UCL

What can AI contribute to neuroscience?

Caswell Barry, UCL.

Novel Computational Paradigms, Cambridge.

Neuroscience

- Encompasses many levels: from labs studying individual neurons to fMRI
- Aim: Understand how the brain generates cognitive phenomena (e.g. how do we remember?)
- Specific questions:
  - what information is present at different points in the brain – how it it represented?
  - how are stimuli encoded and behaviors affected?
  - what computations are performed by neural circuits?

Two problems

- Bottom up - the data problem:
    - fMRI (2GB per brain), electrodes (1GB/minute), microscopy (2GB+/minute)
    - what information is present, how it it encoded, what computations are performed?

What’s encoded?

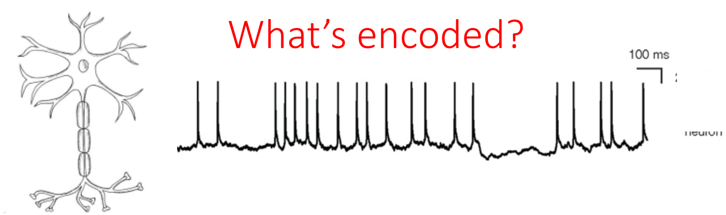

- Top down - the hypothesis & model problem:

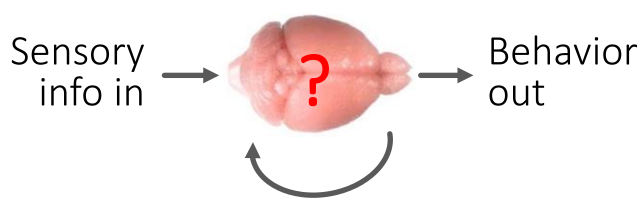

Deep networks (DNNs)

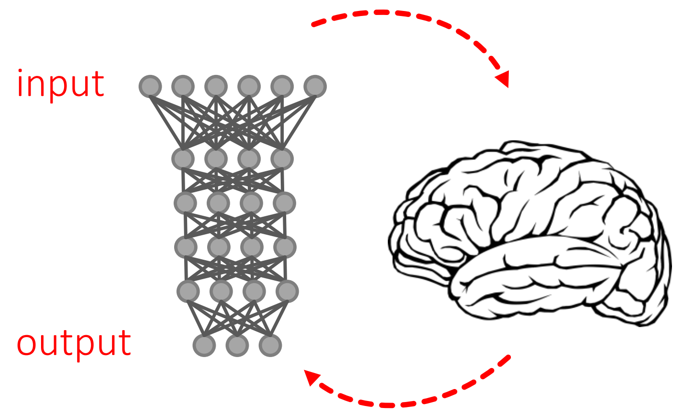

Bottom up

1. Deep learning is extremely good at mapping noisy input data to an output
   - e.g. stimuli to neural data
   - once trained a DNN can be interrogated

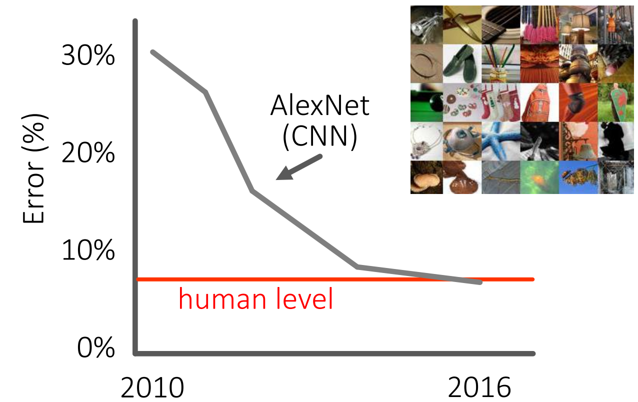

Top down

2. DNNs sometimes solve problems in a similar way to the brain – potentially provides a good model system
   - train to perform similar tasks to animals (e.g. visual recognition, navigation, limb control)

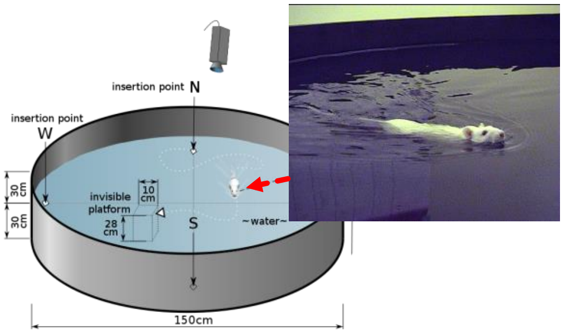

Top down (build a model)

Andrea Banino

Dharshan Kumaran

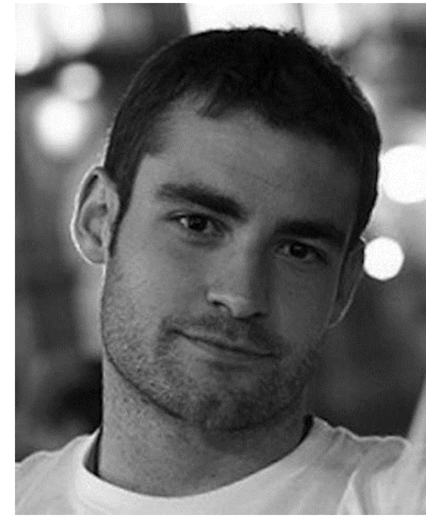

Benigno Uria

DeepMind

UCL CoMPLEX

Banino et al (2018) Vector-based navigation using grid-like representations in artificial agents

Tolman's Cognitive Map (1948)

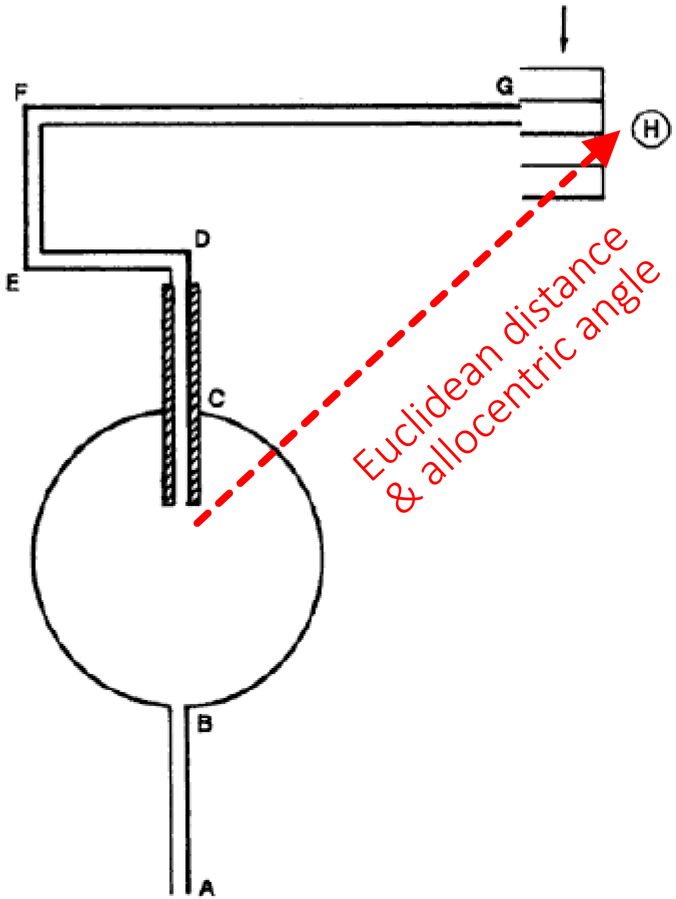

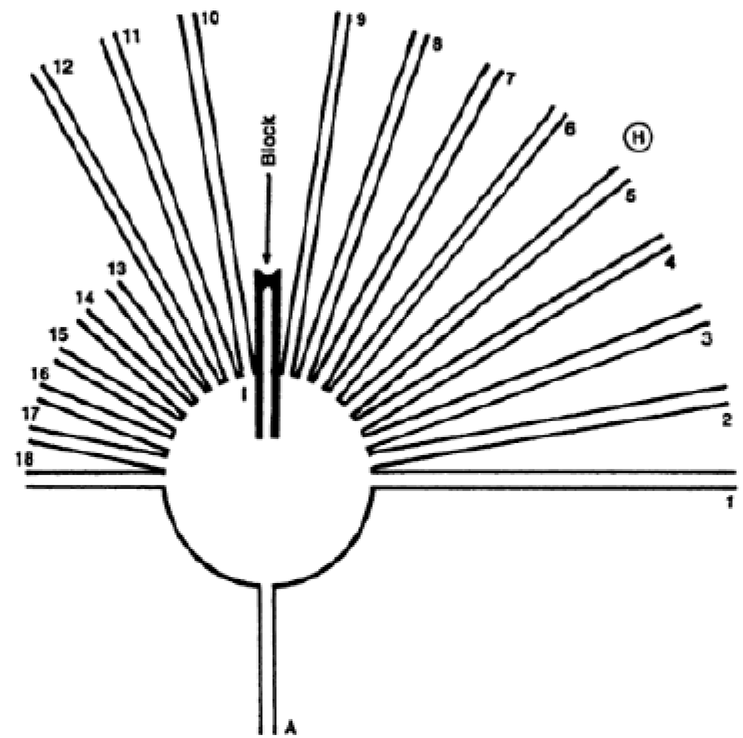

Place cells & grid cells

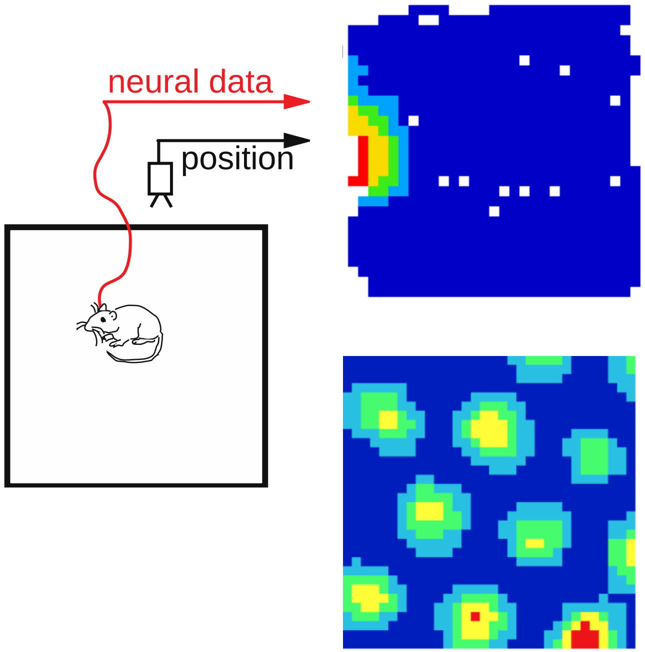

- Stably represent self-location
- Common to mammals
  (& possibly birds)

