Hippocampal mechanisms of memory and cognition

Matthew Wilson
Departments of Brain and Cognitive Sciences and Biology
MIT

The Picower Institute
for learning and memory

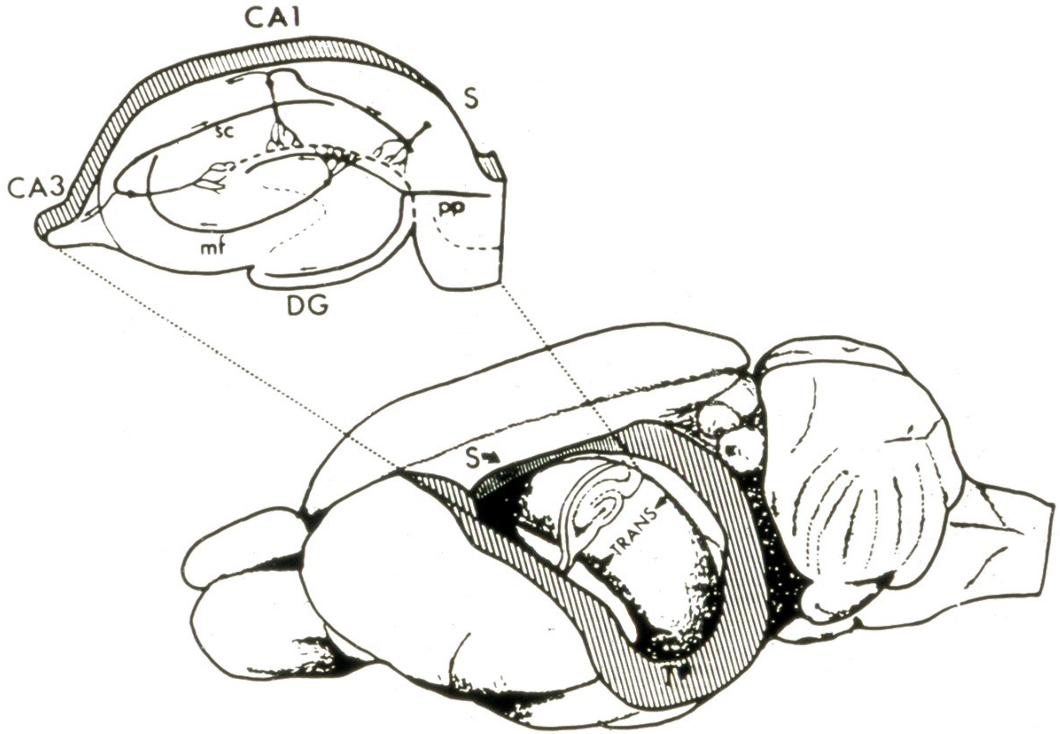

Fig. 2. The position of the hippocampal formation in the rat brain is shown in this drawing of a preparation in which the cortical surface overlying the hippocampus has been removed. The hippocampus is an elongated, C-shaped structure with the long or septotemporal axis running from the septal nuclei rostrally (S) to the temporal cortex (T) ventrocaudally. The short or transverse axis (TRANS) is oriented perpendicular to the septotemporal axis. The major fields of the hippocampal formation (except for the entorhinal cortex) are found in slices taken approximately midway along the septotemporal axis. The slice pictured at top left is a representation of the summary of the major neuronal elements and intrinsic connections of the hippocampal formation as originally illustrated by Andersen *et al.* (see text for details).

Abbreviations: DG, dentate gyrus; mf, mossy fibers; pp, perforant path; sc, Schaffer collaterals.

From Amaral and Witter, 1989

Hippocampus in spatial and episodic memory

- The hippocampus is involved in the formation of episodic memory as well as spatial memory used in navigation.
- Navigation - linkage of spatial locations
- Episodic memory - linkage of events
- Both may depend critically on temporal sequence encoding

Tetrodes simultaneously record the activity of multiple hippocampal cells in awake, behaving rats.

