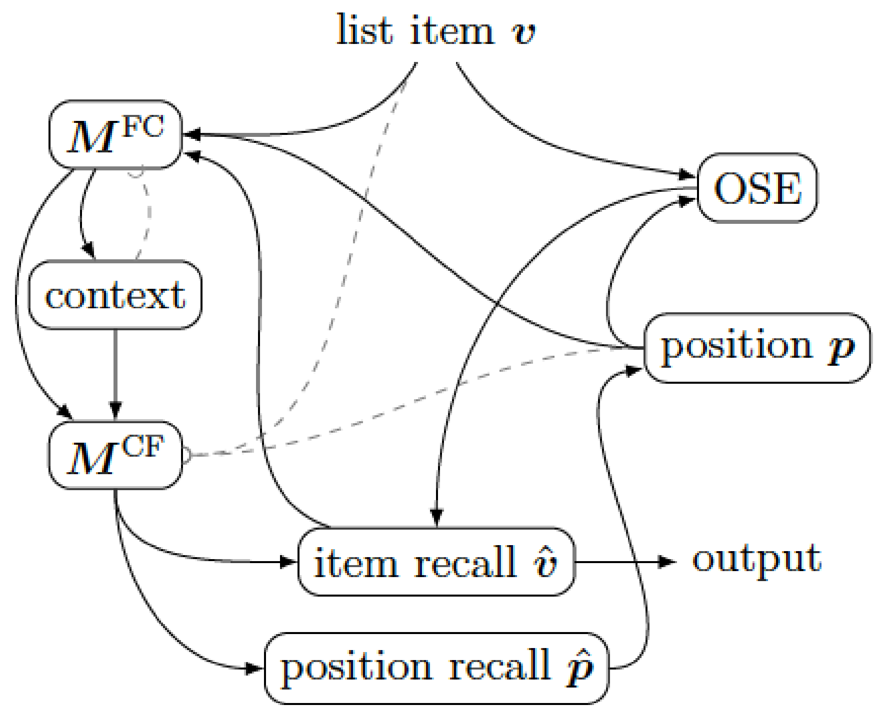

AML Learning Rule

- Association matrix learning (AML)

$$\frac{\mathrm{d}\boldsymbol{W}}{\mathrm{d}t} = \eta \boldsymbol{E}\boldsymbol{v}(t)\left(\boldsymbol{D}^{\top}\boldsymbol{D}\boldsymbol{a}_{\boldsymbol{u}(t)}\right)^{\top}$$

- Derived from PES, but doesn't suffer from catastrophic forgetting
- Also biologically plausible (e.g. hippocampal rapid learning)

