Hippocampal & striatal roles in spatial navigation

AIMS

- Describe how place cells could be used as a spatial memory for the proximity of a goal by synaptic change at the goal location.
- Describe how routes to a goal could be learned by modifying connections between the hippocampus and nucleus accumbens (Brown & Sharp, 1995), including the relevance of the limitations of perceptrons to linearly-separable functions and the problem of temporal credit-assignment

READING

- Brown M A & Sharp P E (1995) `Simulation of spatial-learning in the morris water maze by a neural-network model of the hippocampal-formation and nucleus-accumbens Hippocampus 5 171-188.
- Burgess N, Donnett J G, Jeffery K J & O'Keefe J (1997) `Robotic and neuronal simulation of the hippocampus and rat navigation', Philosophical Transactions of the Royal Society, B 352 1535-1543.
- Bush,D., Barry,C., Manson,D. & Burgess,N. (2015) Using grid cells for navigation. Neuron. 87(3) 507-520

‘Place cells’ encode the rat’s current location

hippocampus

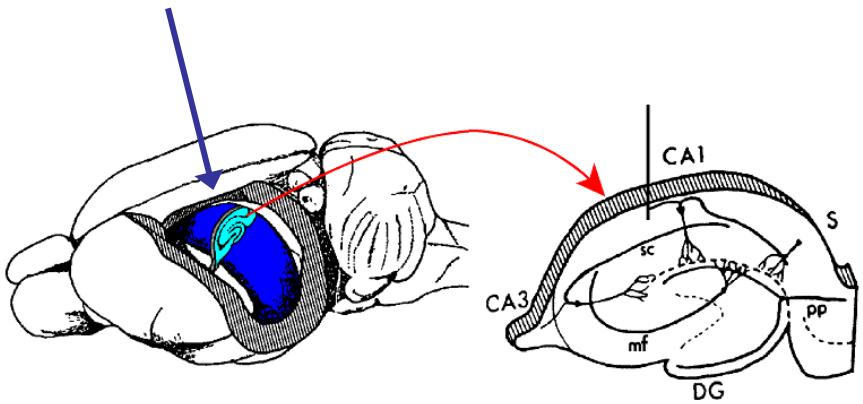

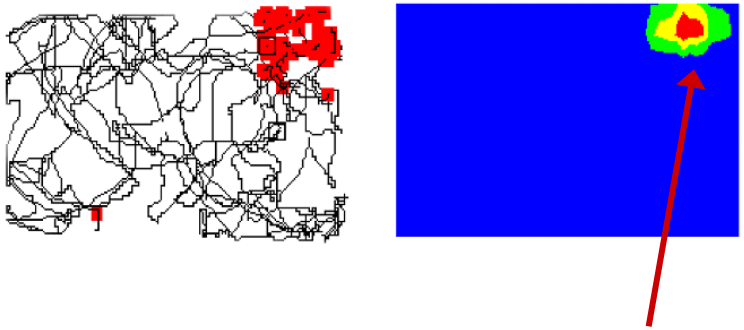

‘place field’

35 SIMULTANEOUSLY RECORDED PLACE CELLS

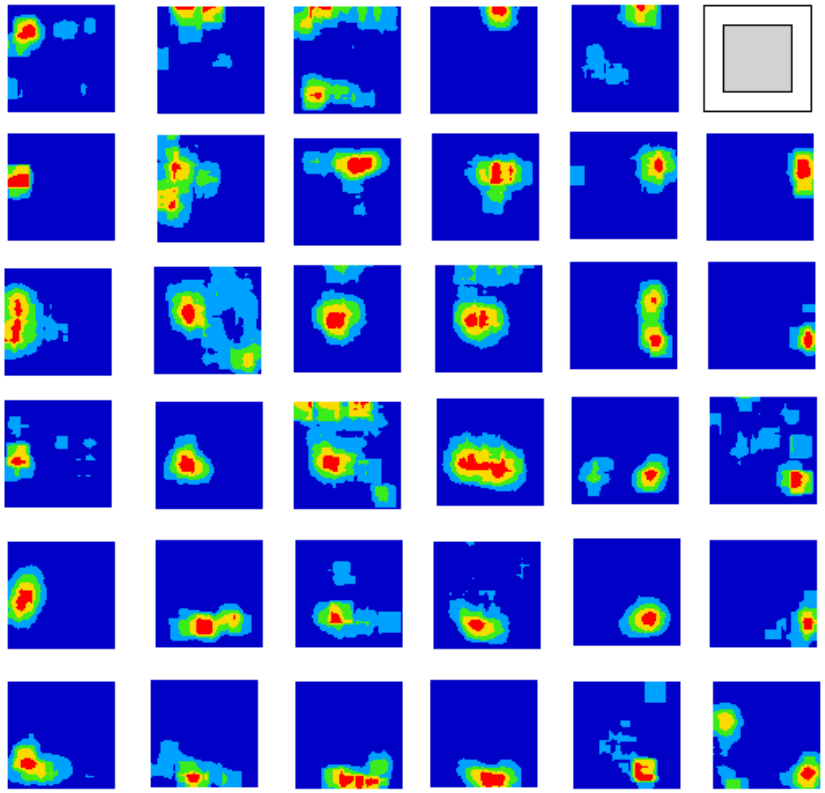

Hippocampus and spatial memory: the water maze

(Morris, 1981)

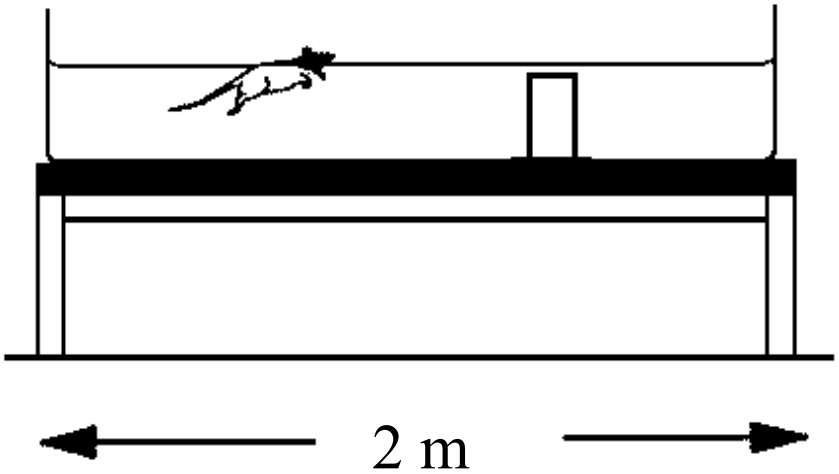

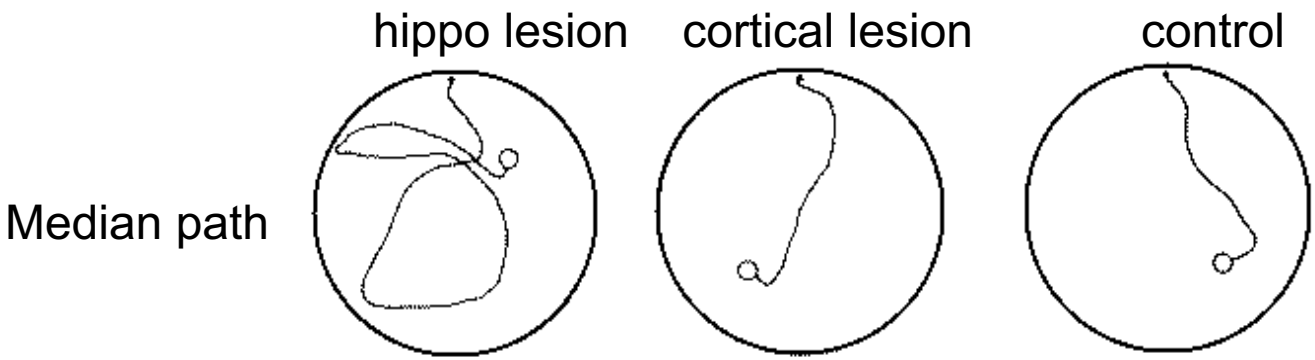

(Morris et al., 1982).

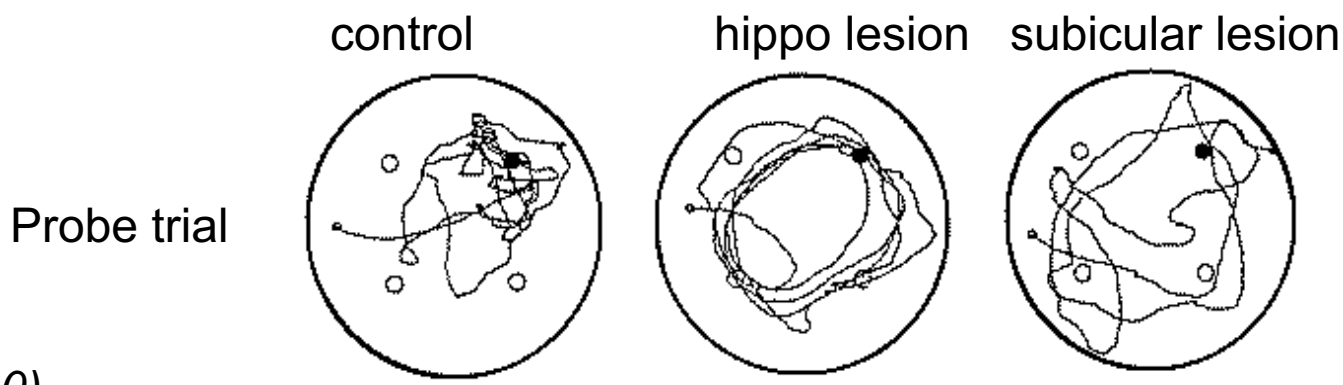

(Morris et al., 1990).

Key properties of hippocampal navigation
(to be explained by a model):

1) Long range
2) Direct, does not require vicarious trial and error
3) Allows short cuts
4) Rapid to compute

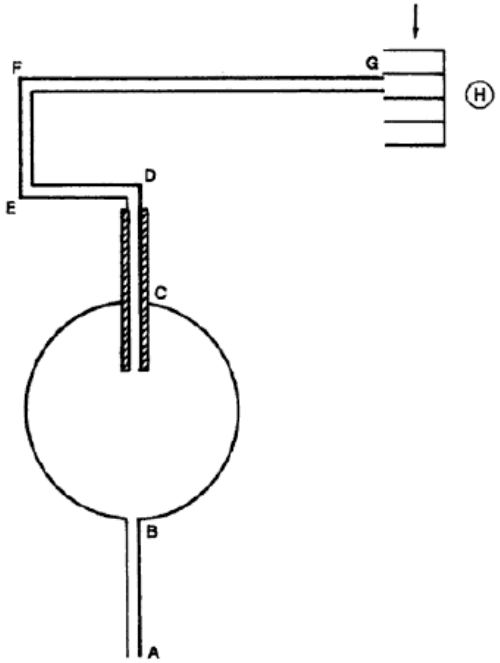

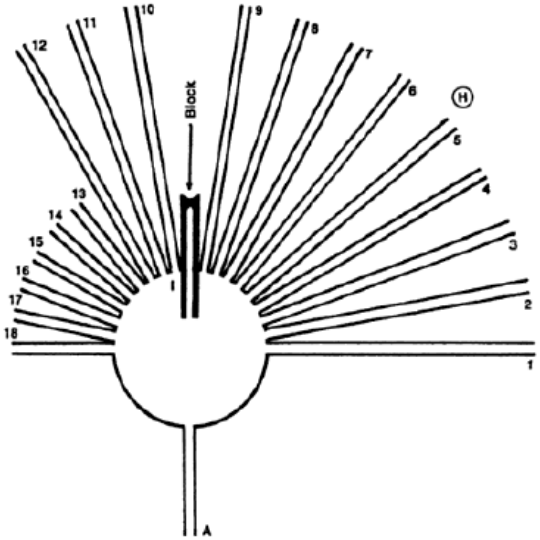

If place cell firing says where you are, how do you know where to go?