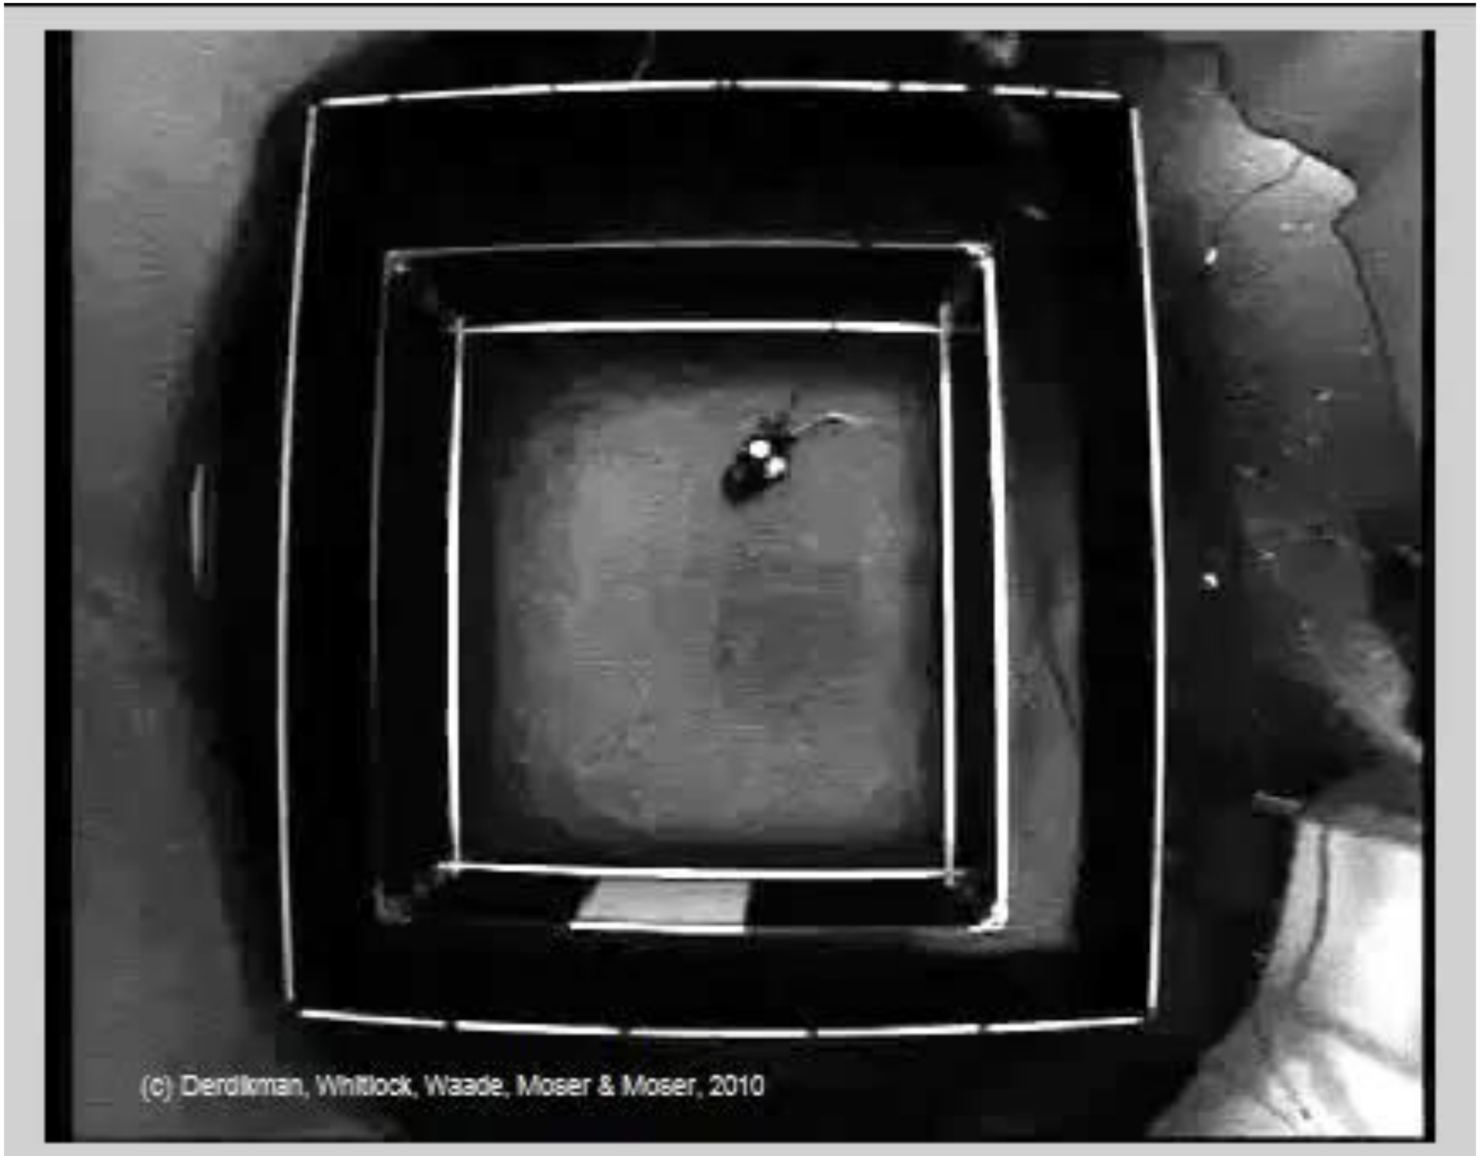
(c) Derdikman, Whitlock, Waade, Moser & Moser, 2010

Place cells & grid cells

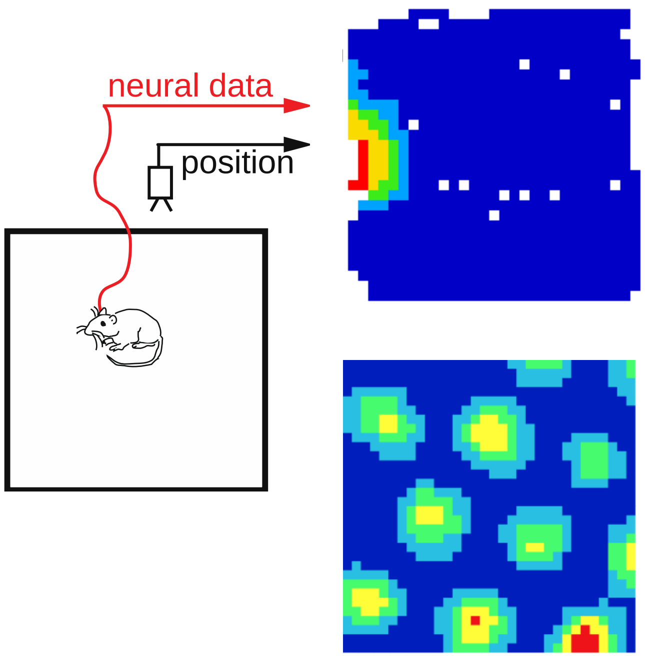

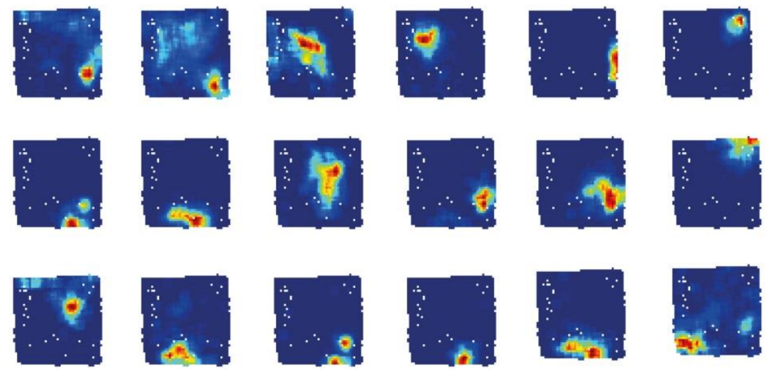

- Stably represent self-location
- Common to mammals (& possibly birds)