AML vs PES Rapid Learning

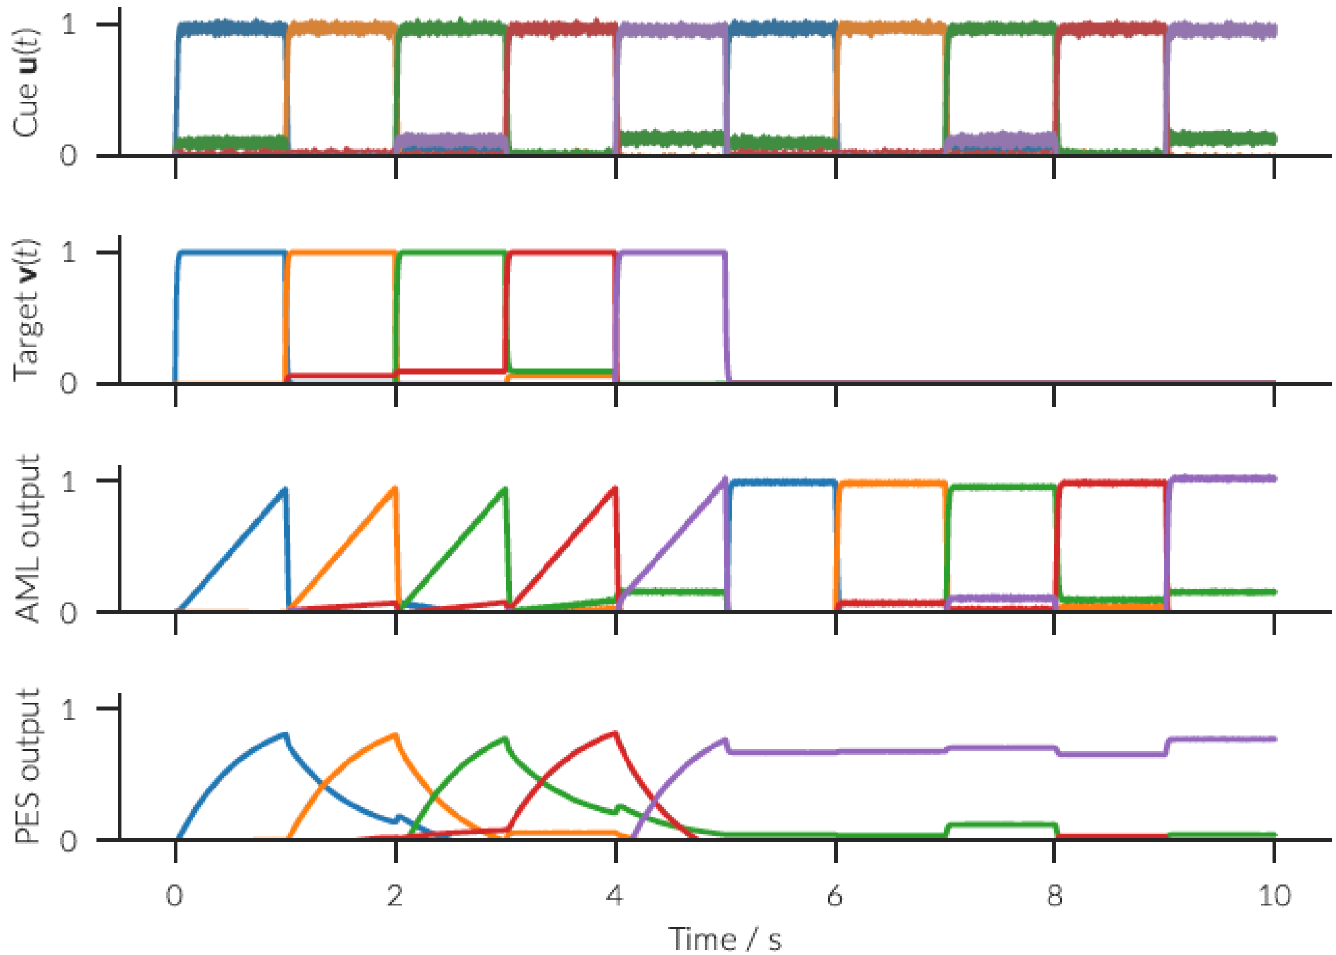

Neural support

- Consistent with human recordings from hippocampus during association learning

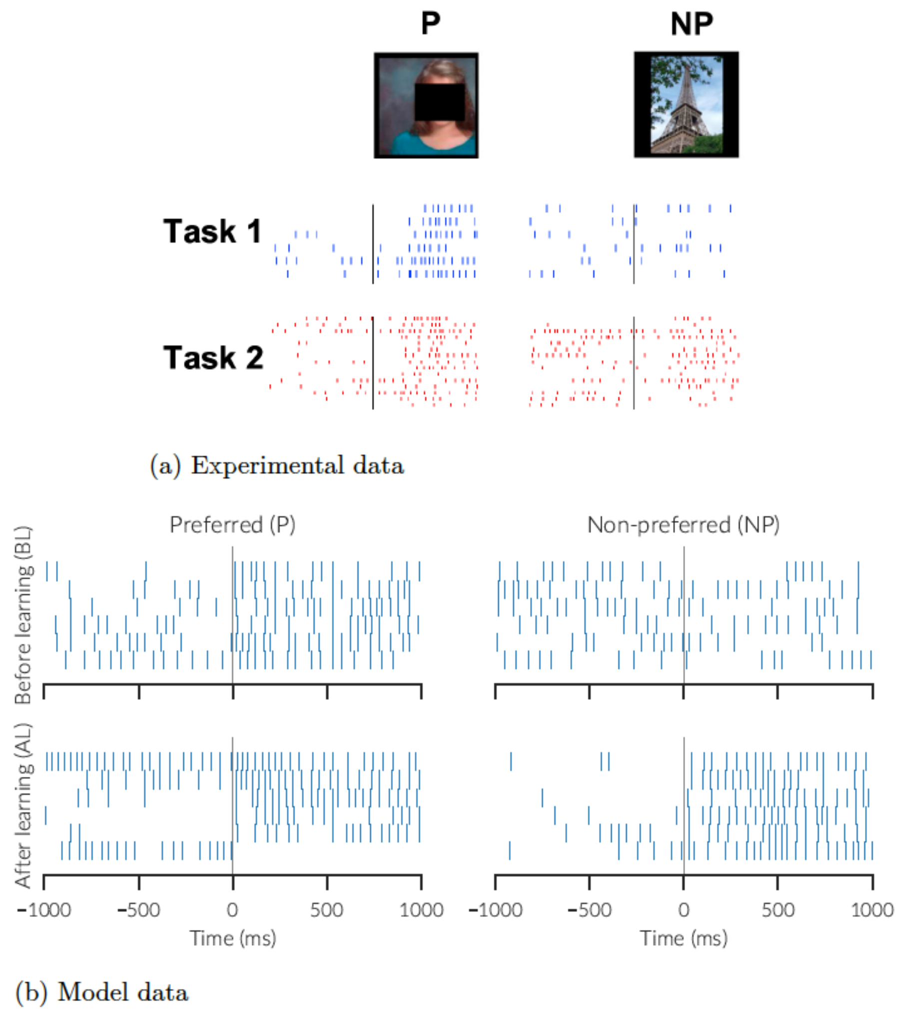

(a) Experimental data

(b) Model data

Overall CUE model

- Three free recall variations
- Better match than past models over this variety
- Explains neural similarity as well (not shown)