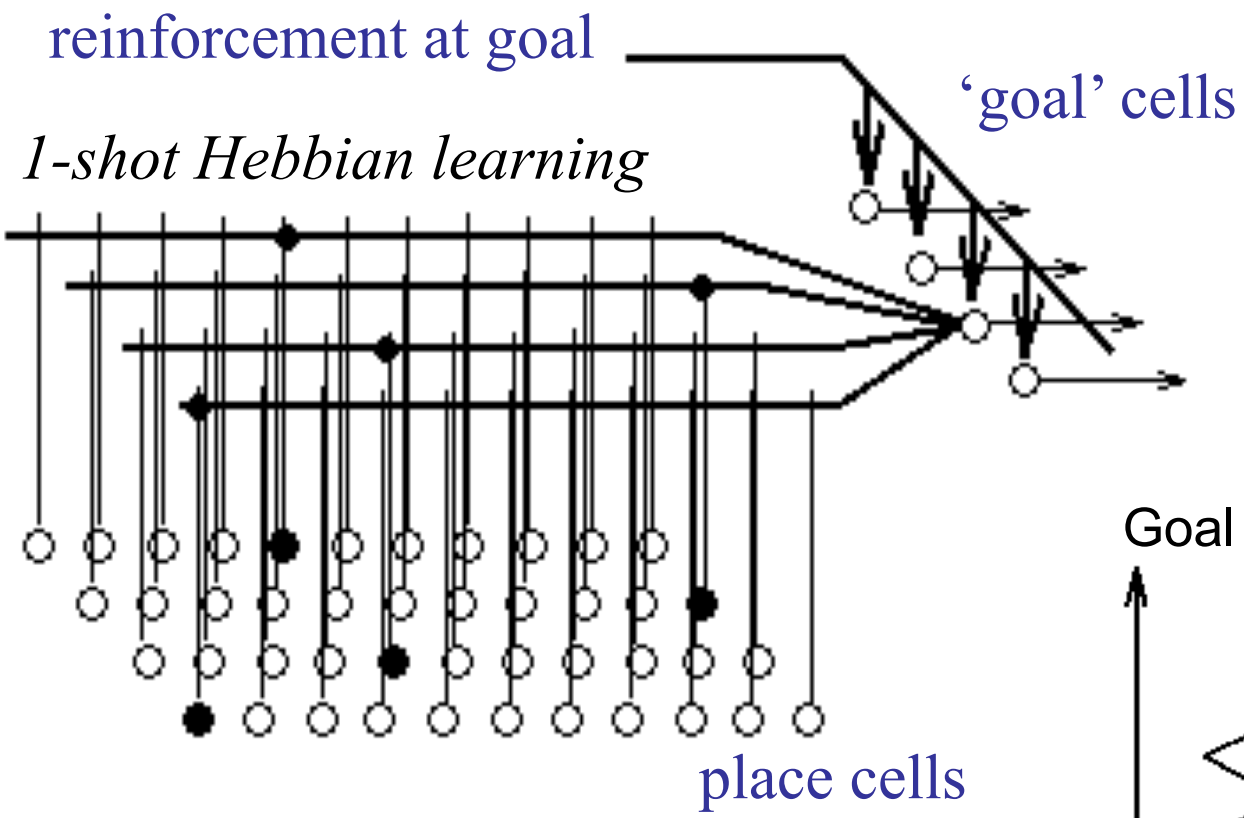

Simple model

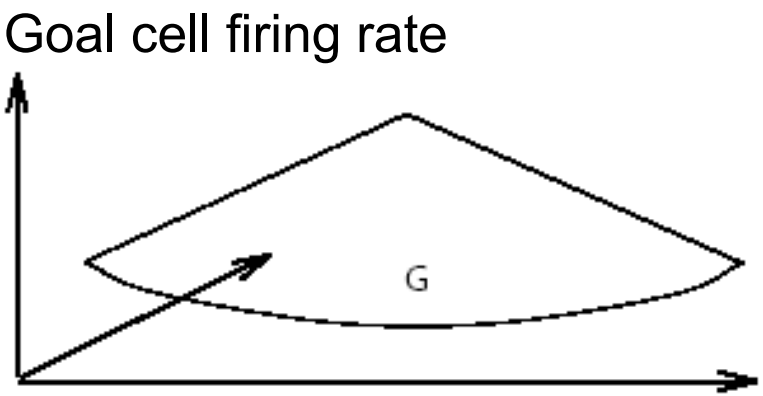

navigation = gradient ascent of firing rate of goal cell

Burgess et al., 1997

Simple model. Evaluation.

- Nice and simple, but too simple – requires ‘hunting’ (aka ‘vicarious trial and error).
- Only works if distance to goal < place cell diameter
- Can be consistent with goal-independent ‘latent learning’ (i.e. rats with prior experience of an environment perform better when the goal is introduced) as improvement of the place cell representation via unsupervised learning.
- A model of navigation should output left/right/ahead motor commands.

A neural circuit for spatial cognition in & around the hippocampus

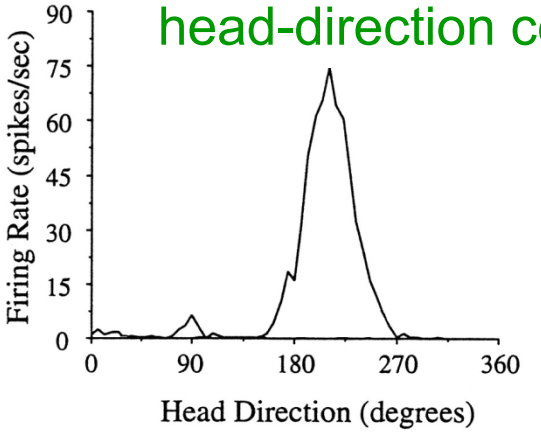

Ranck, 1984; Taube et al 1997

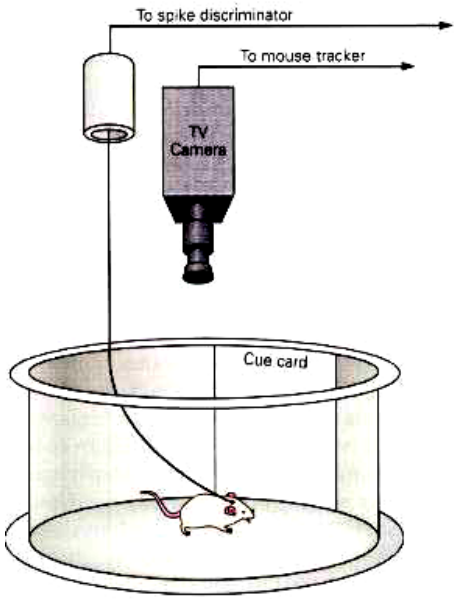

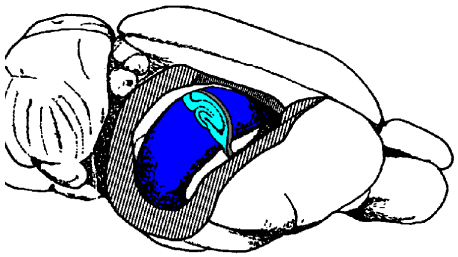

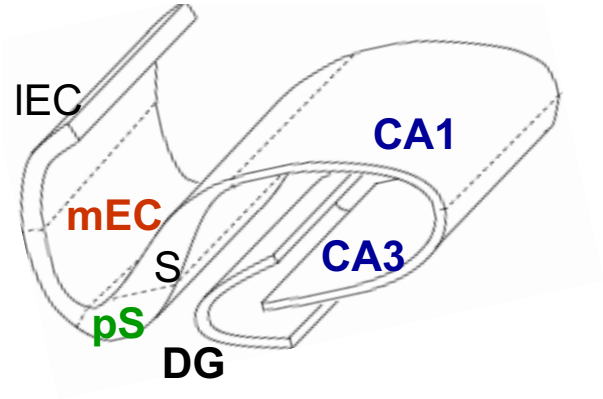

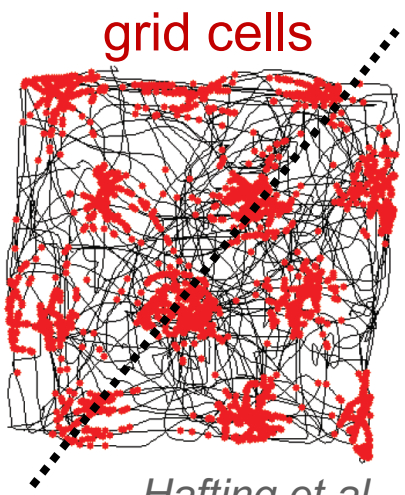

Hafting et al., 2005

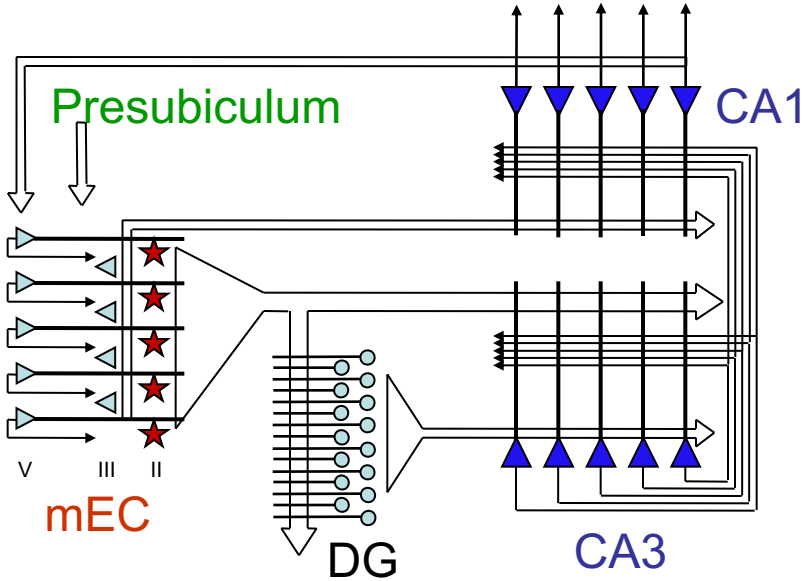

O'Keefe & Dostrovsky 1971

Hippocampus and striatum support ‘place’ & ‘response’ learning respectively

Initial learning of ‘place’, subsequent learning of body turn.

‘Place’ learning is dependent on hippocampus (e.g. Morris et al., 1982; O’Keefe & Nadel, 1978)

‘Response’ learning is dependent on striatum (e.g. Potegal, 1972; Kesner, 1993).

The “Striatum” or “Basal ganglia”: caudate, putatmen, nucleus accumbens, globus pallidus.

Implicated in Parkinson’s disease.