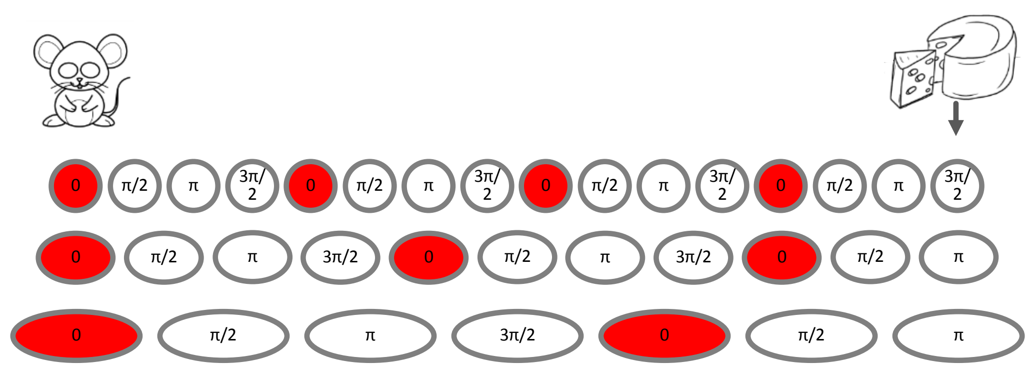

Goal distance = [3π/2, π, π] = +75cm

Aims

1. Test if mammalian-like neural representations emerge in a deep network trained to path integrate

2. Use such a network as a model system on which to conduct experiments

   → Demonstrate that grid cells are an effective basis for vector based navigation

Supervised learning architecture

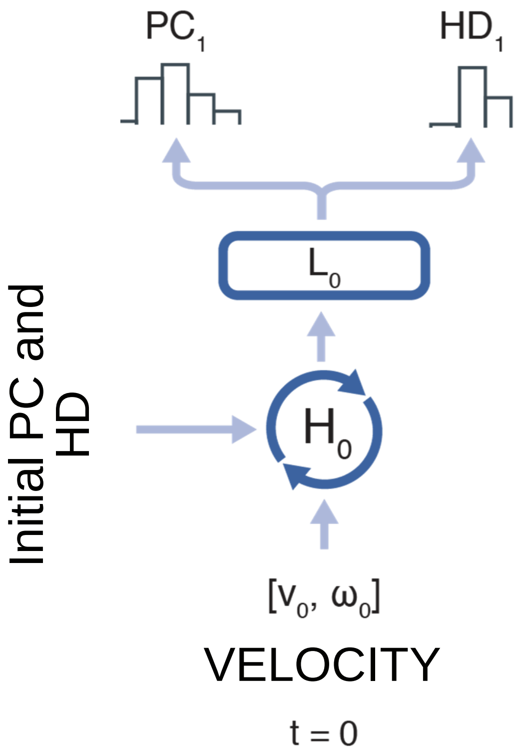

Banino et al (2018) Vector-based navigation using grid-like representations in artificial agents

Supervised learning architecture

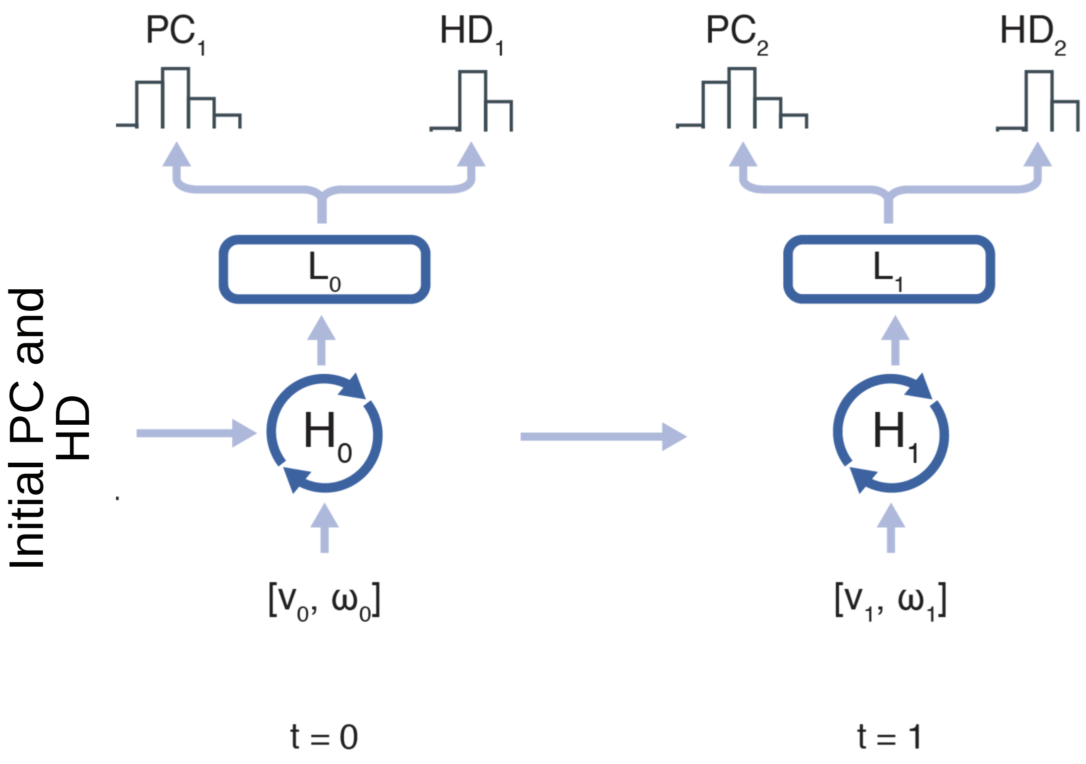

Banino et al (2018) Vector-based navigation using grid-like representations in artificial agents

Supervised learning architecture

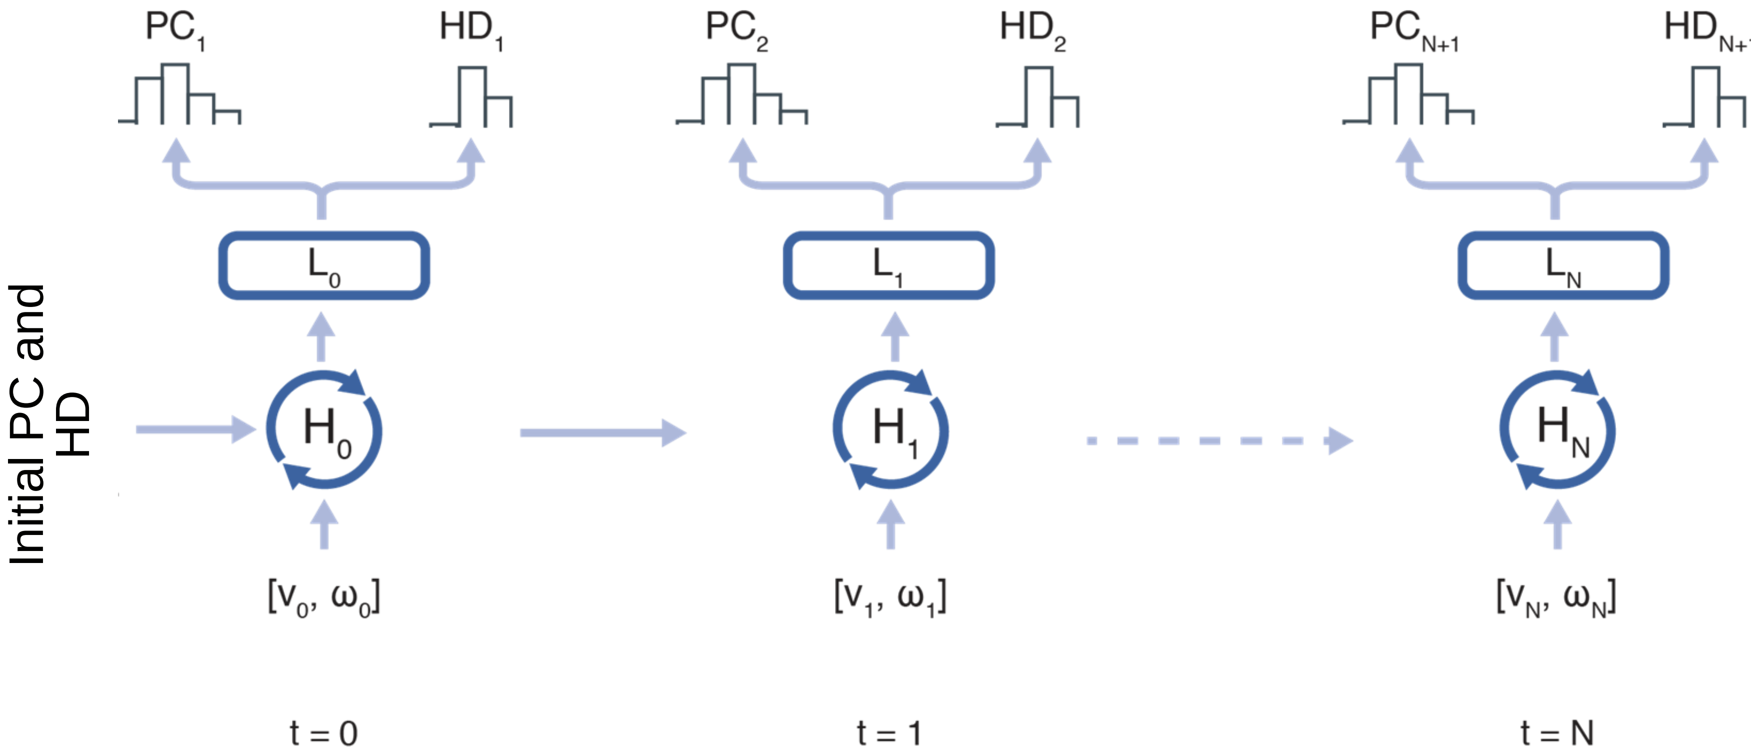

Banino et al (2018) Vector-based navigation using grid-like representations in artificial agents