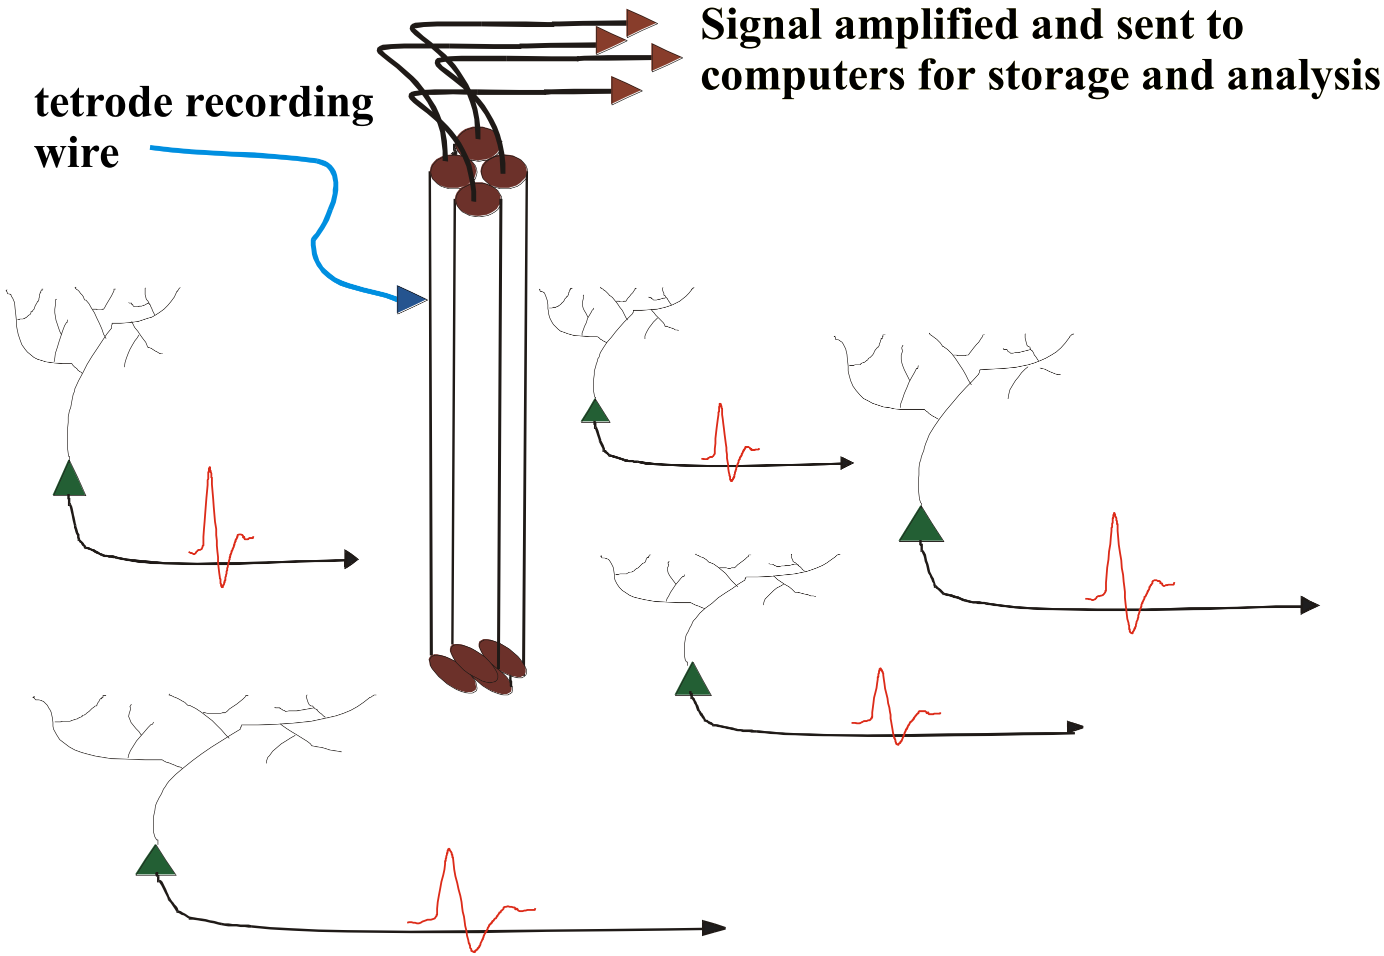

Neural recording device

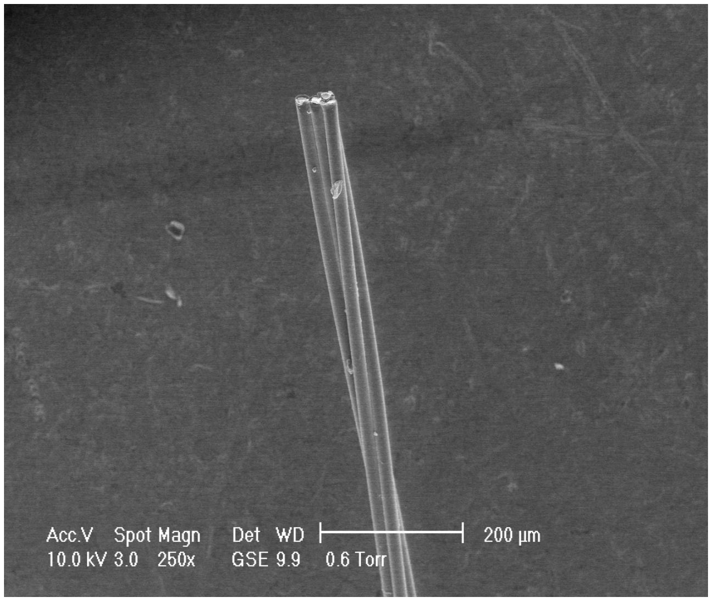

4-channel microwire electrode

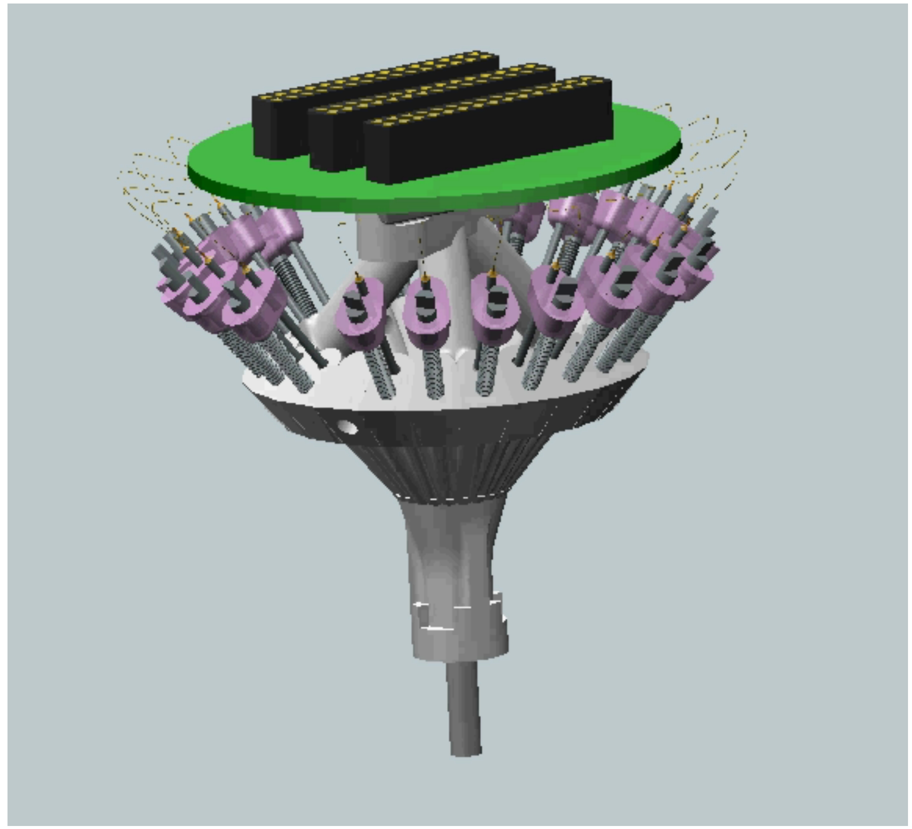

Multiple electrode microdrive array

Spike amplitude clustering

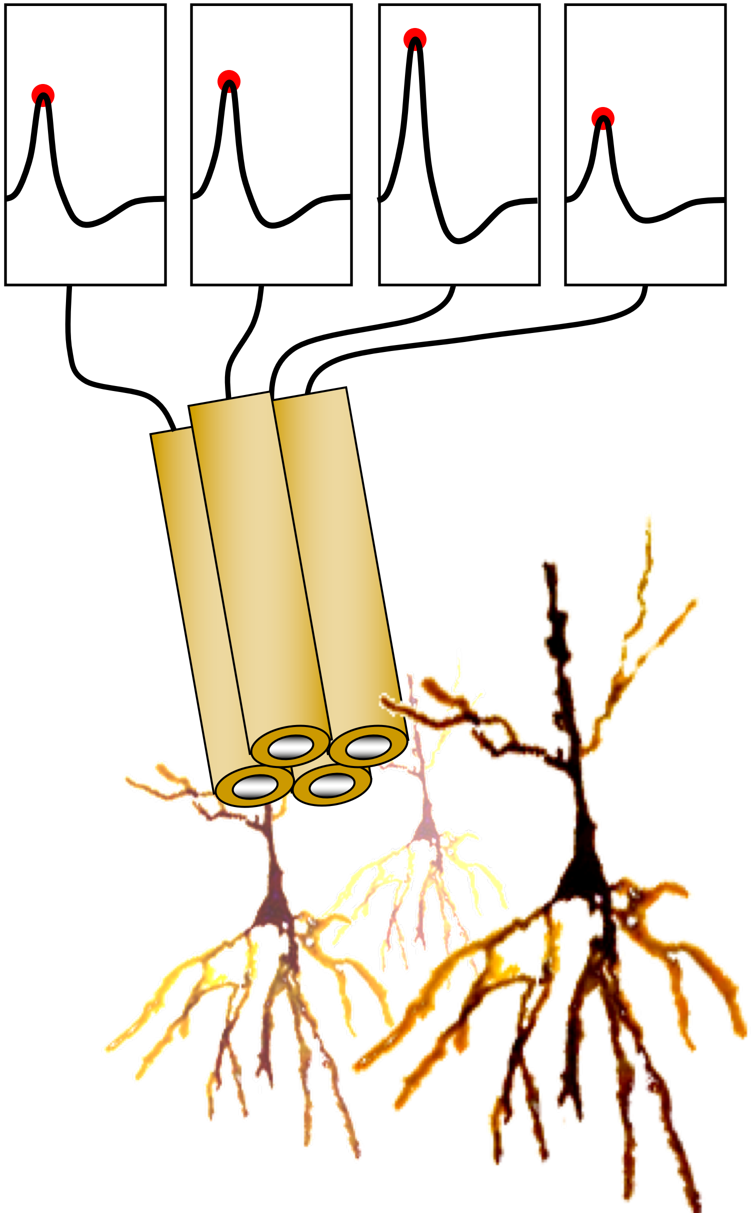

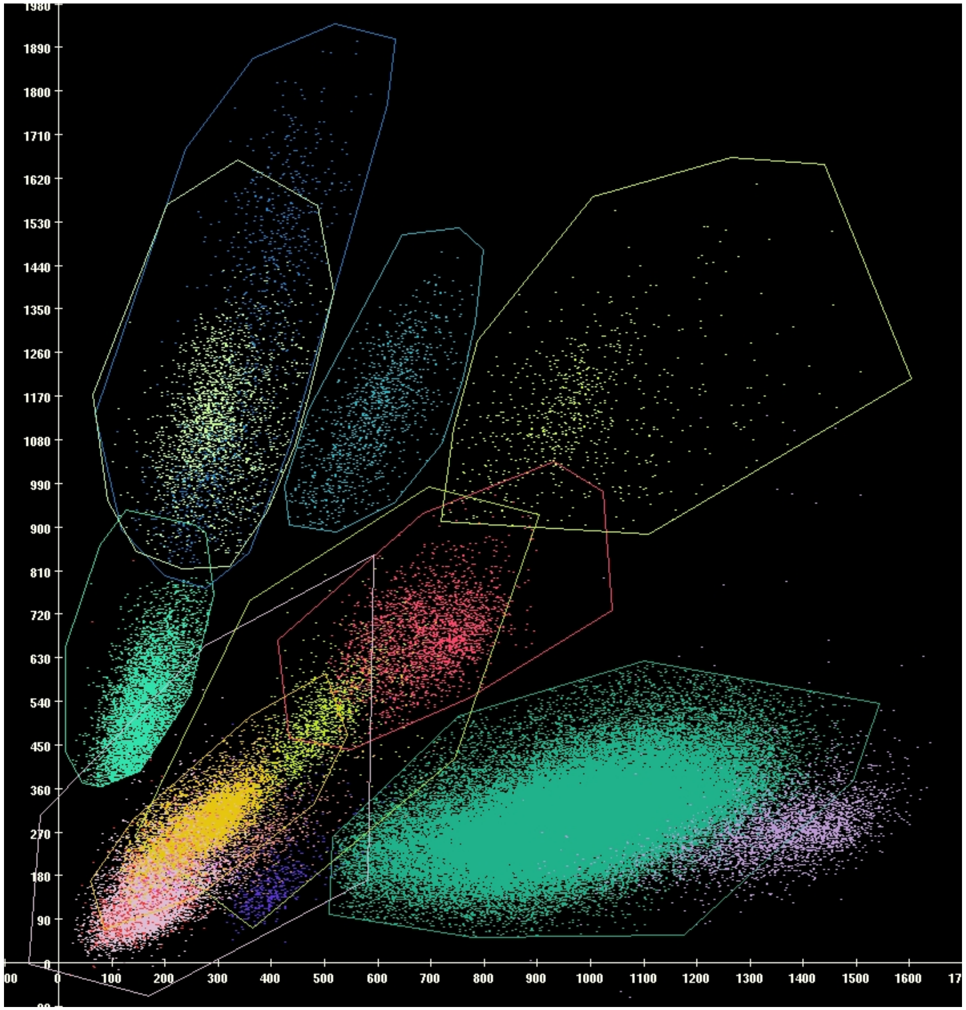

Example of a Simple Spatial Environment

Ensemble Activity in Area CA1 During Spatial Exploration

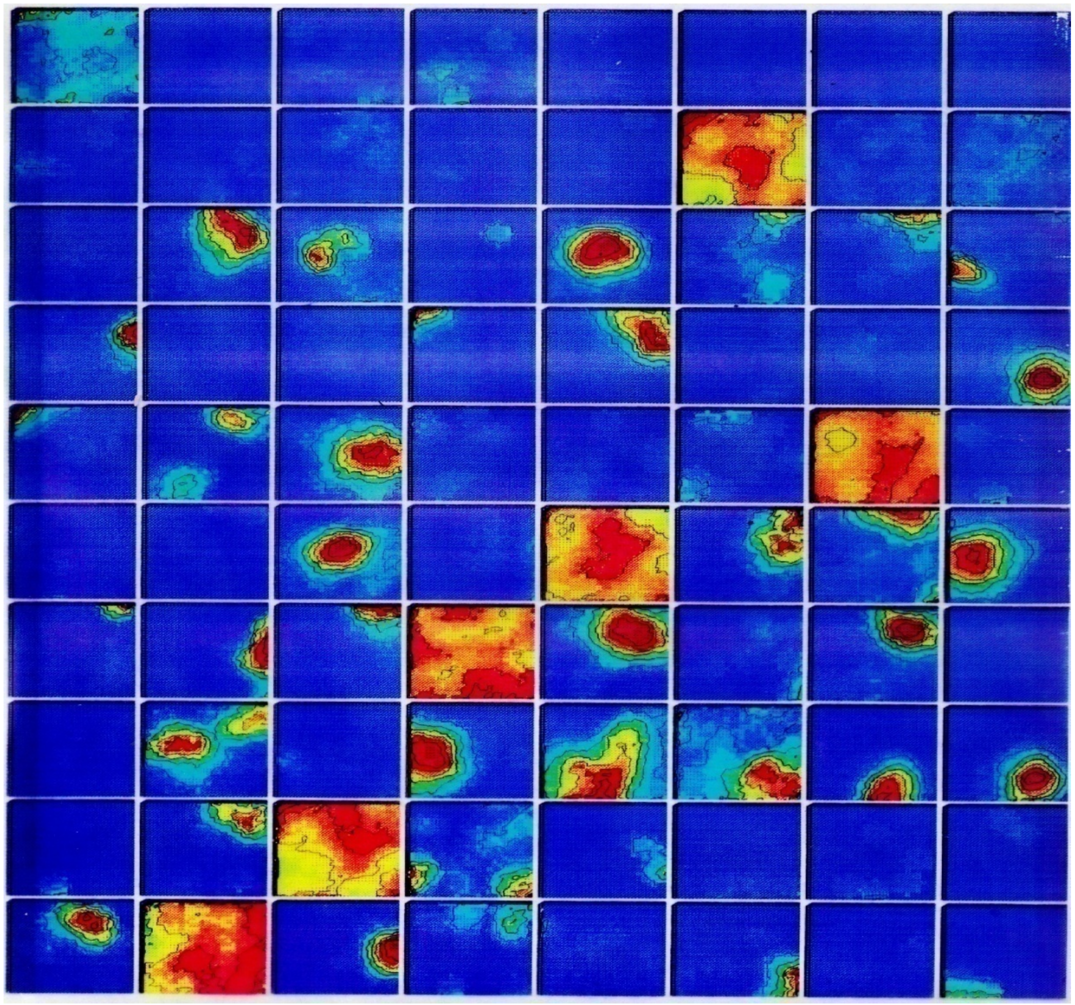

Wilson and McNaughton, *Science*, 1993

Place Fields on Linear Tracks

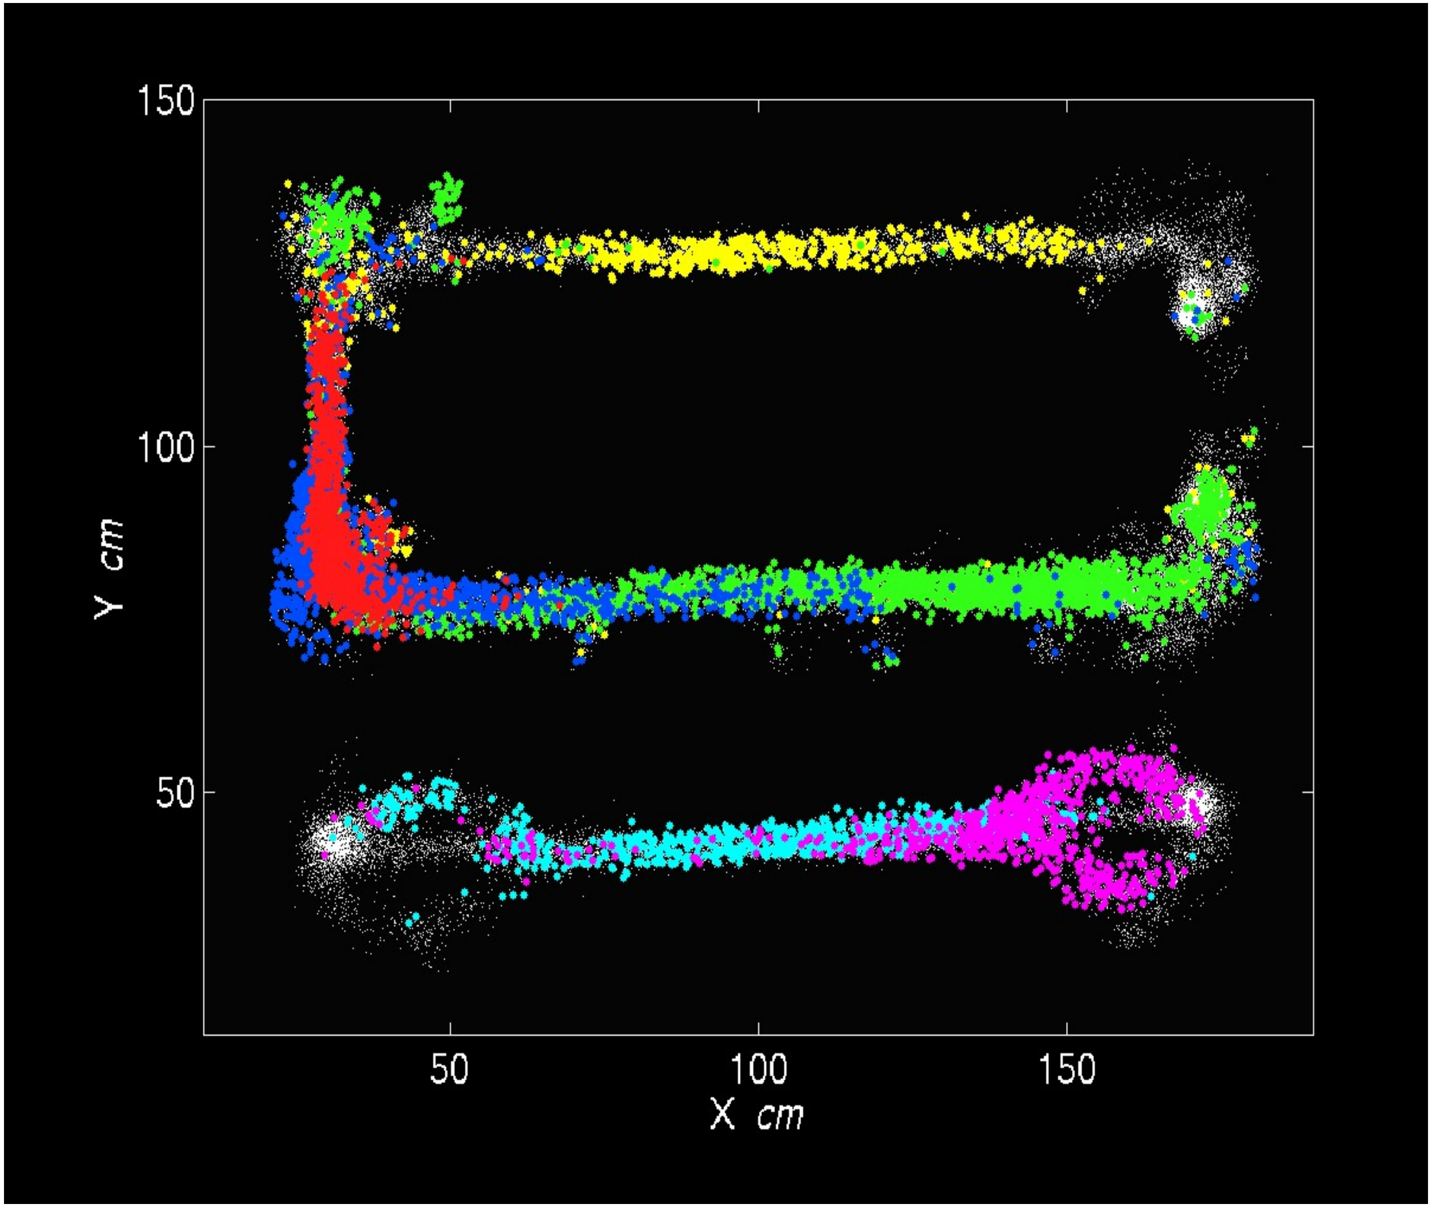

Hippocampal Place Cells

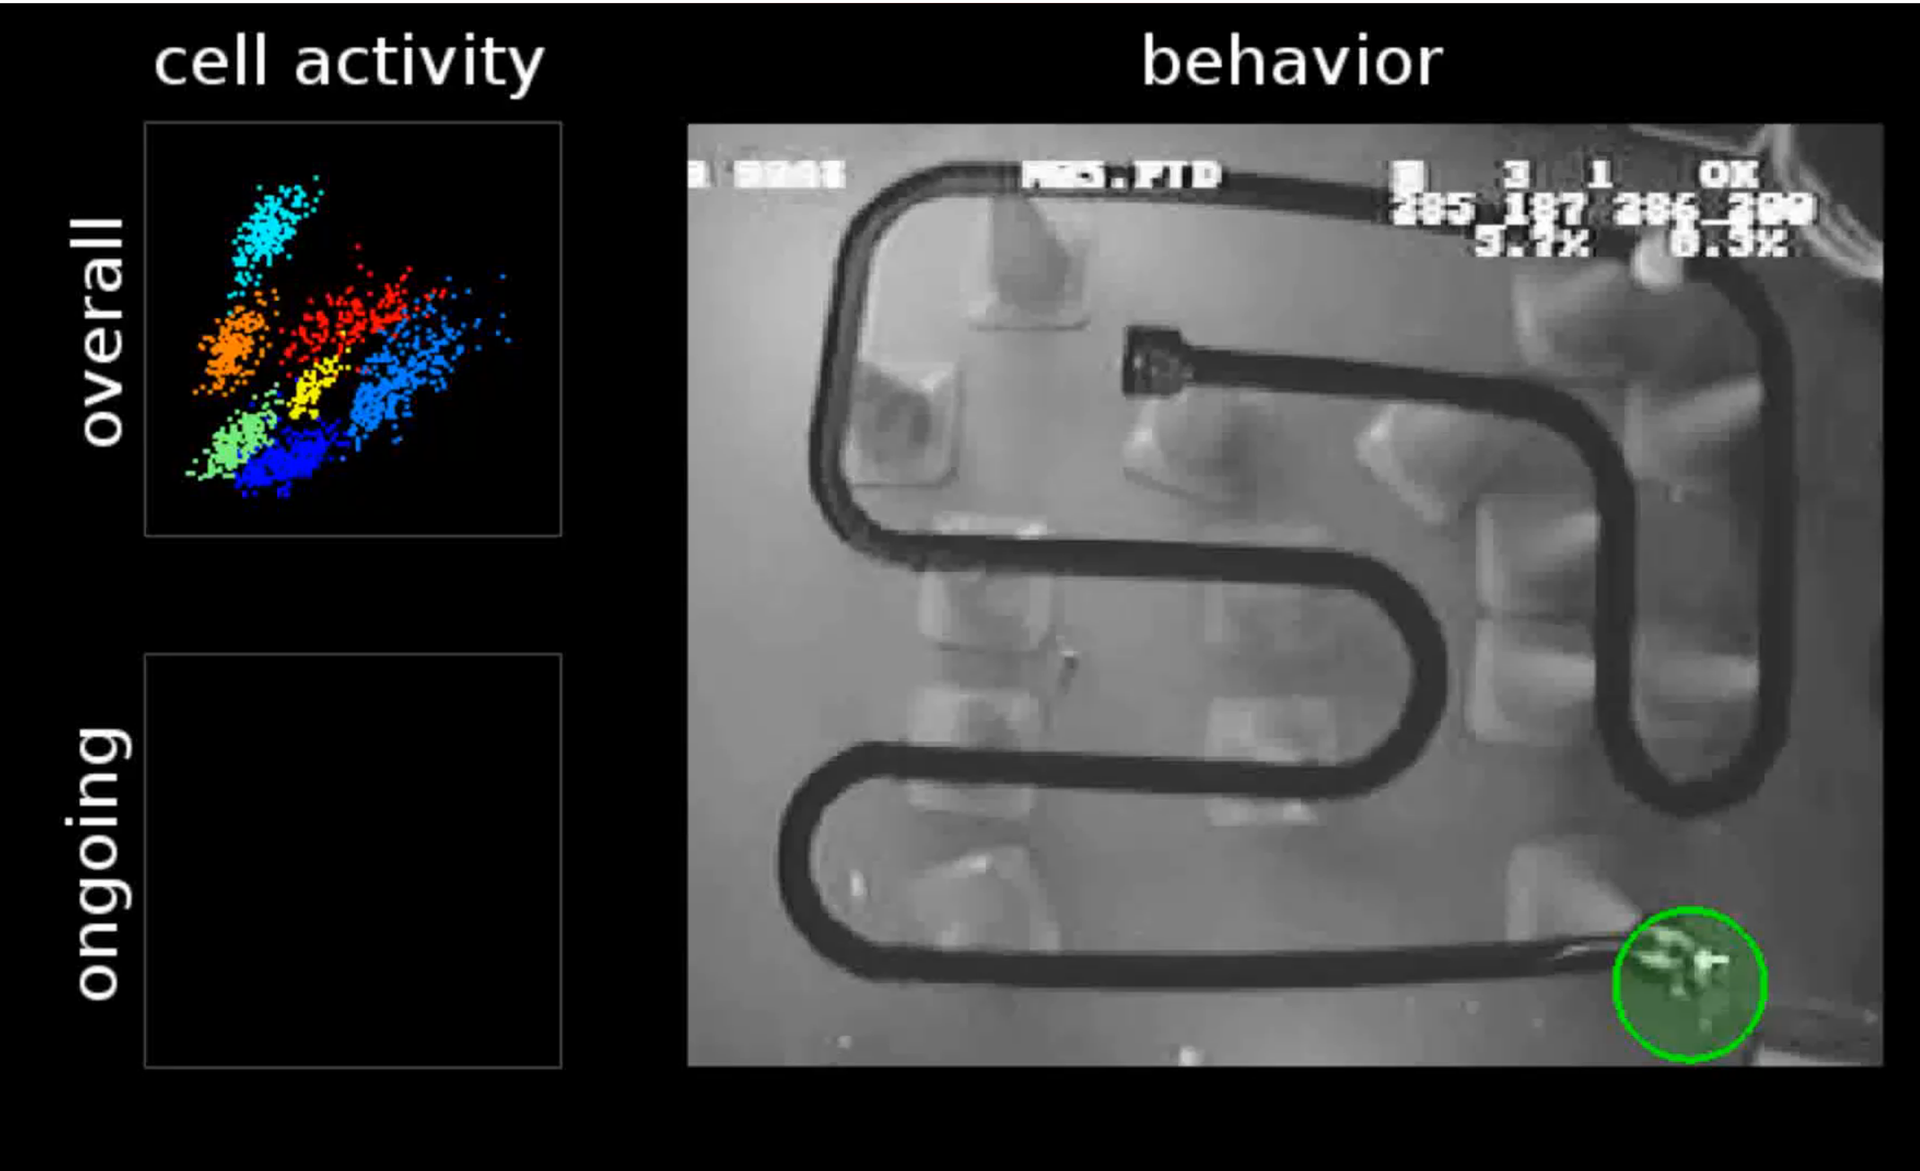

Hippocampal Ensemble Decoding

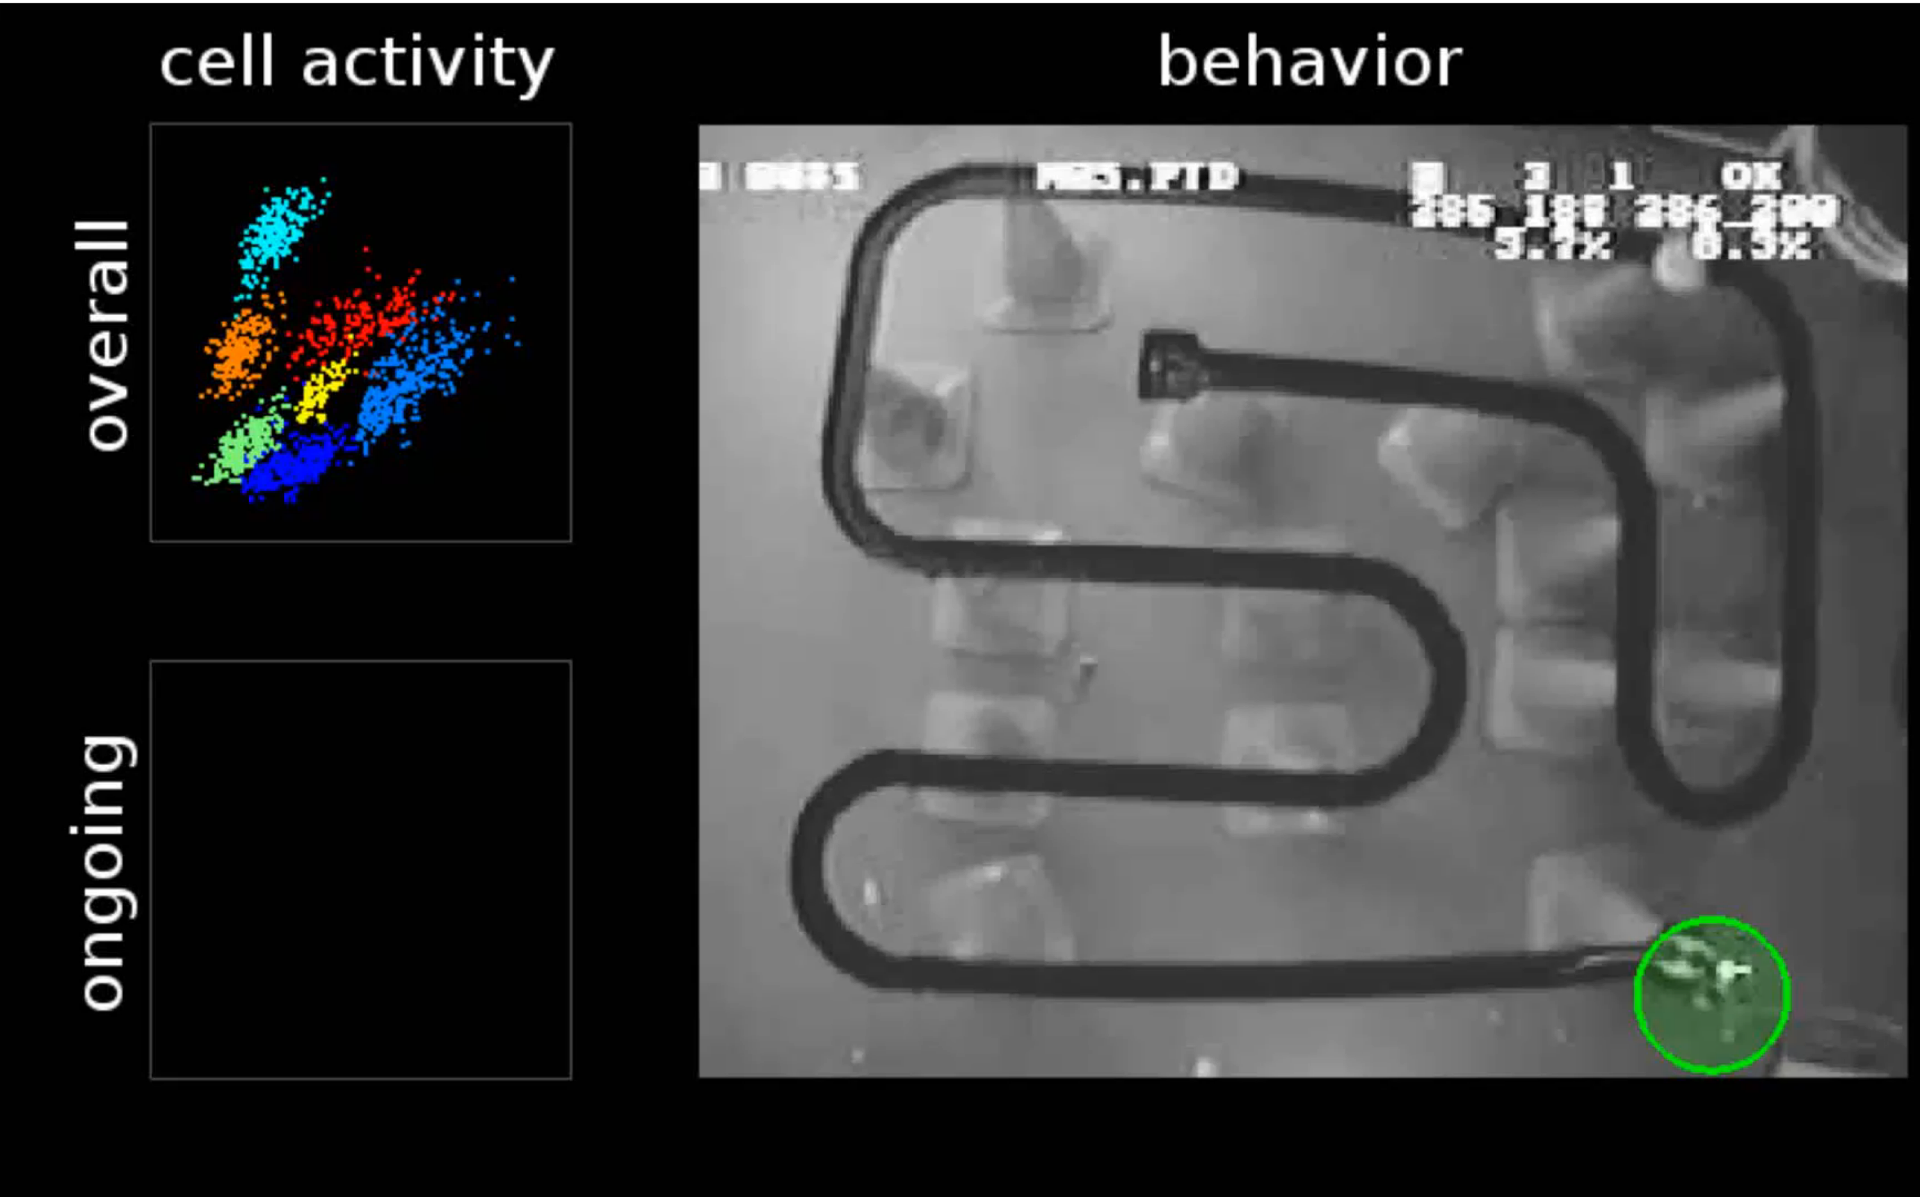

Decoding Sleep Reactivation

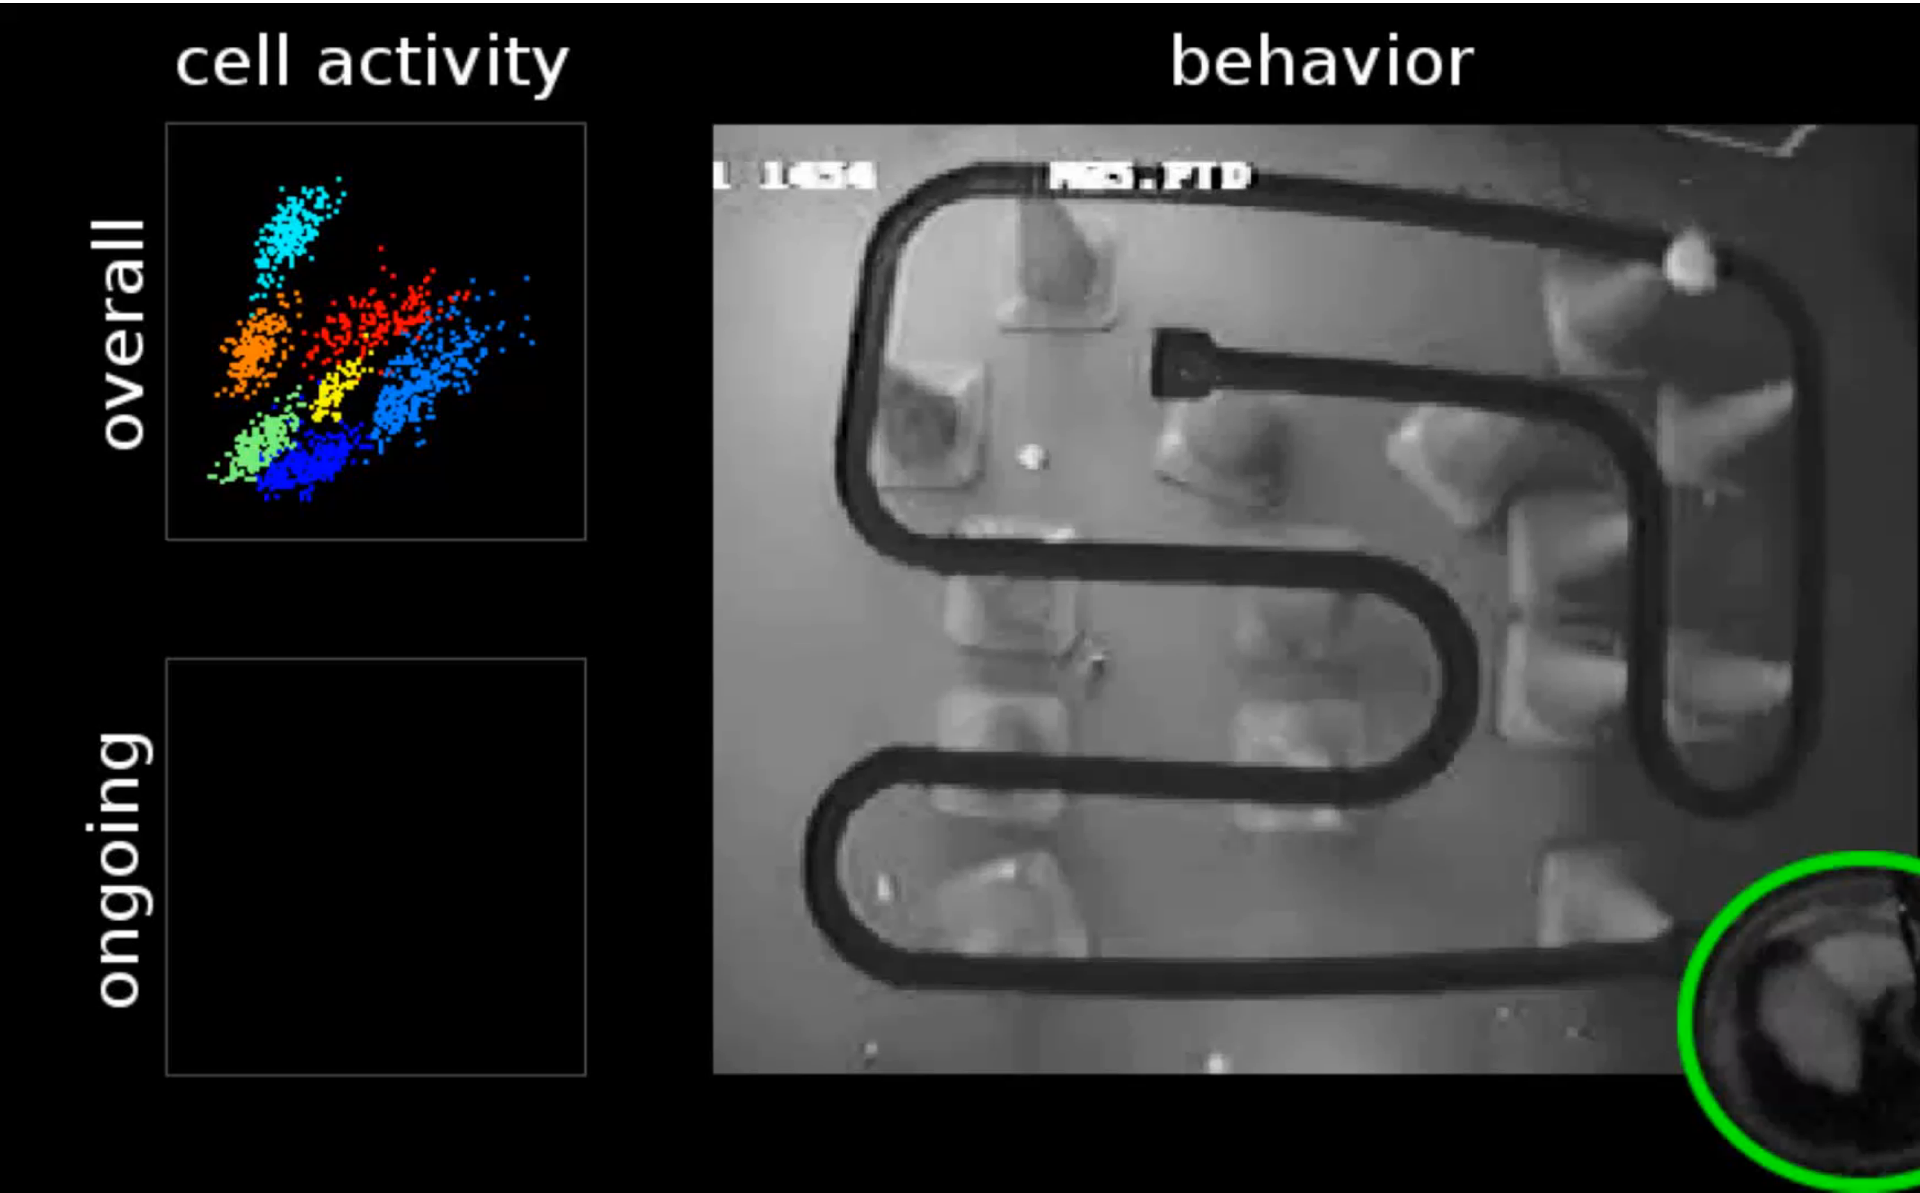

Hippocampus online and offline

Theta rhythm

Sharp wave/ripples

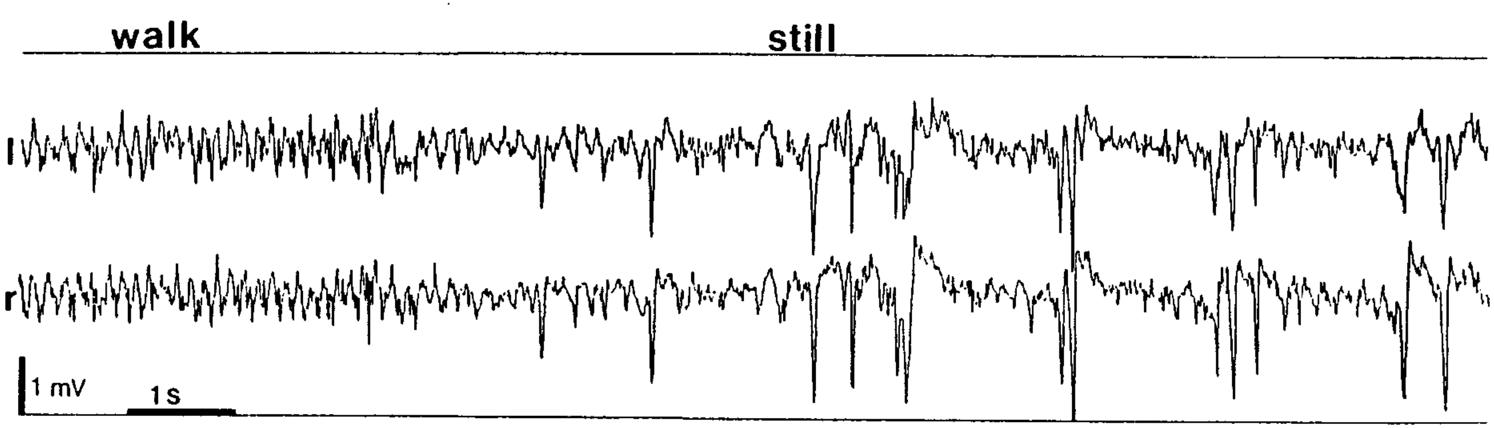

Buzsáki 1989

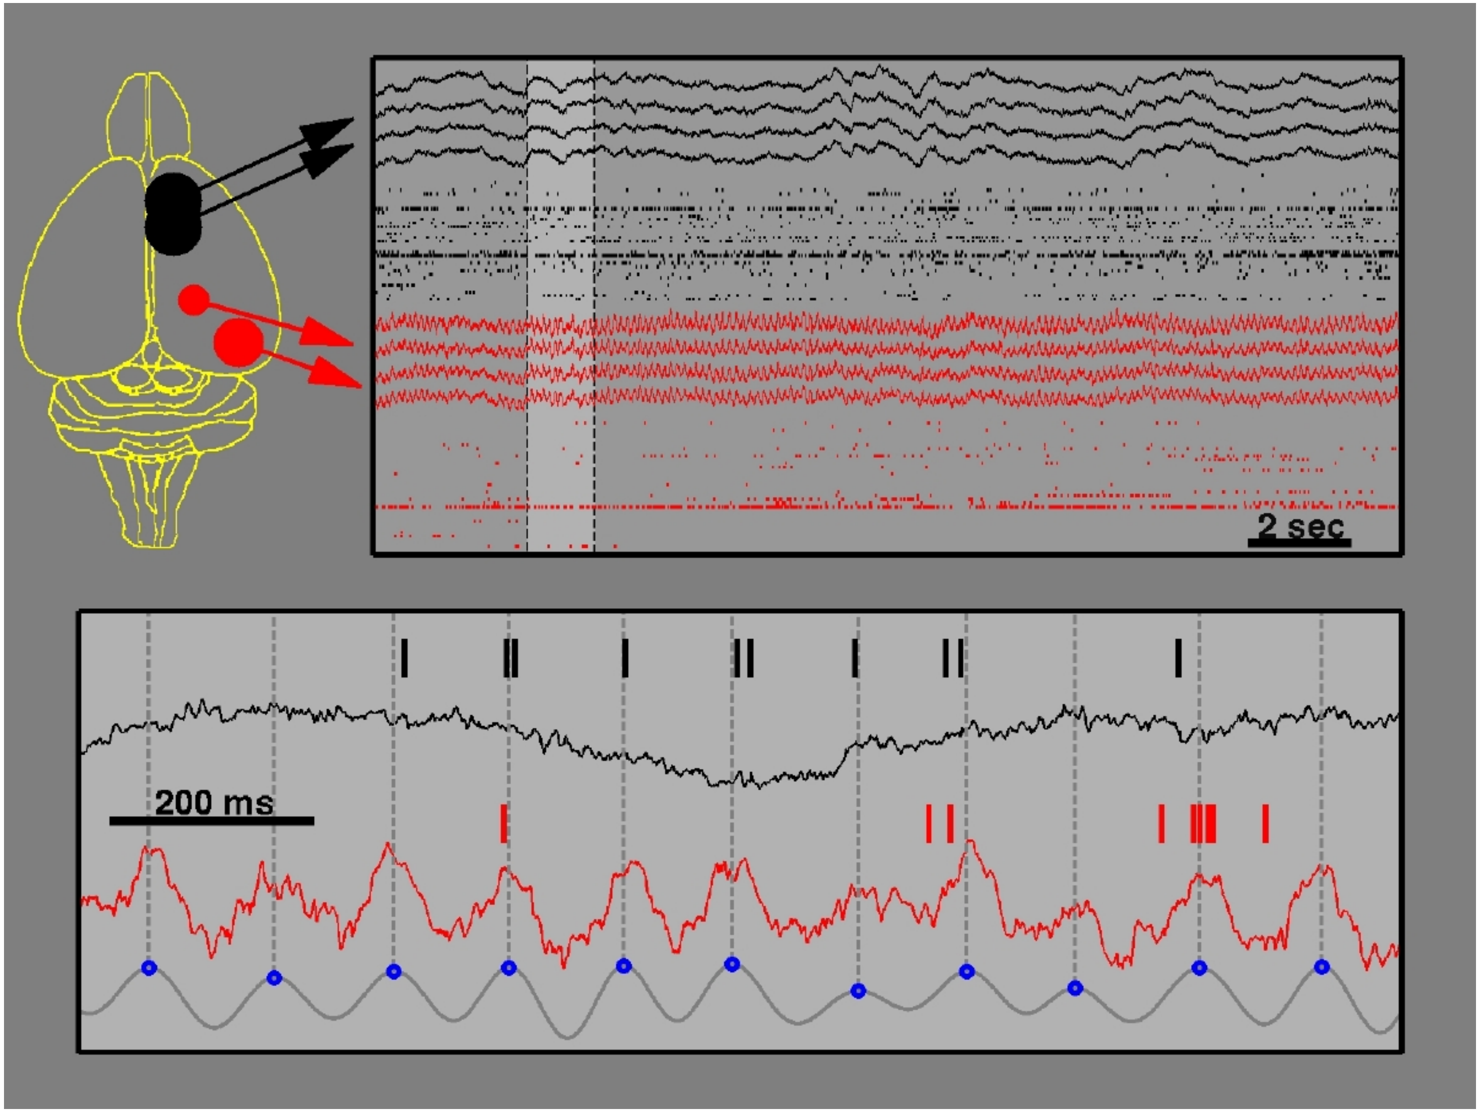

2 sec
200 ms

Interaction of asymmetric excitation with oscillatory variation in inhibition can translate one linear dimension (space) into another (time).

Hippocampal phase precession may be a demonstration of that process.

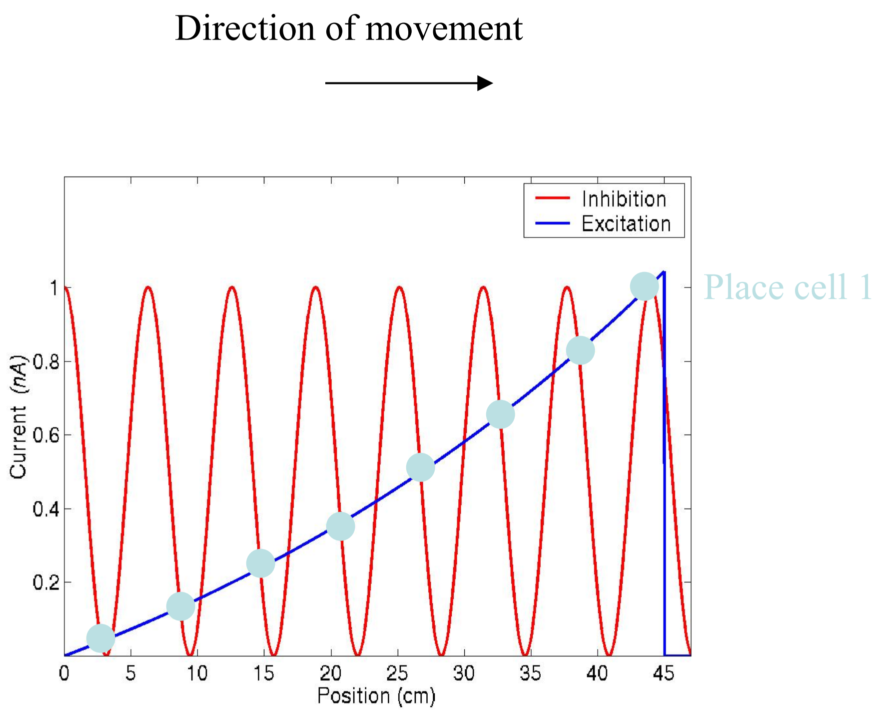

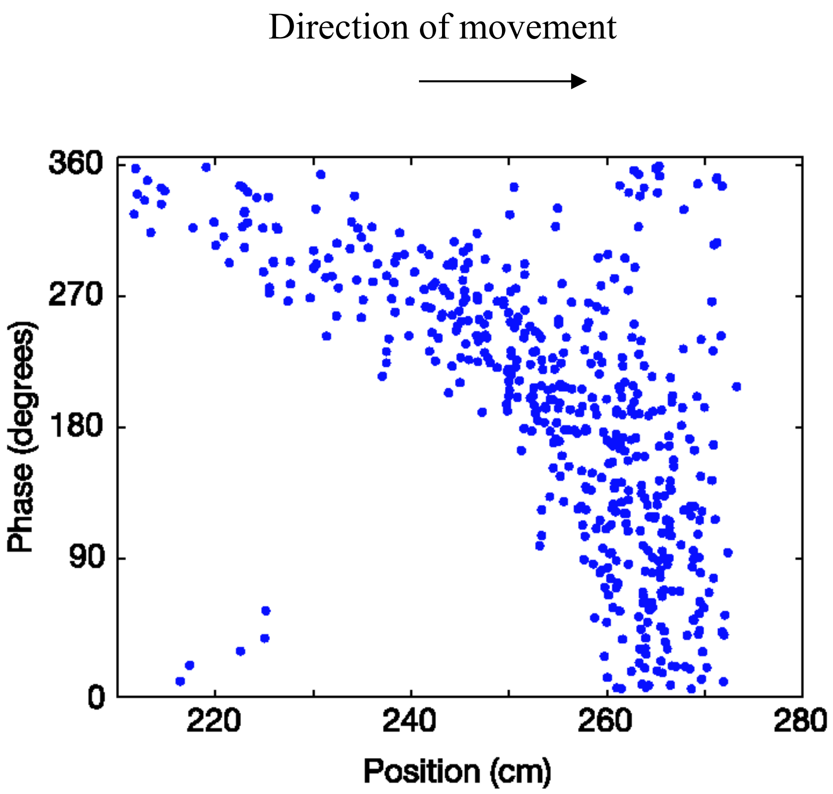

Overlapping asymmetric place fields with oscillatory variation in excitability translate behavioral time relationships to biophysical timescales with preserved temporal order