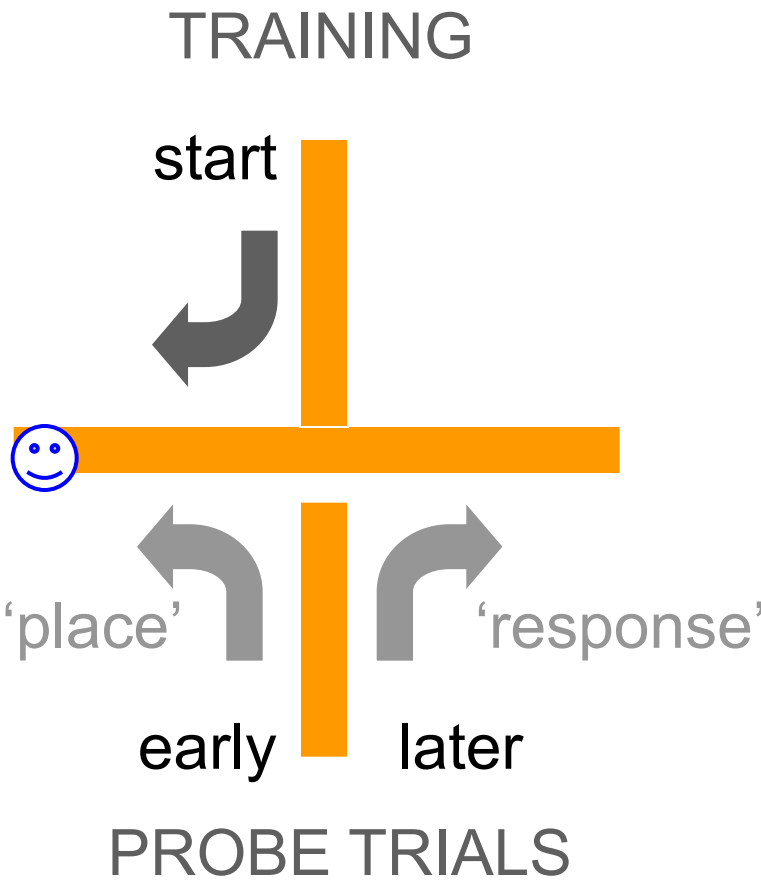

Packard & McGaugh, 1996

Brown & Sharp (1995) naïve model: learning to make correct moves from place and direction info.

First attempt:

Nucleus accumbens (basal ganglia/ motor output).

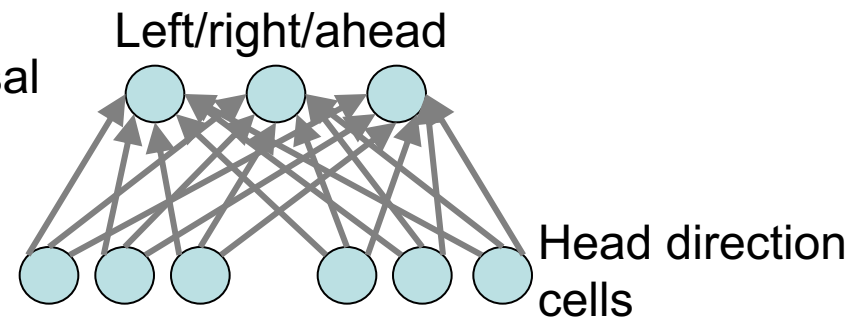

Insufficiently powerful for the same reason the perceptron can't do XOR: navigation is 'non linearly-separable'.

Imagine a simplified network:

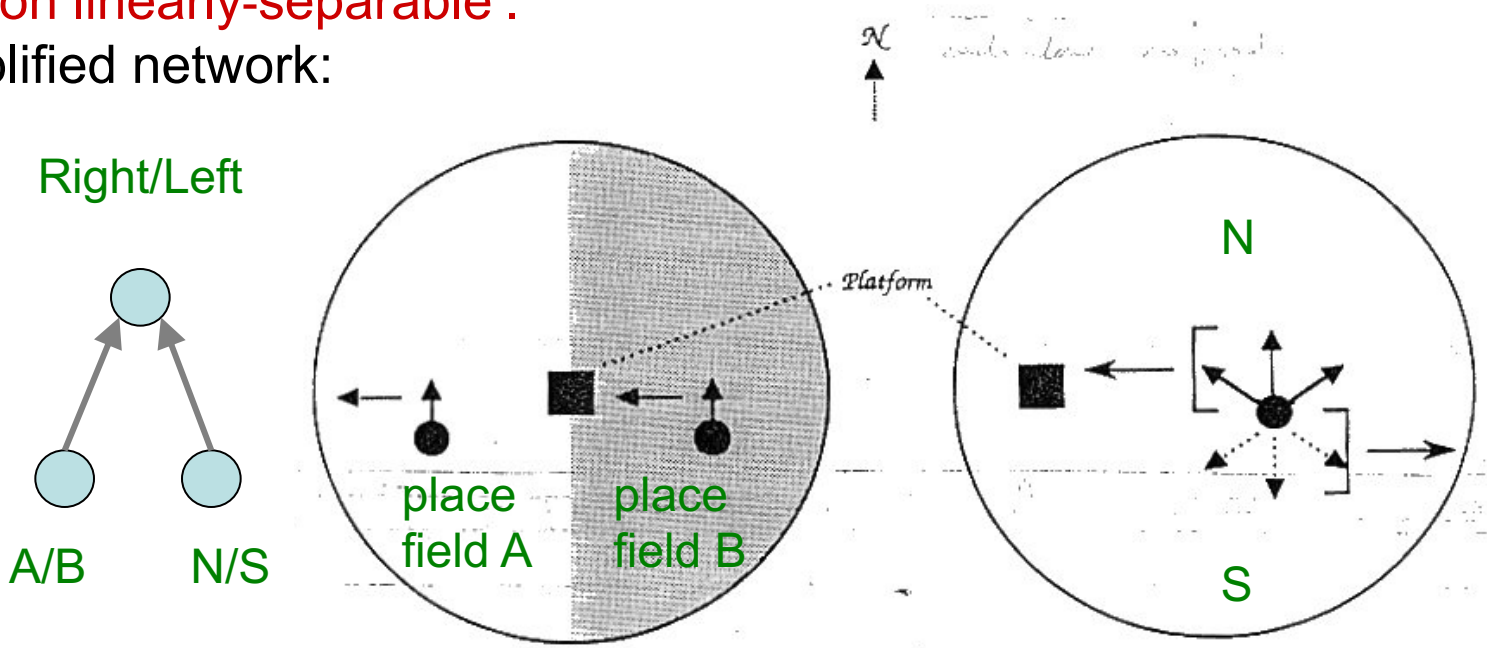

Brown & Sharp (1995) naïve model: learning to make correct moves from place and direction info.

First attempt:

Nucleus accumbens (basal ganglia/ motor output).

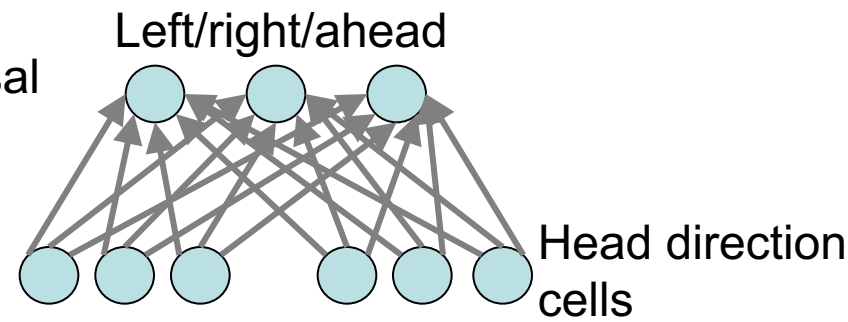

Insufficiently powerful for the same reason the perceptron can't do XOR: navigation is 'non linearly-separable'.

Imagine a simplified network:

		A	B
		0	1
N	0	**0**	**1**
S	1	**1**	**0**

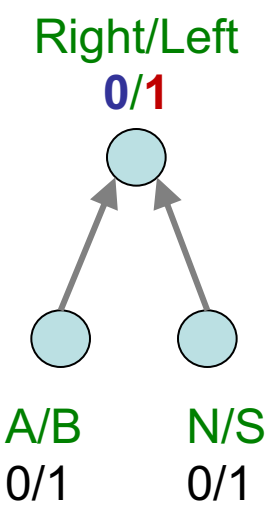

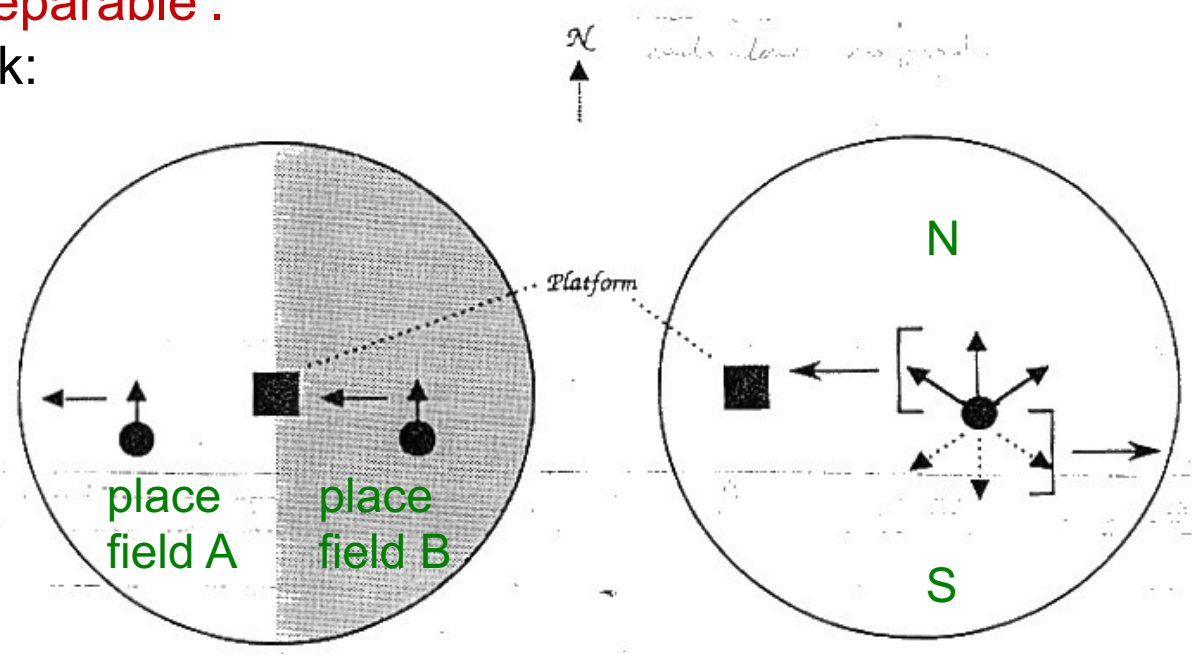

Brown & Sharp (1995) model

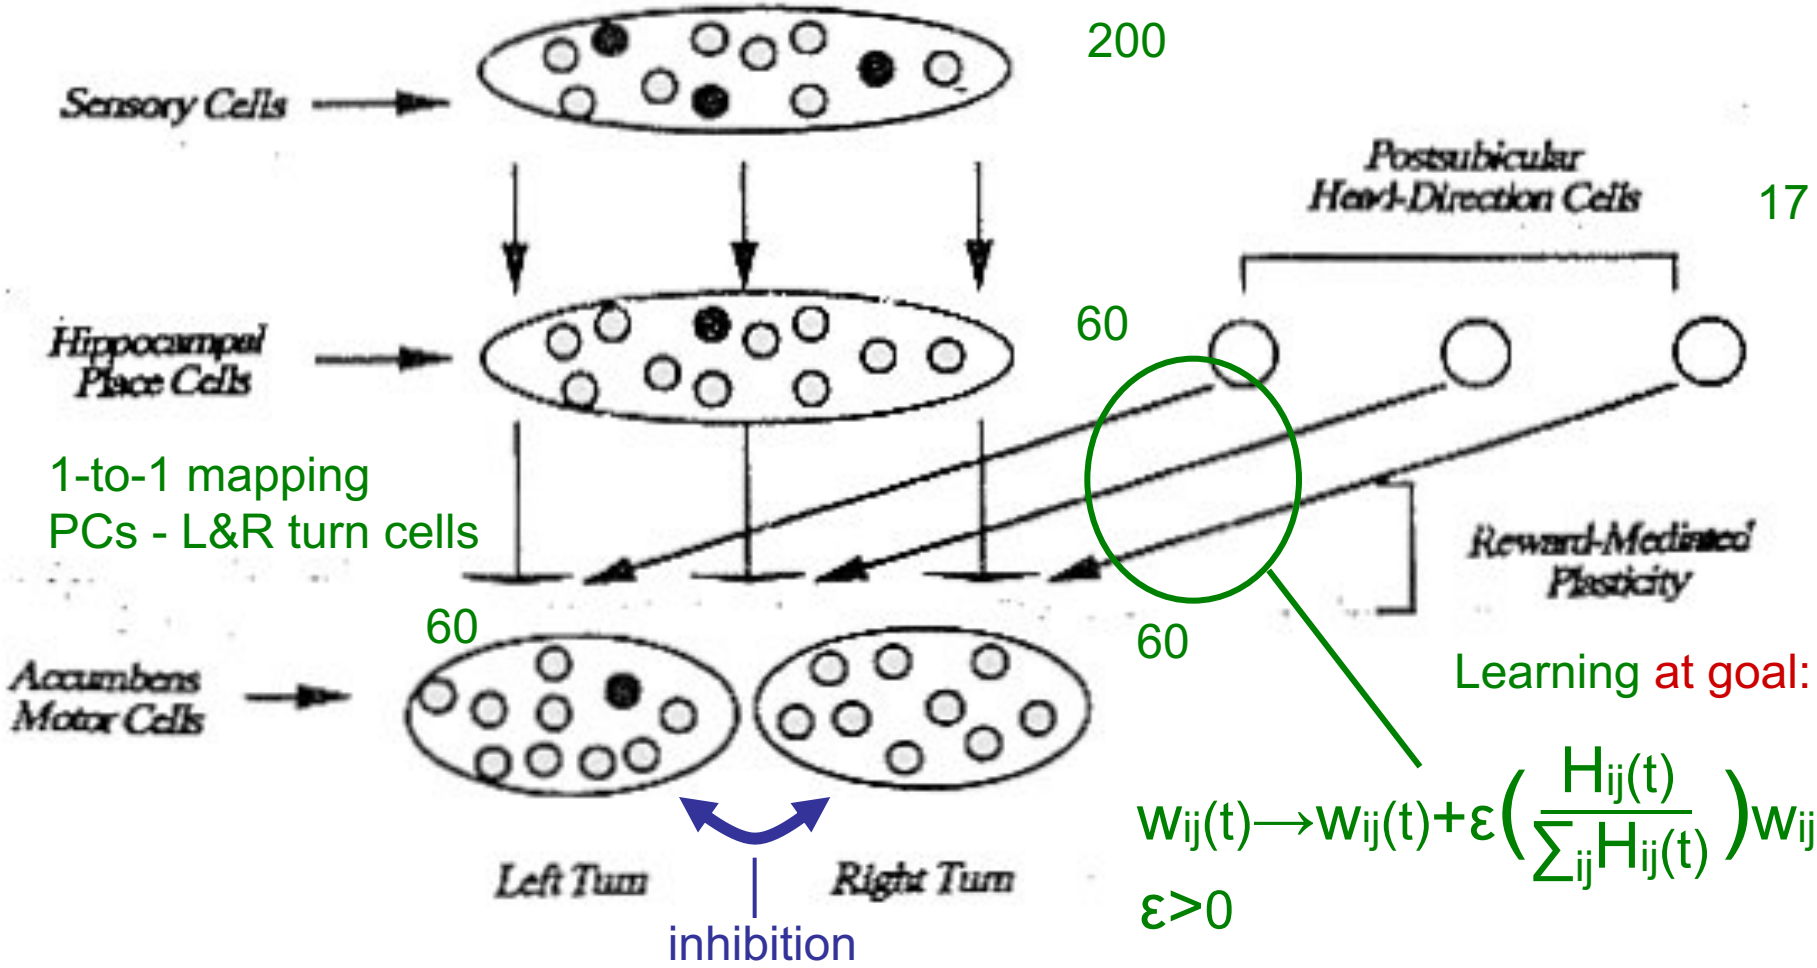

Learning at goal:

$$w_{ij}(t) \rightarrow w_{ij}(t) + \varepsilon \left(\frac{H_{ij}(t)}{\sum_{ij} H_{ij}(t)} \right) w_{ij}(t)$$

$\varepsilon > 0$

Recency-weighted cumulative 'Hebbian' activity: $H_{ij}(t+1) = r\, H_{ij}(t) + B_{ij}(t)$,

$B_{ij}(t) = 1$ if $x_j(t)$ and $x_i(t)$ are both active (0 otherwise),

$0<r<1$ gives the rate of decay, changes normalised by $\sum_{ij} H_{ij}(t)$ term.

Brown & Sharp (1995) model. Performance.

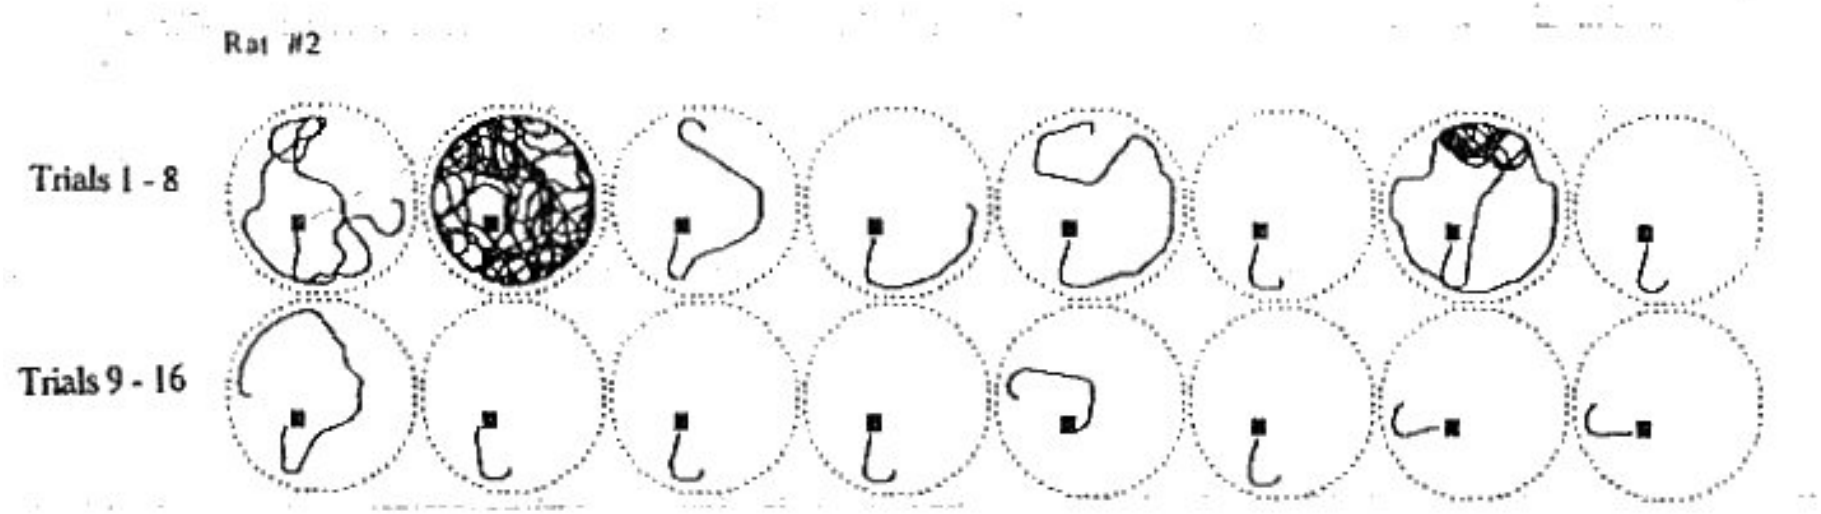

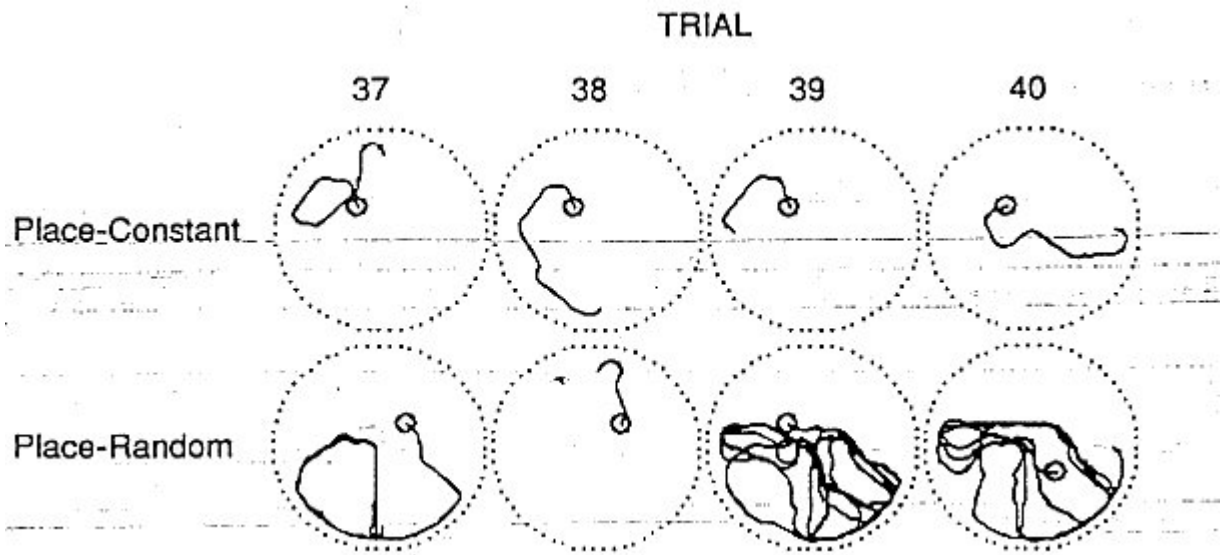

Brown & Sharp (1995) model. Main points.

1. Each place cell allows 1 left and 1 right turn cell to be active – separate representation of L & R for each place makes the task easy (linearly separable). One group (L or R) inhibits the other.
2. Learning occurs in head-direction to turn cells (i.e. whether to go L or R at a given place, according to your direction)
3. Learning, when at goal, according to recency-weighted cumulative amount of Hebbian activity (H_{ij}) at the synapse:

$$w_{ij}(t+1) \longrightarrow w_{ij}(t) + \varepsilon \left(\frac{H_{ij}(t)}{\sum_{ij} H_{ij}(t)} \right) w_{ij}(t)$$

$$H_{ij}(t+1) = r\, H_{ij}(t) + B_{ij}(t),$$

$$B_{ij}(t) = 1 \text{ if } x_j(t) \,\&\, x_i(t) \text{ both active}$$
$$0 \text{ otherwise,}$$

i.e. recent direction/place/turn combinations are reinforced if they have lead to the goal.