Path integration task

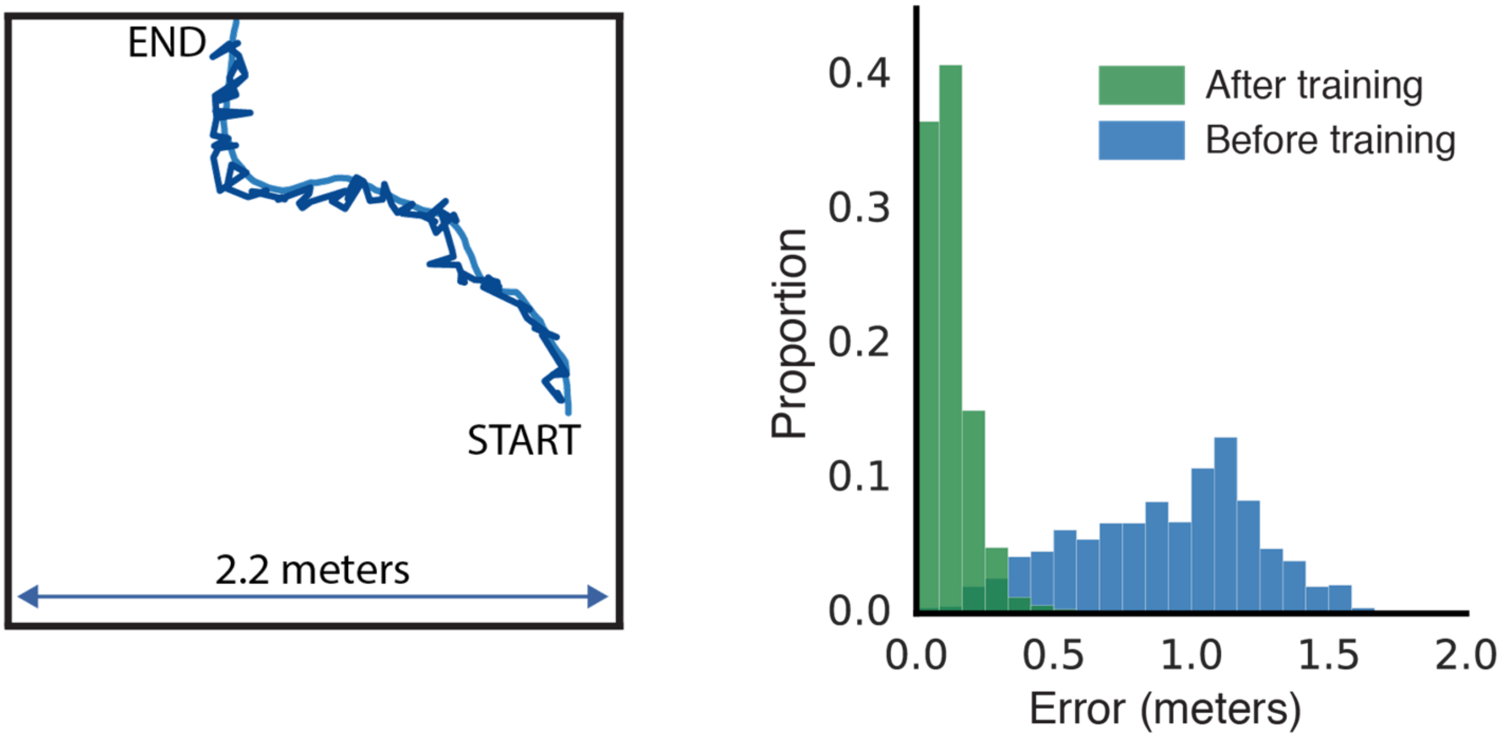

Analysis of linear layer

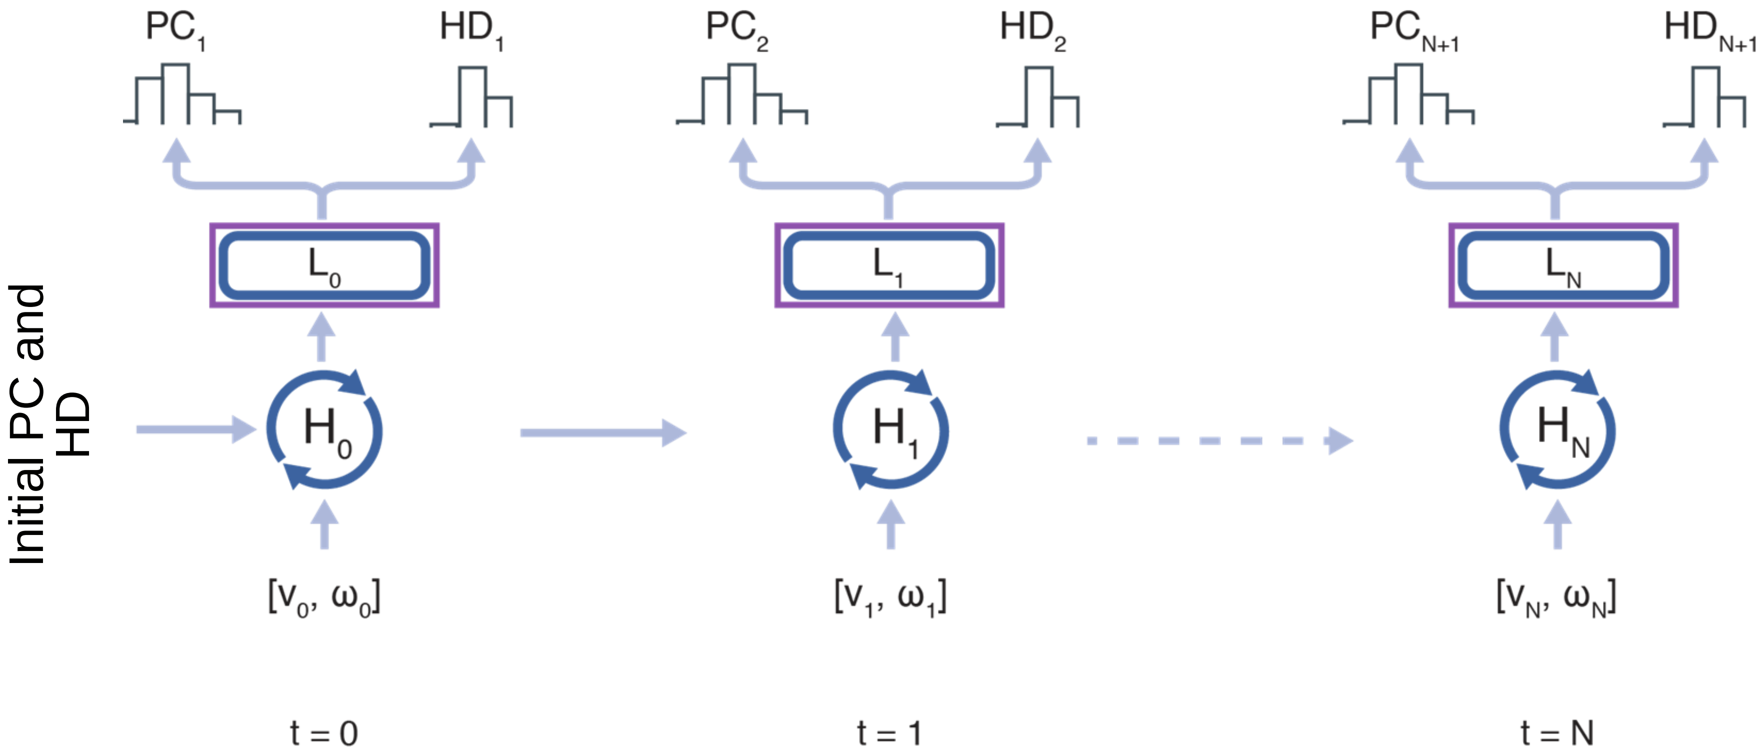

Linear layer activations

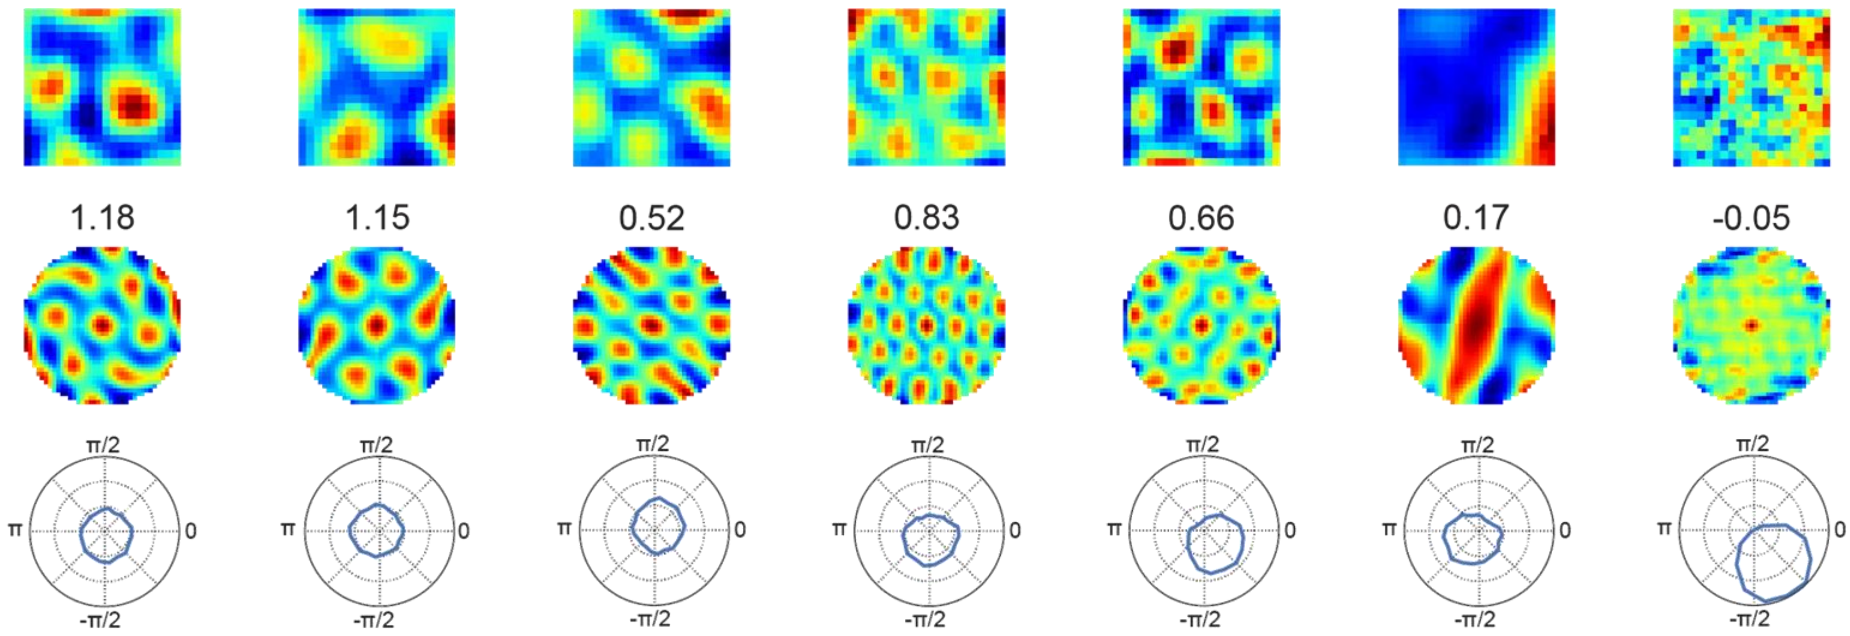

Linear layer: properties

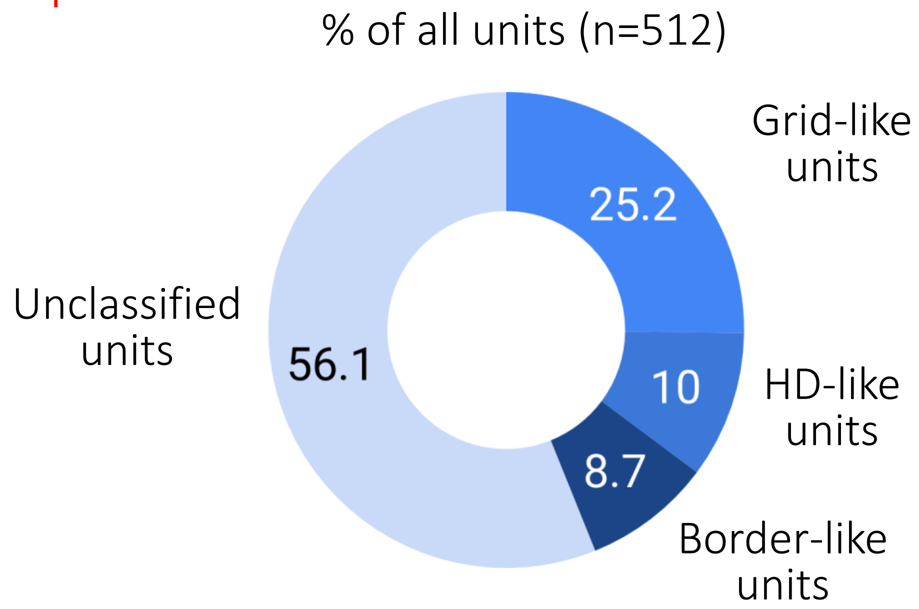