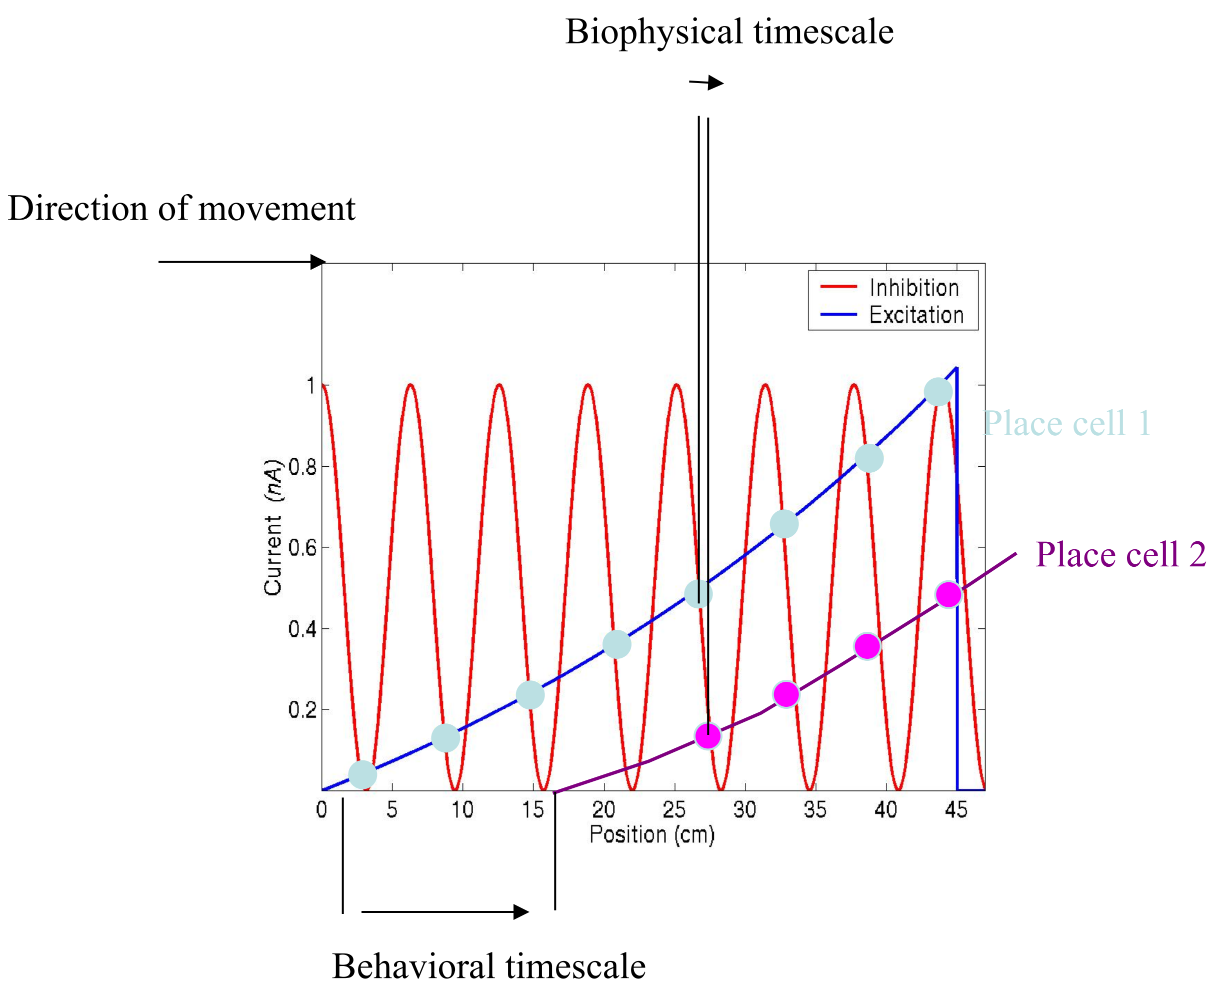

Hippocampal theta sequences

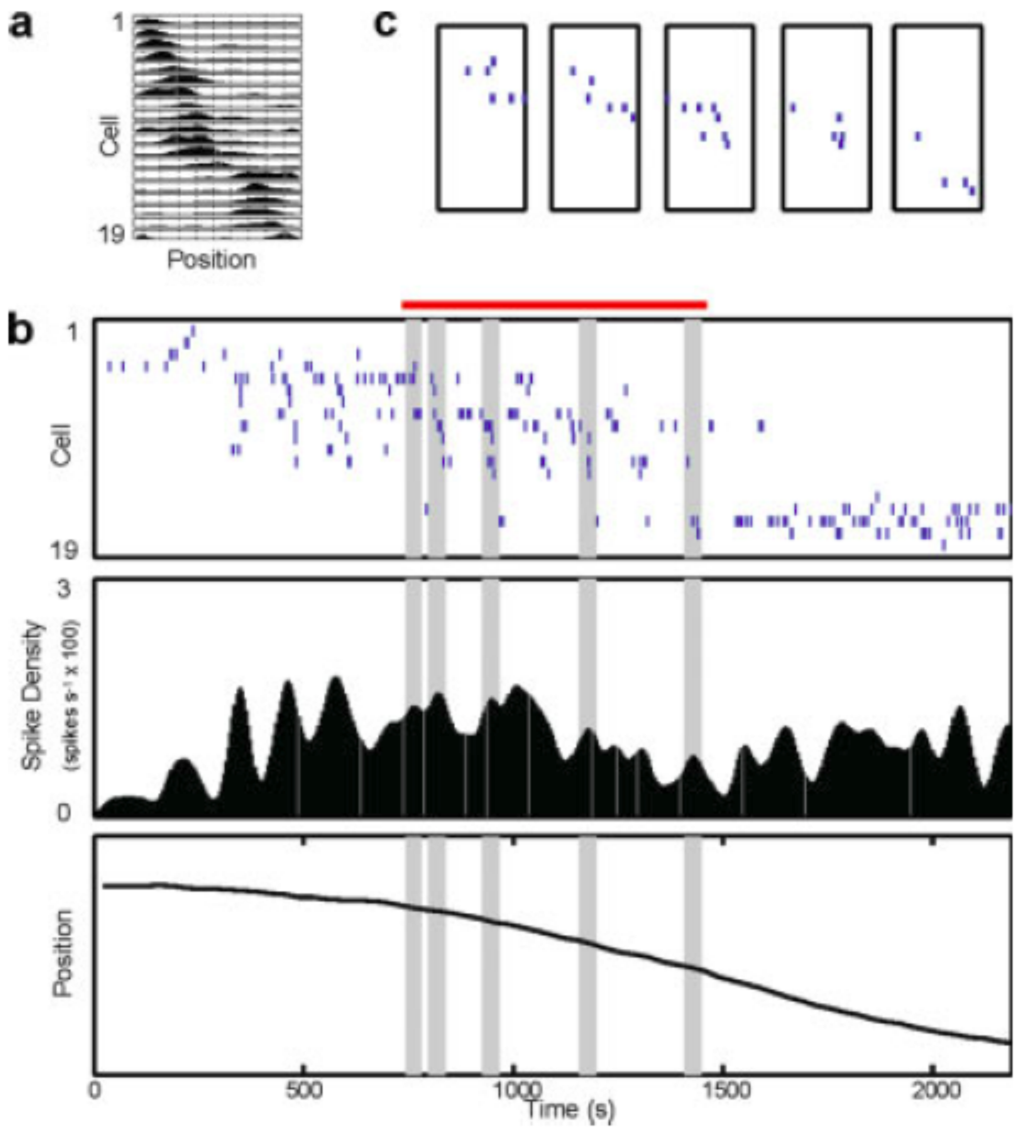

Foster and Wilson, 2007

Hippocampal spatial representations are encoded as sequences during behavior

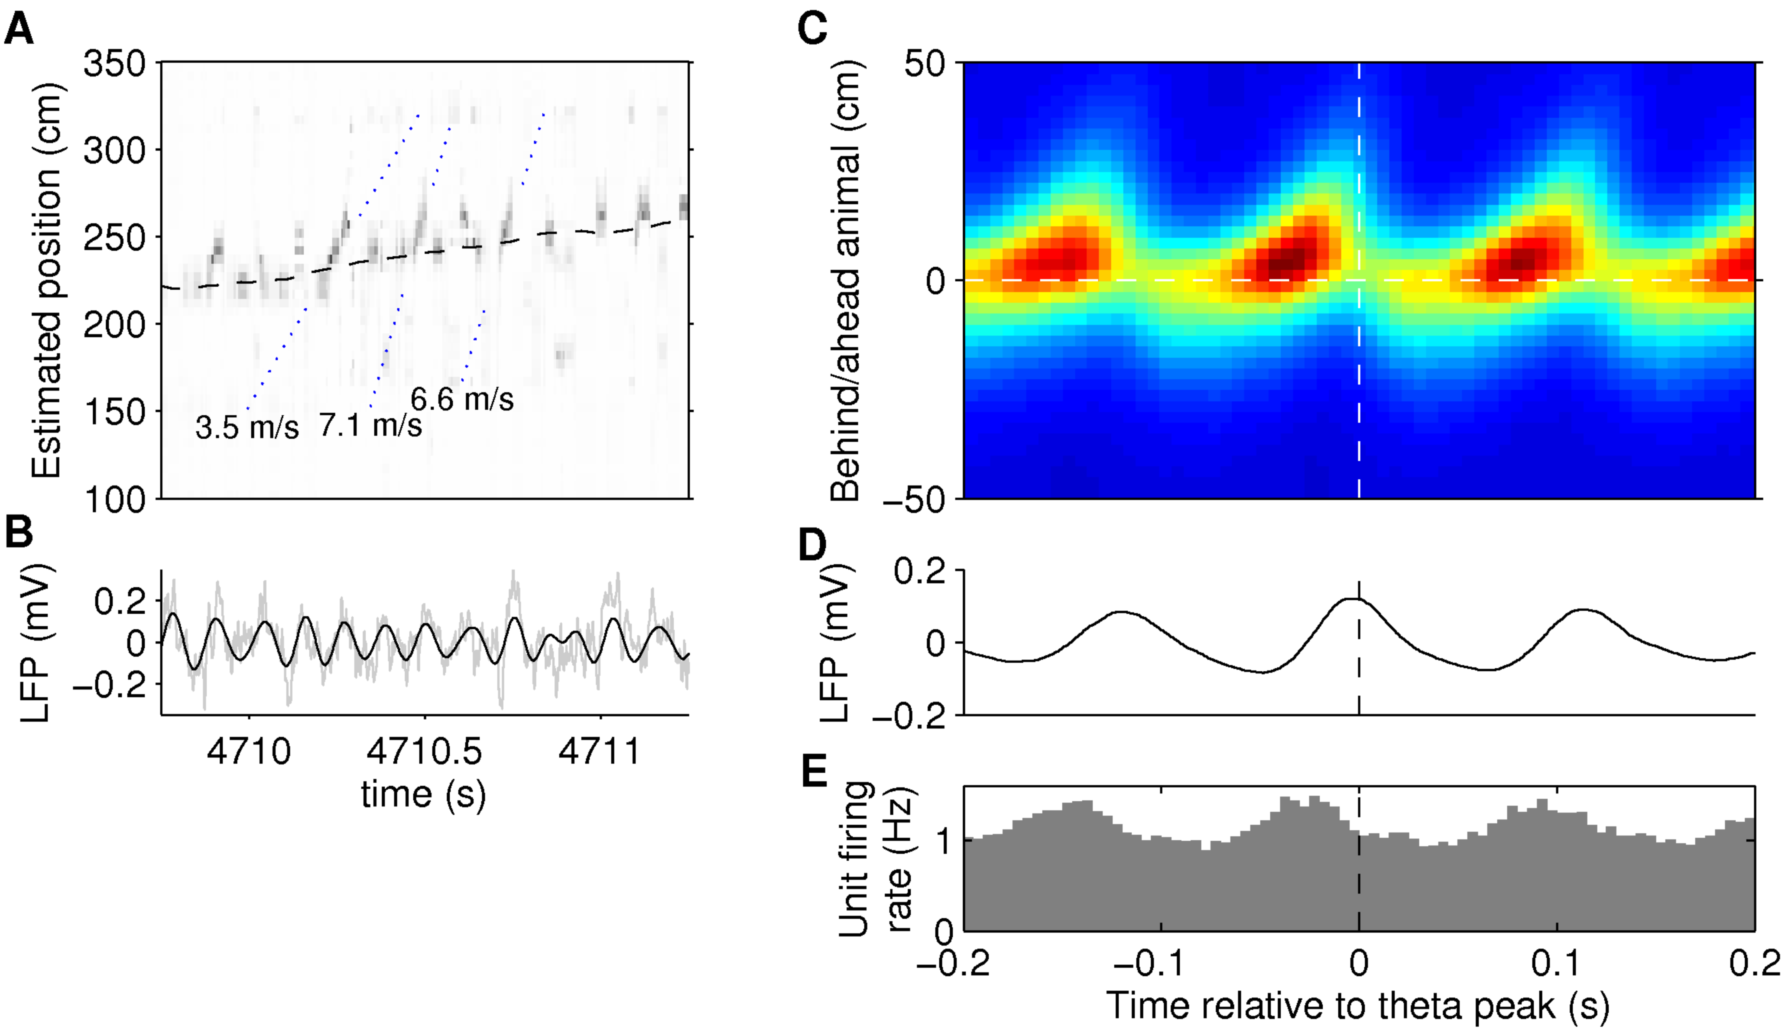

Davidson, Kloosterman, and Wilson, *Neuron*, 2009

HPC-PFC: functionally connected

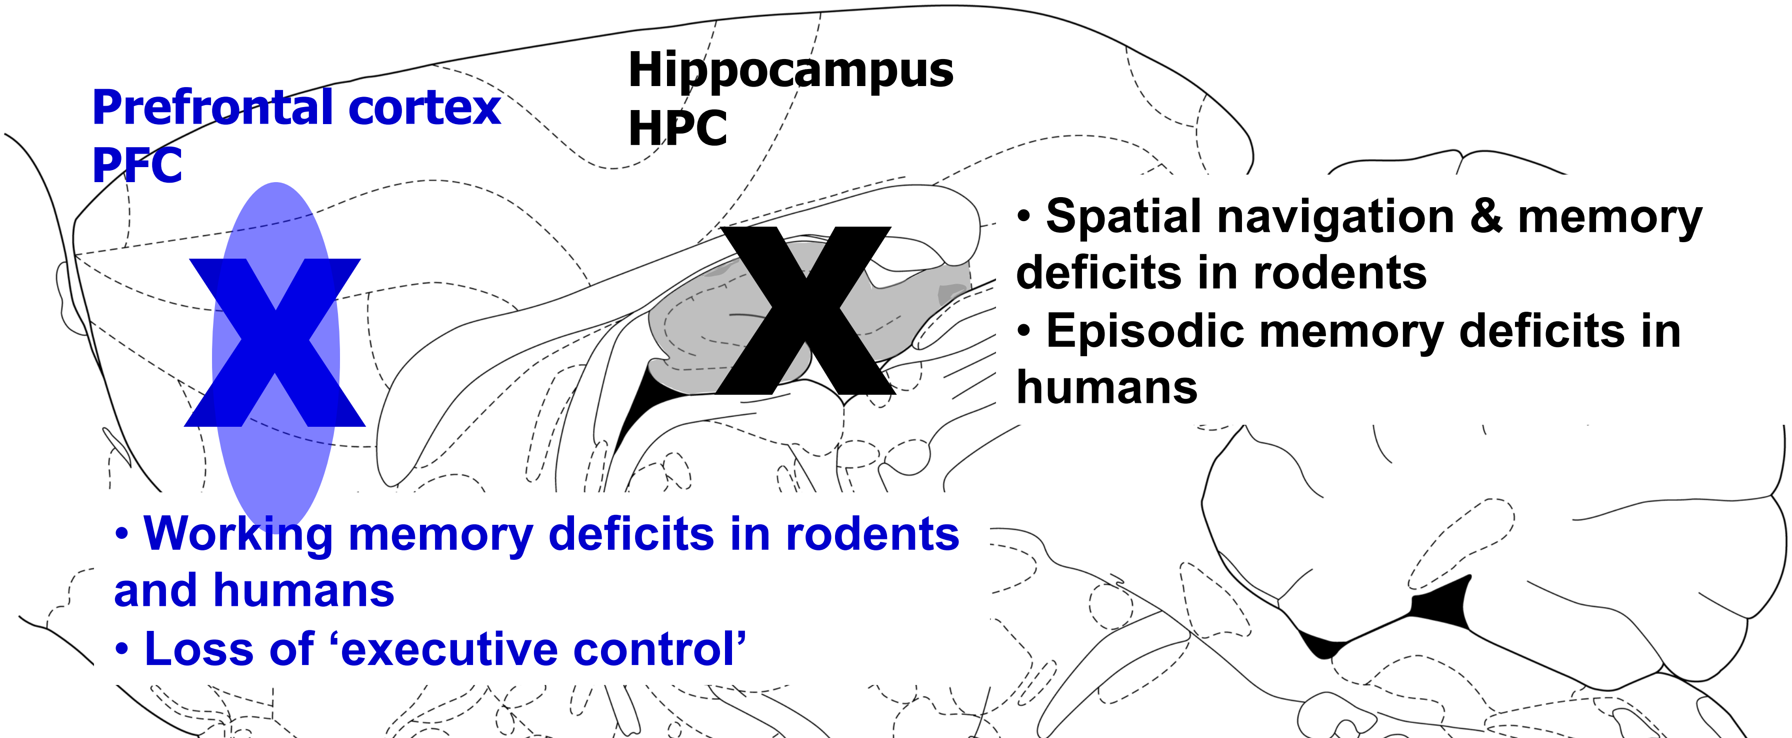

The hippocampus: **encoding and recognising spatial context**

The prefrontal cortex: **integrating the cues of current context (held on-line in working memory) to control appropriate behaviour**

HPC-PFC: functionally connected during spatial working memory tasks

Hippocampus

Prefrontal c

PFC

X

ation & memory

nts

nory deficits in

• Working r

and human

• Loss of 'e

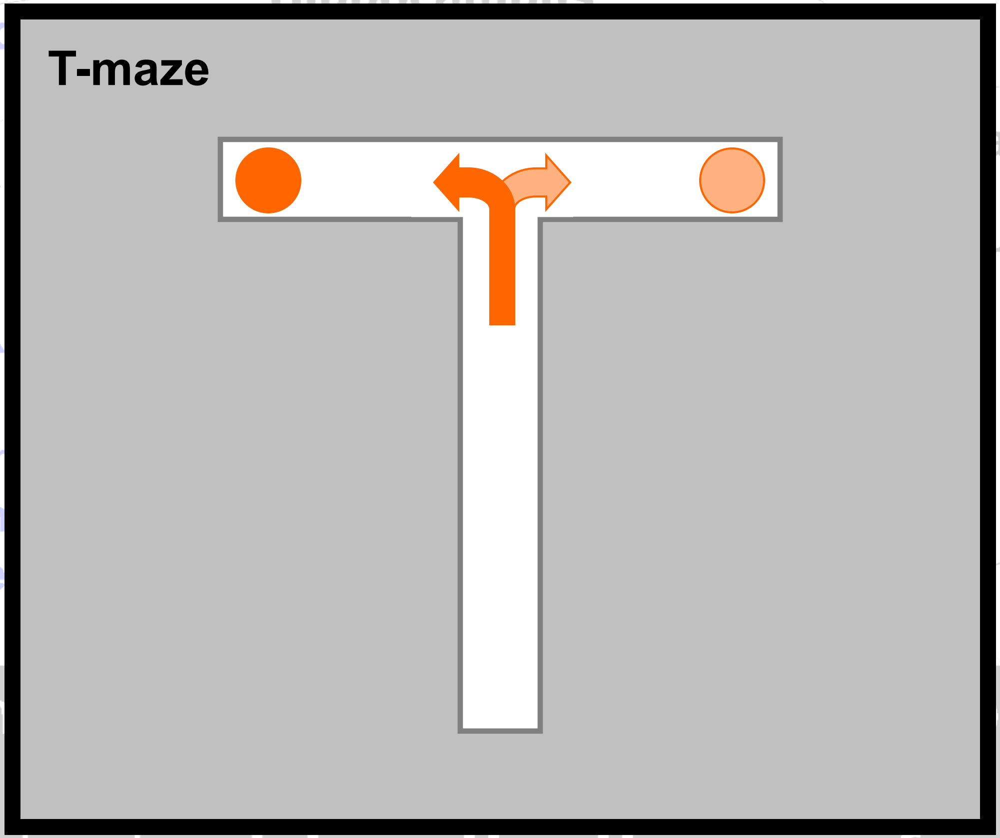

The hippocan atial context

The prefrontal cortex: **integrating the cues of current context (held on-line in working memory) to control appropriate behaviour**

HPC-PFC: individual electrophysiologies

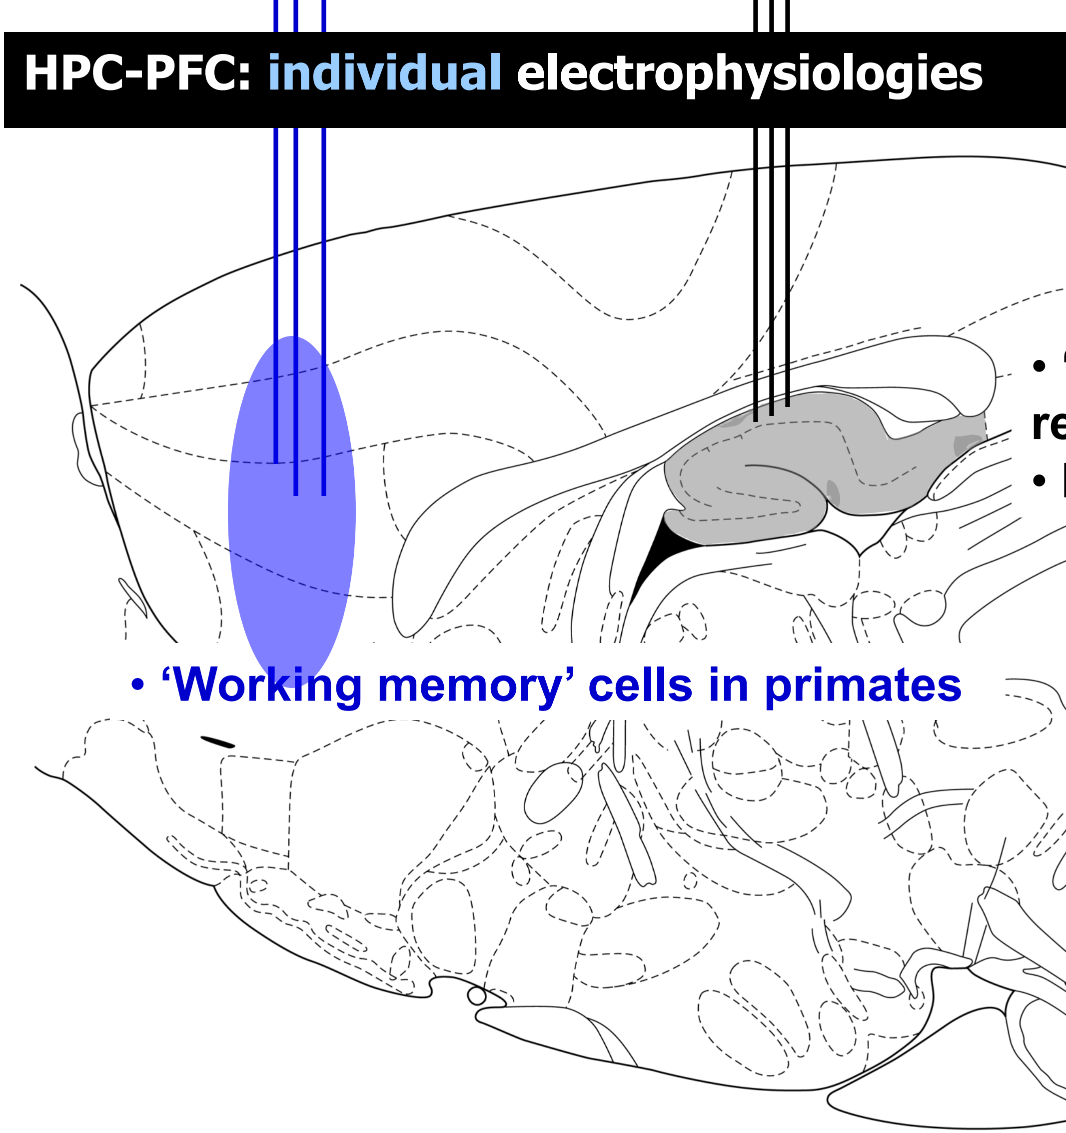

- ‘Place cells’: neurons with spatial receptive fields (‘place fields’)
- Rodents, primates, humans

- ‘Working memory’ cells in primates

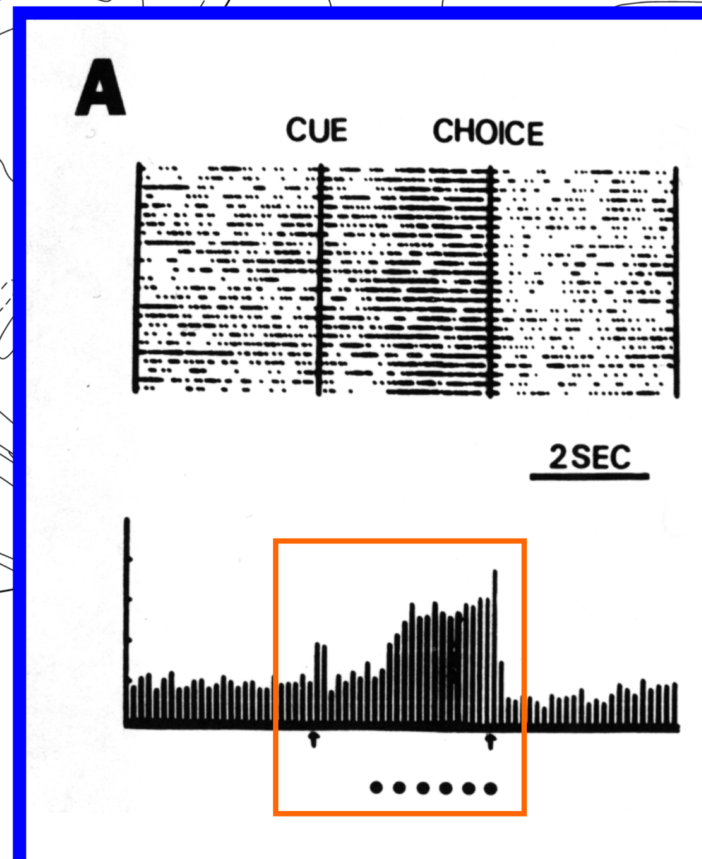

from Niki (1974) *Brain Res*. **70**, 346-349

Multiple units from multiple electrodes in multiple sites

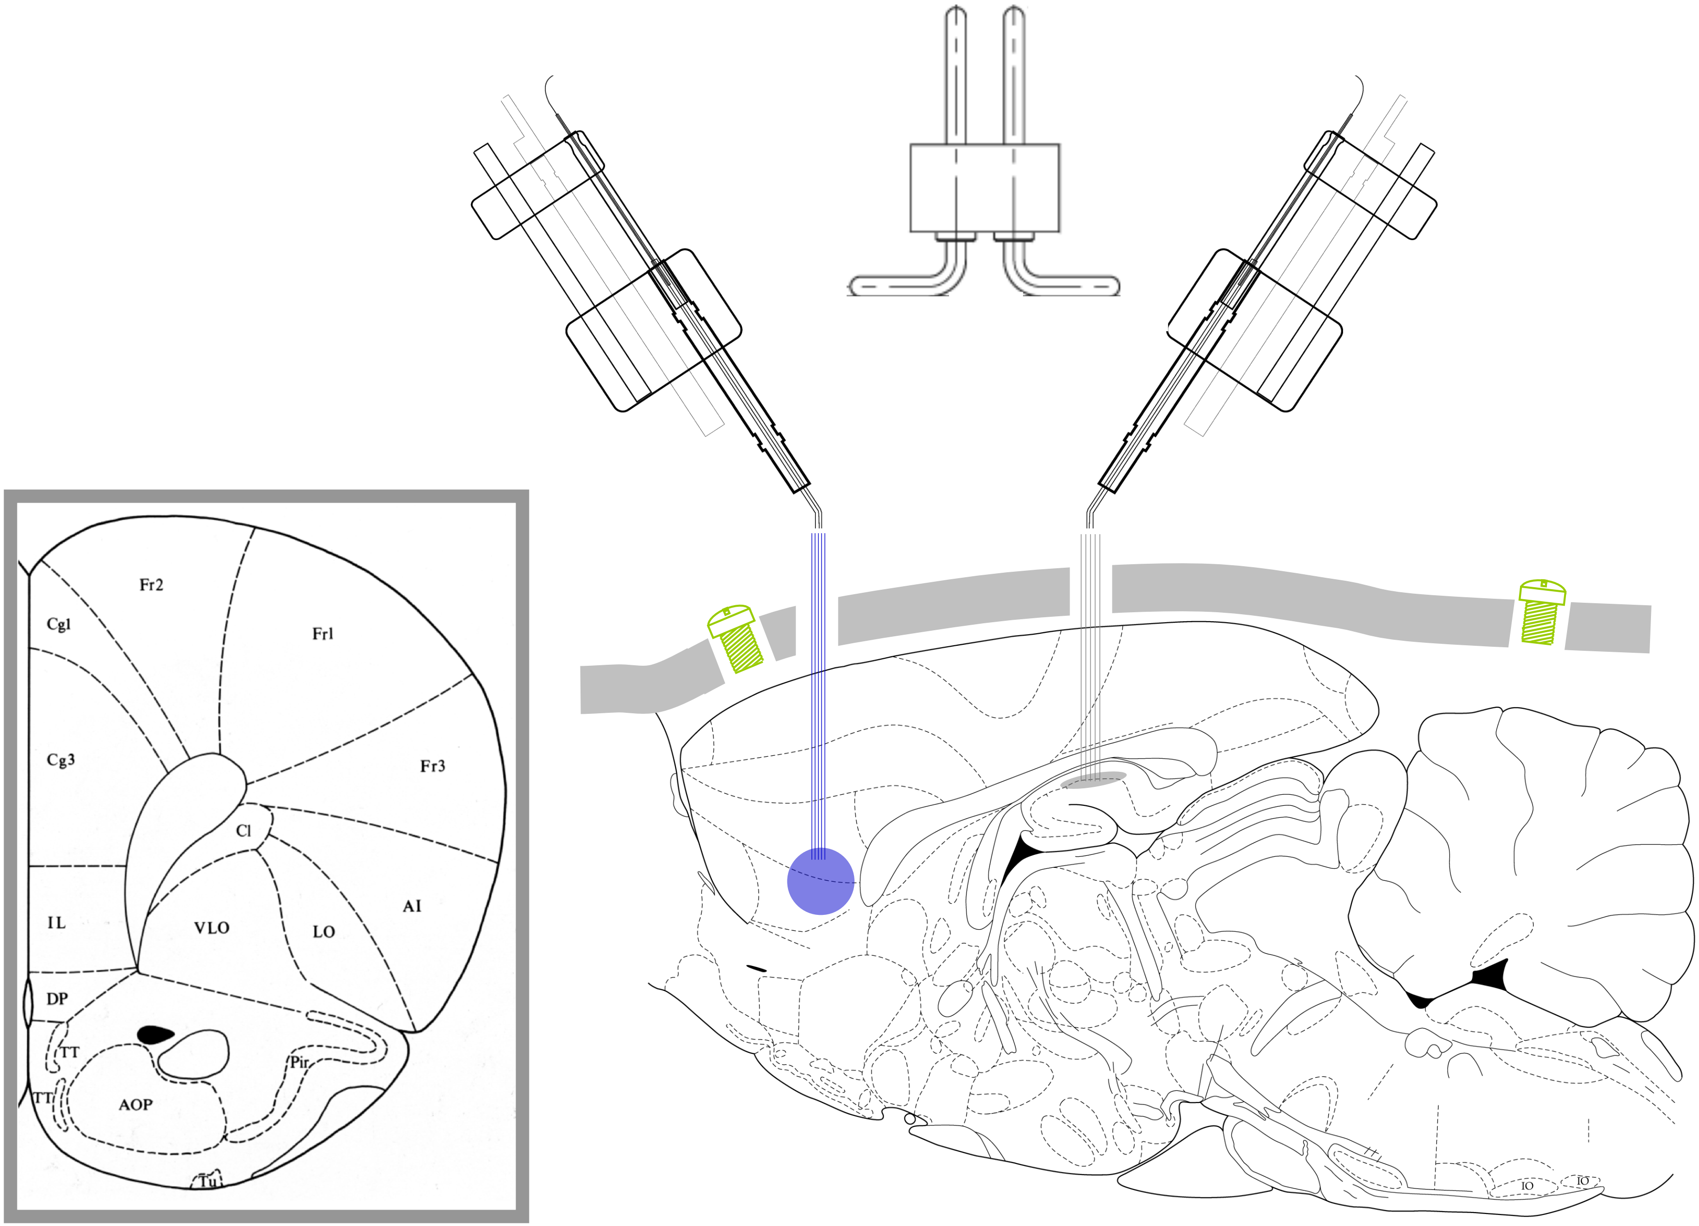

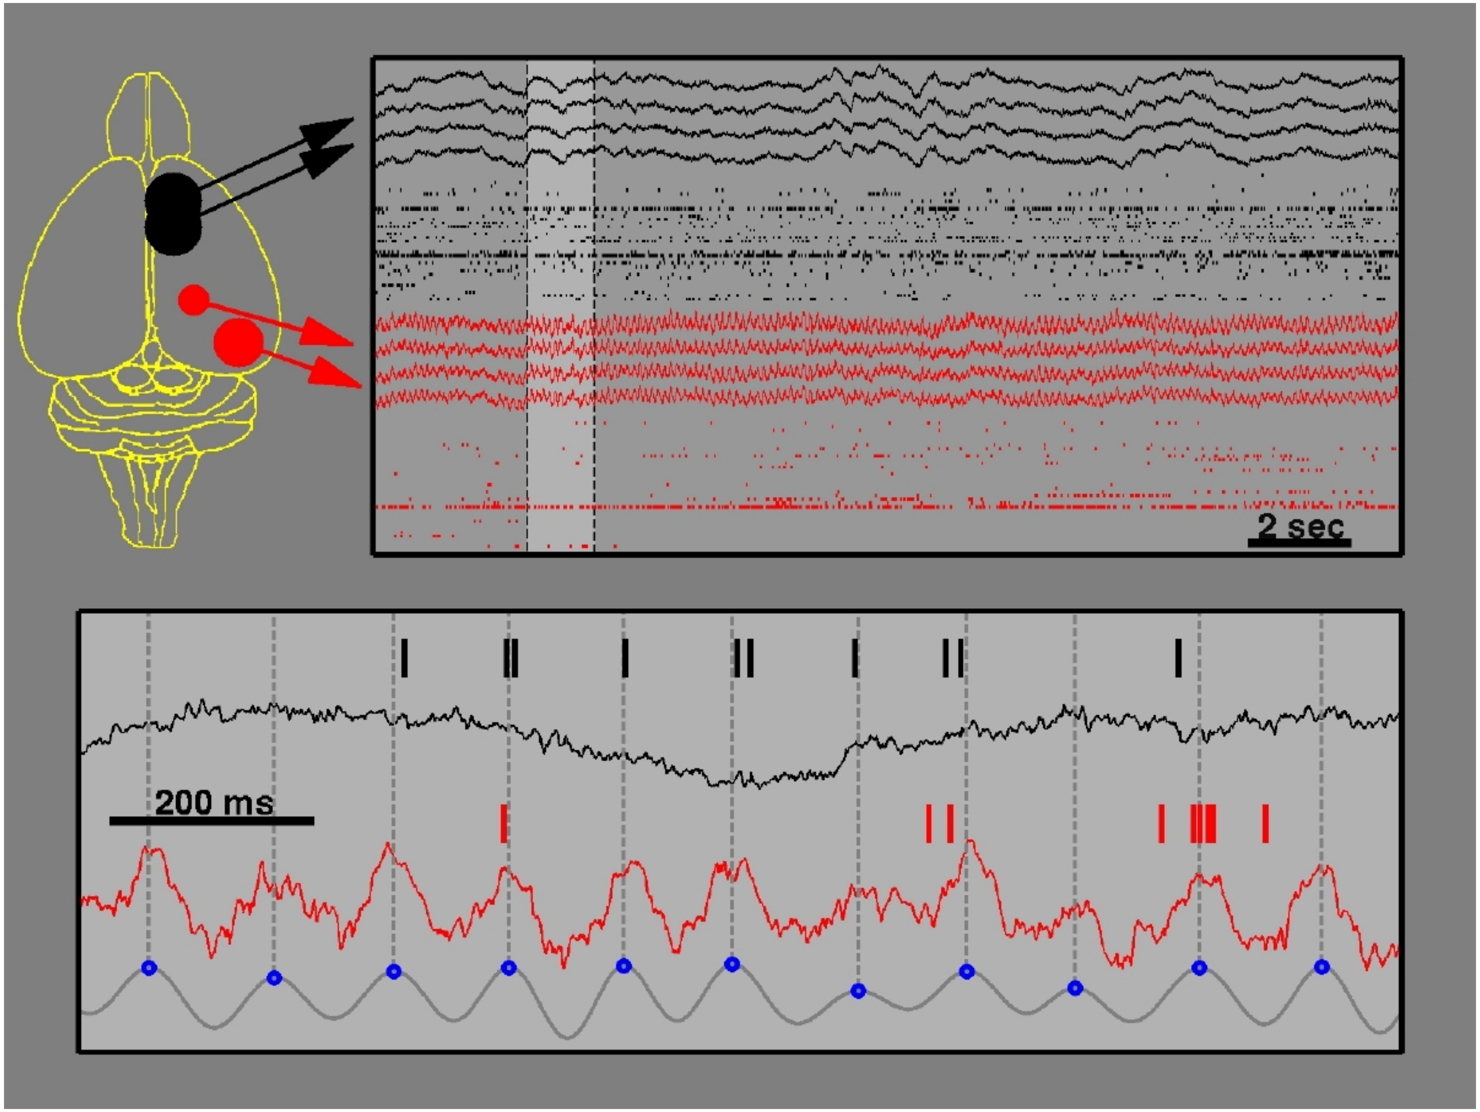
2 sec
200 ms

Data

Behaviour & Position

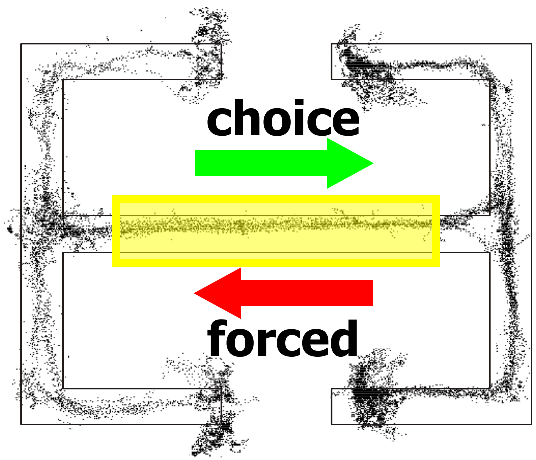

Extracellular Action Potentials (spikes)

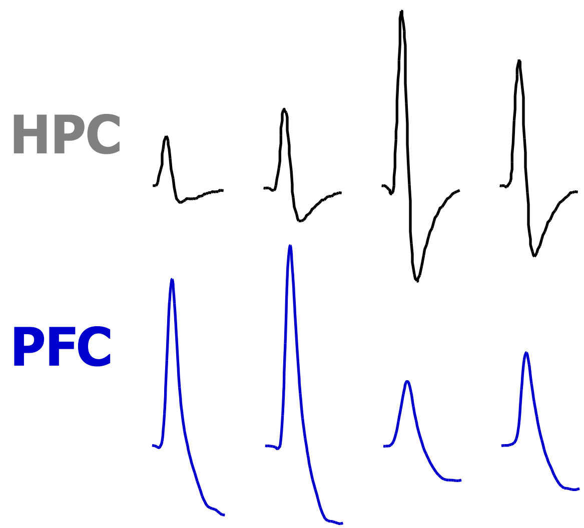

Local Field Potentials (LFP)