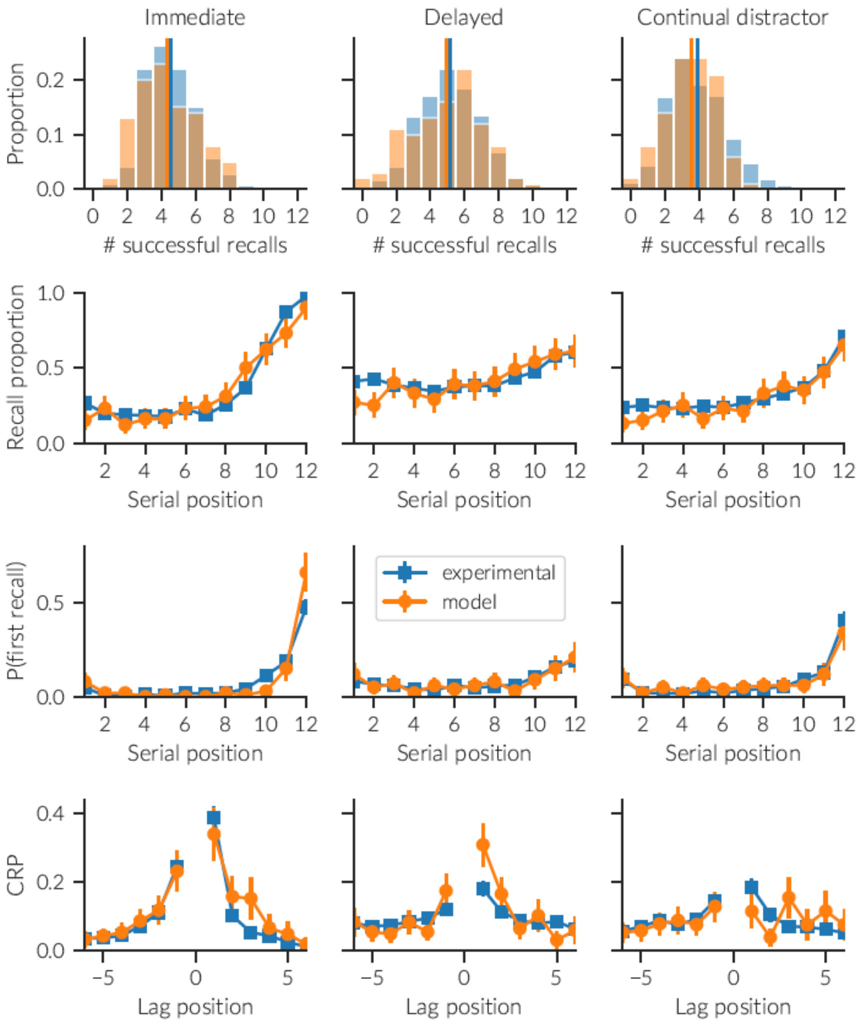

Continuous Time Processing

- Legendre Memory Units (LMUs)
  - Optimal temporal compression
  - 10^6x more accurate while compressing 10^4x more data than LSTMs
  - Time cells (hippo)

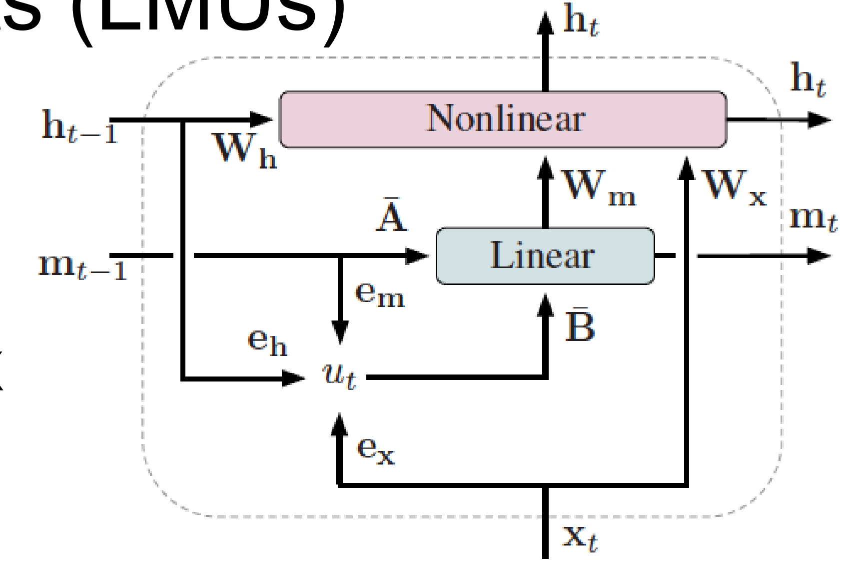

- New proposal ff-LMU
  - Purely feedforward training (recurrent inference)
  - 11-60x faster training than LSTMs; scales like ff-net (e.g. transformers)

ML Results Summary

- Beating state of the art on benchmarks for RNNs
  - psMNIST
  - QQP
- Parameter efficient
  - First transformer to beat LMU on QQP uses 50,000x more parameters
  - 36-1000x fewer parameters than LSTMs on 4 tested datasets (with higher accuracy)

Neuro Results Summary

(with DeJong et al.)