4. Good performance requires many runs to the goal via different routes: noise in turn cell activations help sample different routes.

Brown & Sharp (1995) model. Evaluation.

- Consistent with ‘response’ (body-turn responses) in basal ganglia versus ‘place’ representations in hippocampus.

- The model performs complex ‘stimulus-response’ learning – generalises to new situations (place/direction combinations) only insofaras the same place and head-direction cells fire for similar places and directions. May not calculate totally novel shortcuts.

- Solves the *temporal credit assignment* problem by using the recency-weighted cumulative Hebbian term. (Reinforcement learning is a more principled solution.)

- Solves the non linearly-separable nature of navigation by providing an expanded representation (place x turn).

The problem with place cells

Place cells do not allow for the direct calculation of heading vectors between points that are separated by distances exceeding the size of the largest place field. At distances greater than this there is no overlap between the place cell vector at the current and goal location.

Grid cells

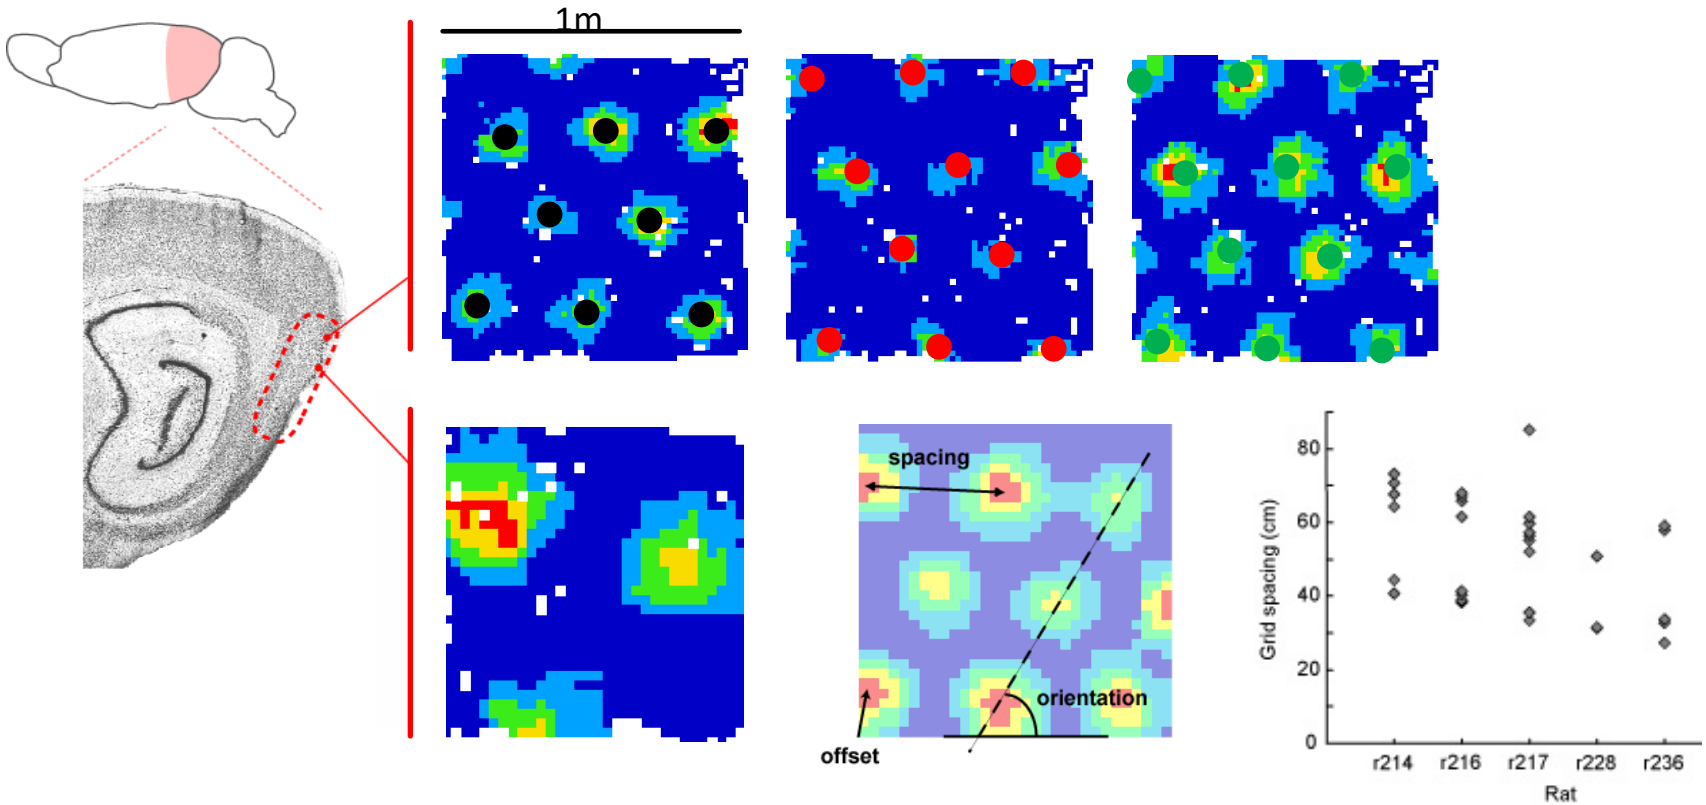

- Multiple regularly spaced fields tile the environment
- Co-recorded cells typically have same scale & orientation
- Independent modules distributed along the entorhinal cortex

Hafting et al. (2005); Barry et al. (2007); Yoon et al (2013)

Using grids to recover the vector between two points

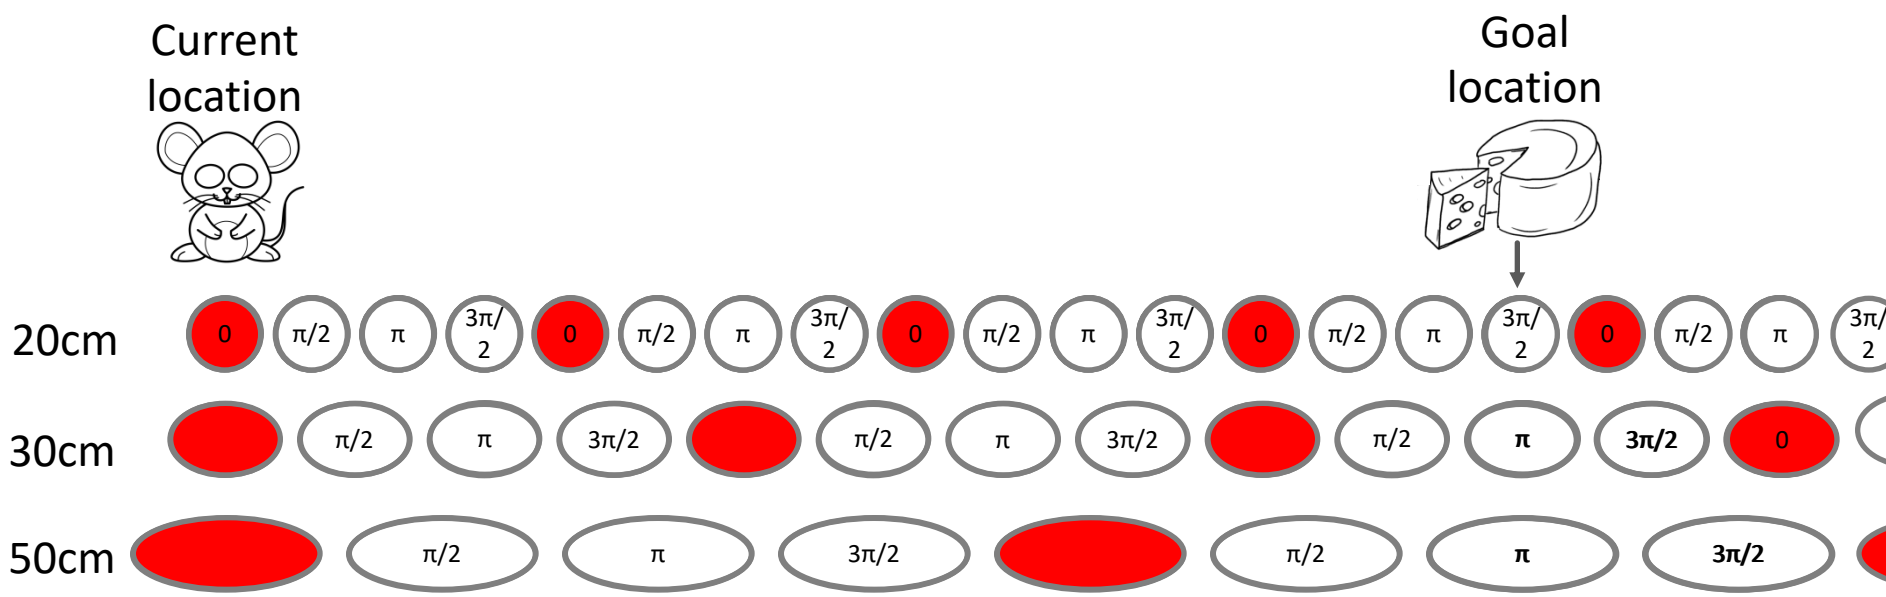

Goal distance = 75cm [3π/2, π, π]

- The vector between two points in grid space is unique within the capacity of the network & evolves continuously
- This process can be extended to 2D (or more)
- Accuracy depends on regular grid patterns
- Several neural implementations possible: sequential search, direct look-up