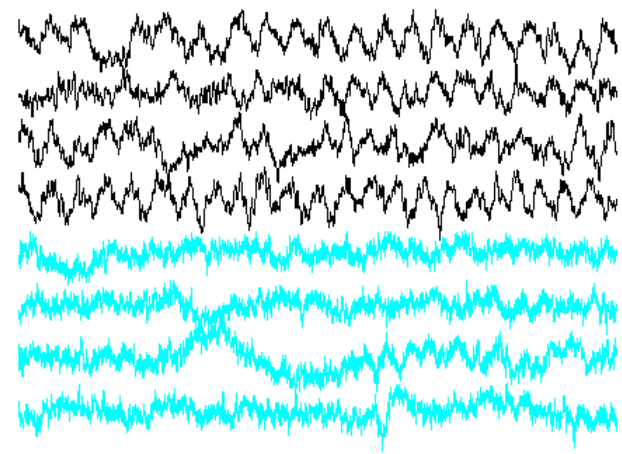

Interactions: spikes vs. LFP

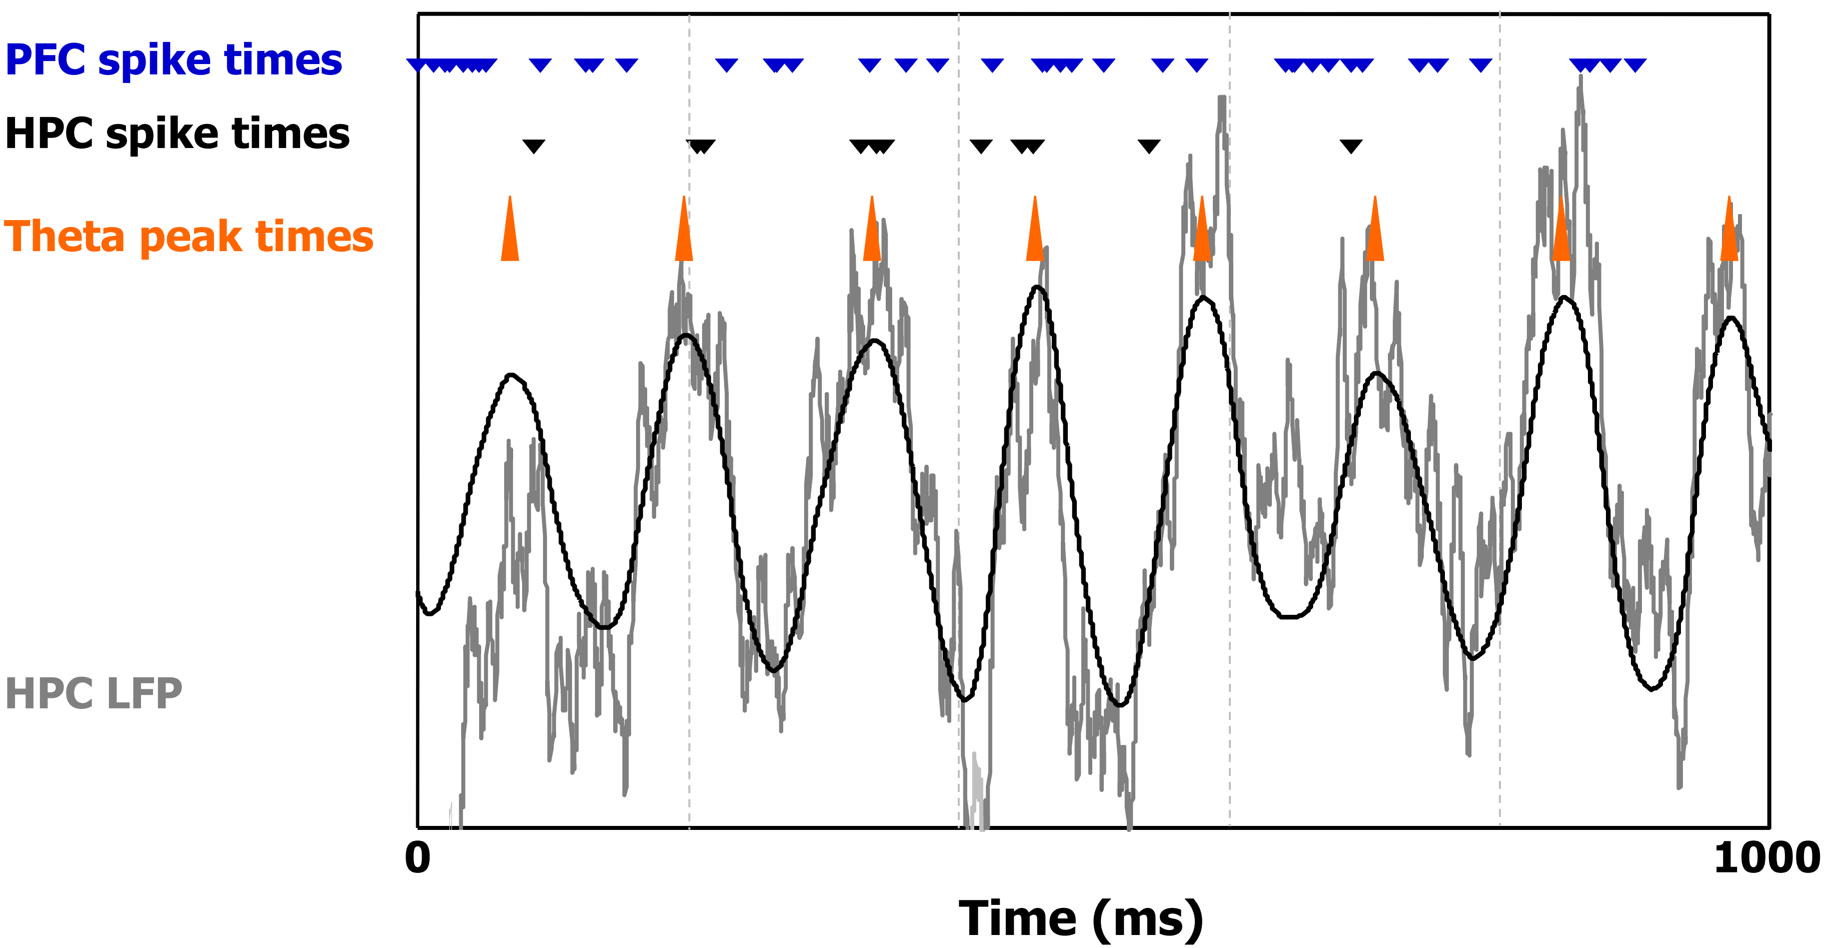

Enhanced theta-phase locking during ‘ correct choice’

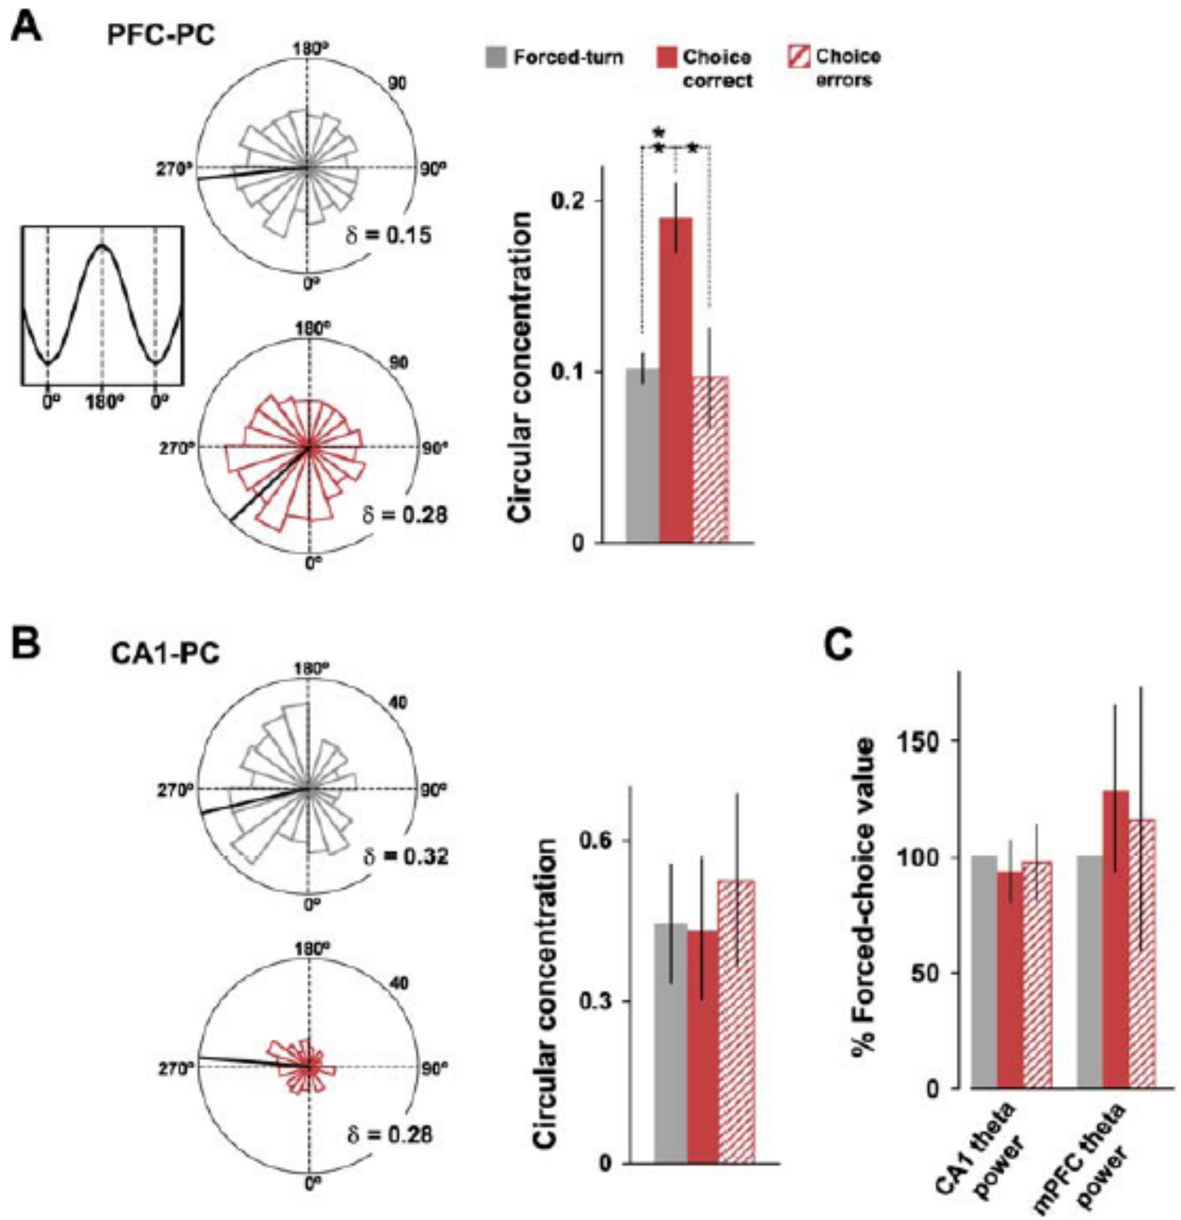

Jones and Wilson, *PLoS Biol.*, 2005

LFP *vs*. LFP: Coherence

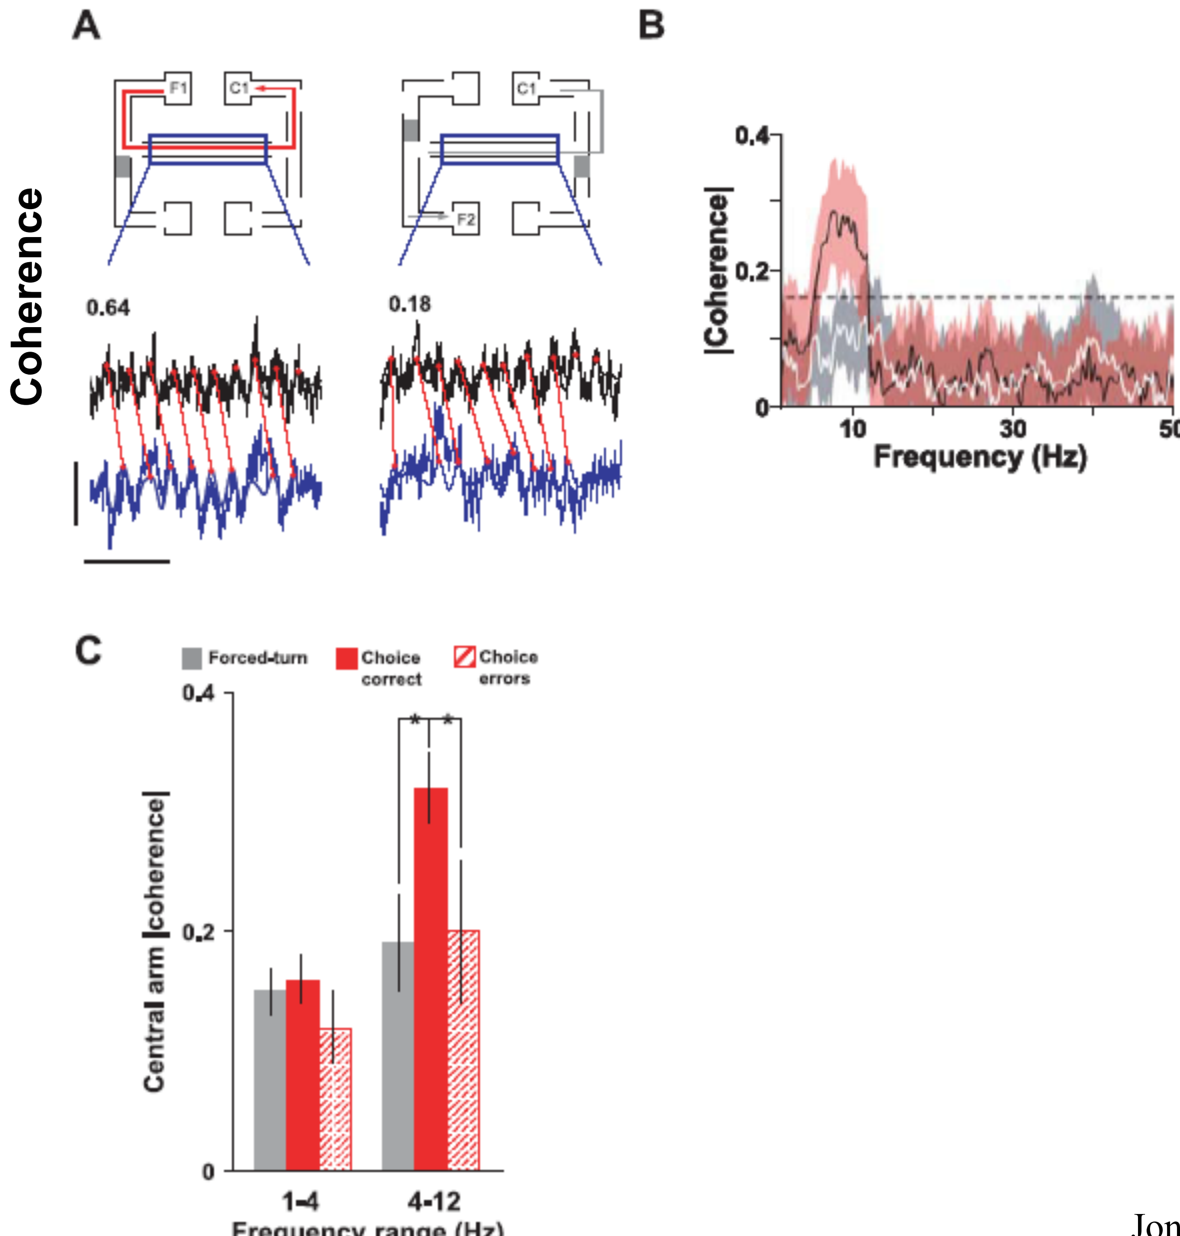

Jones, M. et al. *PLoS* 2005

Optogenetic manipulation of hippocampal inhibitory cells

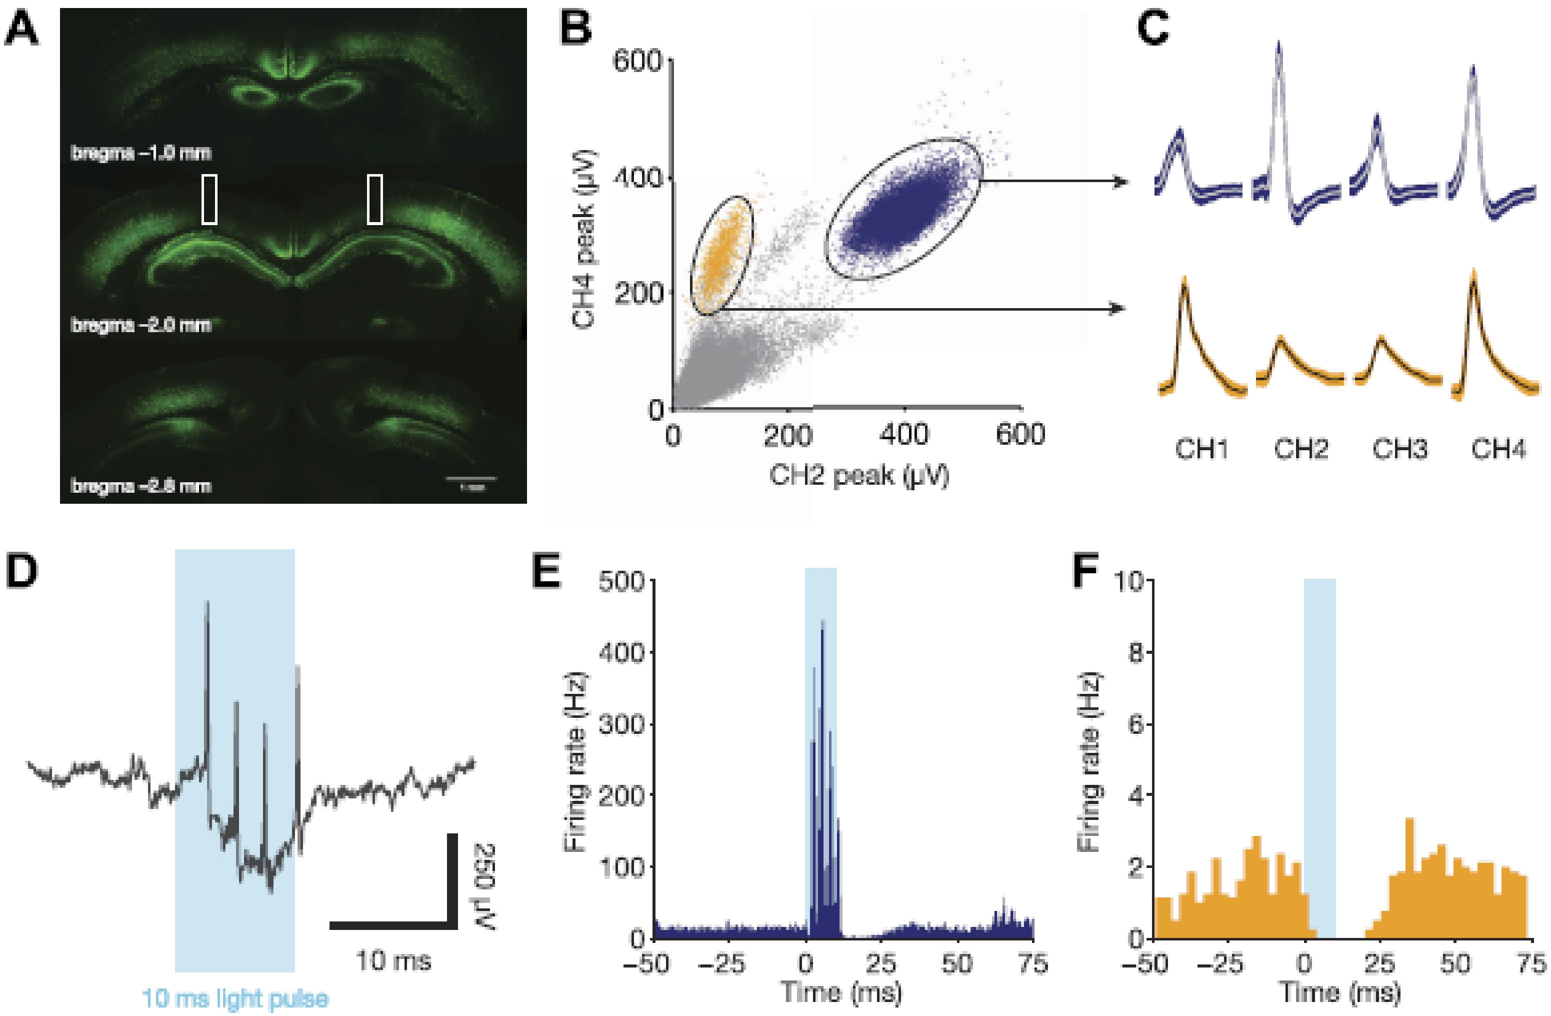

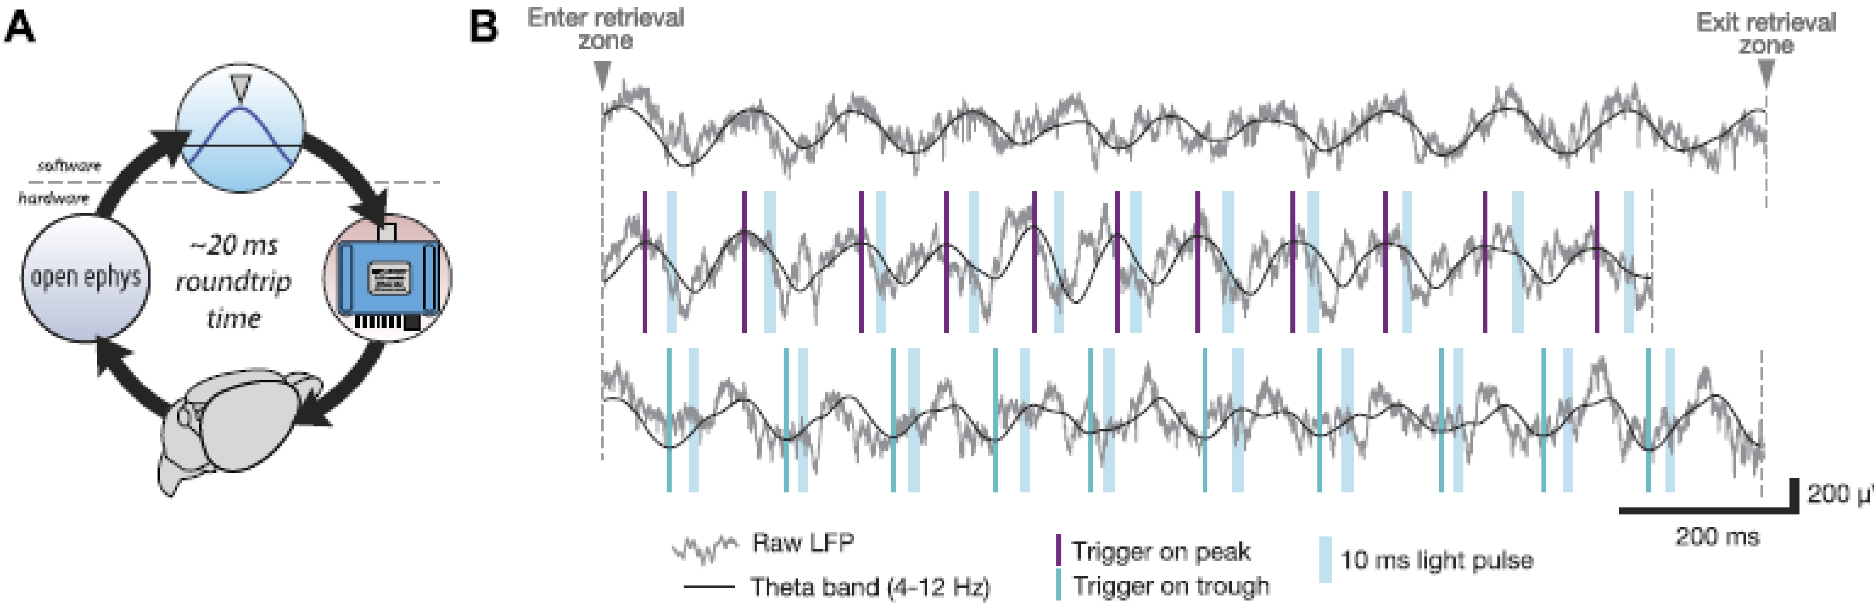

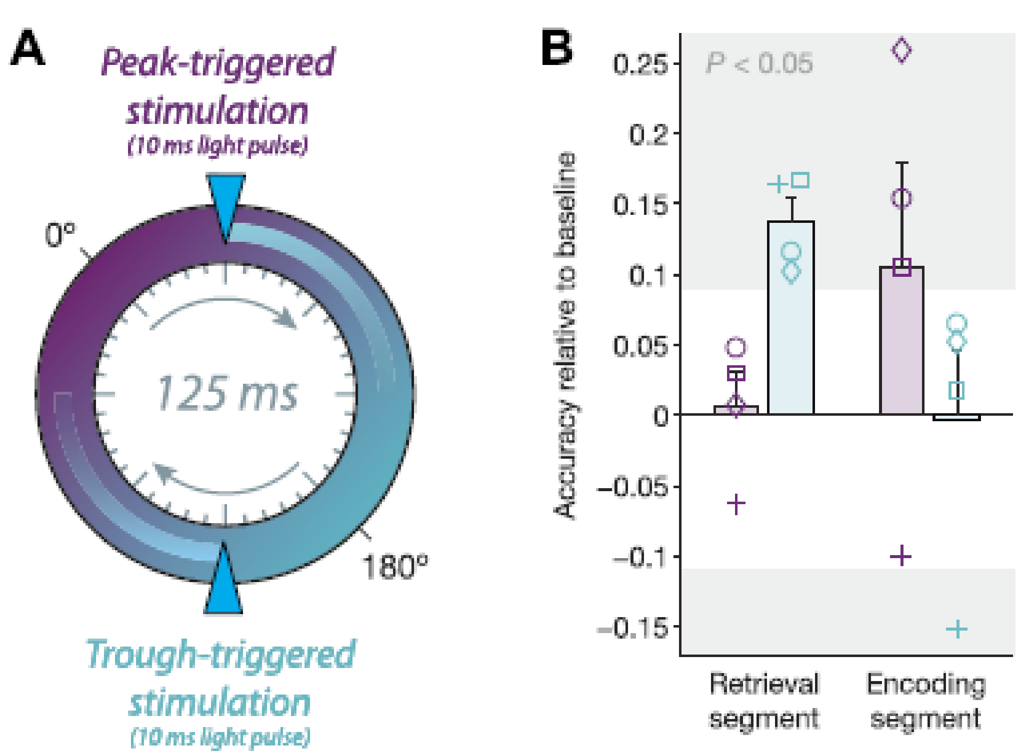

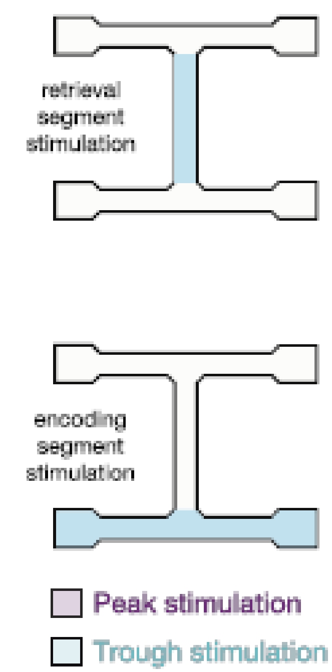

Siegle and Wilson, *eLIFE*, 2014

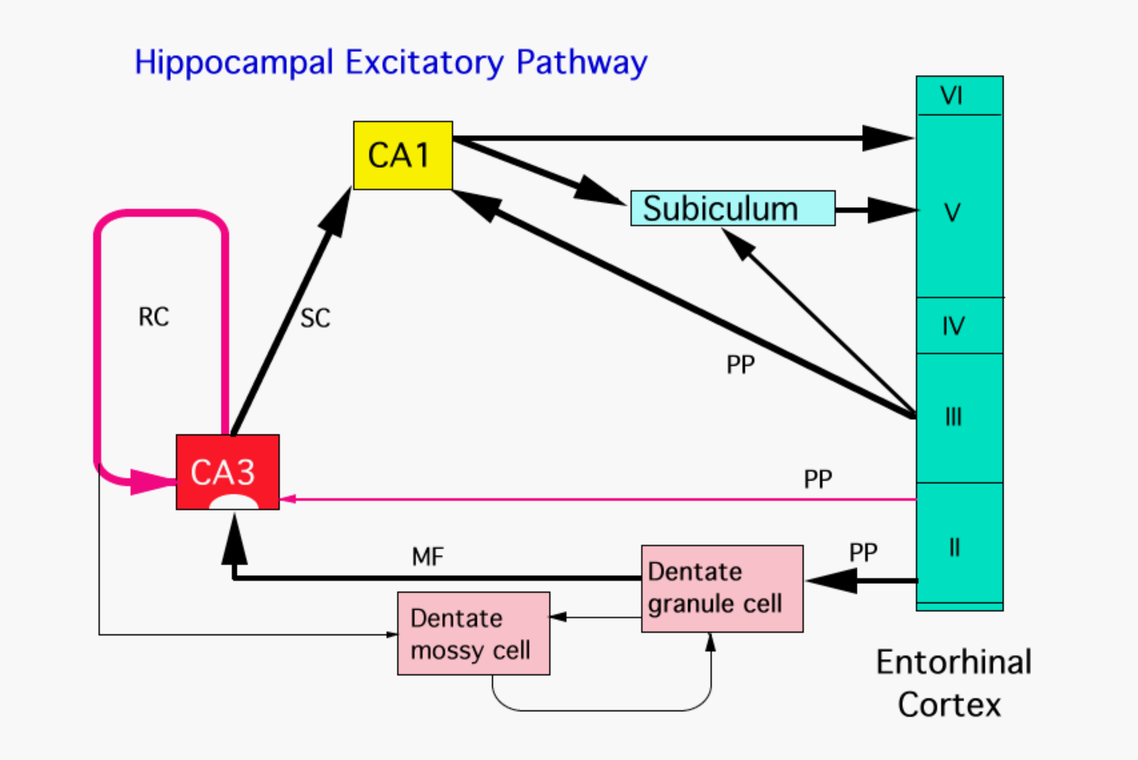
Hippocampal Excitatory Pathway
CA1
Subiculum
VI
V
IV
III
II
RC
SC
PP
PP
PP
CA3
MF
Dentate granule cell
Dentate mossy cell
Entorhinal Cortex

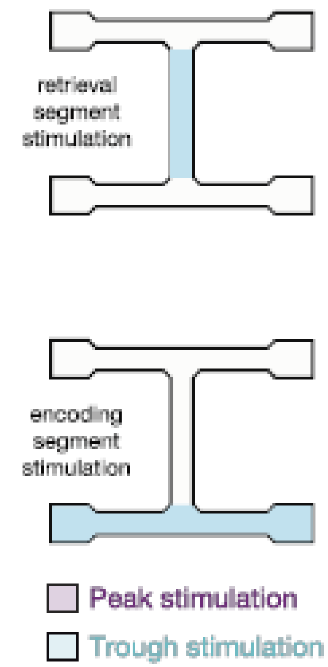
retrieval segment stimulation
encoding segment stimulation
Peak stimulation
Trough stimulation

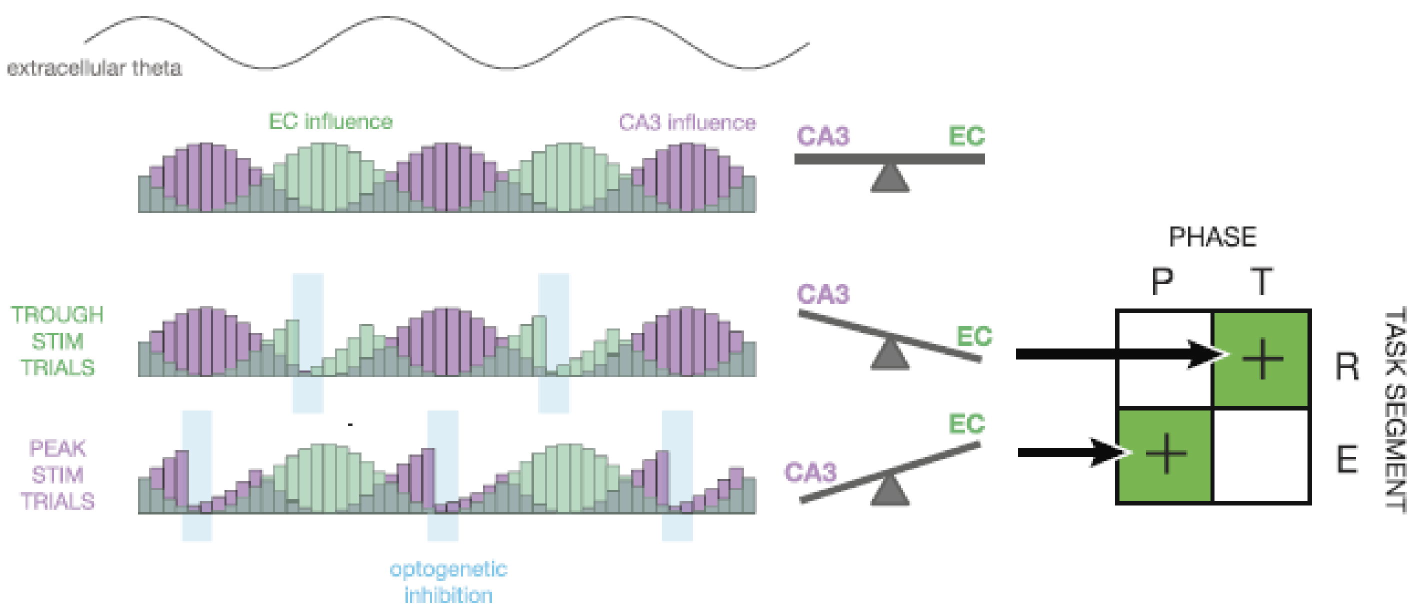
extracellular theta
EC influence
CA3 influence
CA3
EC
TROUGH STIM TRIALS
CA3
EC
PEAK STIM TRIALS
EC
CA3
optogenetic inhibition
PHASE
P
T
R
E
TASK SEGMENT