- Captures known violation of the scalar property (not captured by past models)
- Matches Wang et al data

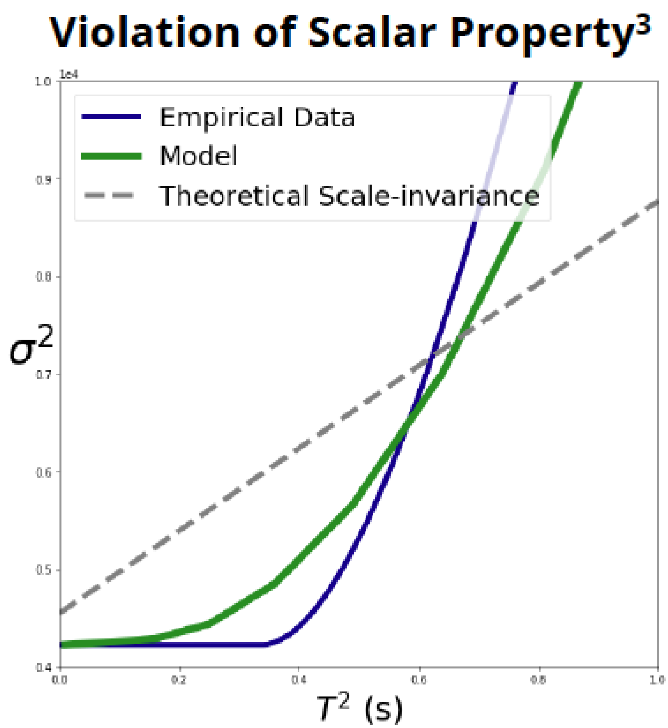

Neural Responses

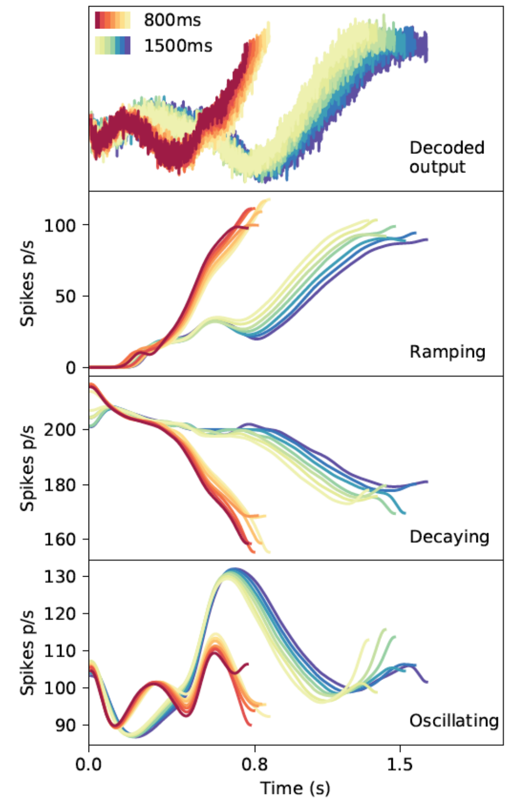

Spatial Semantic Pointers

- Semantic pointers represent standard discrete structures (lists, trees, etc.)
- SSPs allow recurrent convolutions to have fractional powers

$$B^k = \underbrace{B \circledast B \circledast \ldots \circledast B}_{B \text{ appears } k \text{ times}}$$

- Compute fractional k in Fourier space

$$B^k = \mathcal{F}^{-1}\left\{\mathcal{F}\left\{B\right\}^k\right\}, \quad k \in \mathbb{R}$$

Spatial Semantic Pointers

- Represent continuous space (Clifford torus)