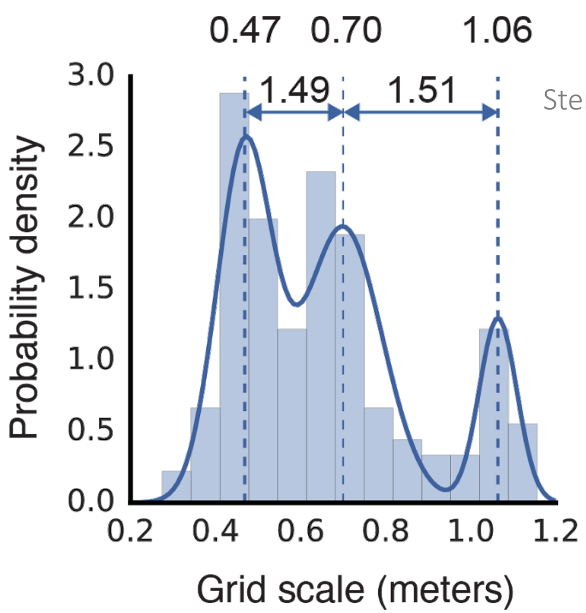

Stemmler et al. (2015)

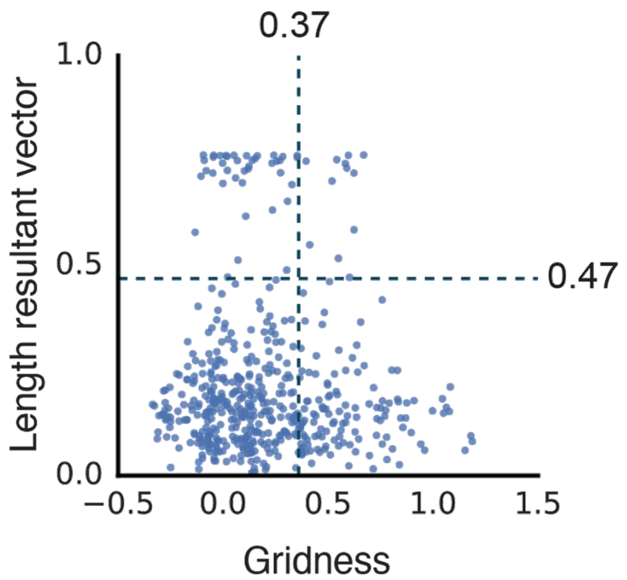

Grid cell agent: architecture

“Morris Water Maze”

Grid cell agent: architecture

“Morris Water Maze”

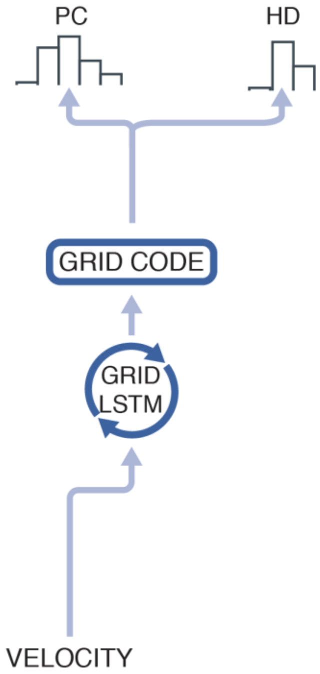

Grid cell agent: architecture

"Morris Water Maze"