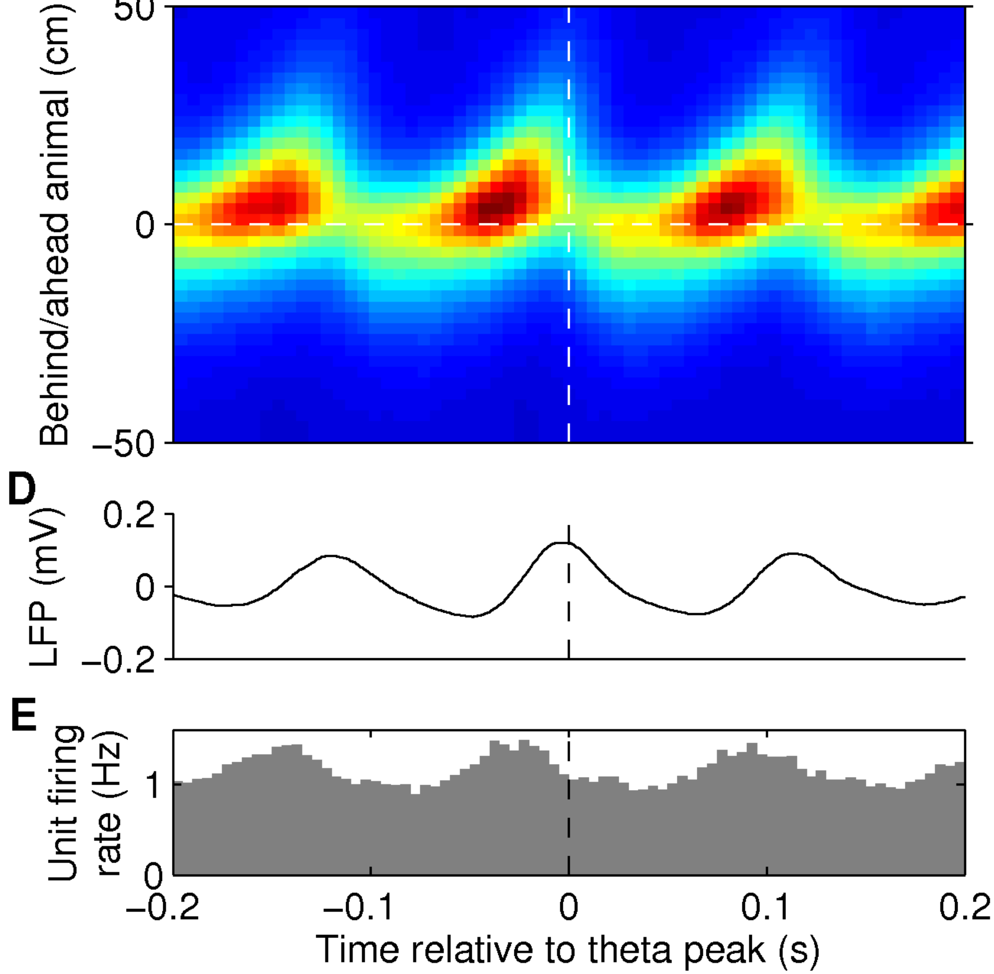

Behind/ahead animal (cm)
50
0
−50
D
LFP (mV)
0.2
0
−0.2
E
Unit firing rate (Hz)
1
0
−0.2
−0.1
0
0.1
0.2
Time relative to theta peak (s)

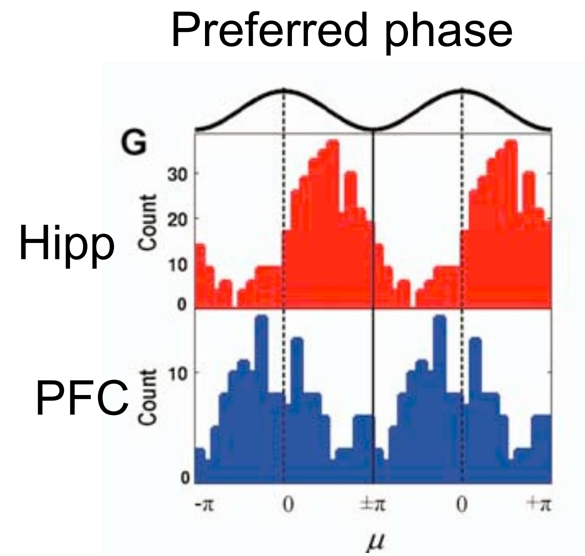

Preferred phase
G
Hipp
PFC
Count
30
20
10
0
-π
0
±π
0
+π
μ

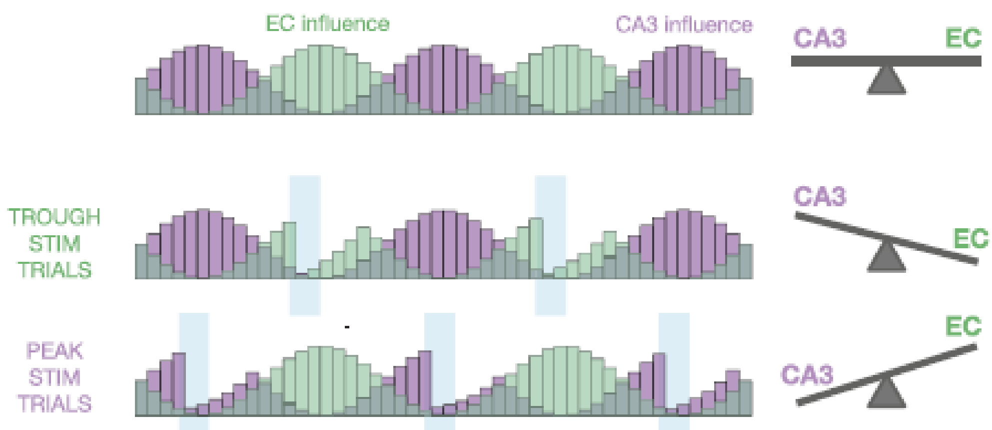

EC influence
CA3 influence
CA3
EC
TROUGH STIM TRIALS
PEAK STIM TRIALS

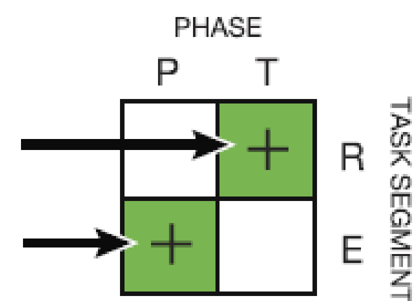

PHASE
P
T
R
E
TASK SEGMENT

Role of Sleep in Memory

- Sleep allows examination of memory independent of behavior.
- The formation of lasting memories may involve the communication of information between brain areas during sleep.
- Broadly identify two stages of non-REM sleep –(NREM) and rapid eye movement sleep (REM).

Experimental design

Compressed Run sequences are expressed in hippocampus during nREM sleep

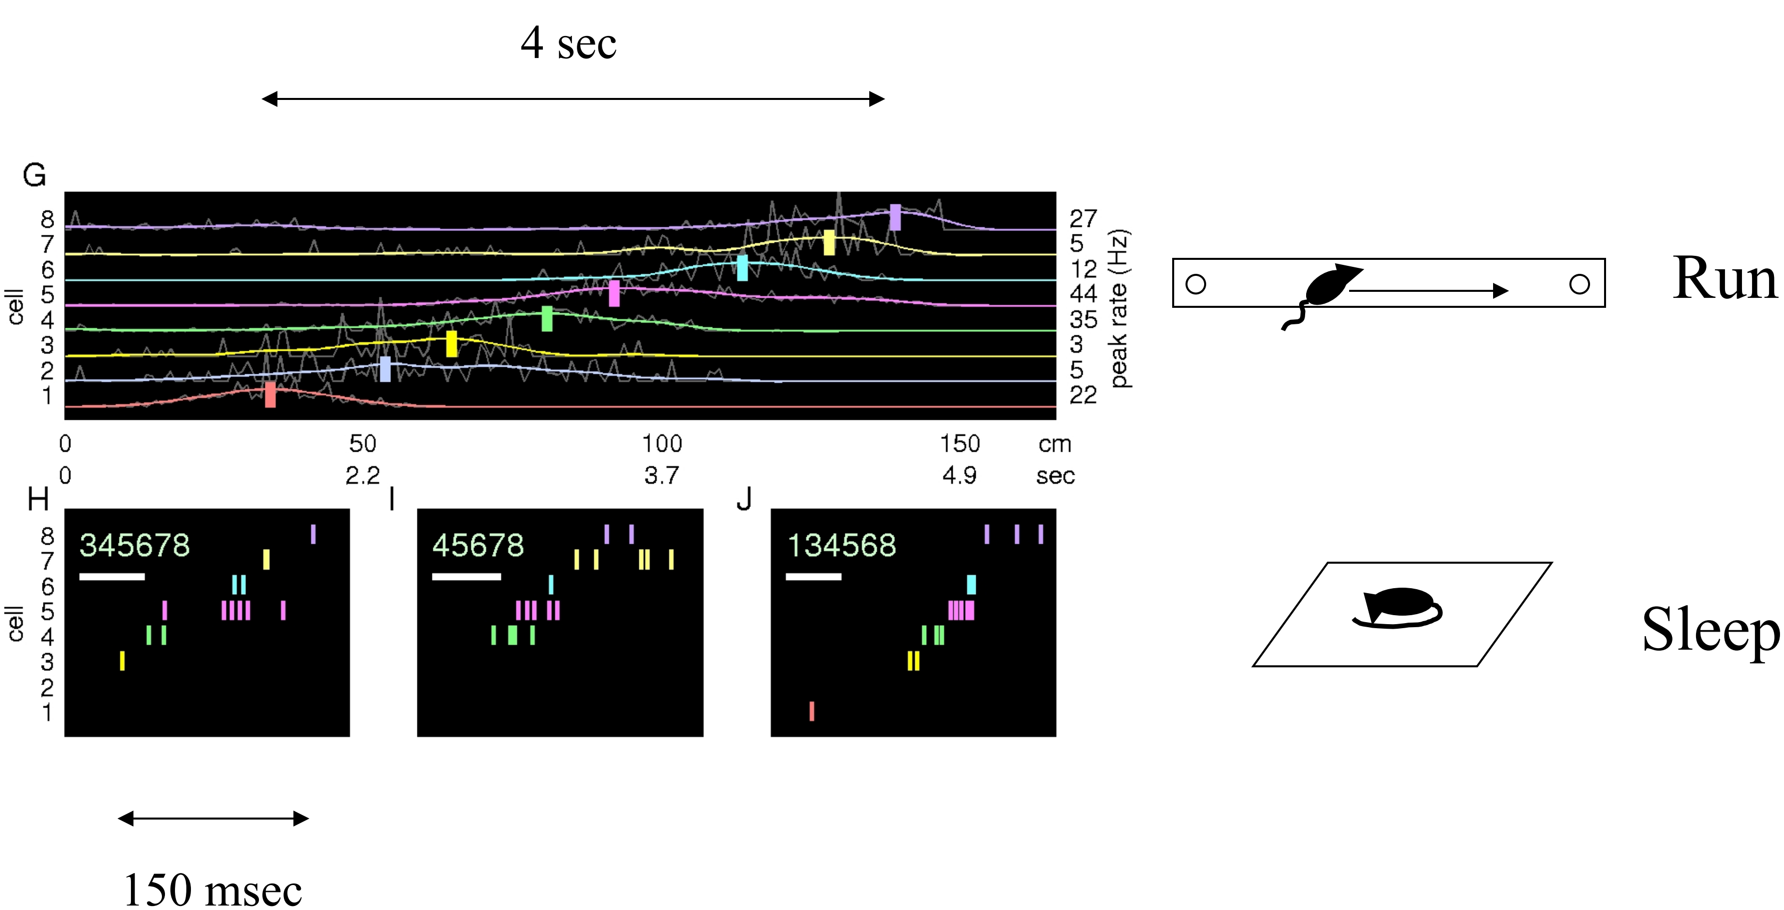

Lee and Wilson, *Neuron*, 2002

Sequences are re-expressed during CA1 ripple events

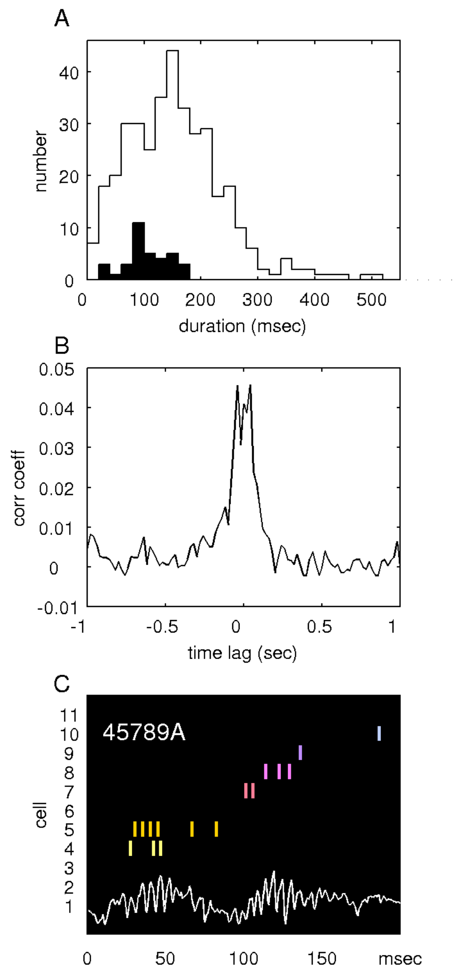

Duration of low probability sequences

Correlation of low probability sequences and ripples

Example of a low probability sequence and a ripple event

Overlapping asymmetric place fields with oscillatory variation in excitability translate behavioral time relationships to biophysical timescales with preserved temporal order

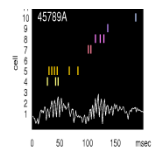

Direction of movement

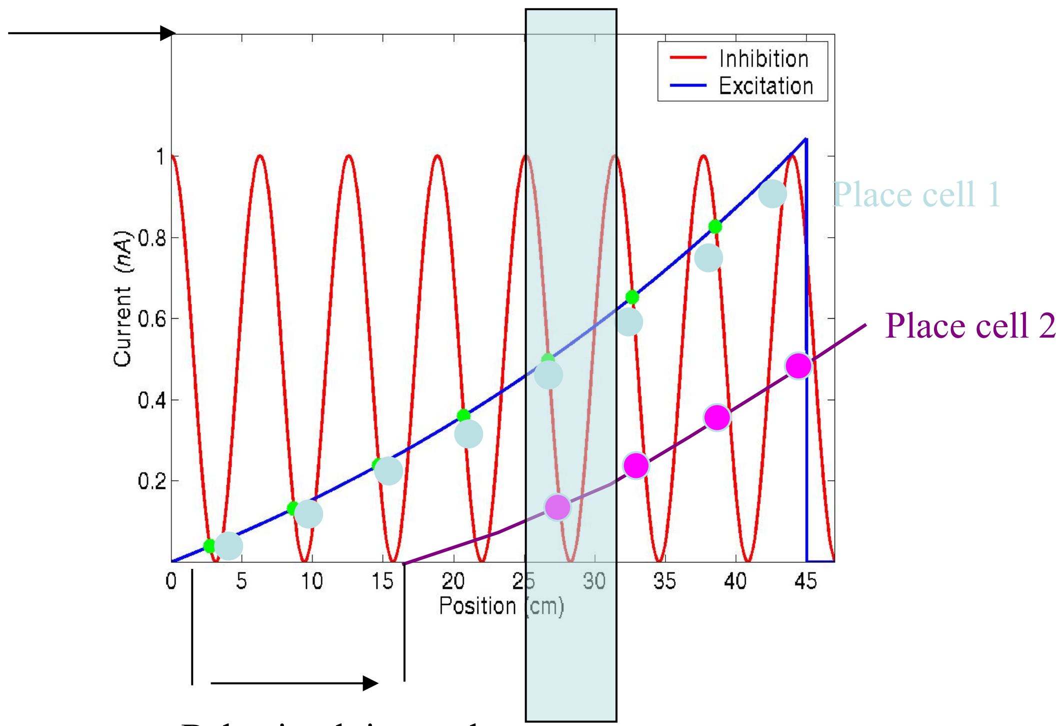

Behavioral timescale

Are there signatures of memory reactivation in the neocortex during hippocampal reactivation

- Simultaneously record in the hippocampus and primary and secondary visual cortex during spatial behavior.
- Look for reactivation in both structures during sleep.

A

Experimental Design:

PRE (1-2hrs) — RUN (20-40mins) — POST (1-2hrs)

C

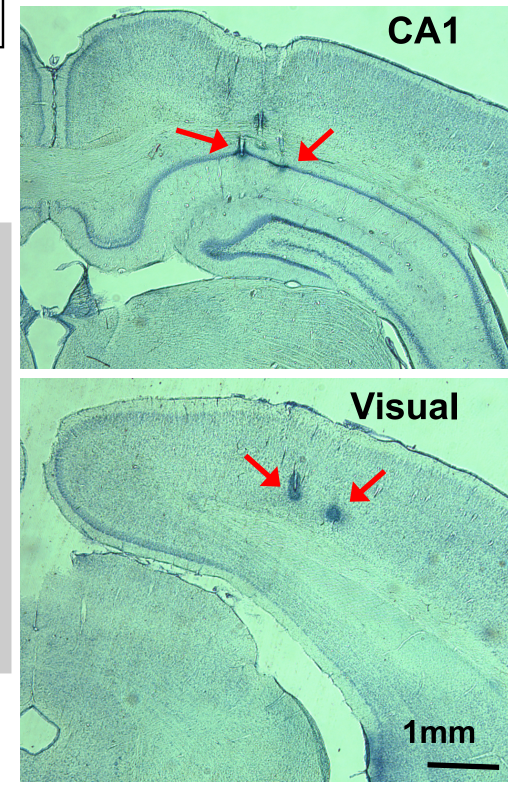

B

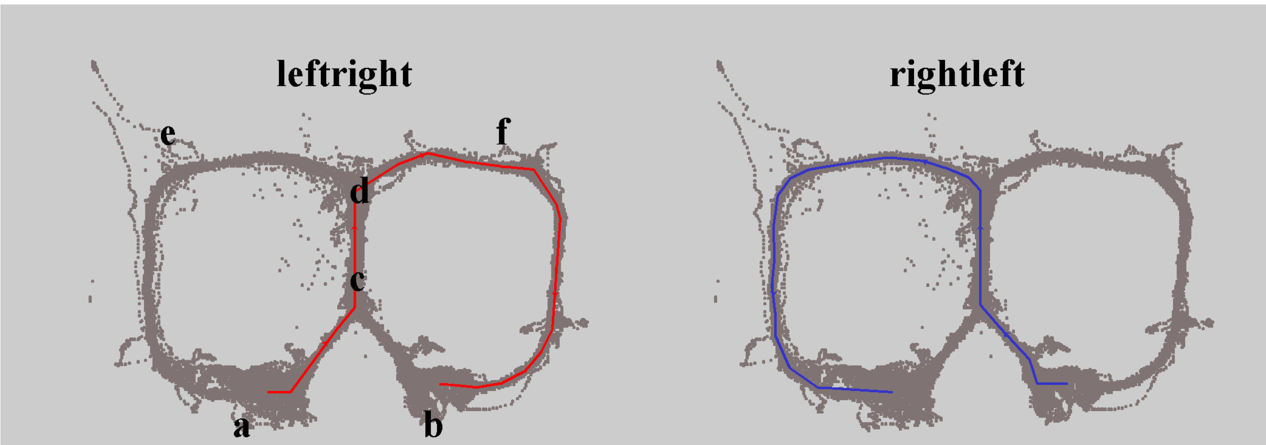

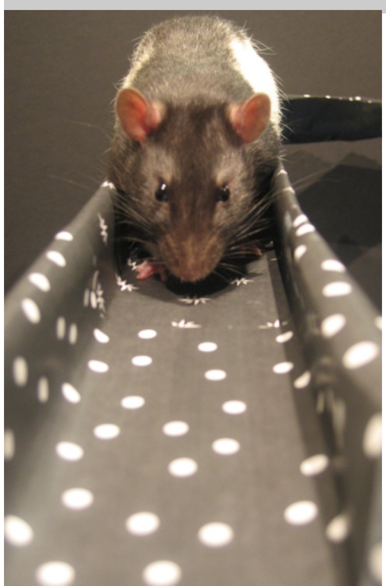

1. Intra-maze local cues, no prominent distal cues
2. Well trained animals: alternation task
3. Recording sites: visual cortex (Occ1, Occ2) and CA1
4. Sleep states (SWS, REM, Wake, Int) classified using EMG and hippocampal EEG

Sequence memory reactivation in hippocampus and visual cortex

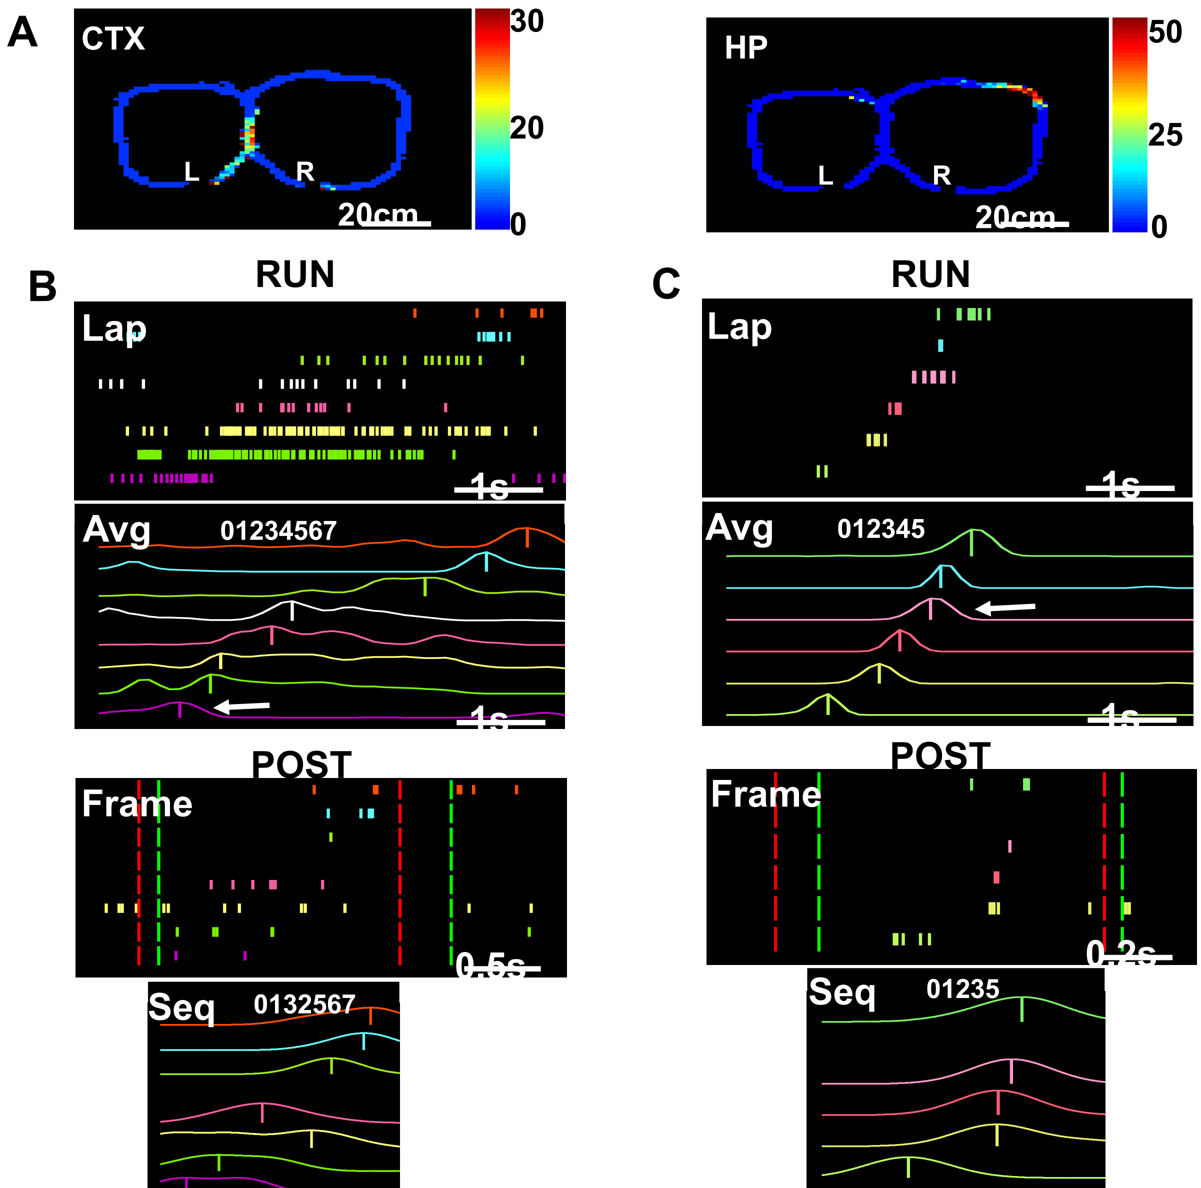

Ji and Wilson, *Nature Neuroscience*, 2007

Reactivation occurs during activity frames correlated with the slow oscillation

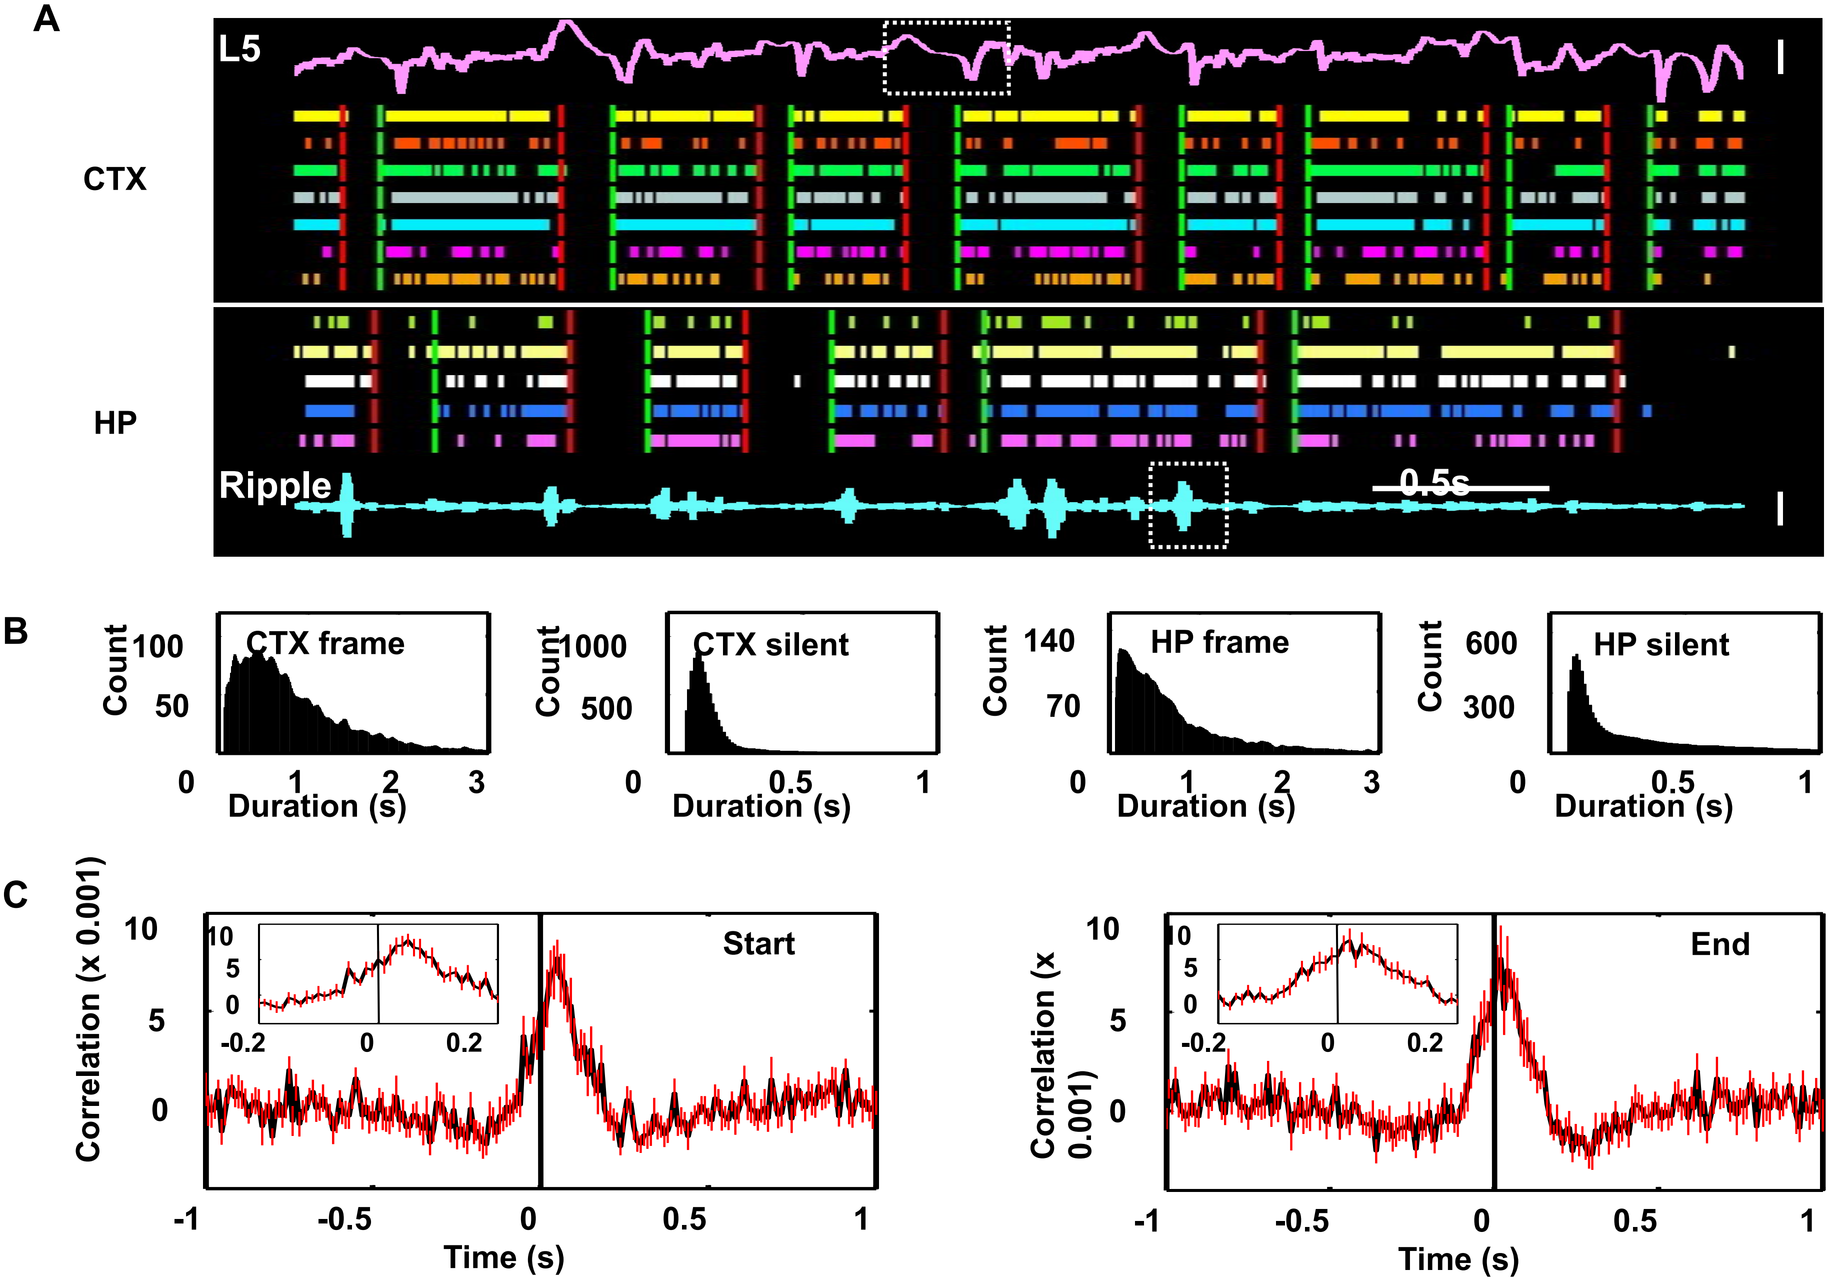

Can we influence memory reactivation during sleep?