$$S(x,y) = X^x \circledast Y^y$$

- Bind
  objects at locations

$$M = \sum_{i=1}^{m} OBJ_i \circledast S_i$$

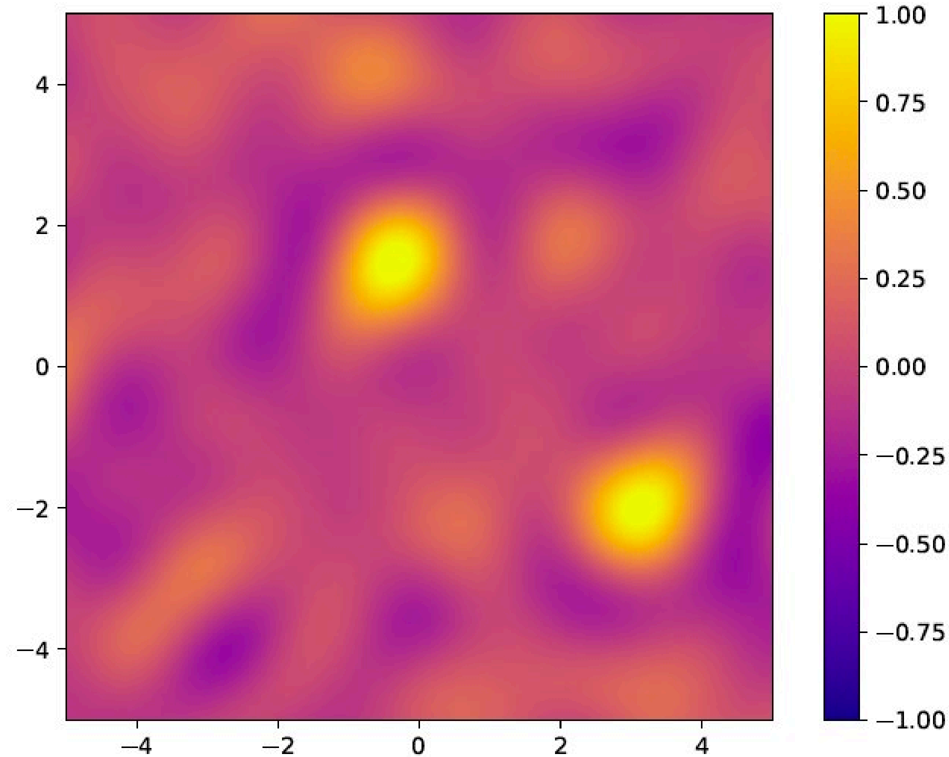

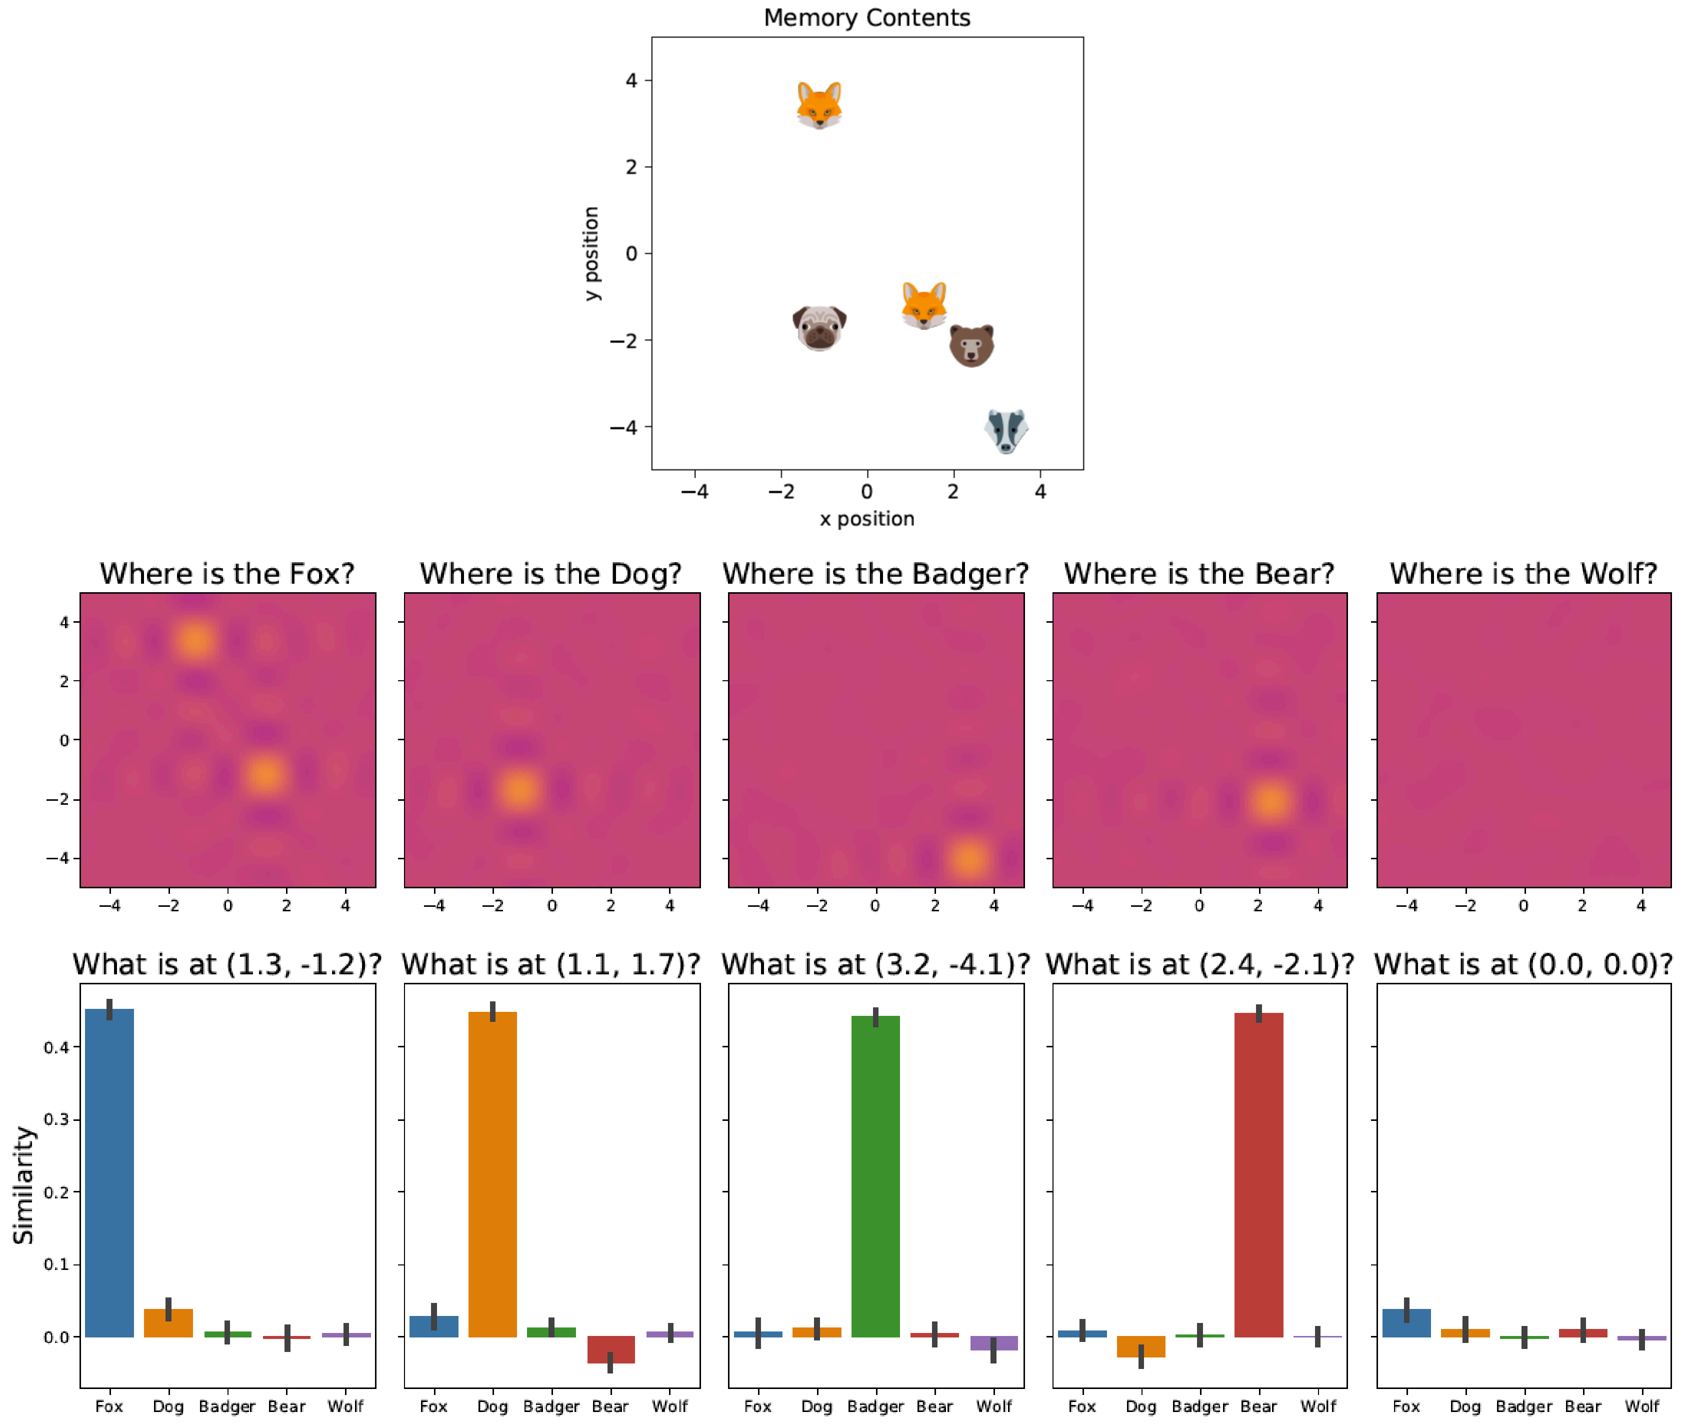
Memory Contents
y position
x position
Where is the Fox?
Where is the Dog?
Where is the Badger?
Where is the Bear?
Where is the Wolf?
What is at (1.3, -1.2)?
What is at (1.1, 1.7)?
What is at (3.2, -4.1)?
What is at (2.4, -2.1)?
What is at (0.0, 0.0)?
Similarity
Fox
Dog
Badger
Bear
Wolf