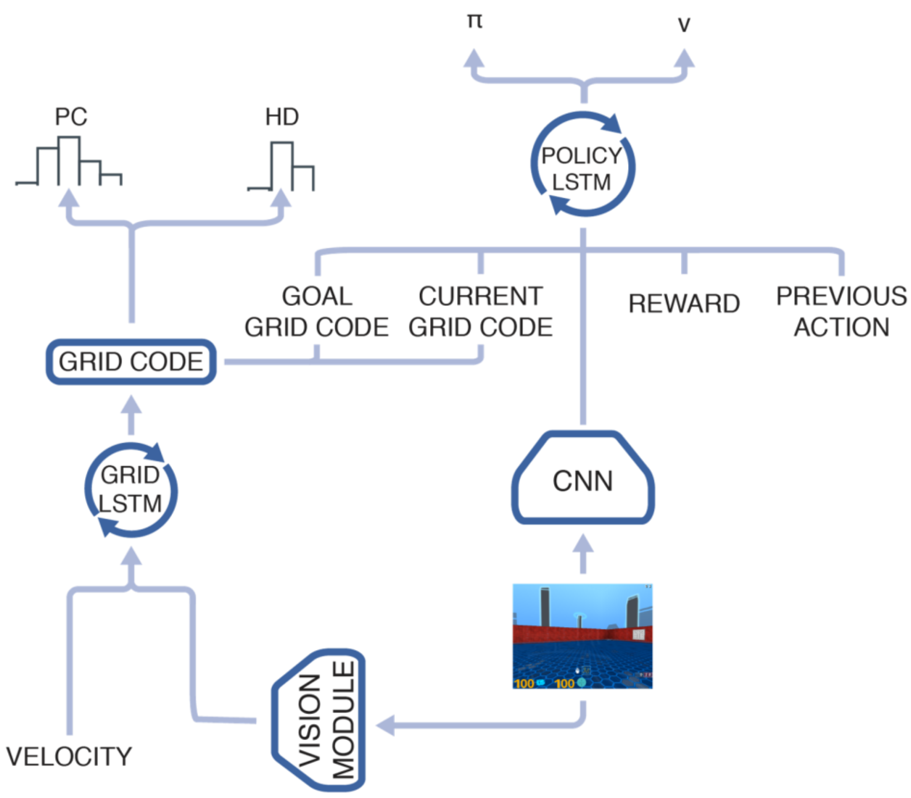

Goal: maximise expected cumulative discounted future reward

$$G_t = \mathbb{E}[\sum_{j=1}^{\infty} \gamma^{j-1} R_{t+j}]$$

Morris Water Maze

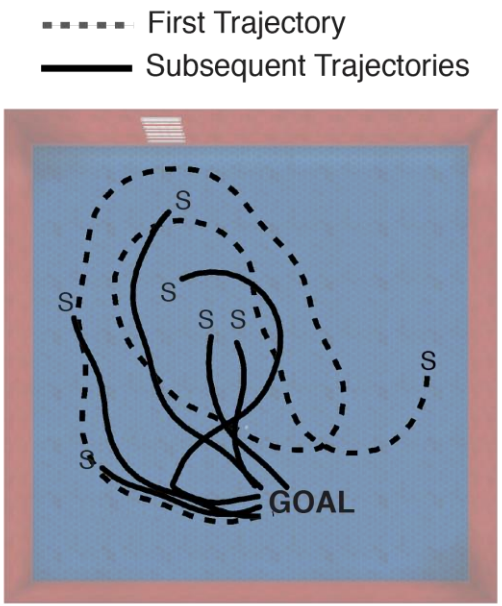

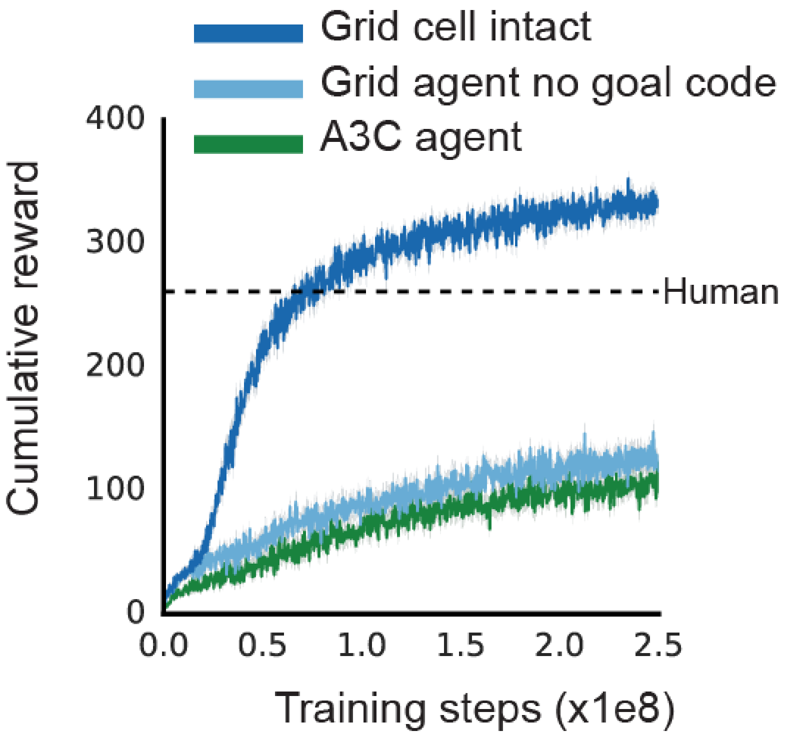

Shortcut linearized sunburst

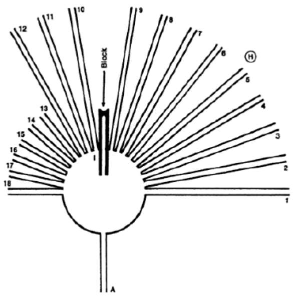

First trajectory

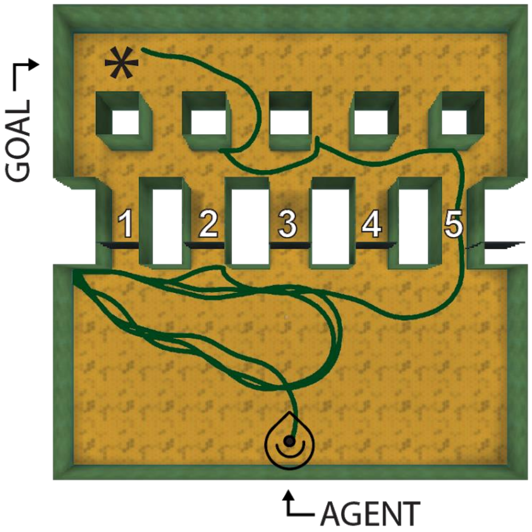

Shortcut linearized sunburst

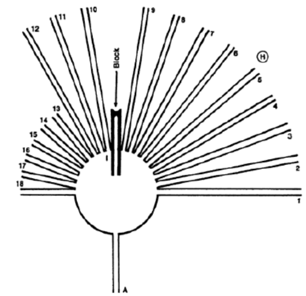

First trajectory

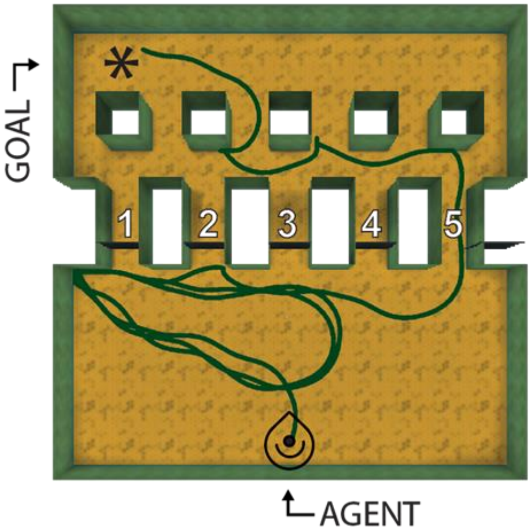

Subsequent trajectories

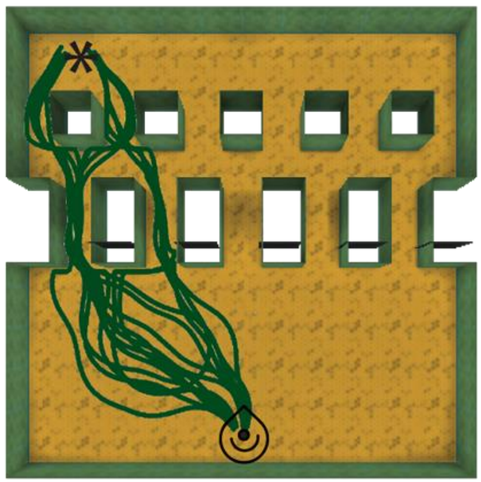

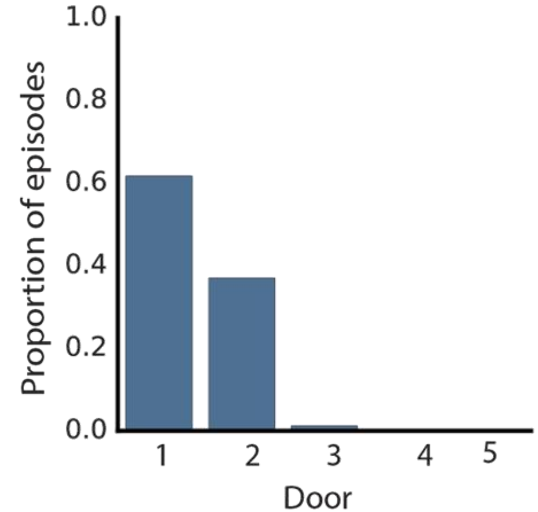

Complex maze: stochastic doors

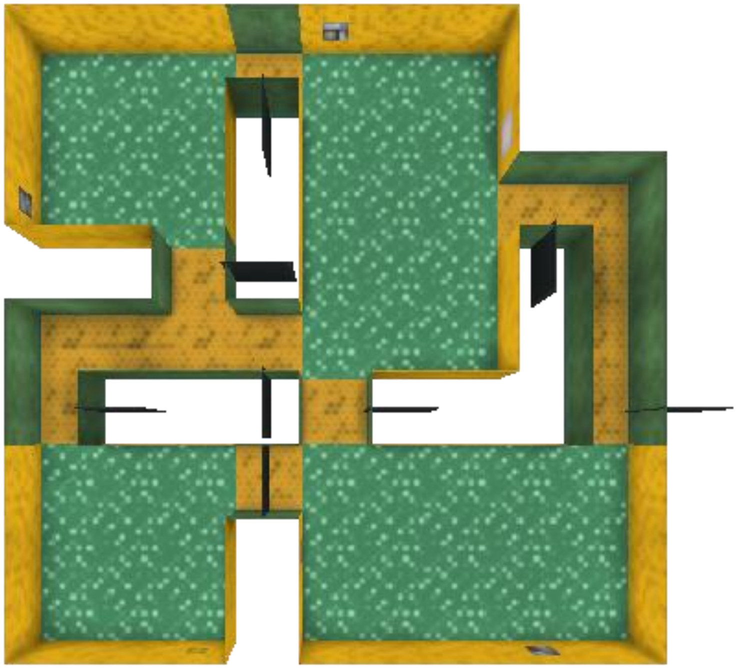

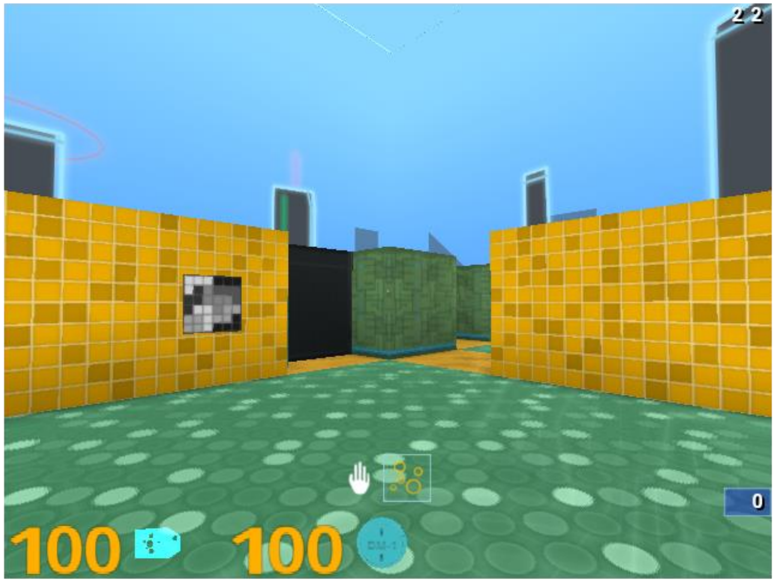

A novel maze configuration (colours, wall position, goal location) is generate for each episode