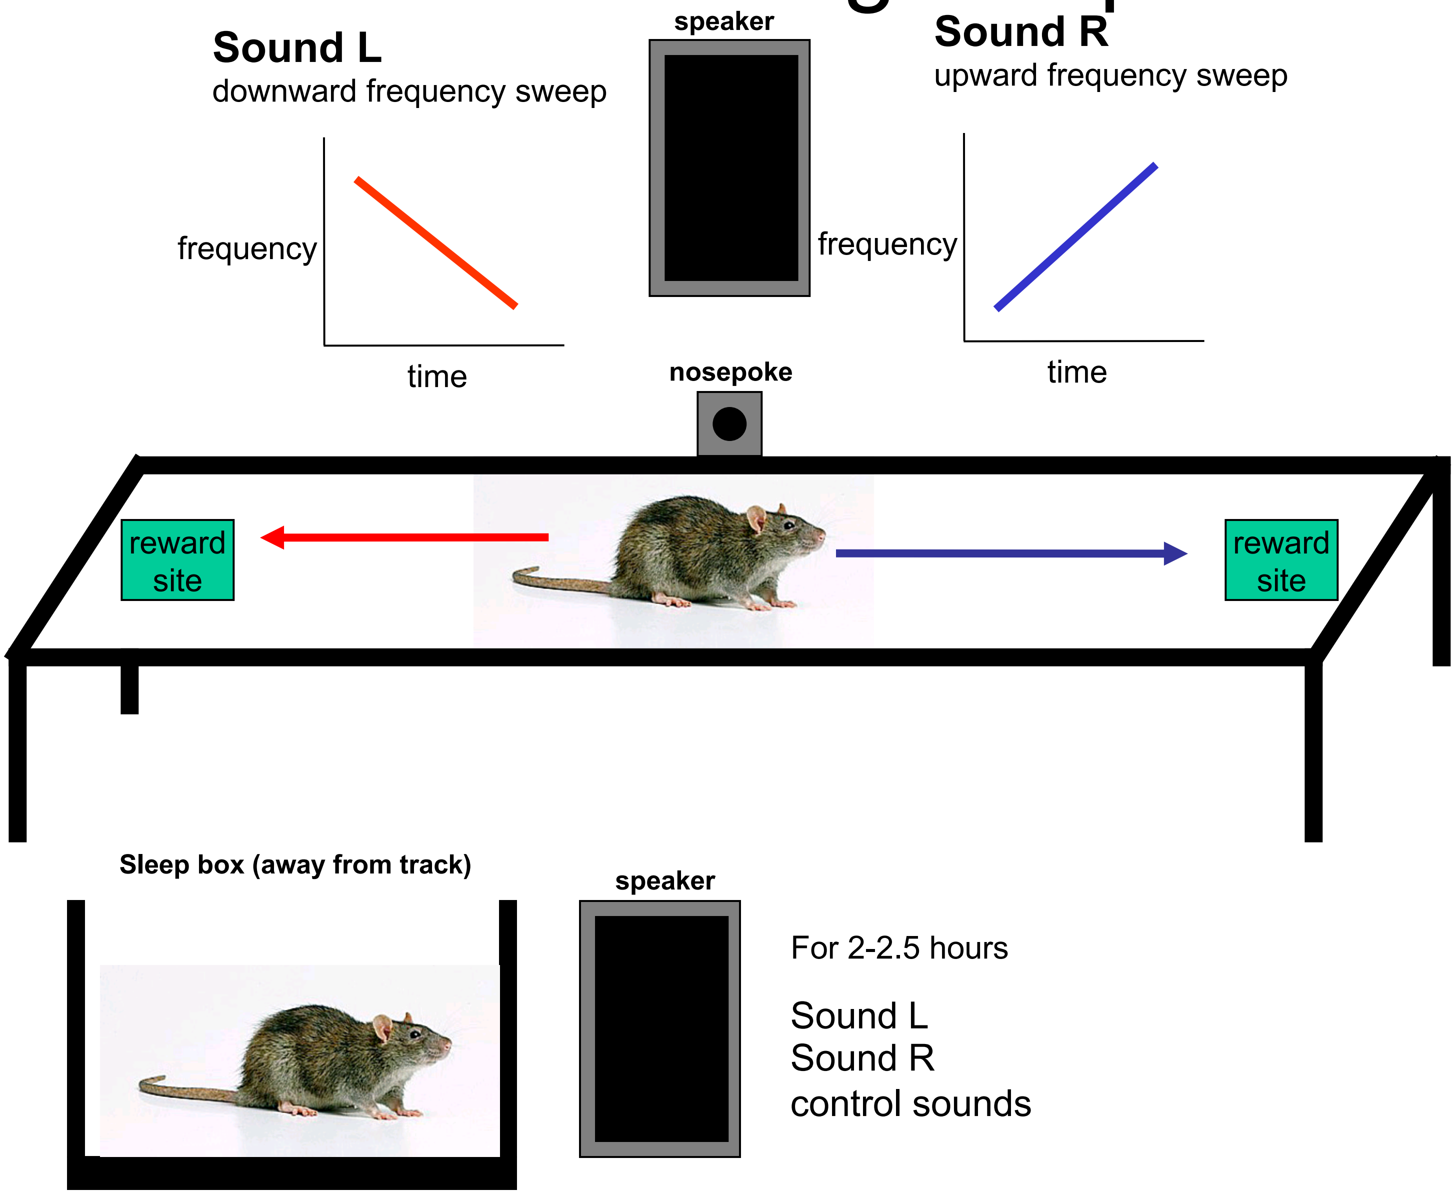

Behavioral task design

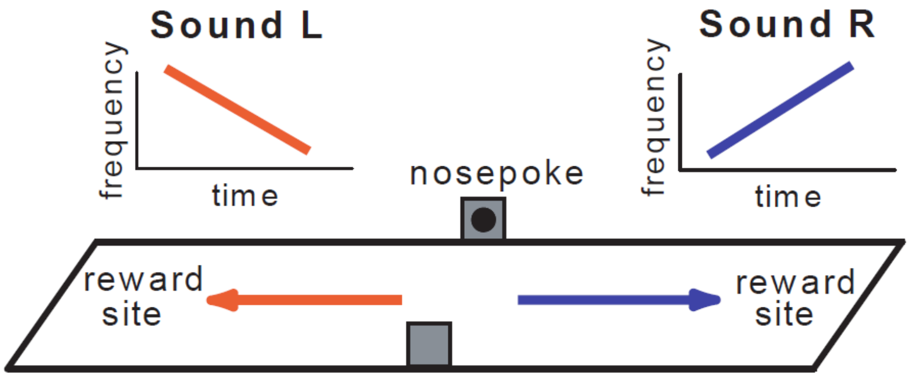

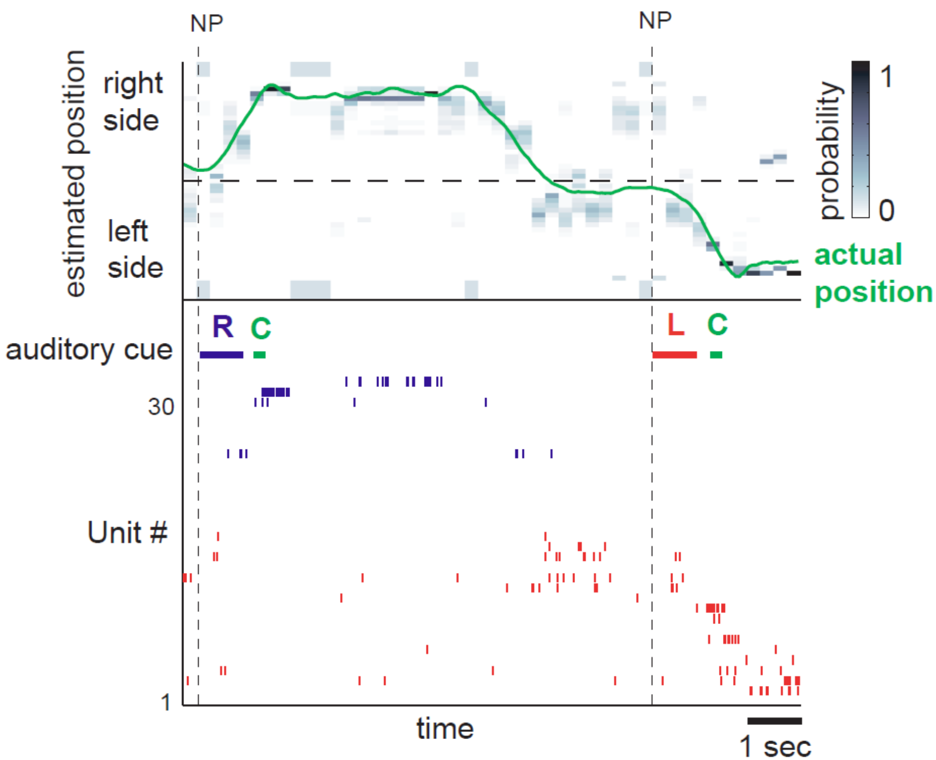

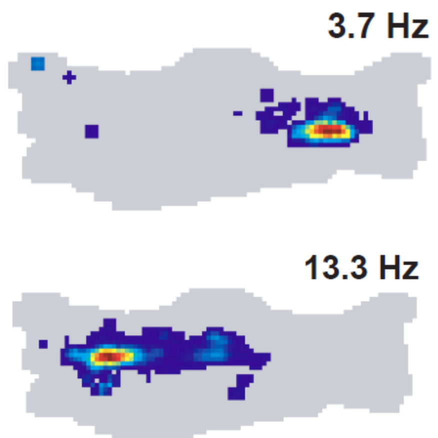

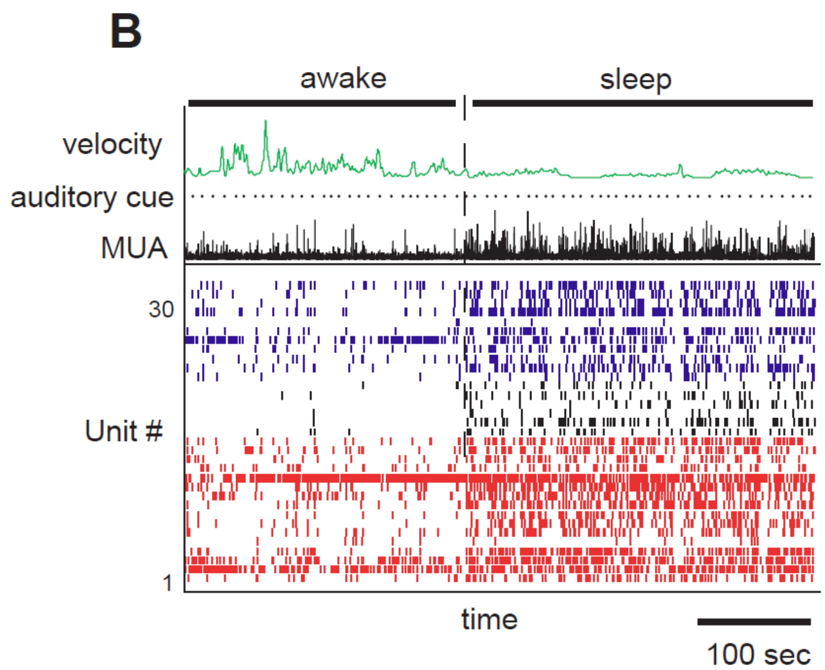

Do task-related sounds bias the content of future replay?

Hypothesis:

Sound R- place cells with **right-sided** place fields are more active during replay

Sound L- place cells with **left-sided** place fields are more active during replay

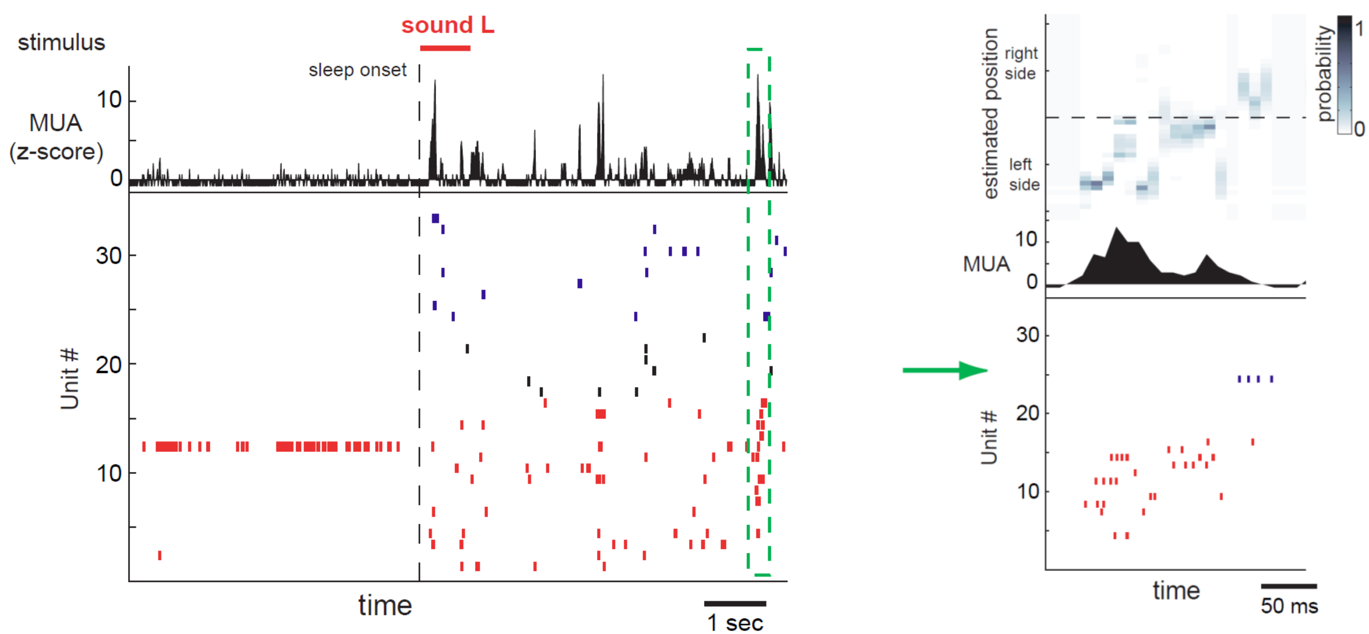

Bias observed in individual place cell responses

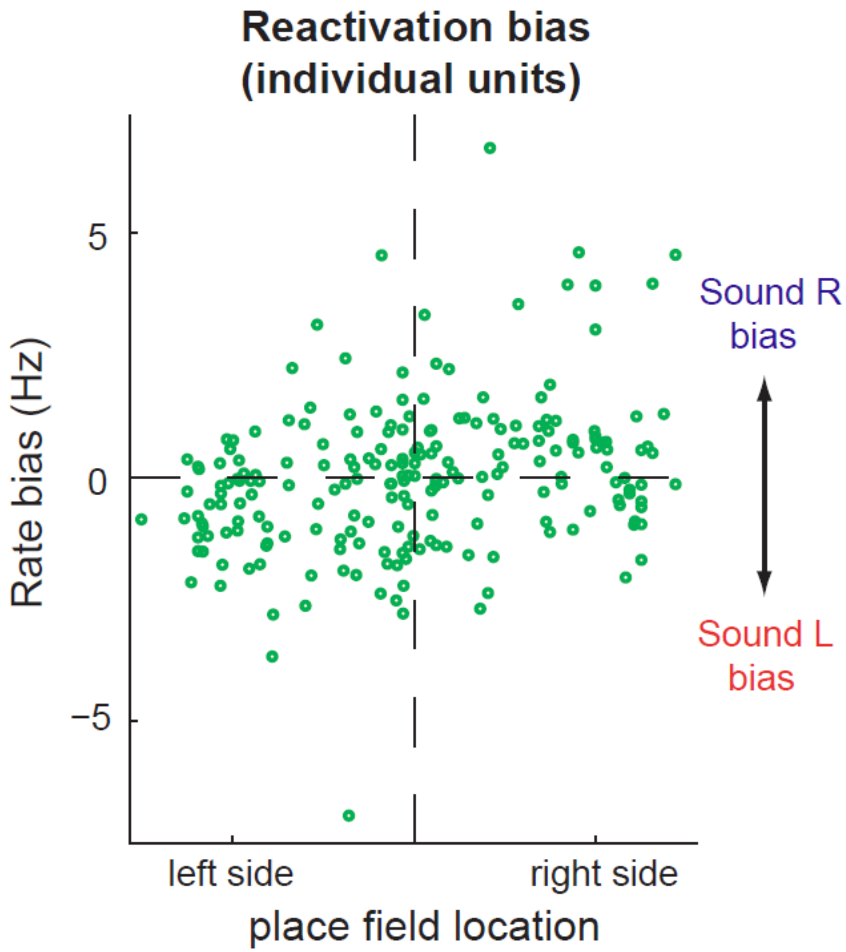

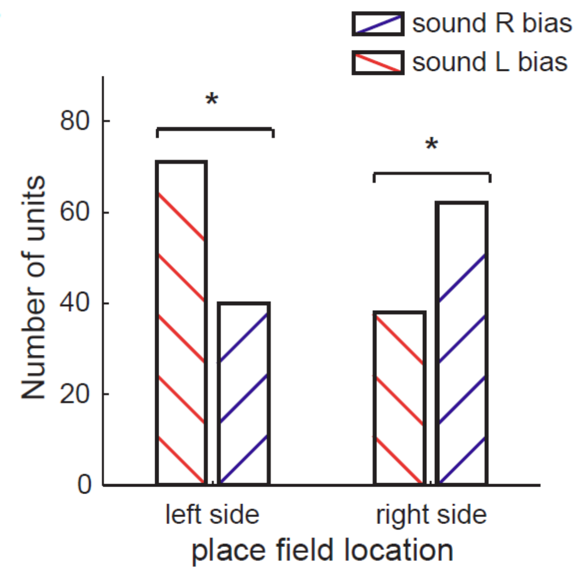

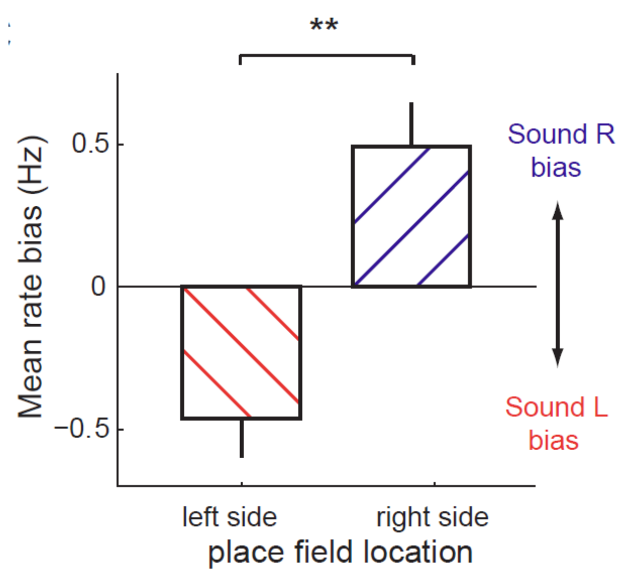

Bendor and Wilson, *Nature Neuroscience*, 2012

Bias is maintained after initial cueing

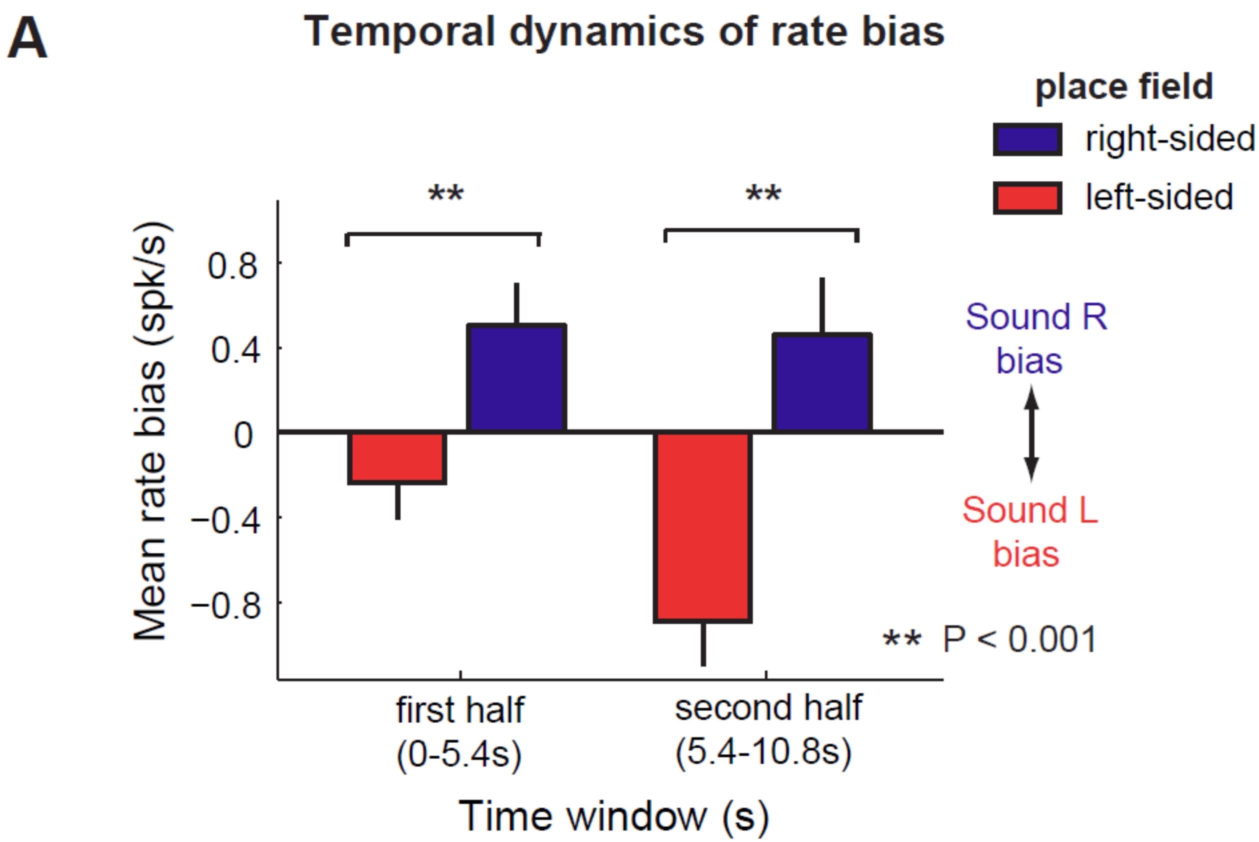

Bendor and Wilson, *Nature Neuroscience*, 2012

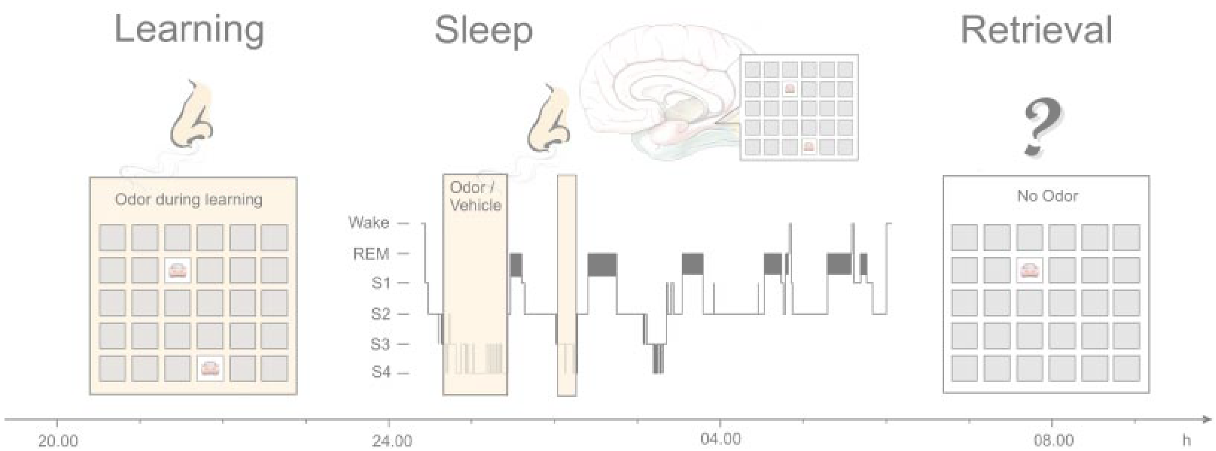

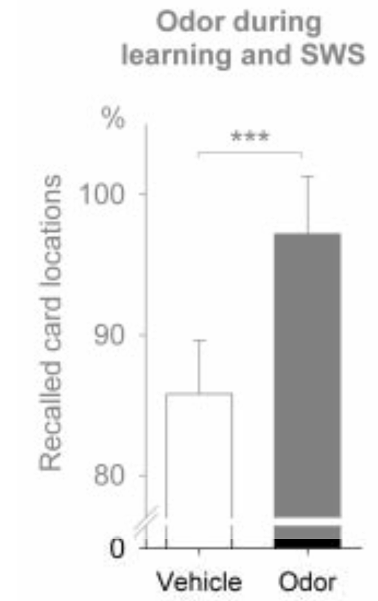

Rasch et al. 2007

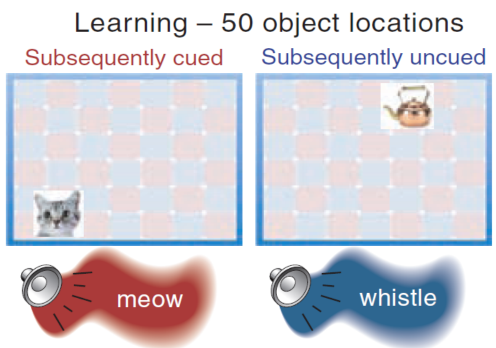

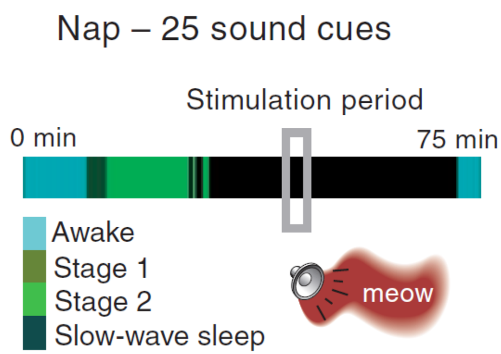

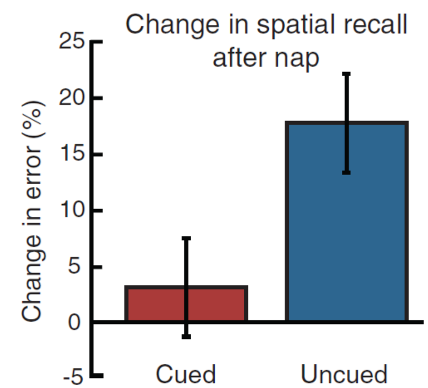

Rudoy et al. 2009

stimulation → Cortex

Bias pre-replay brain state

Cortex ⇄ Hippocampus

Bias which memories are transferred

Hippocampus online and offline

Theta rhythm

Sharp wave/ripples

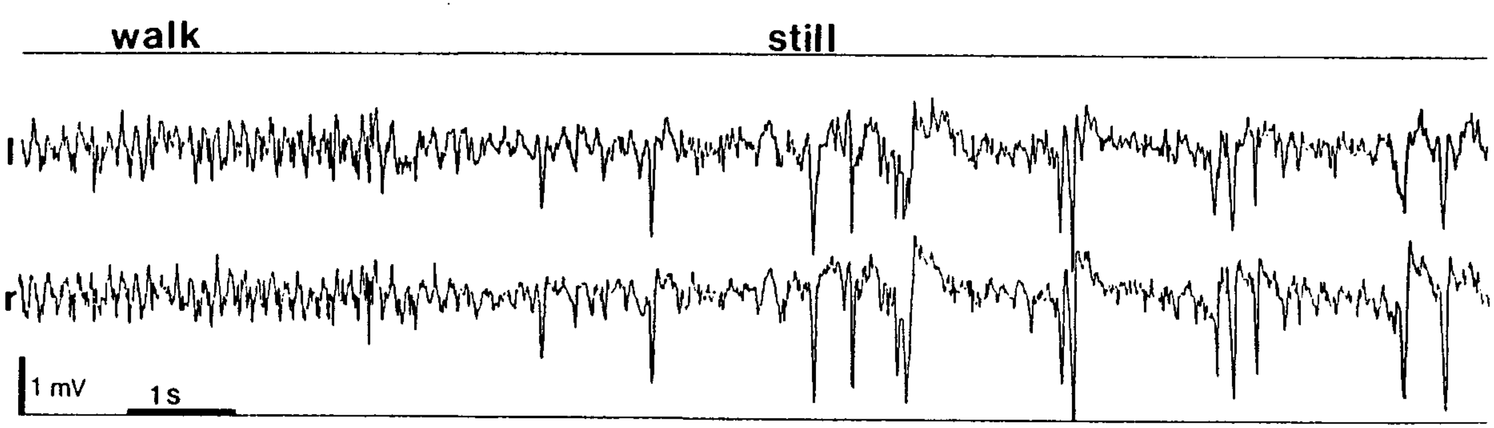

Buzsáki 1989

Hippocampal activity during quiet wakefulness

- During awake behavior, there are periods of quiet wakefulness that have EEG that is similar to NREM consisting of brief bursts of activity modulated by high frequency “ripple” oscillations.
- Is there structure to the patterns of multiple single neuron activity during this state?

Does sequence reactivation occur during quiet wakefulness?

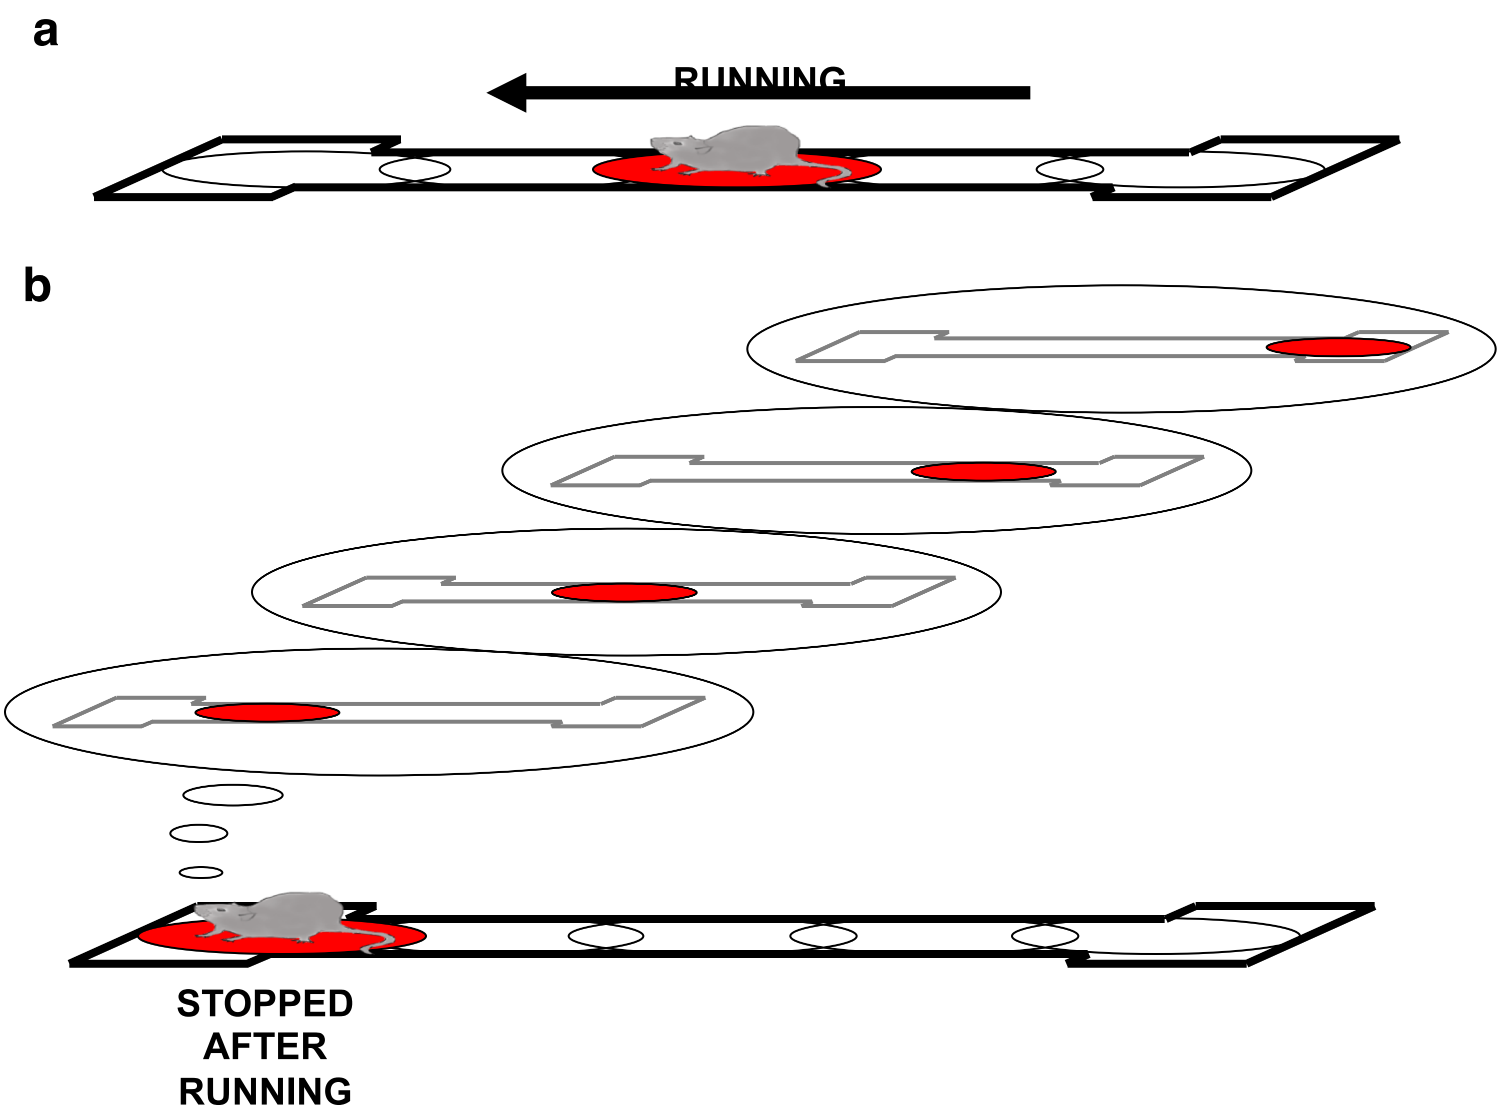

Memory of recent experience replayed in reverse-time order

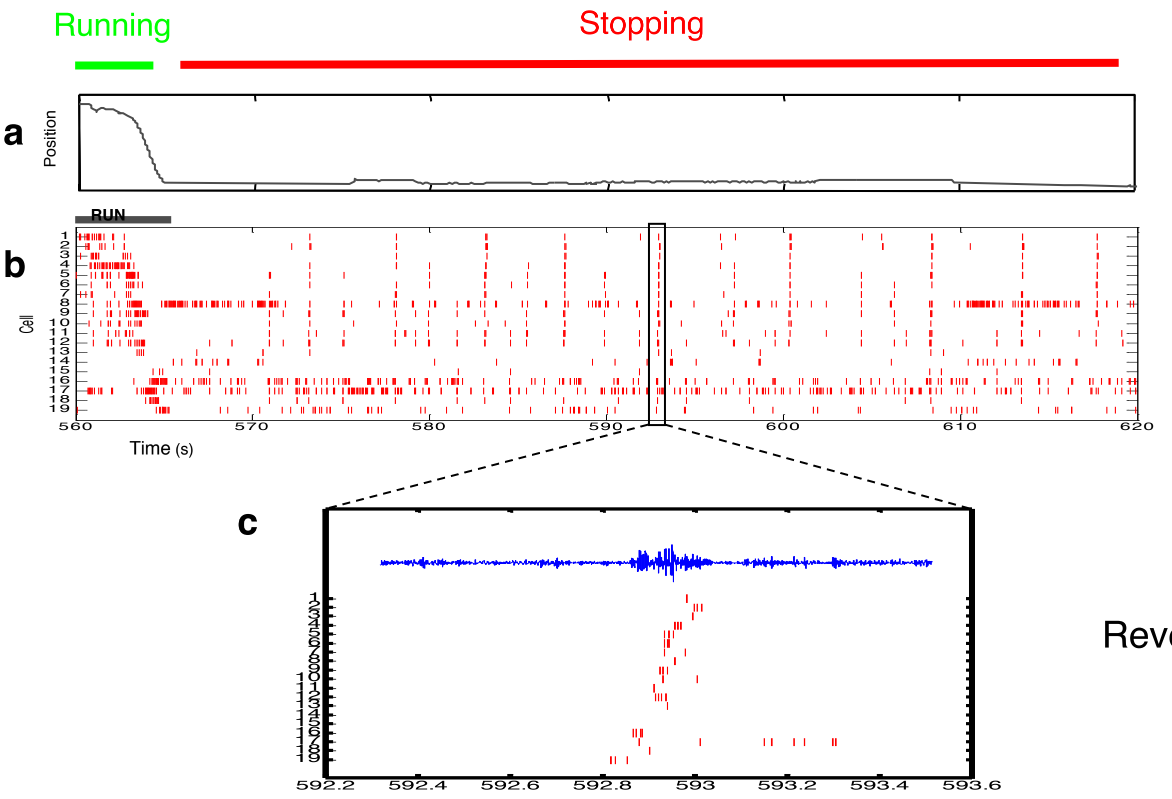

Position vs. time

Hippocampal place-cell activity vs. time

Reverse-time sequence replay during hippocampal ripples

Foster and Wilson, *Nature*, 2006

Learning sequences of actions

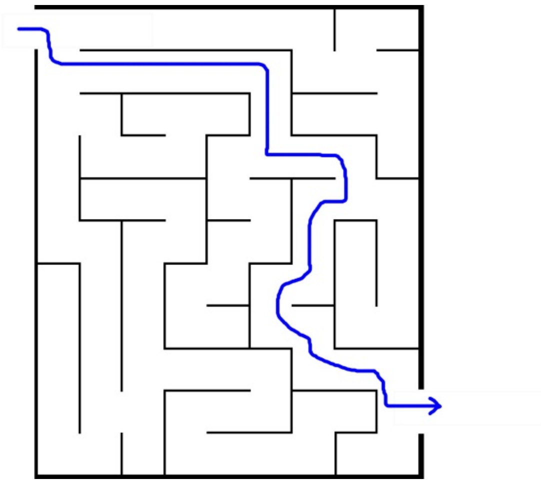

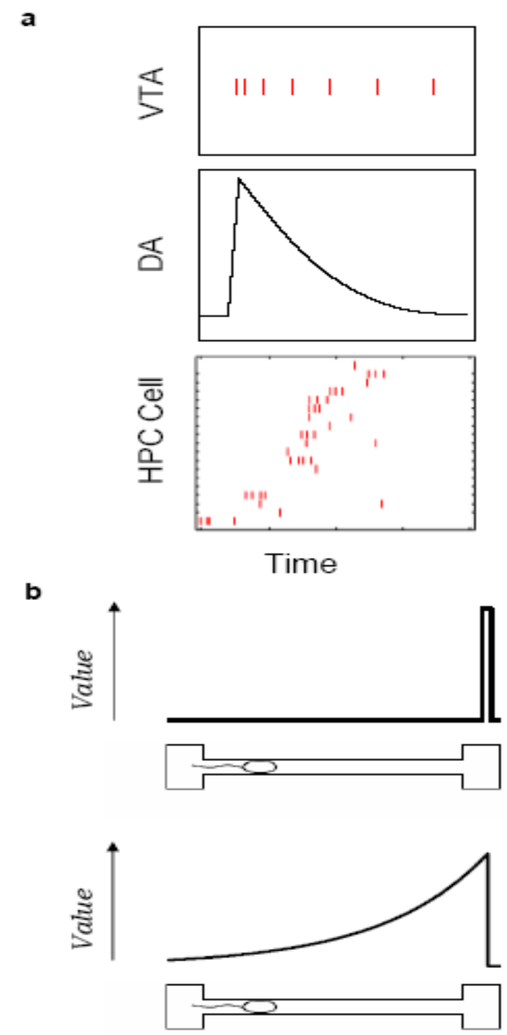

Temporal credit assignment

Dopamine unit activity could differentially weight the content of hippocampal sequences, propagating value information from the rewarded location backwards along the incoming trajectory.