Grid Cells

- We found a method to choose X, Y to give grid cells $S(x,y) = X^x \circledast Y^y$
- And combine them to get place cells (with standard NEF decoders)

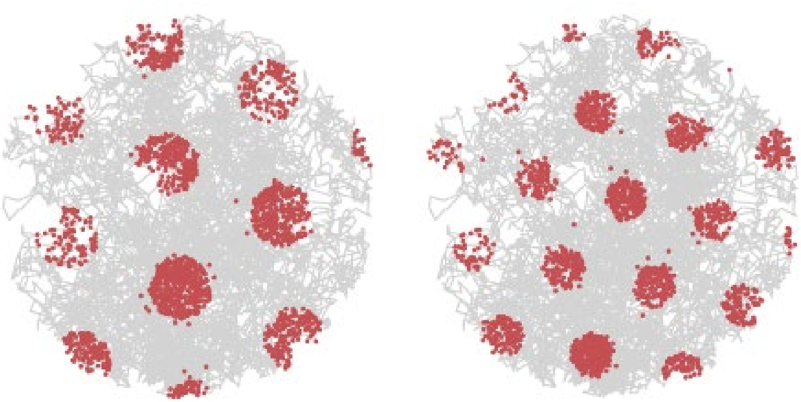

SSP Grid cells

Place cells

Hexagonal SSPs

- They are better on a wide variety of tasks

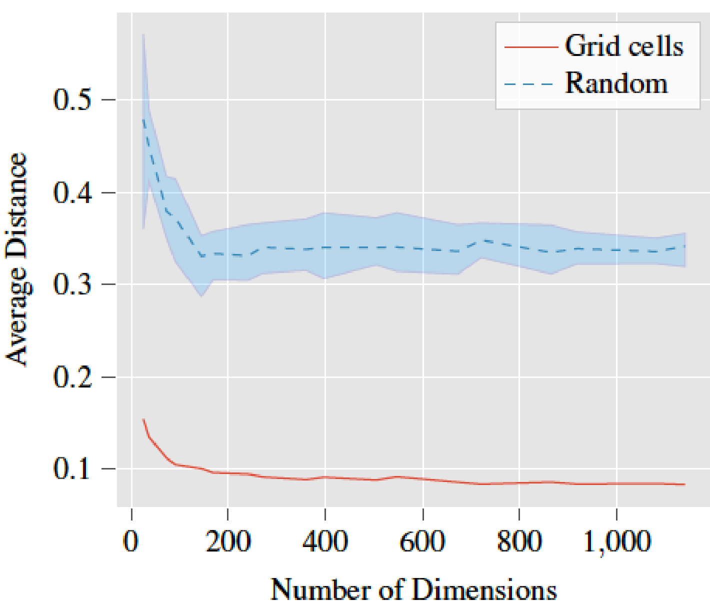

Accuracy for representing locations

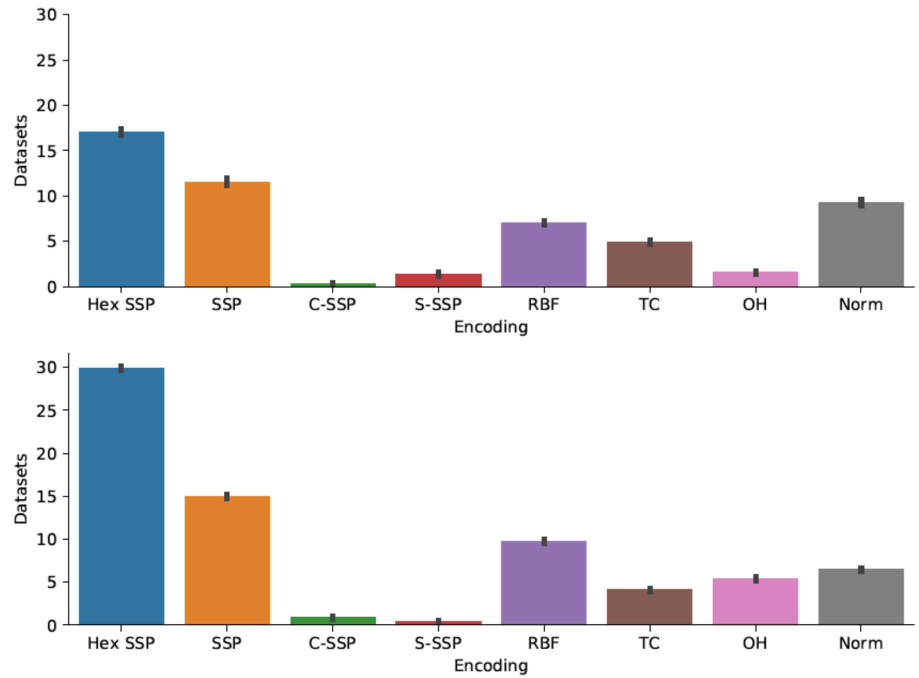

Arbitrary ML problems

Lots more…

- Generalizes better outside of the trained space than other repns
- Has higher capacity than other repns
- Works well with RL (fast convergence)
- Built a full location memory system with navigation – one-shot/episodic learning
- PhD thesis

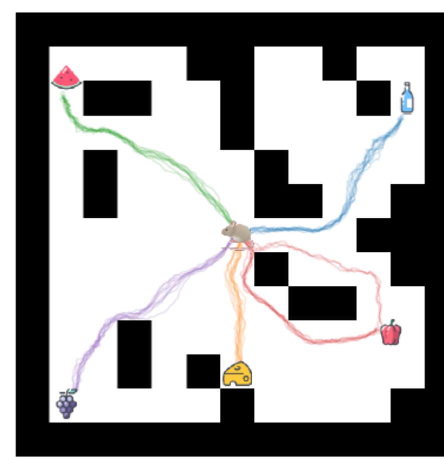

Conclusion

- Exploring new hippocampus inspired representations for both ML and psych models
- Specifically,
  - Improved spatial representations (critical to hippocampus)
  - Improved continuous temporal representation
  - Improved integration of long-term and WM
- Integrated system with short term episodic memory