Complex maze: analysis

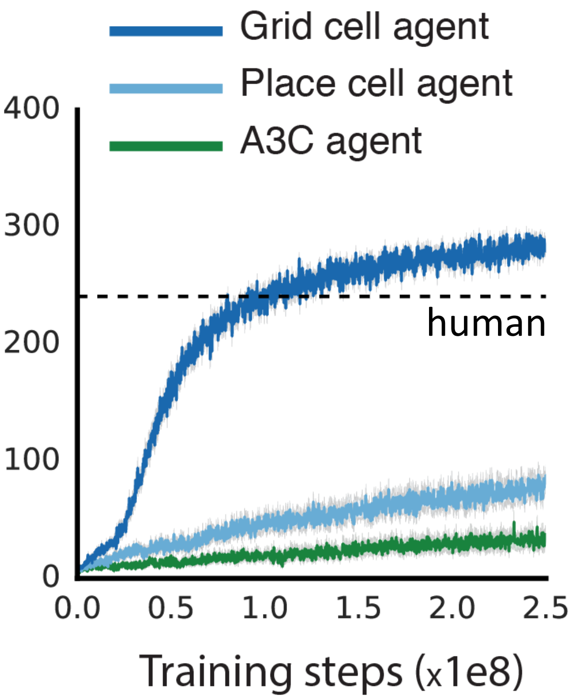

Multivariate decoding

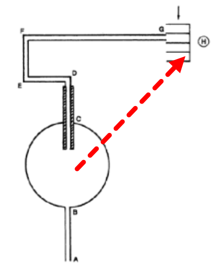

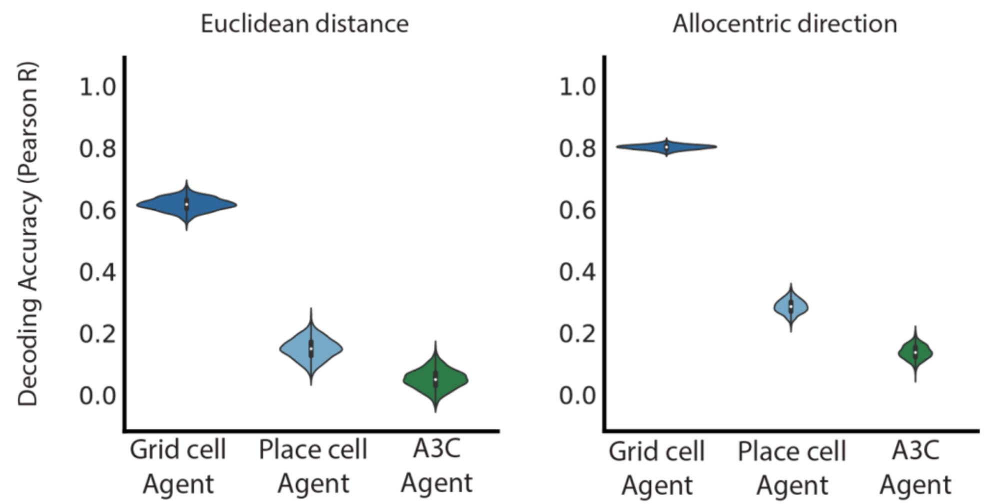

Bottom up (data focused)

Decoding wide-band neural data with DNNs

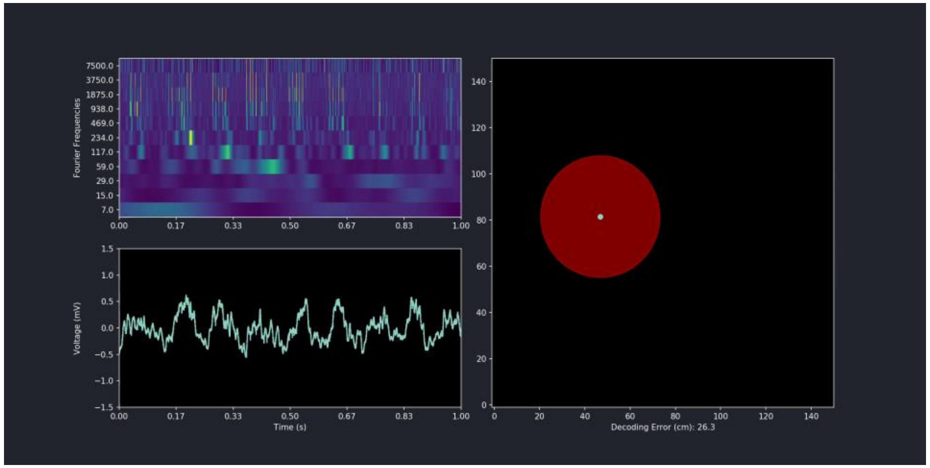

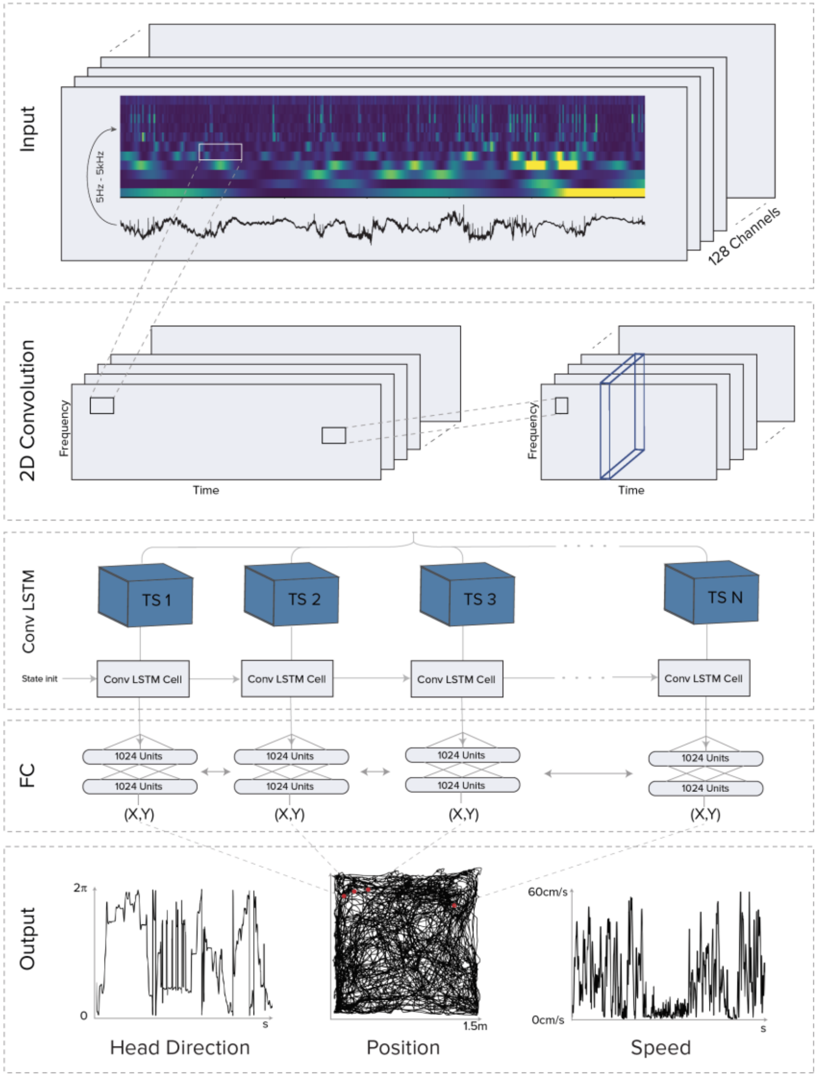

DNN comfortably outperforms standard decoding methods

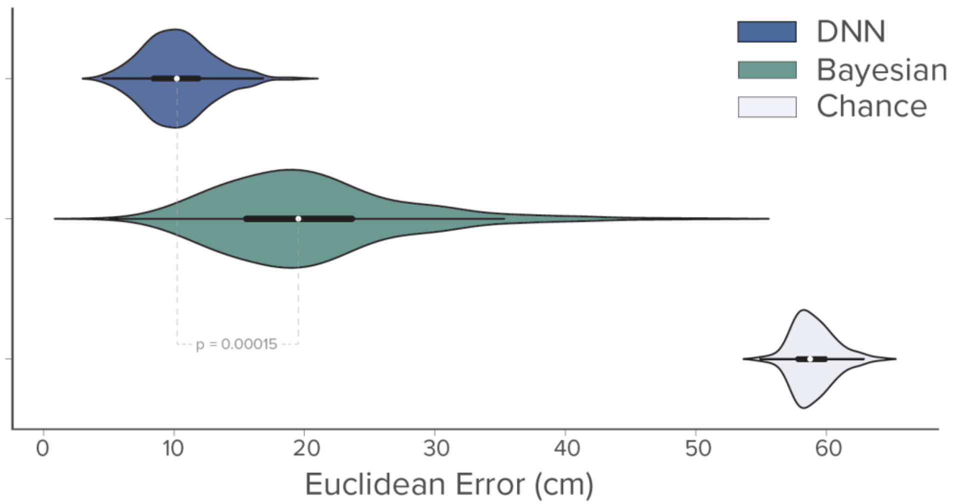

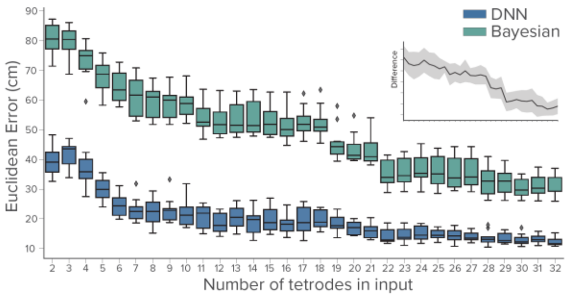

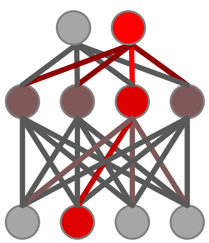

Gradient analysis

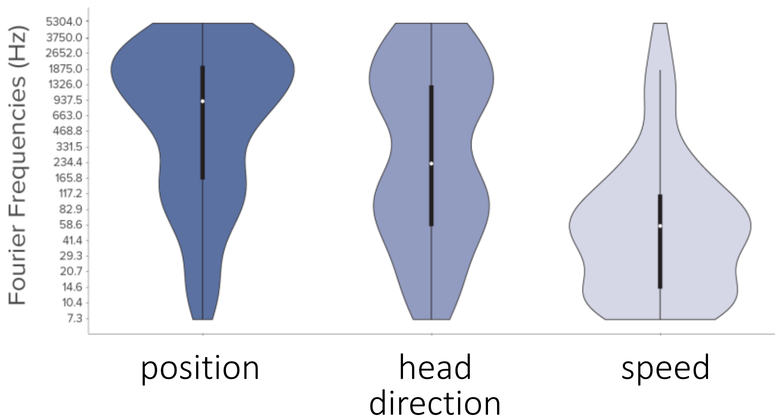

Conclusions

- Grid-like units emerge spontaneously when performing self localization and match many properties of mammalian spatially modulated neurons