Bush, Barry, Manson & Burgess (2015) Neuron

Supervised learning architecture

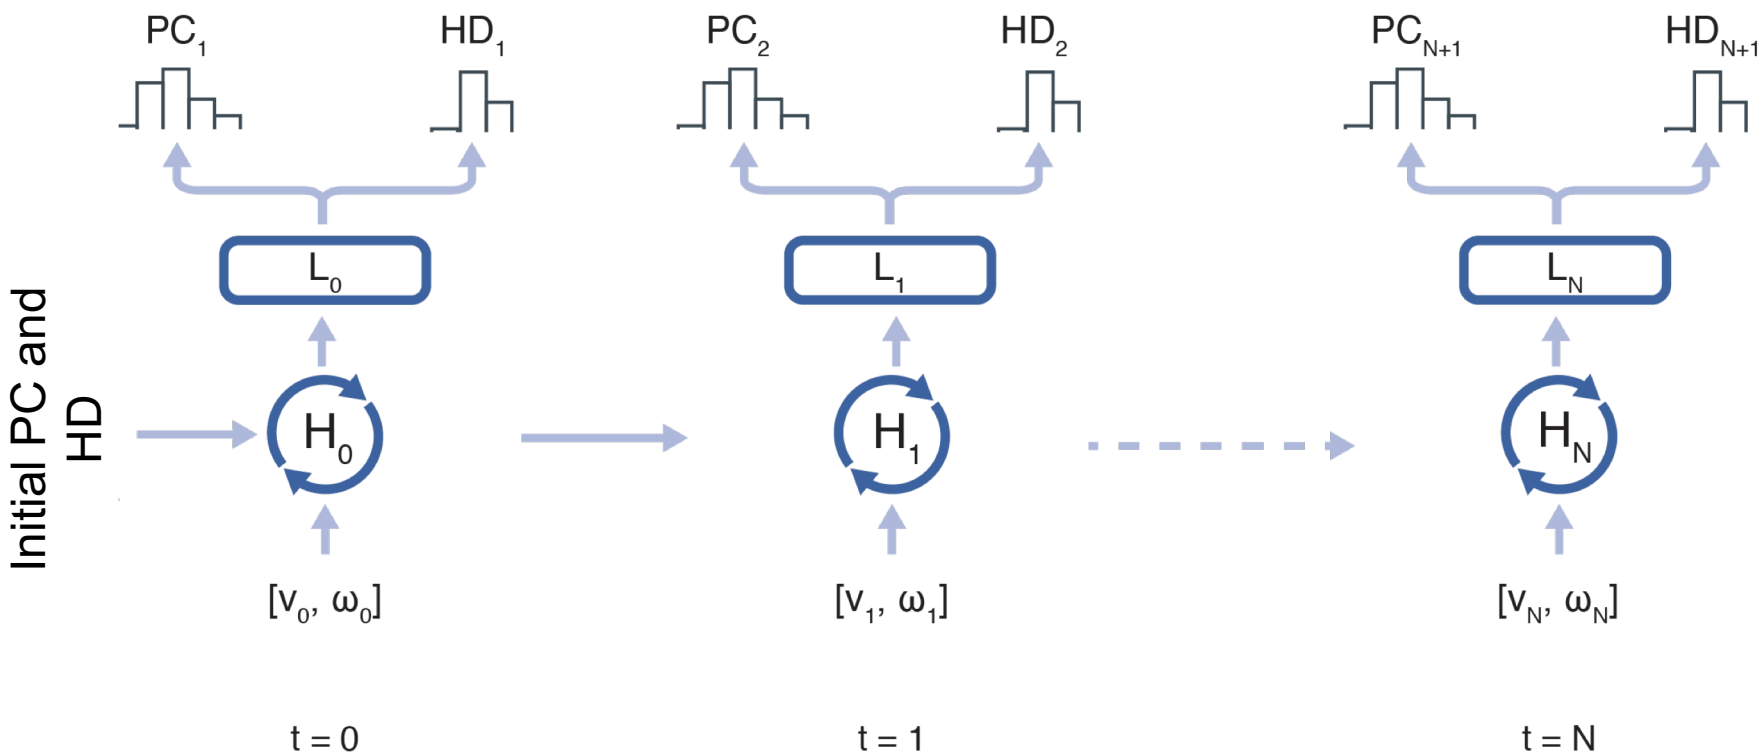

Banino et al (2018) Vector-based navigation using grid-like representations in artificial agents. Nature