List of Publications/Awards Attributed to the Grant

- *Best cognitive modeling prize - Language* Peter Blouw and Chris Eliasmith. Inferential role semantics for natural language. In Thora Tenbrink, Glenn Gunzelmann, Andrew Howes and Eddy Davelaar, editors, *Proceedings of the 39th Annual Conference of the Cognitive Science Society, 142–147*. Philadelphia, Pennsylvania, 2017. Cognitive Science Society.
- Peter Duggins, Terrence C. Stewart, Xuan Choo, and Chris Eliasmith. Effects of guanfacine and phenylephrine on a spiking neuron model of working memory. *Topics in Cognitive Science*, 2017.
- Jan Gosmann, Aaron R. Voelker, and Chris Eliasmith. A spiking independent accumulator model for winner-take-all computation. In *Proceedings of the 39th Annual Conference of the Cognitive Science Society*. London, UK, 2017. Cognitive Science Society.
- Jan Gosmann and Chris Eliasmith. Automatic optimization of the computation graph in the nengo neural network simulator. *Frontiers in Neuroinformatics*, 11:33, 2017.
- Ivana Kajić, Jan Gosmann, Brent Komer, Ryan W. Orr, Terrence C. Stewart, and Chris Eliasmith. A biologically constrained model of semantic memory search. In *Proceedings of the 39th Annual Conference of the Cognitive Science Society*. London, UK, 2017. Cognitive Science Society.

List of Publications/Awards Attributed to the Grant

- Ivana Kajić, Jan Gosmann, Terrence C. Stewart, Thomas Wennekers, and Chris Eliasmith. A spiking neuron model of word associations for the remote associates test. *Frontiers in Psychology*, 8:99, 2017.
- Daniel Rasmussen, Aaron R. Voelker, and Chris Eliasmith. A neural model of hierarchical reinforcement learning. *PLoS ONE*, 12(7):1–39, 2017.
- Peter Blouw and Chris Eliasmith (2018) Using Neural Networks to Generate Inferential Roles for Natural Language. *Frontiers in Psychology*.
- Terrence C. Stewart, Sverrir Thorgeirsson, and Chris Eliasmith. Supervised learning of action selection in cognitive spiking neuron models. In 40th Annual Conference of the Cognitive Science Society, 1086–1091. Cognitive Science Society, 2018.
- Sverrir Thorgeirsson, Terrence C. Stewart, and Chris Eliasmith. Analysis of learning action selection parameters in a neural cognitive model. In International Conference on Cognitive Modelling. 2018.
- Sverrir Thorgeirsson, Brent Komer, and Chris Eliasmith. Incorporating an adaptive learning rate in a neural model of action selection. In Cognitive Computational Neuroscience. Philadelphia, USA, 2018.
- Jan Gosmann and Chris Eliasmith. Vector-derived transformation binding: an improved binding operation for deep symbol-like processing in neural networks. Neural Computation, 31(5):849-869, 05 2019.

List of Publications/Awards Attributed to the Grant

- Aaron R. Voelker and Chris Eliasmith. Improving spiking dynamical networks: accurate delays, higher-order synapses, and time cells. Neural Computation, 30(3):569-609, 03 2018
- Brent Komer, Terrence C. Stewart, Aaron R. Voelker, and Chris Eliasmith. A neural representation of continuous space using fractional binding. In 41st Annual Meeting of the Cognitive Science Society. Montreal, QC, 2019. Cognitive Science Society.
- Joost de Jong, Aaron R. Voelker, Hedderik van Rijn, Terrence C. Stewart, and Chris Eliasmith. Flexible timing with delay networks – the scalar property and neural scaling. In 17th Annual Meeting of the International Conference on Cognitive Modelling (ICCM). 2019.
- Peter Duggins and Chris Eliasmith. A spiking neuron model of pharmacologically-biased fear conditioning in the amygdala. In SfN Abstracts. Chicago USA, 2019.
- Aaron R. Voelker, Ivana Kajić, and Chris Eliasmith. Legendre memory units: continuous-time representation in recurrent neural networks. In Advances in Neural Information Processing Systems, 15544–15553. 2019.
- Thomas Lu, Aaron R. Voelker, Brent Komer, and Chris Eliasmith. Representing spatial relations with fractional binding. In 41st Annual Meeting of the Cognitive Science Society. Montreal, QC, 2019. Cognitive Science Society.

List of Publications/Awards Attributed to the Grant

- Andreas Stöckel, Terrence C. Stewart, and Chris Eliasmith. Connecting biological detail with neural computation: application to the cerebellar granule-golgi microcircuit. In 18th Annual Meeting of the International Conference on Cognitive Modelling. Toronto, ON, 2020. Society for Mathematical Psychology. ***Best paper award
- Andreas Stöckel, Terrence C. Stewart, and Chris Eliasmith. A biologically plausible spiking neural model of eyeblink conditioning in the cerebellum. In 42nd Annual Meeting of the Cognitive Science Society, 1614–1620. Toronto, ON, 2020. Cognitive Science Society.
- Brent Komer and Chris Eliasmith. Efficient navigation using a scalable, biologically inspired spatial representation. In 42nd Annual Meeting of the Cognitive Science Society. Toronto, ON, 2020. Cognitive Science Society.
- Nicole Sandra-Yaffa Dumont and Chris Eliasmith. Accurate representation for spatial cognition using grid cells. In 42nd Annual Meeting of the Cognitive Science Society, 2367–2373. Toronto, ON, 2020. Cognitive Science Society.
- Jan Gosmann and Chris Eliasmith. CUE: a unified spiking neuron model of short-term and long-term memory. Psychological Review, 2020.