Foster and Wilson, Nature 440: 680-683, 2006

Hippocampal place-cell sequences depict future paths to remembered goals
Brad E. Pfeiffer & David J. Foster
Nature, 2013

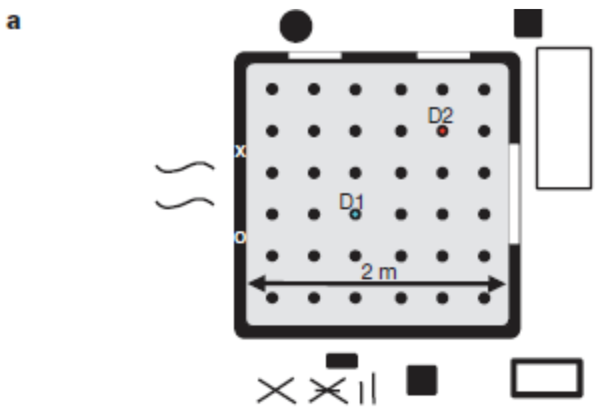

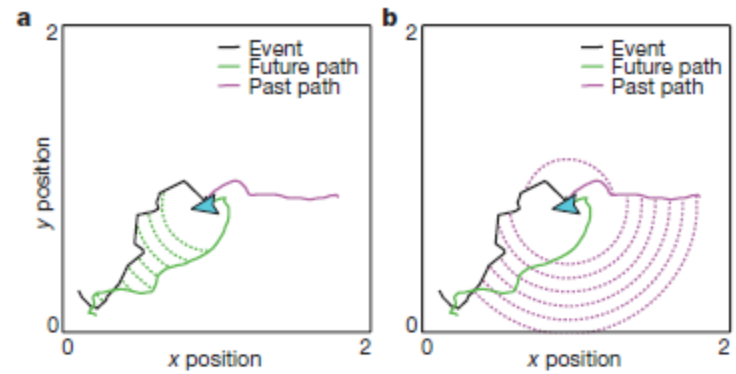

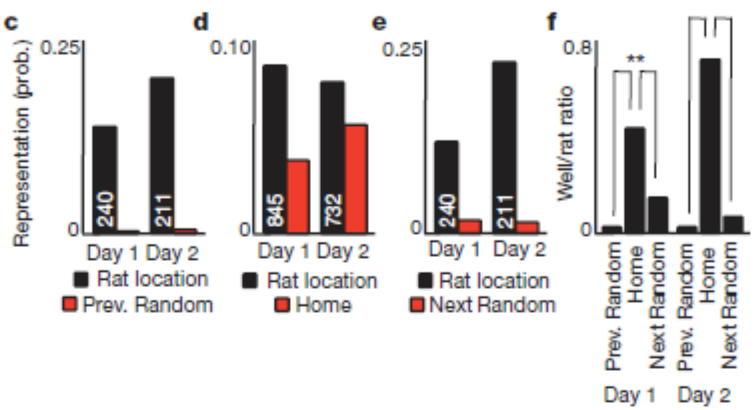

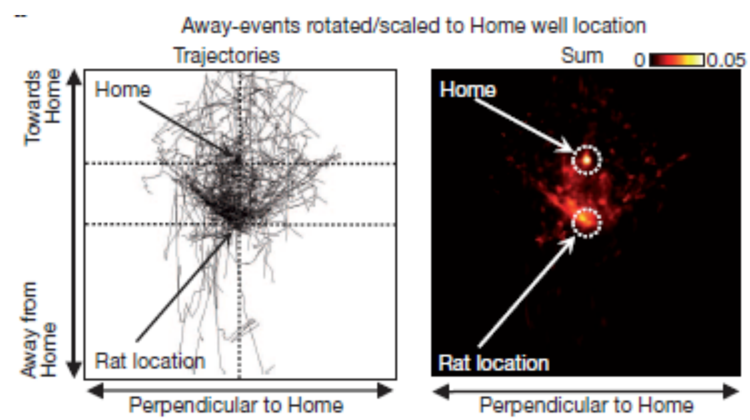

Long behavioral sequences on a 10m track

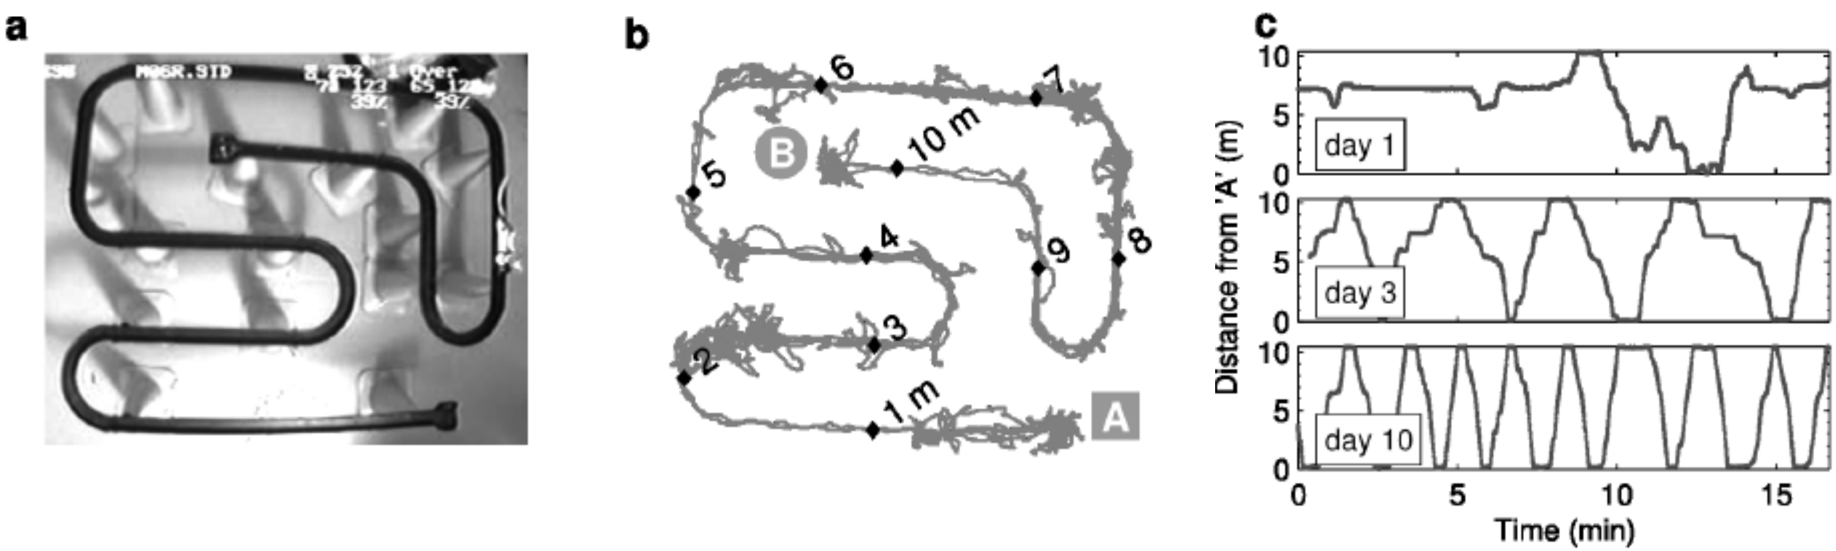

Davidson, Kloosterman and Wilson, *Neuron*, 2009

Reconstruction of extended sequence replay during quiet wakefulness

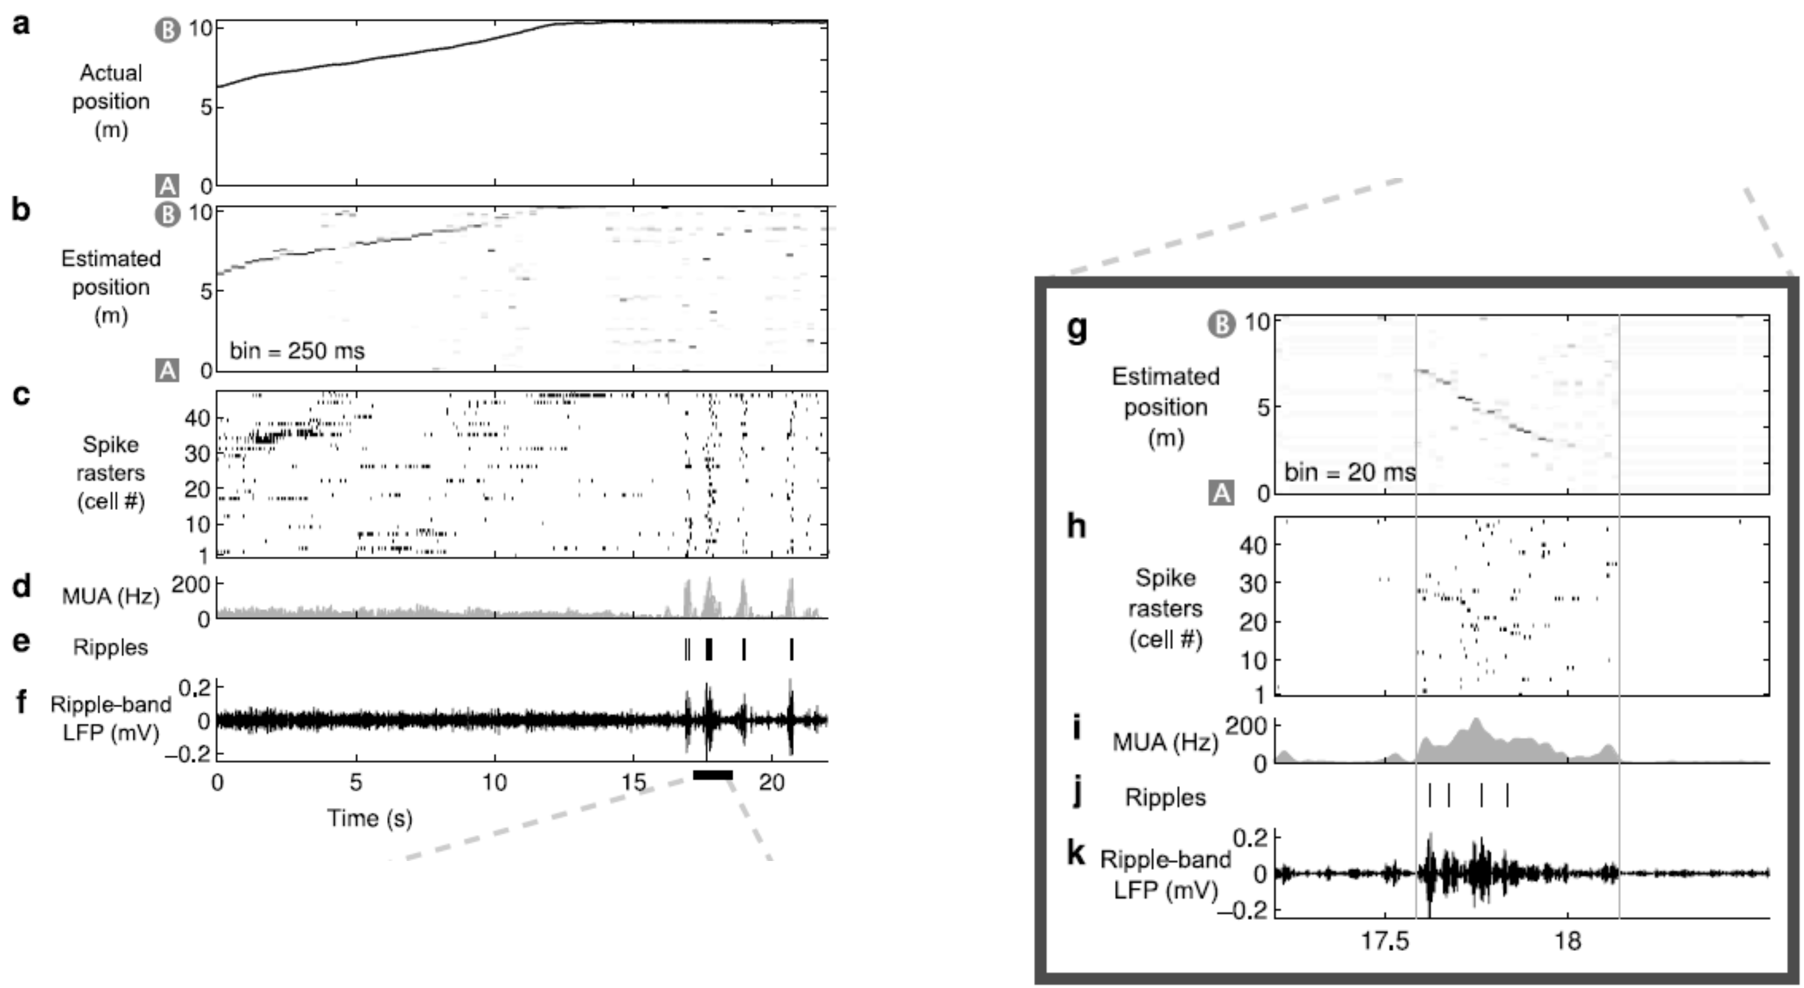

Davidson, Kloosterman and Wilson, *Neuron*, 2009

Forward Replay from A to B

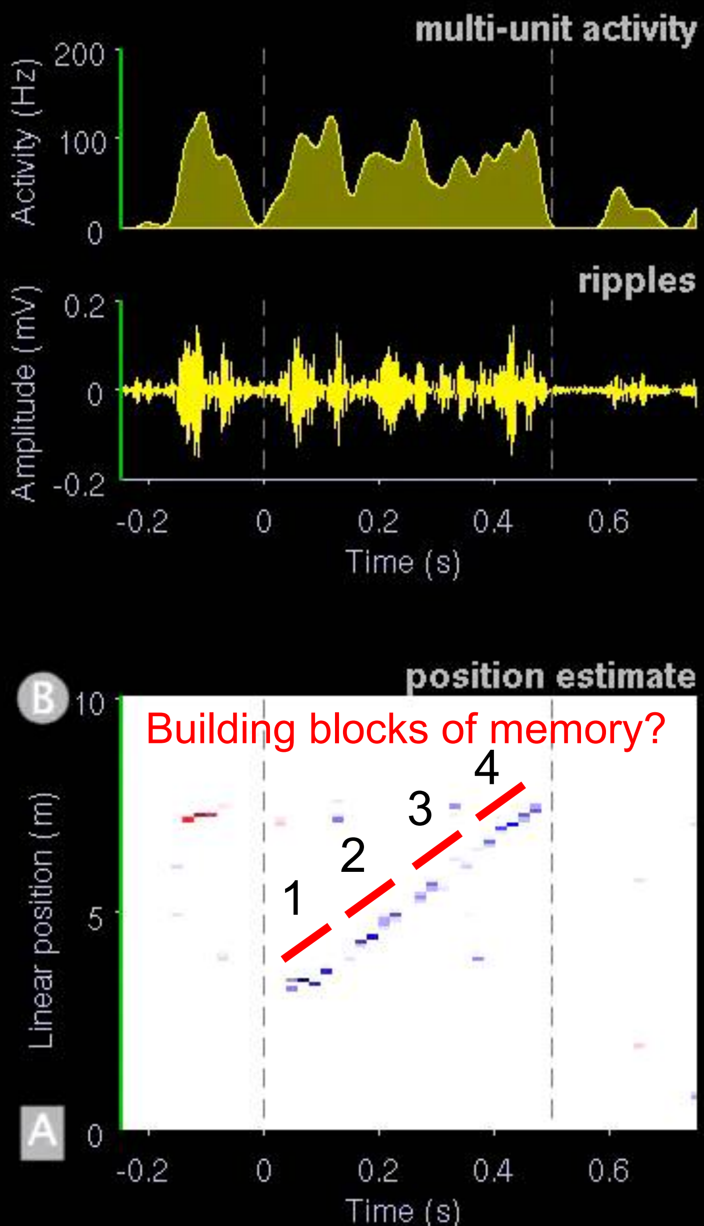

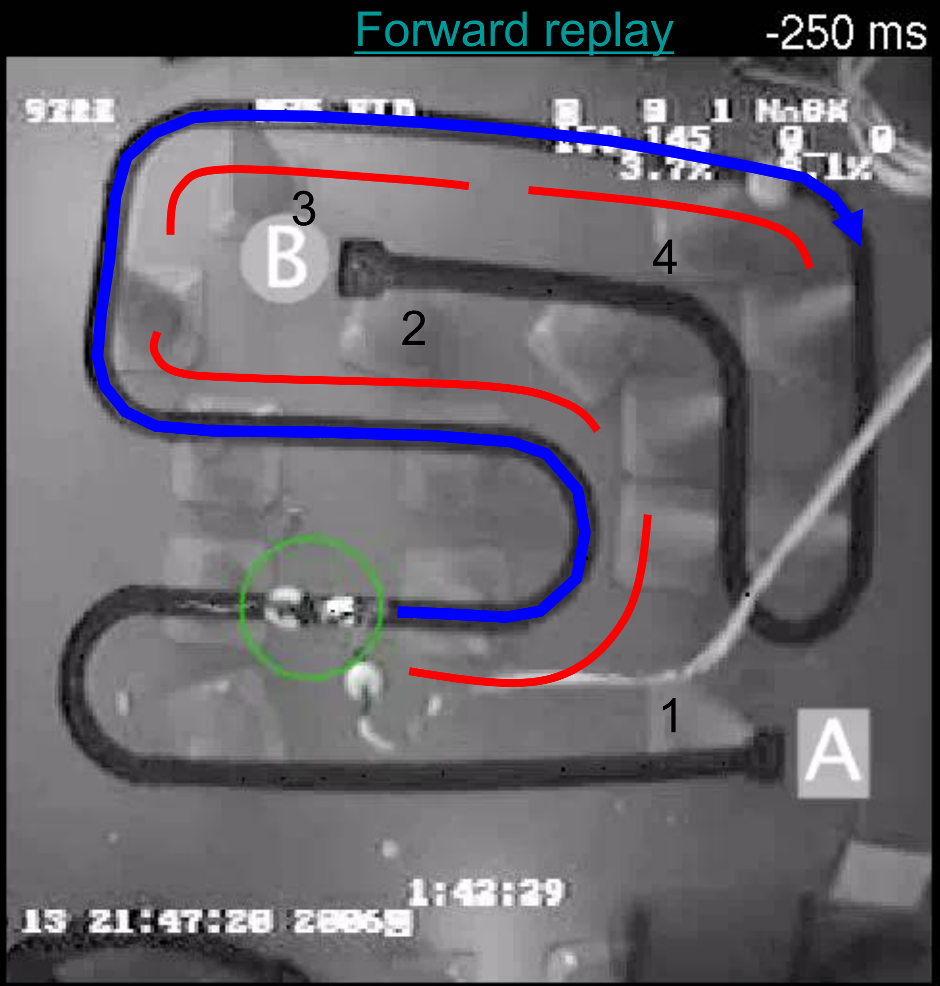

Forward Replay from A to B

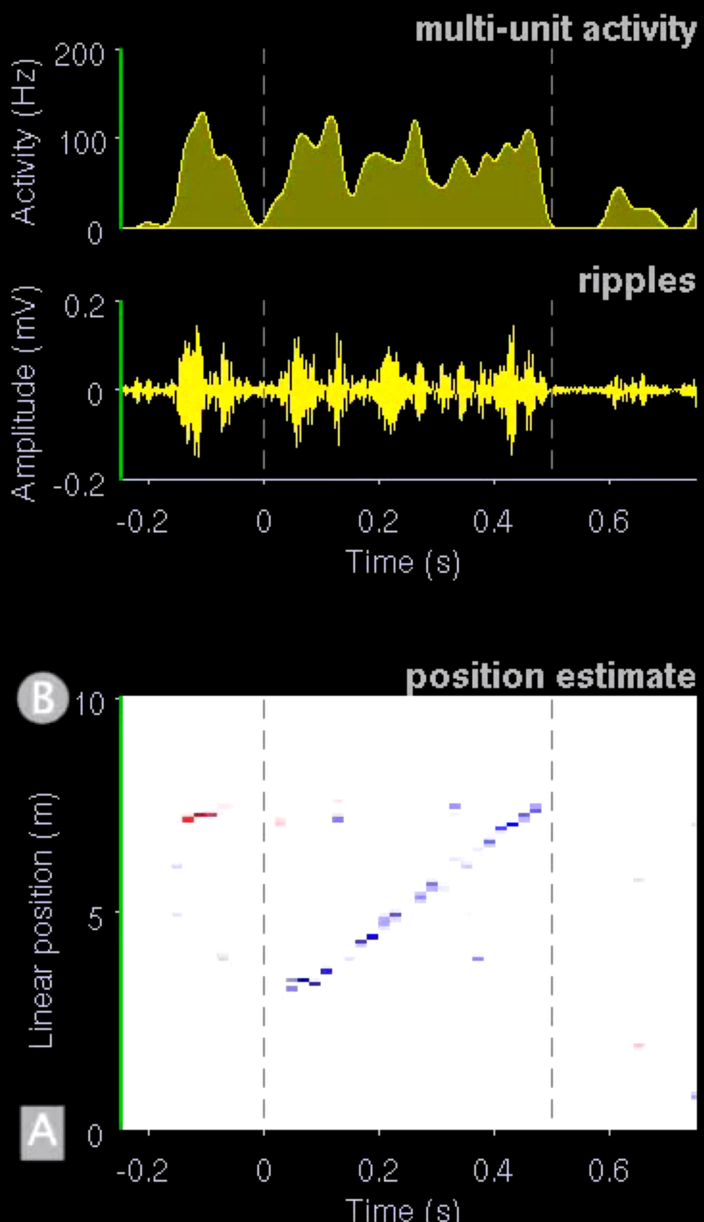

-250 ms

Extended replay spans multiple ripple events

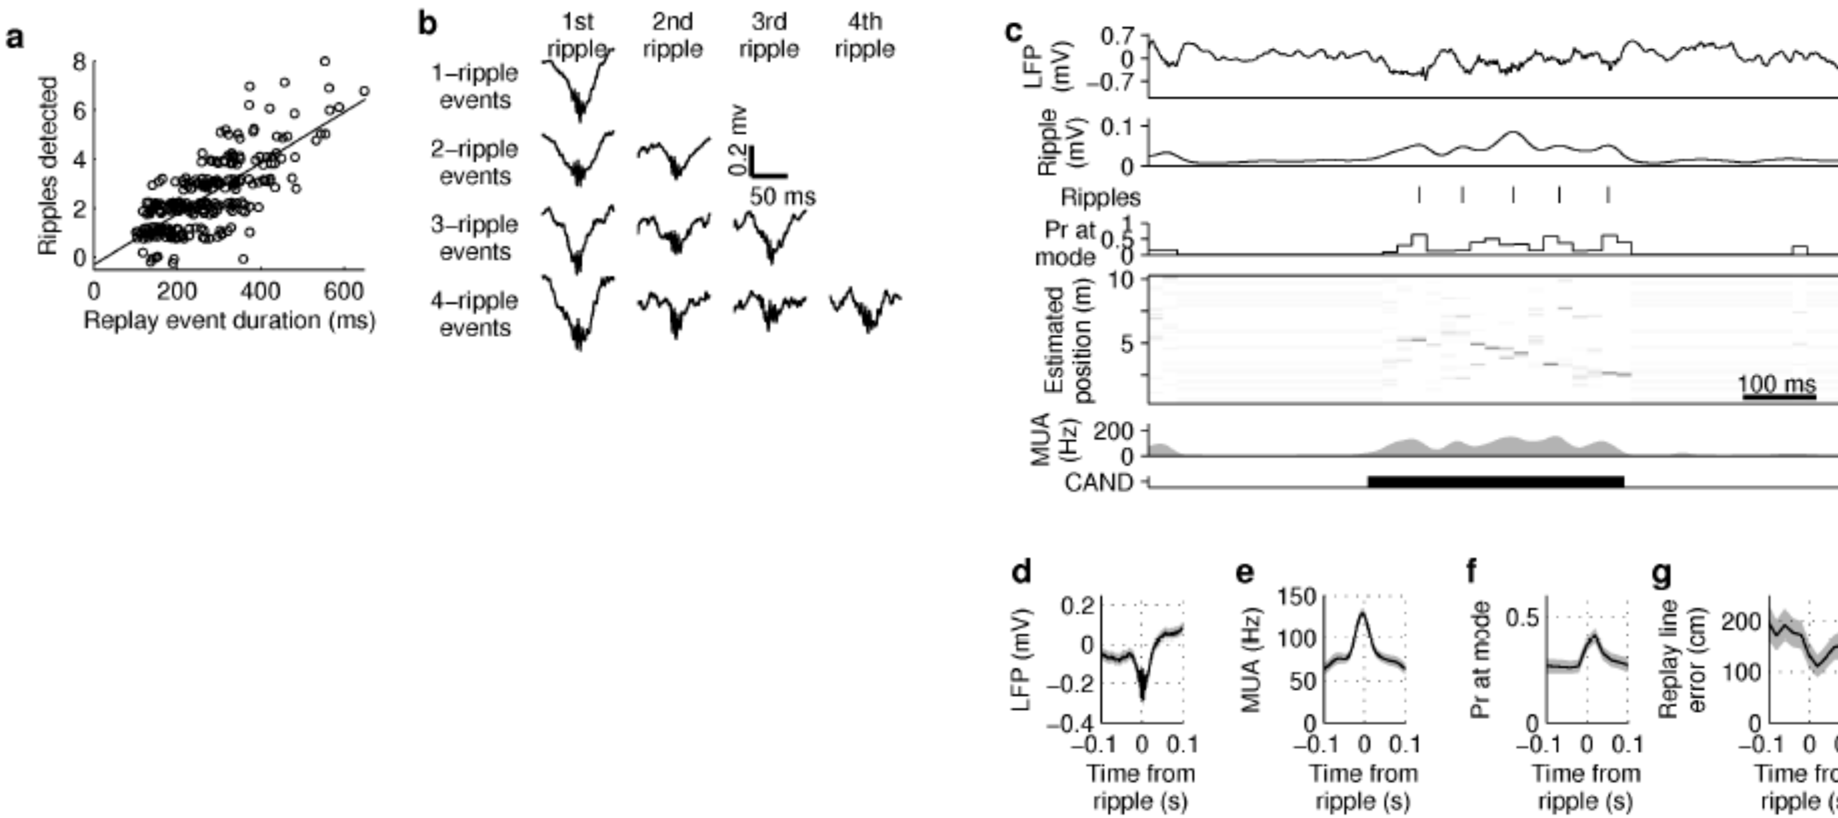

Davidson, Kloosterman and Wilson, *Neuron*, 2009

Extended replay has a characteristic speed

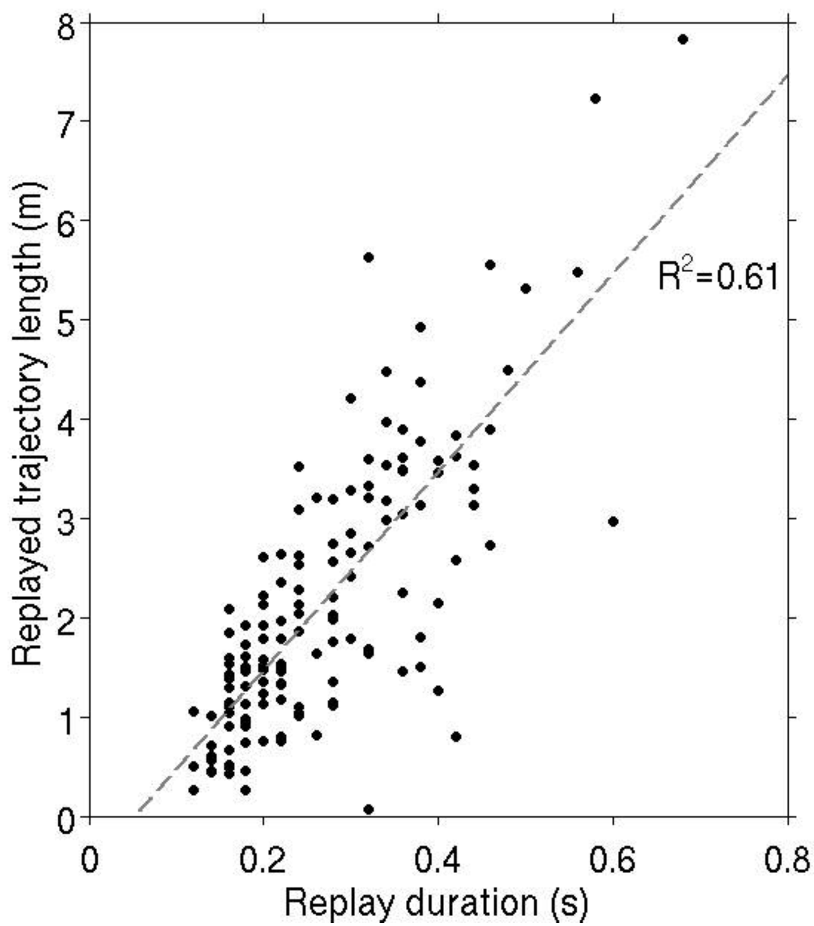

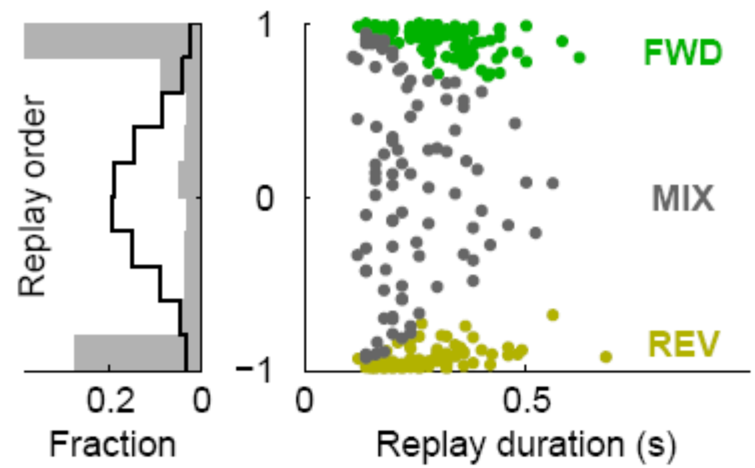

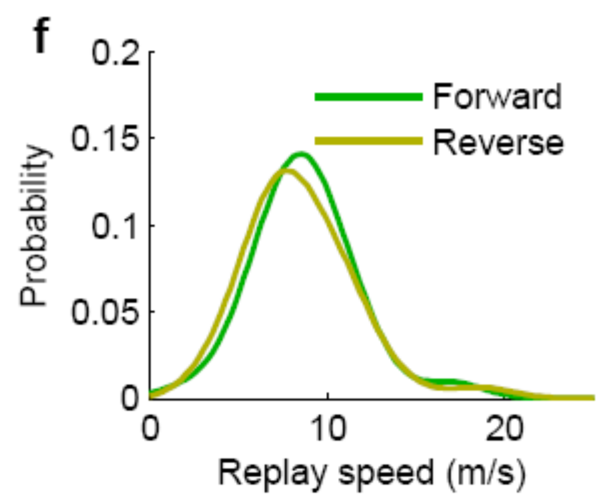

Davidson, Kloosterman and Wilson, *Neuron*, 2009

Single ripple sequences are at same scale as theta sequences

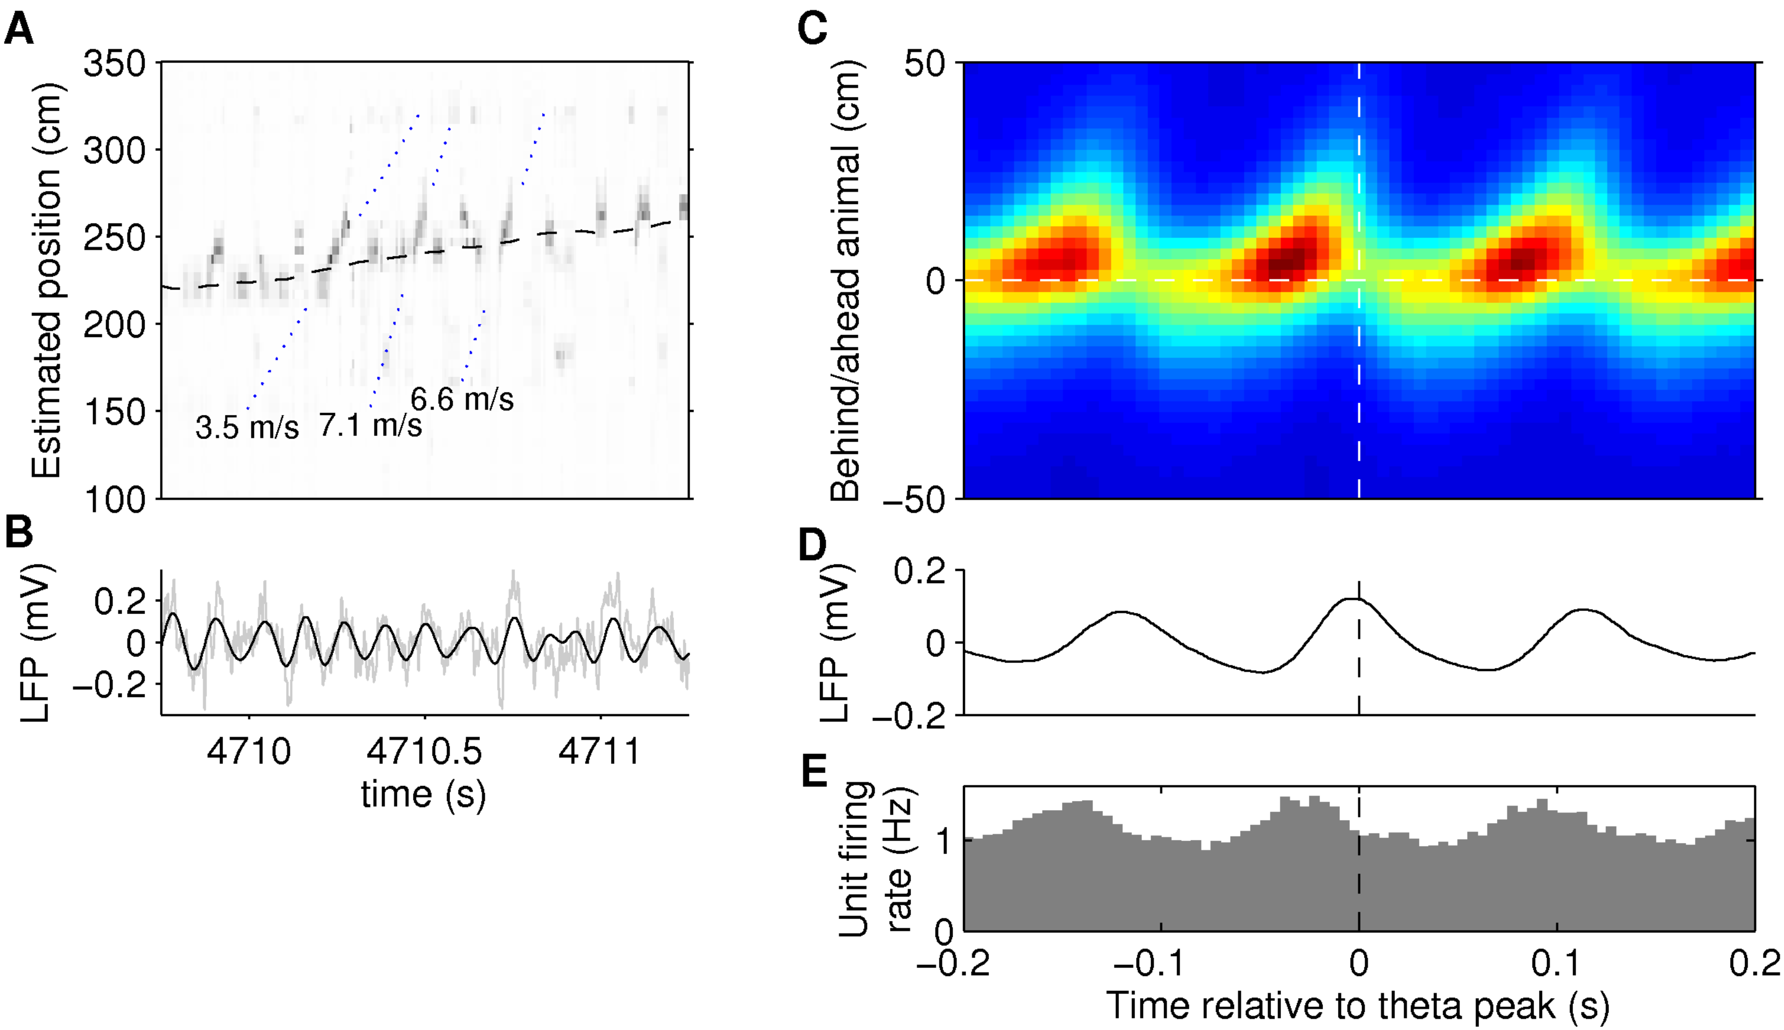

Davidson, Kloosterman, and Wilson, *Neuron*, 2009

Learning sequences of actions

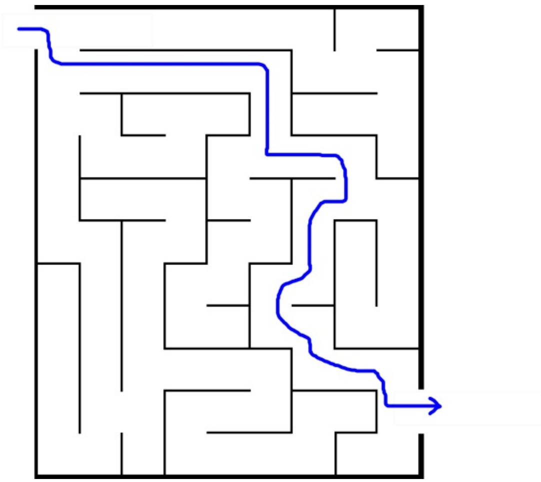