Path integration task

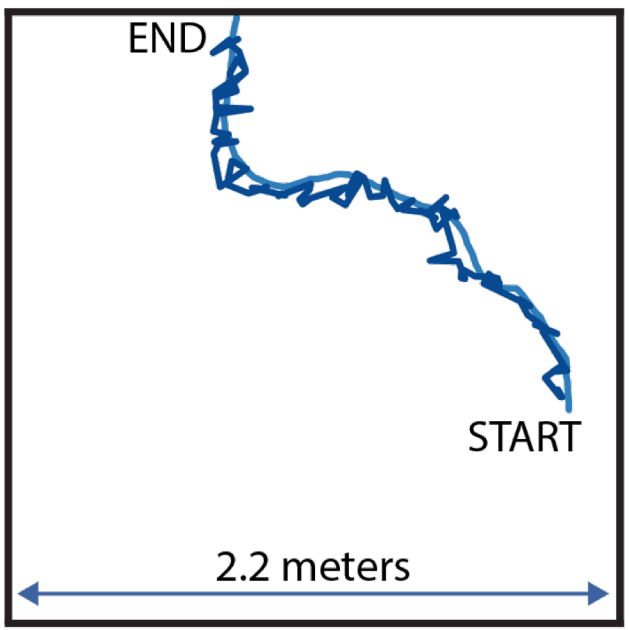

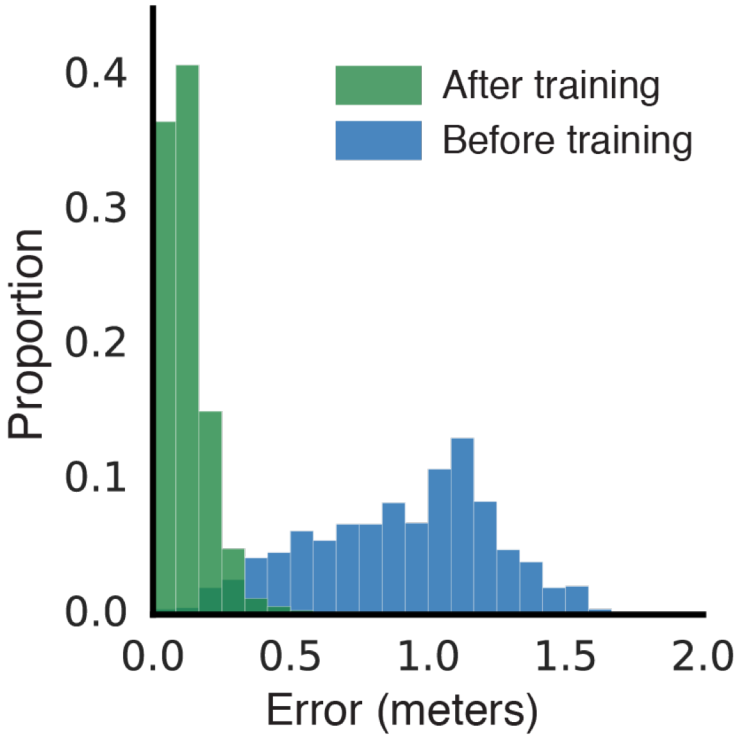

Analysis of linear layer

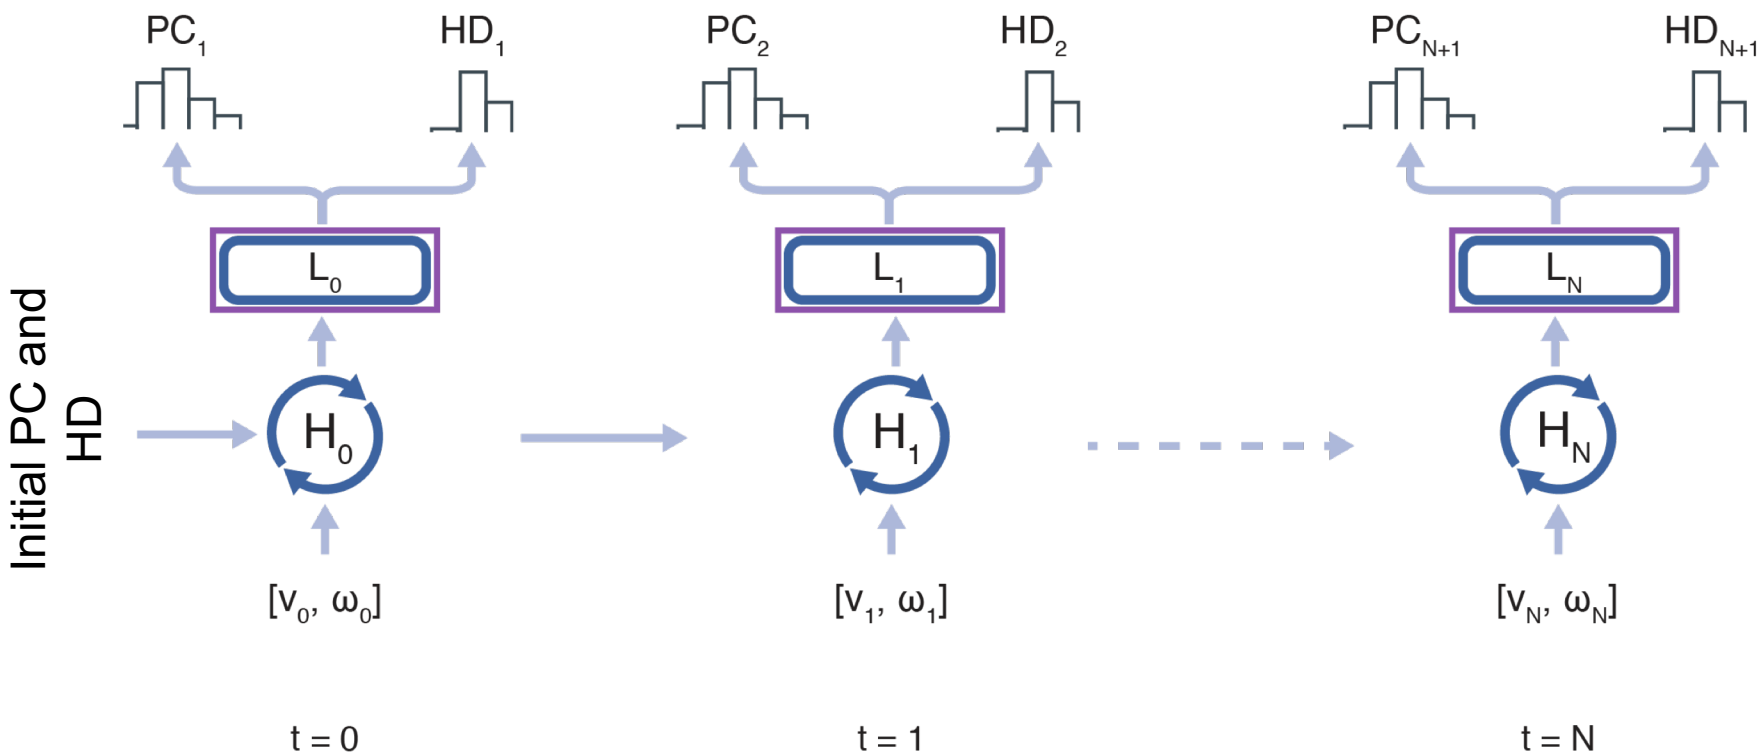

Linear layer activations – grid-like

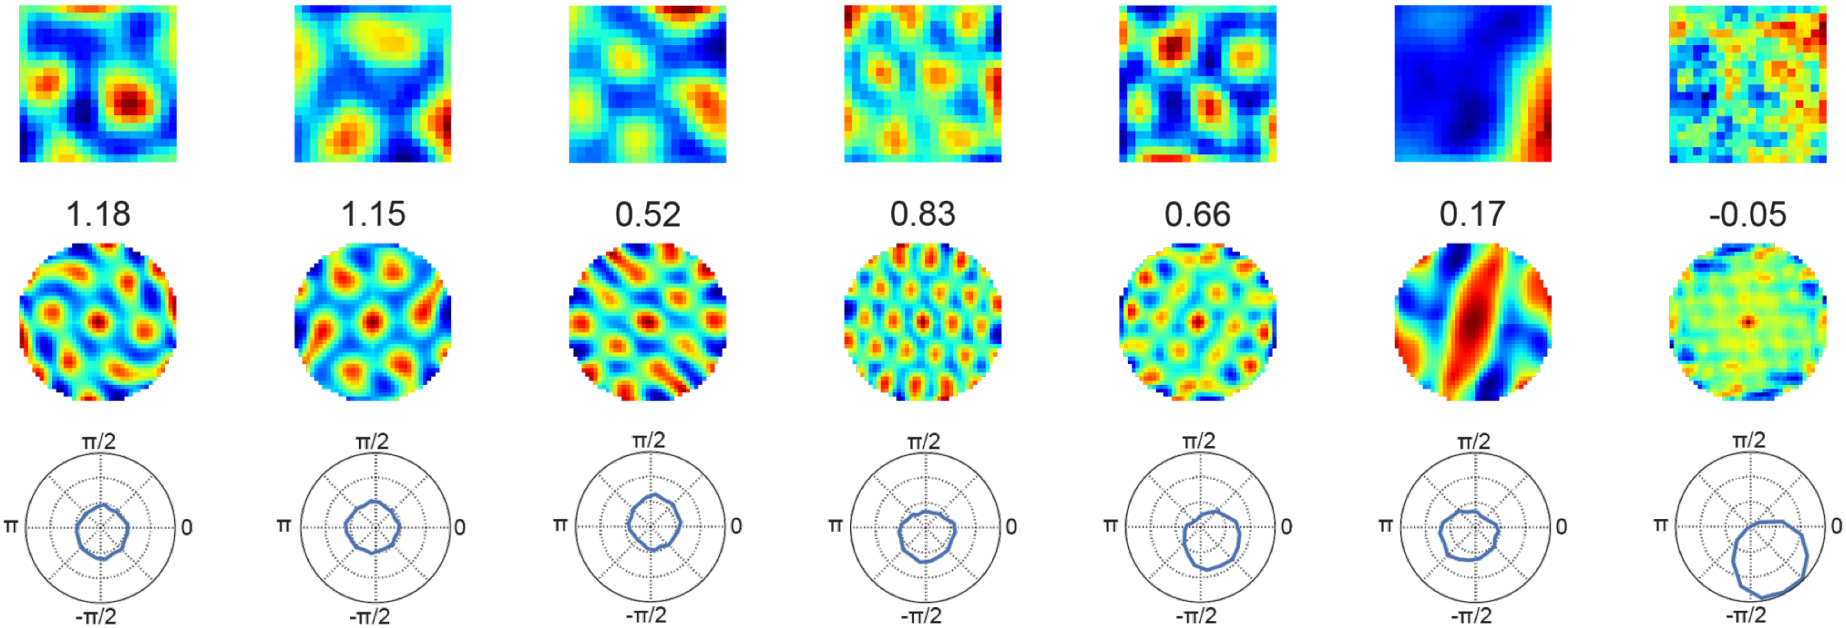

Linear layer properties

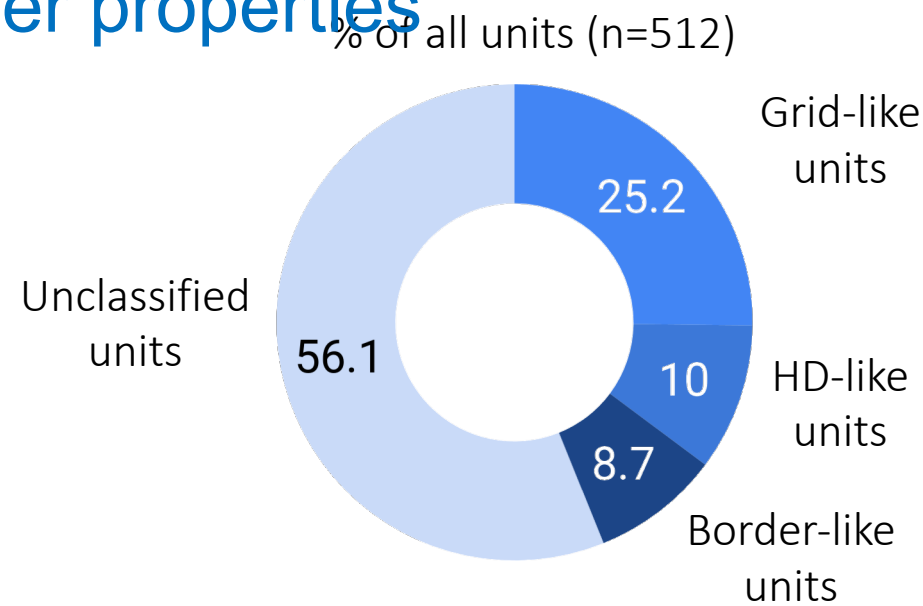

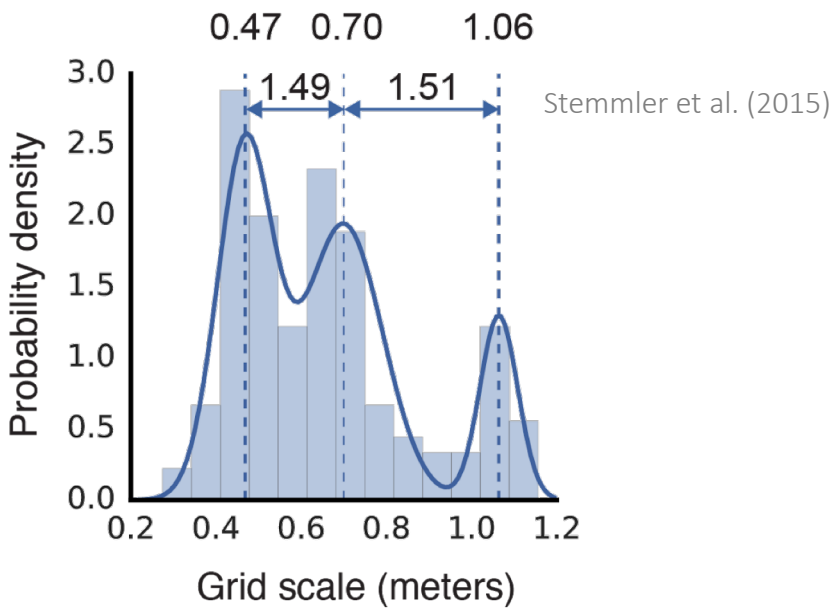

Stemmler et al. (2015)

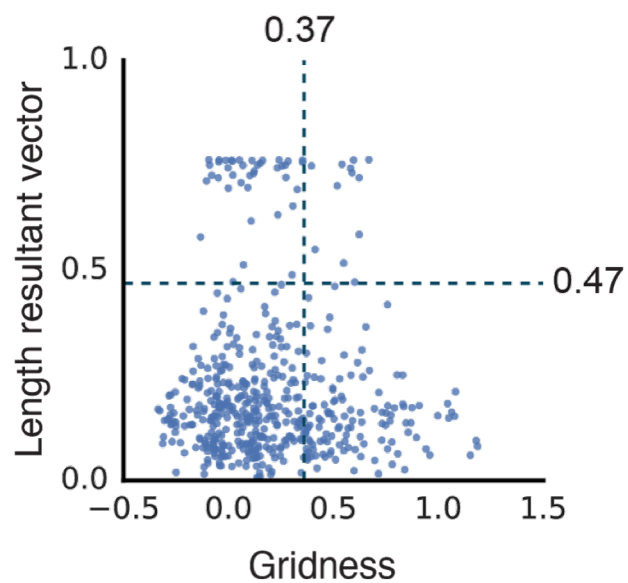

RL: grid cell agent

"Morris Water Maze"

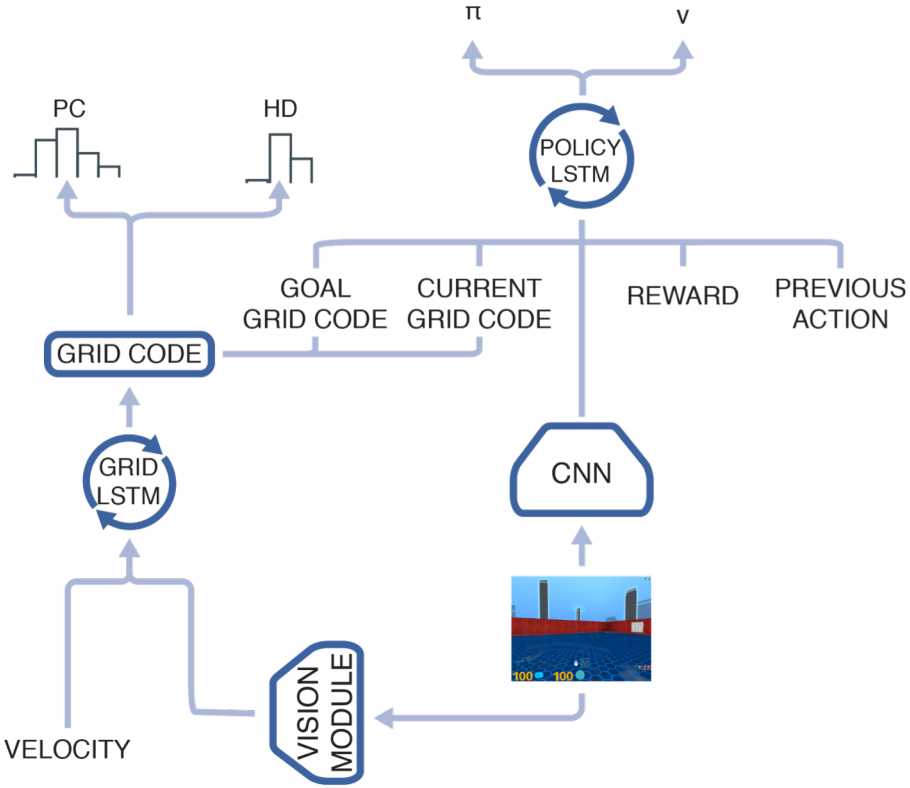

Goal: maximise expected cumulative discounted future reward

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

Performance in ‘water maze’