- Emergent grid-like representations provides a Euclidean spatial metric and associated vector operations
  - supporting proficient navigation

- DNNs provide a powerful tool for interrogating neural codes
  - understand what is encode, how, & when

Thanks

- Andrea Banino
- Beni Uria
- Dharsh Kumaran

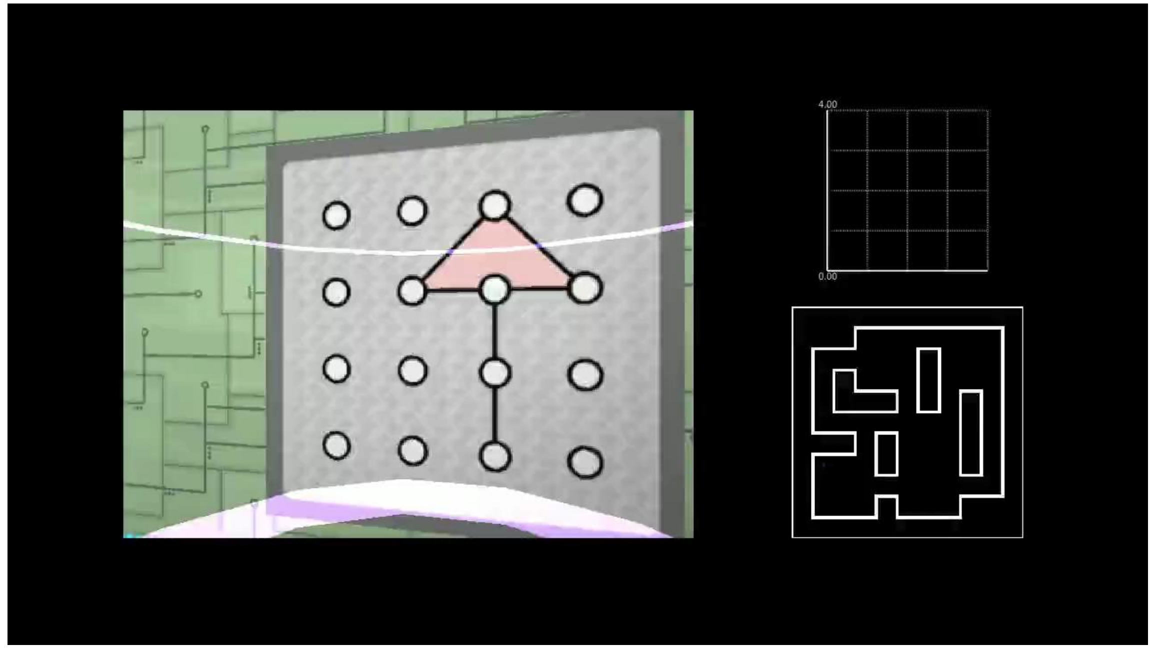

- Sander Tanni
- Markus Frey
- Christian Doeller

wellcome

DeepMind

NVIDIA

UCL CoMPLEX

THE ROYAL SOCIETY

Can we decode location from human data?

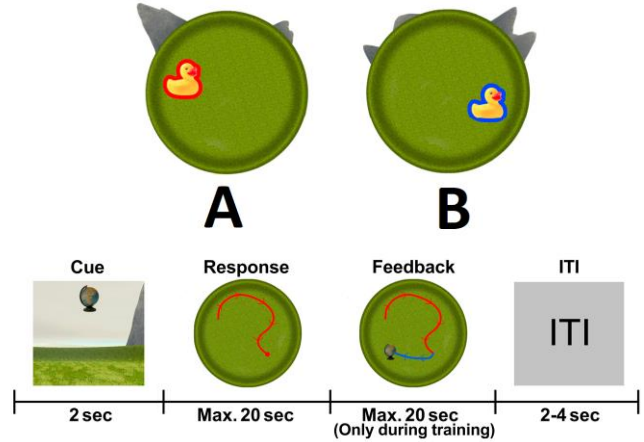

- 18 subjects perform a simple spatial memory task in the scanner
  - learn location of objects in two environments that are disambiguated by back ground
  - ~35 minutes of data per person

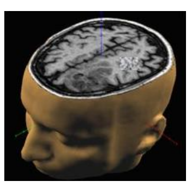

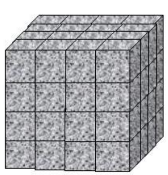

Hippocampal voxels

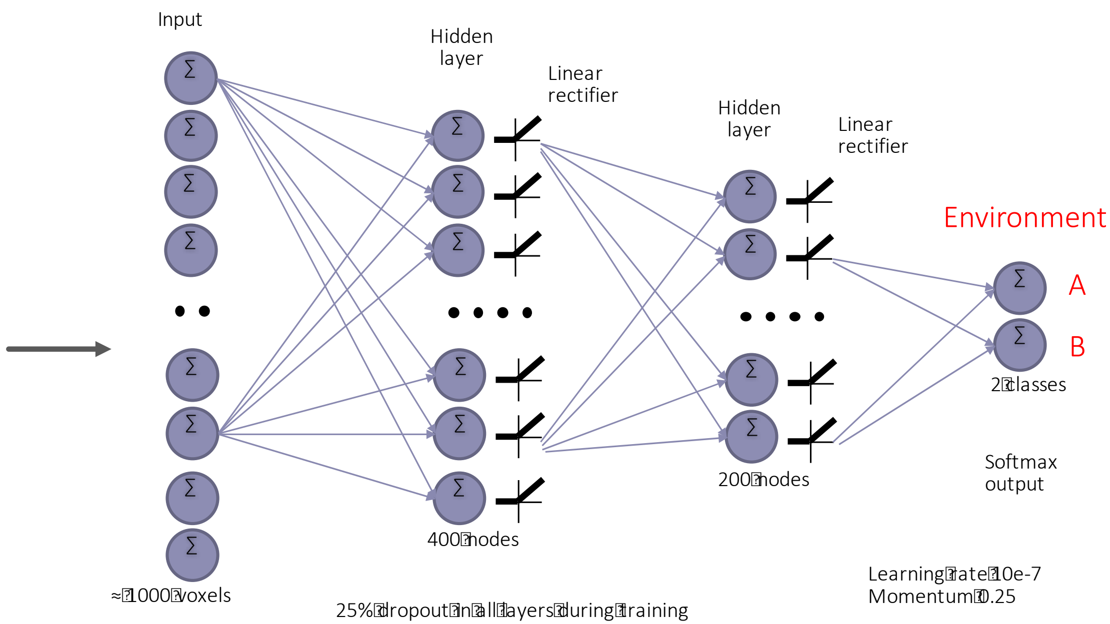

Tanni, S. & Barry, C. (in prep)

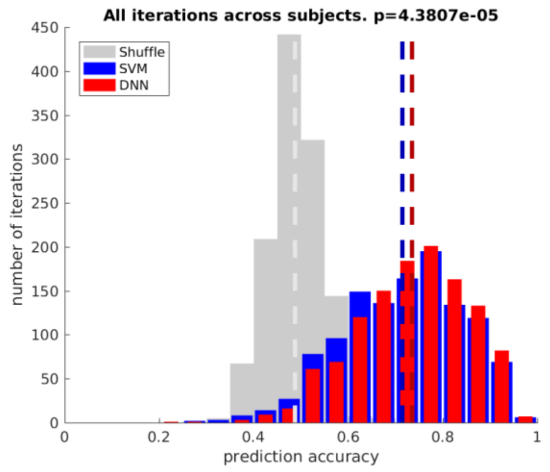

- Based solely on hippocampal voxels:
  - performance exceeds SVM (just)
  - 73% correct vs 71%
- But categorization is hard:
  - hippocampus is a small deep structure
  - subjects might not be accurate themselves – don't know what max score is

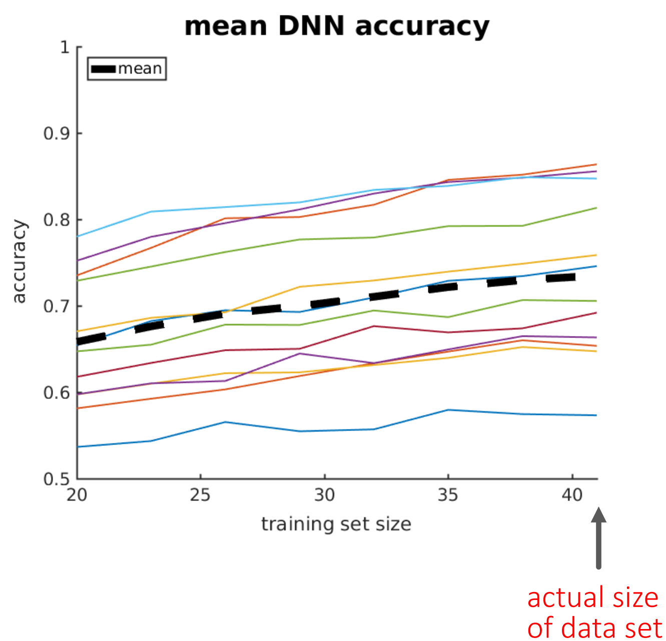

- Do we have enough data?
  - subsampled data indicates not
  - with more data DNN performance increases further above SVM (e.g. 1 hour plus)
- We need to investigate methods for augmenting the existing data
- Next steps:
  - explore which voxels are informative
  - how information is distributed & encoded etc.

Tanni, S. & Barry, C. (in prep)