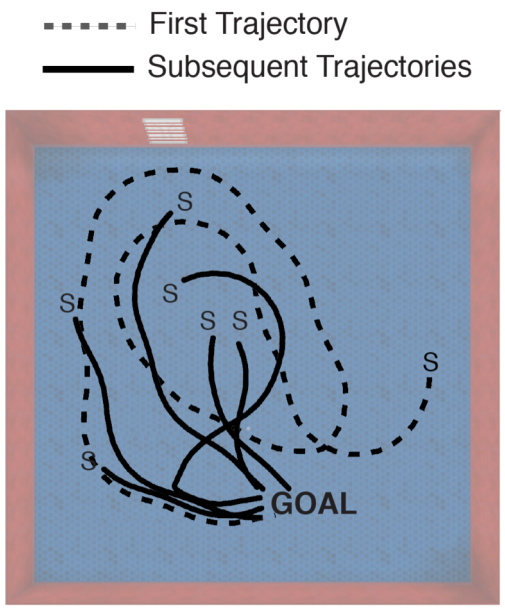

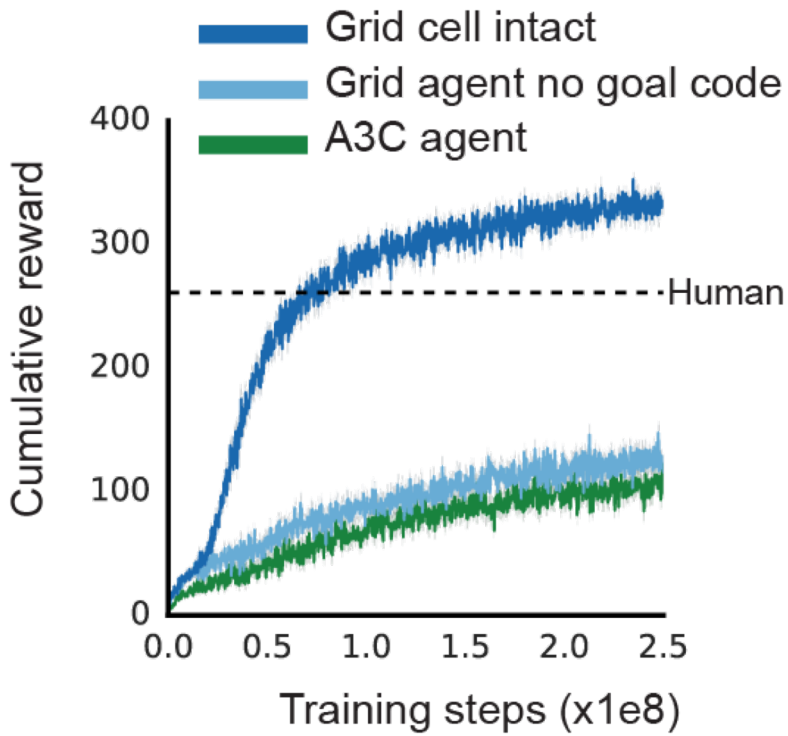

Performance in linear sunburst

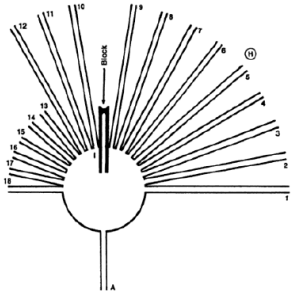

First trajectory

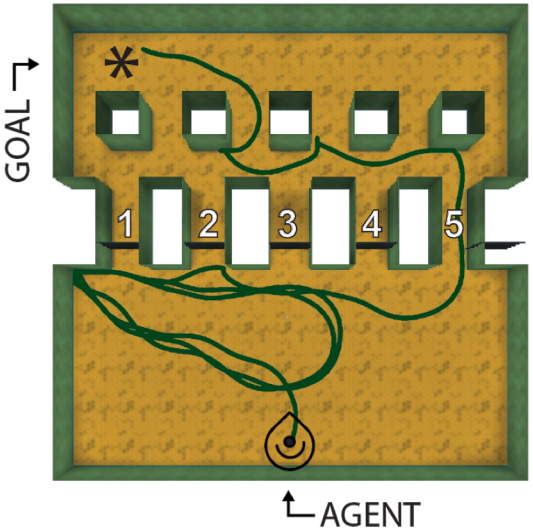

Subsequent trajectories

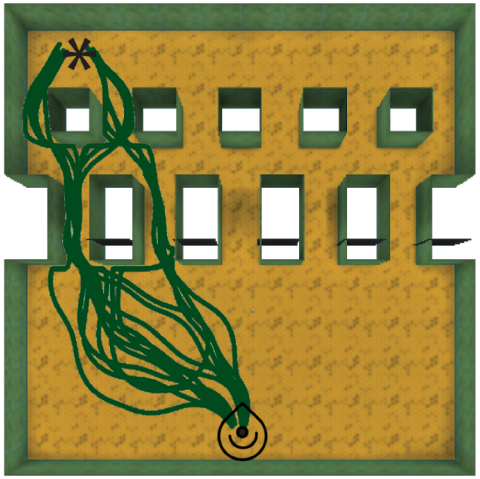

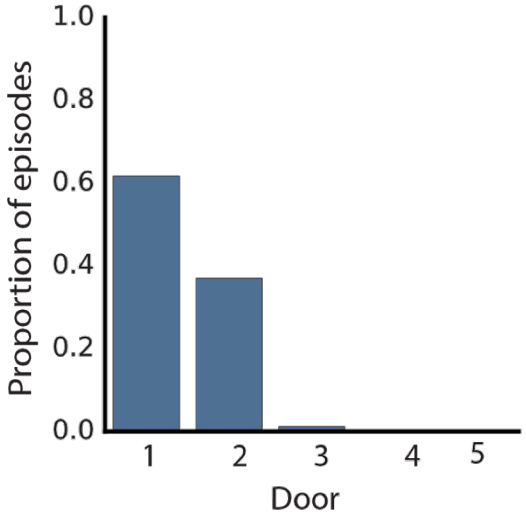

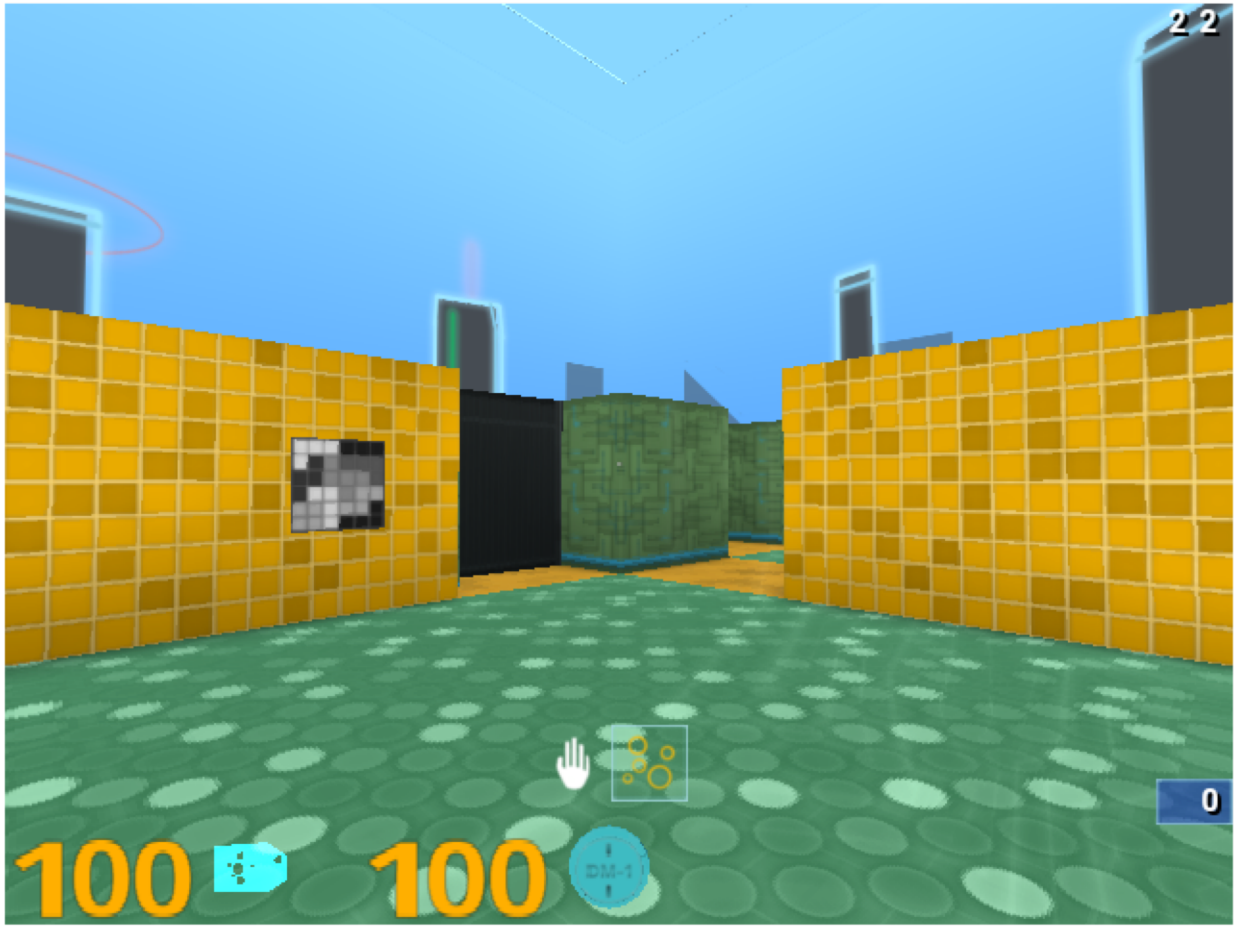

A novel maze configuration (colours, wall position, goal location) is generate for each episode

Complex maze analysis

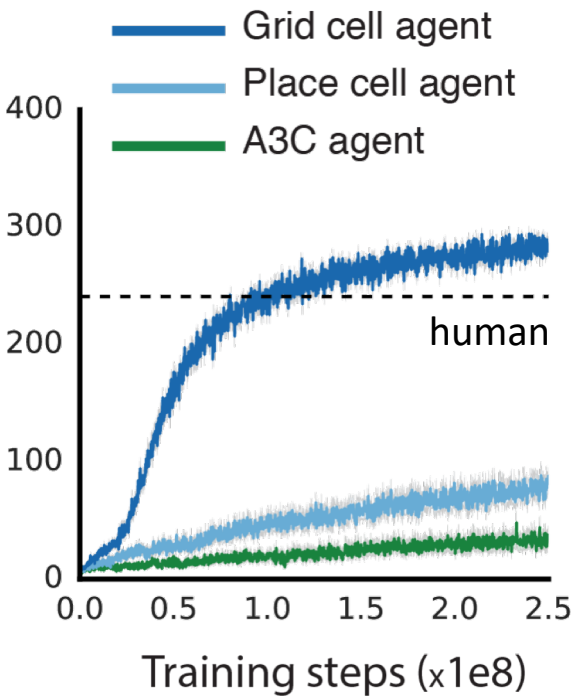

Multivariate decoding

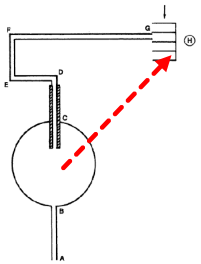

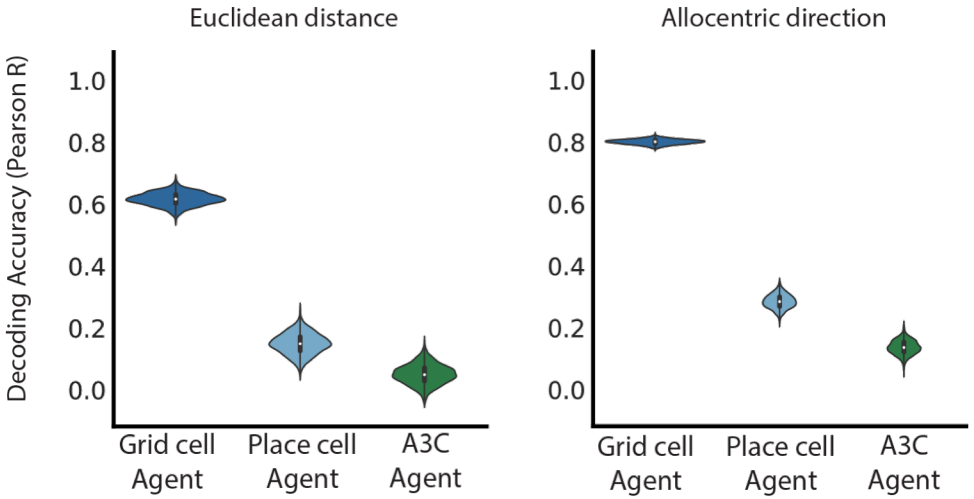

Grid cell based model. Evaluation.

- Grid cells provide a spatial metric that spans the environment. Hence vector between any two points can be calculated allowing direct navigation.

- Does not require exploration of region between current location and goal – so novel shortcuts can be taken. But probably does require that goal has been explored.

- Range covered is similar to the capacity of the grid system as a whole, which is belived to greatly exceed the scale of the individual grid cells (upper limit being the least common multiple of the grid scales)

- Number of possible mechanisms by which vector might be calculated, some plausible some less so