Temporal credit assignment

Dopamine unit activity could differentially weight the content of hippocampal sequences, propagating value information from the rewarded location backwards along the incoming trajectory.

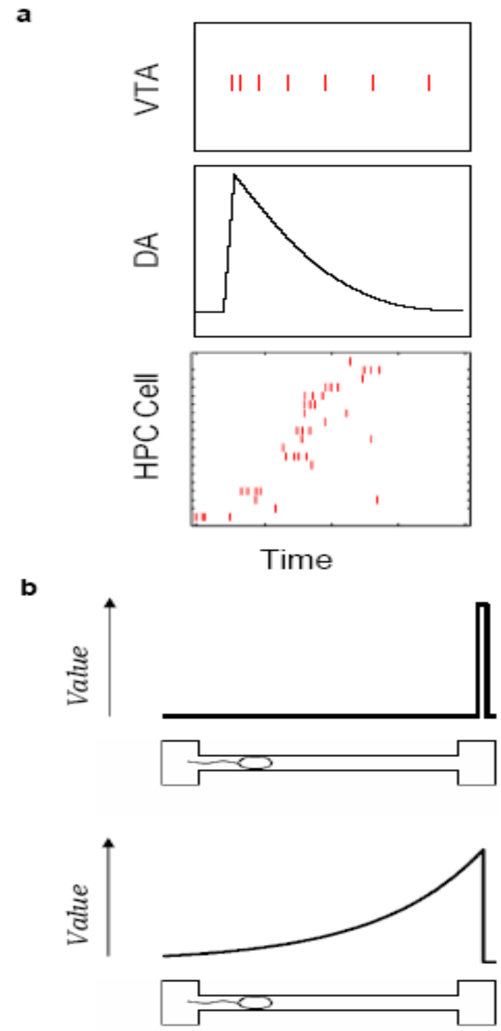

Foster and Wilson, Nature 440: 680-683, 2006

Hippocampal place-cell sequences depict future paths to remembered goals
Brad E. Pfeiffer & David J. Foster
Nature, 2013

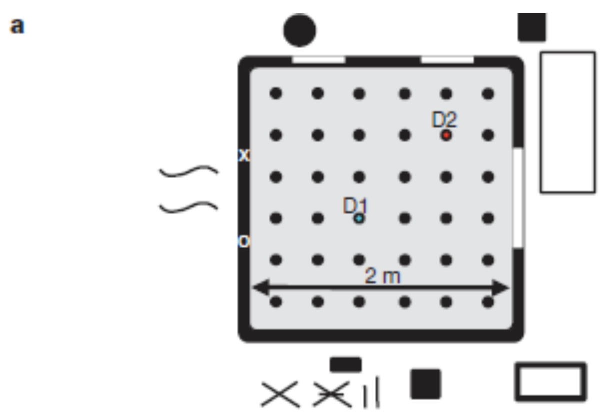

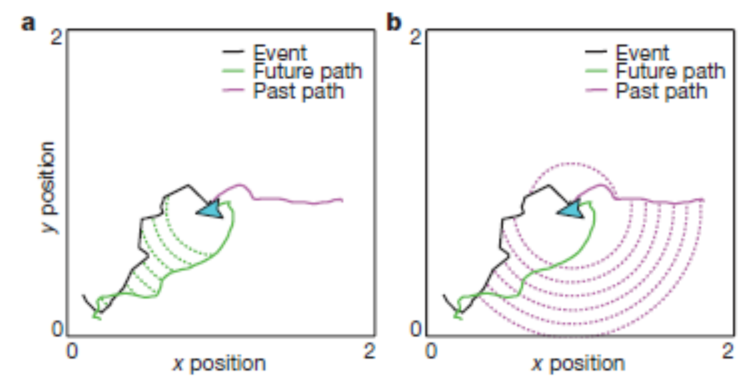

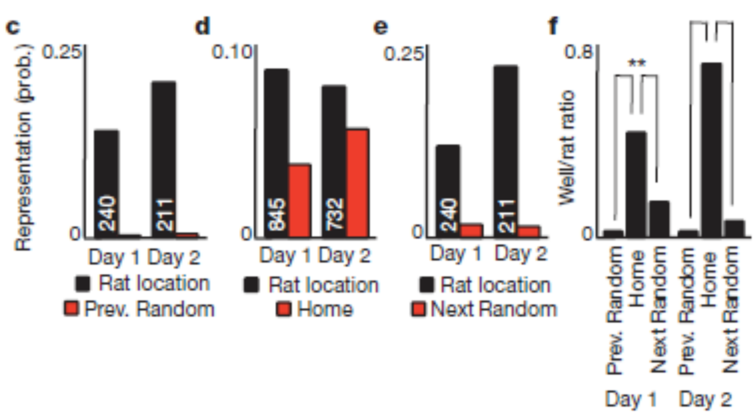

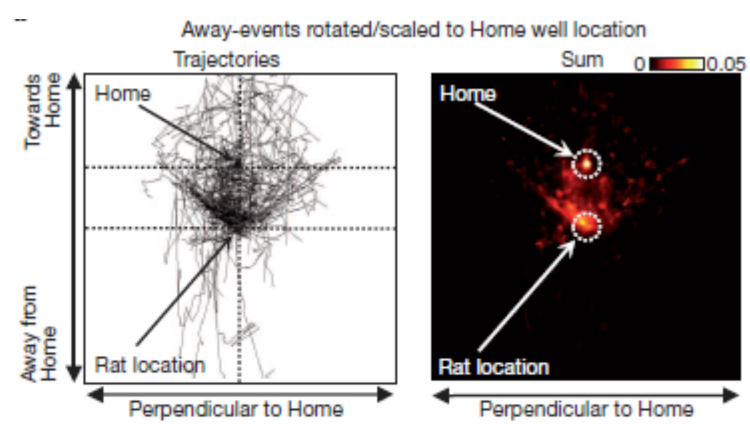

Dopamine cell representations

- unexpected reward
- predictors of reward
- errors in the prediction of reward

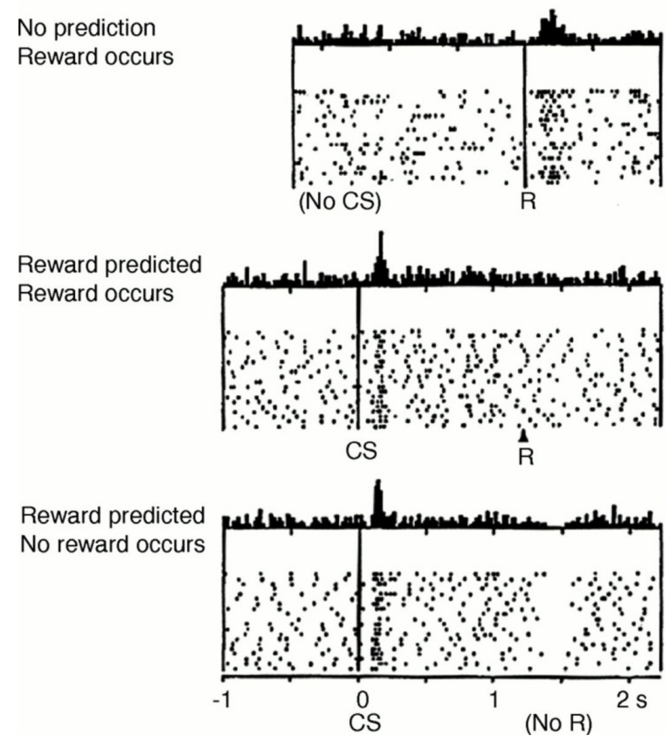

Reward prediction error

Current reward – Expected reward

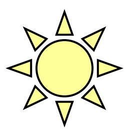

An error signal
to teach target
brain regions

Schultz, Dayan, Montague. Science 1997

VTA Hippocampus co-recordings

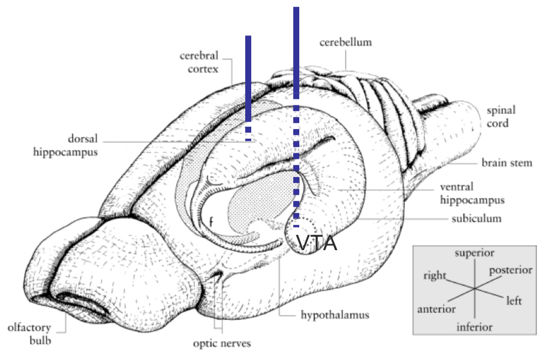

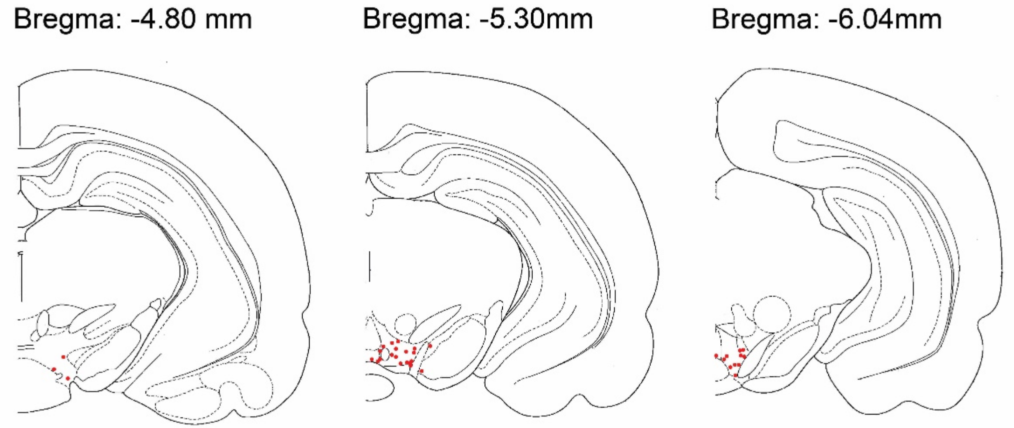

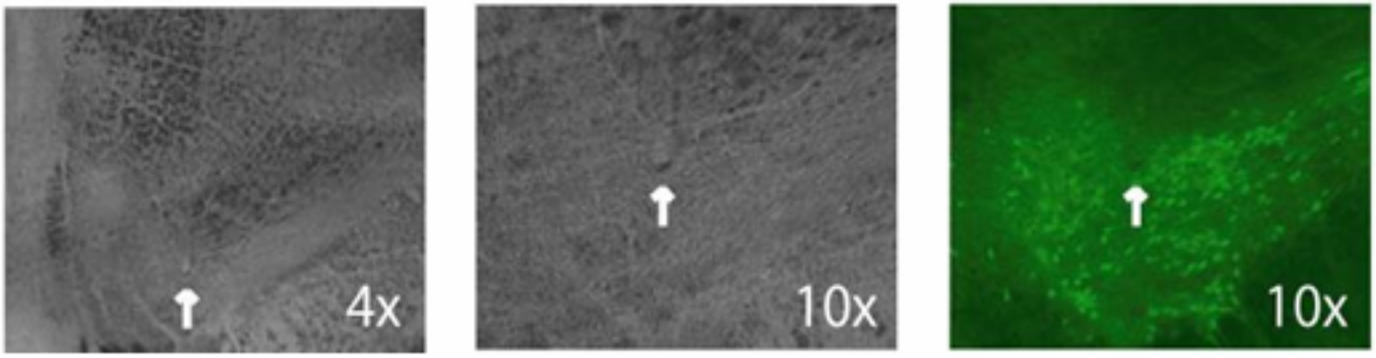

Light microscopy

Anti-TH Ab

Spatial working memory task

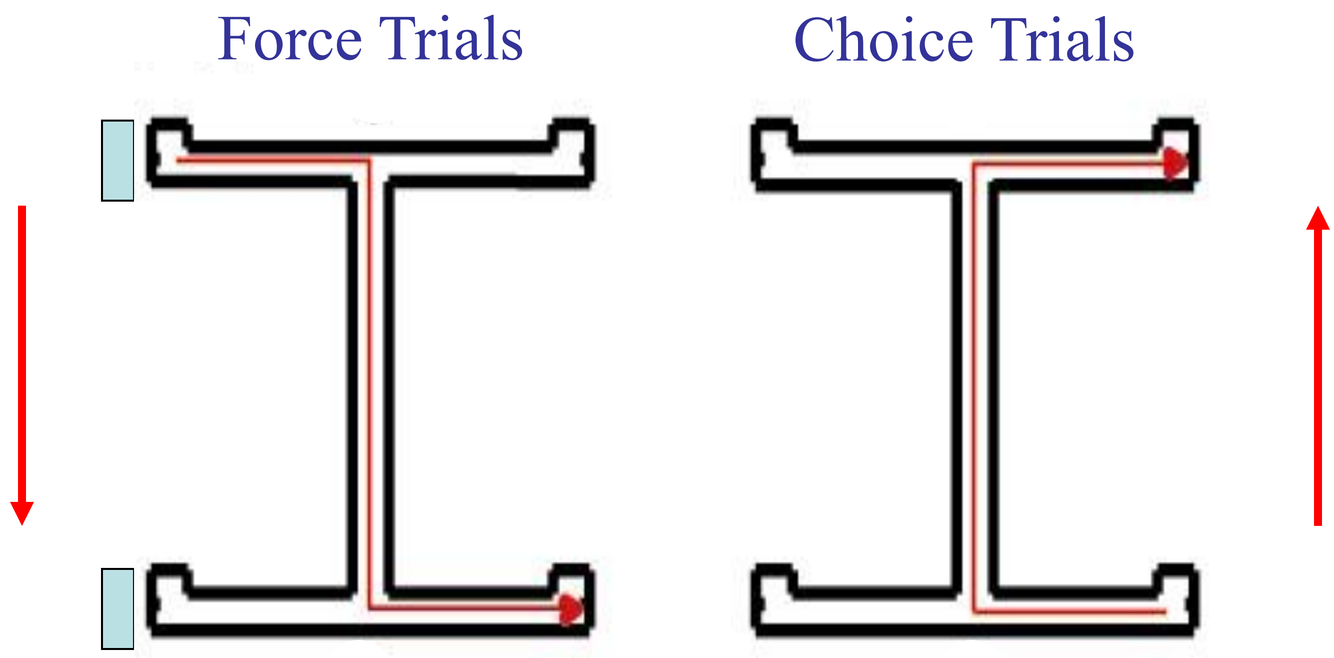

Task contingency associated VTA unit activity

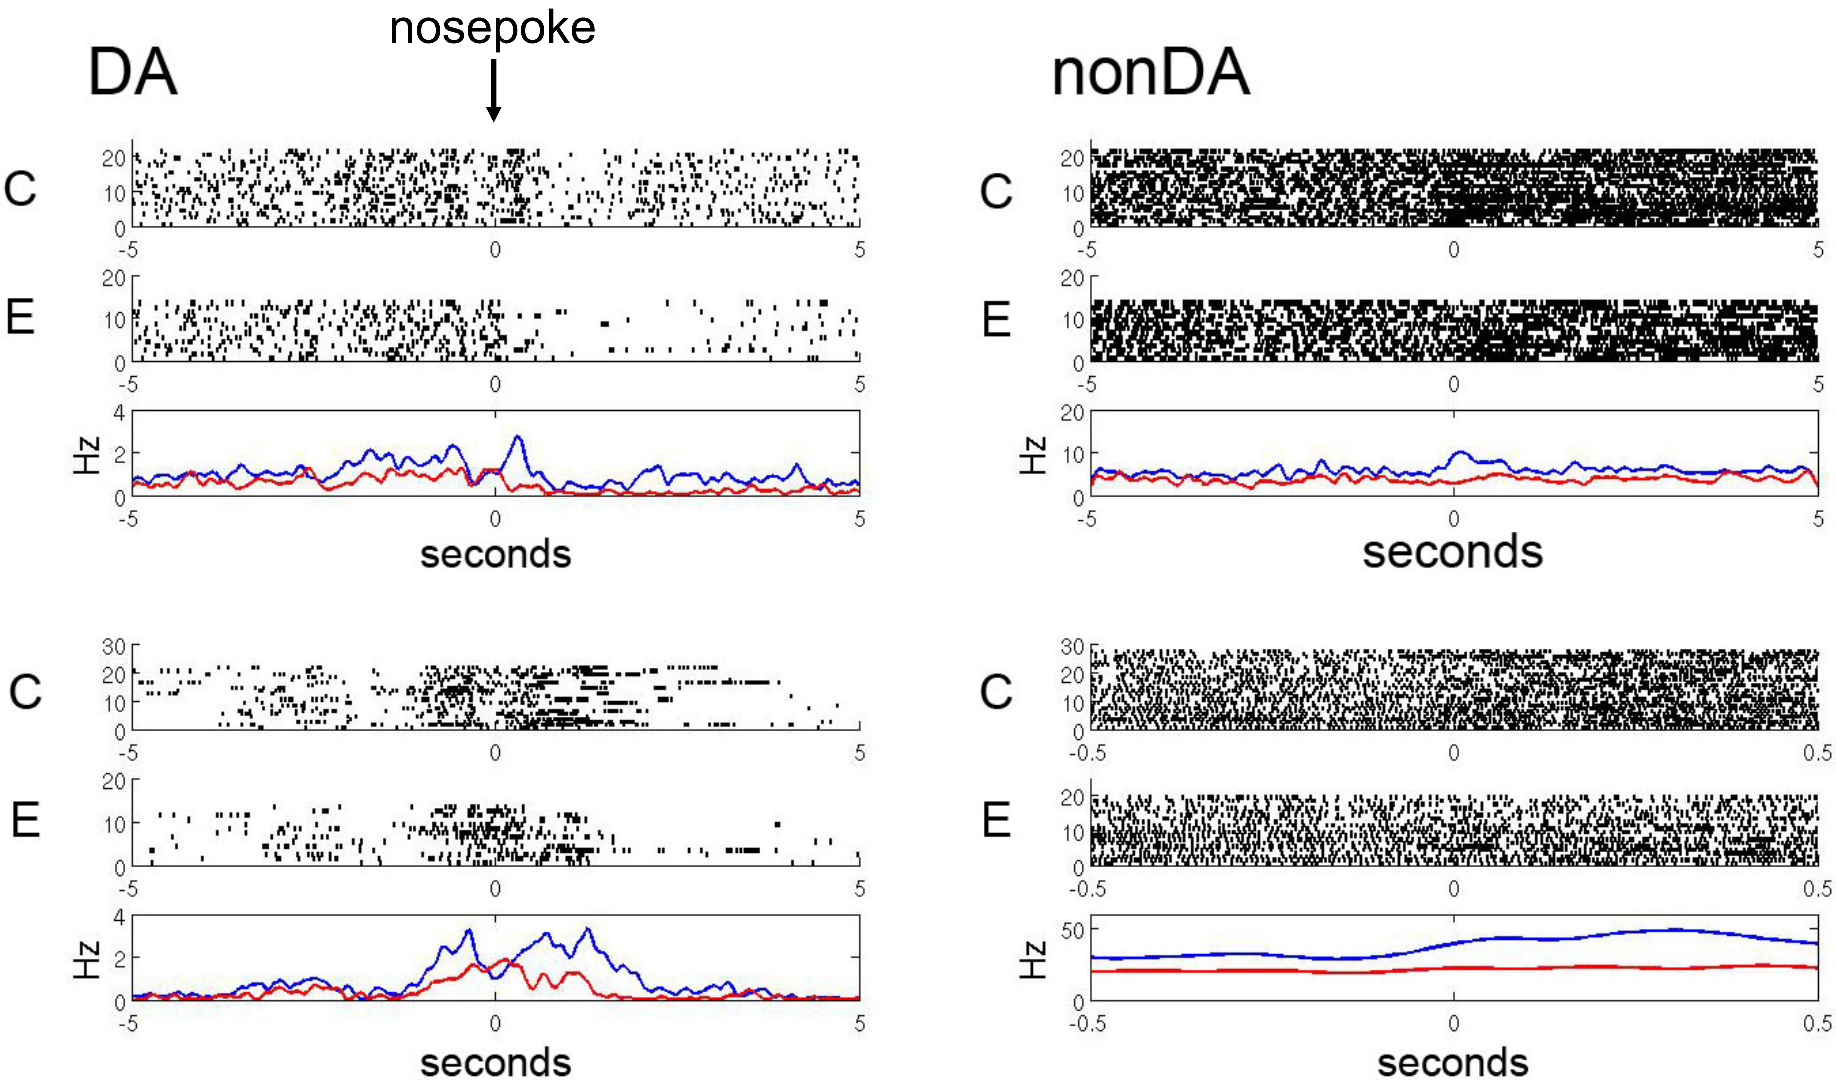

VTA unit activity during RUN

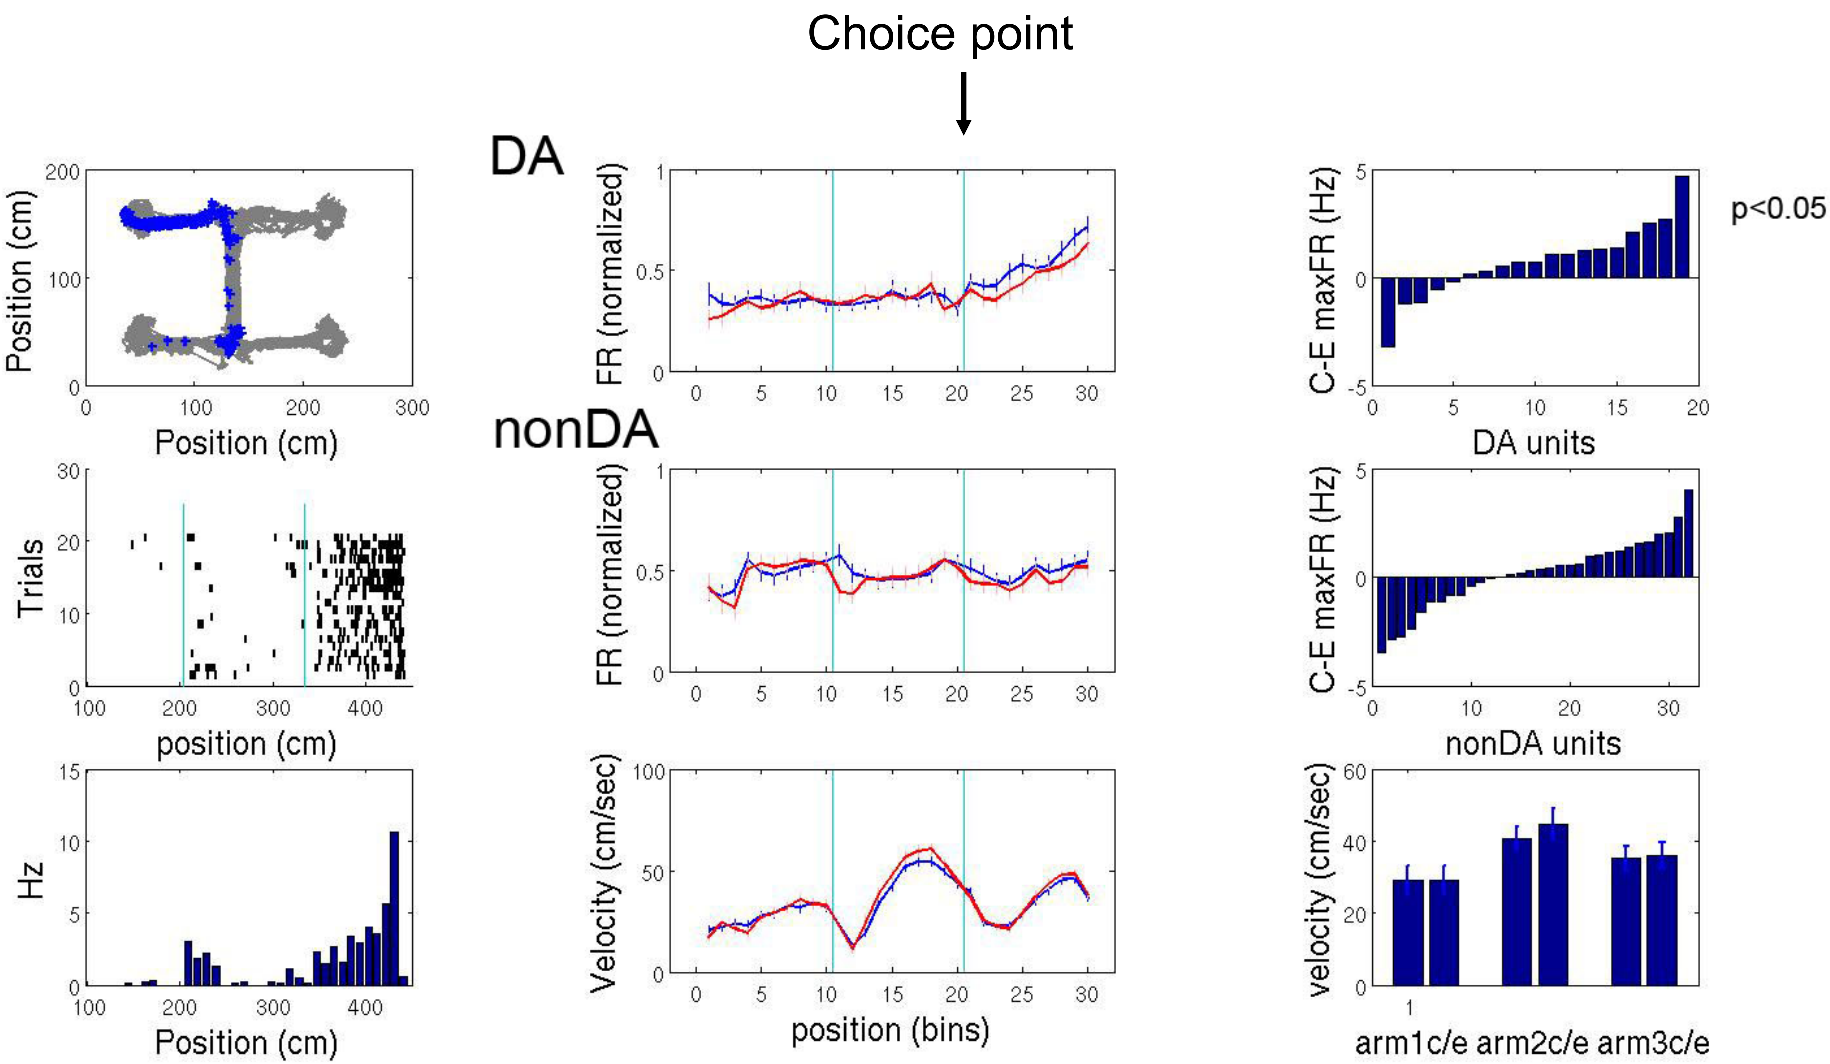

HC ripple bursts modulate DA unit activity

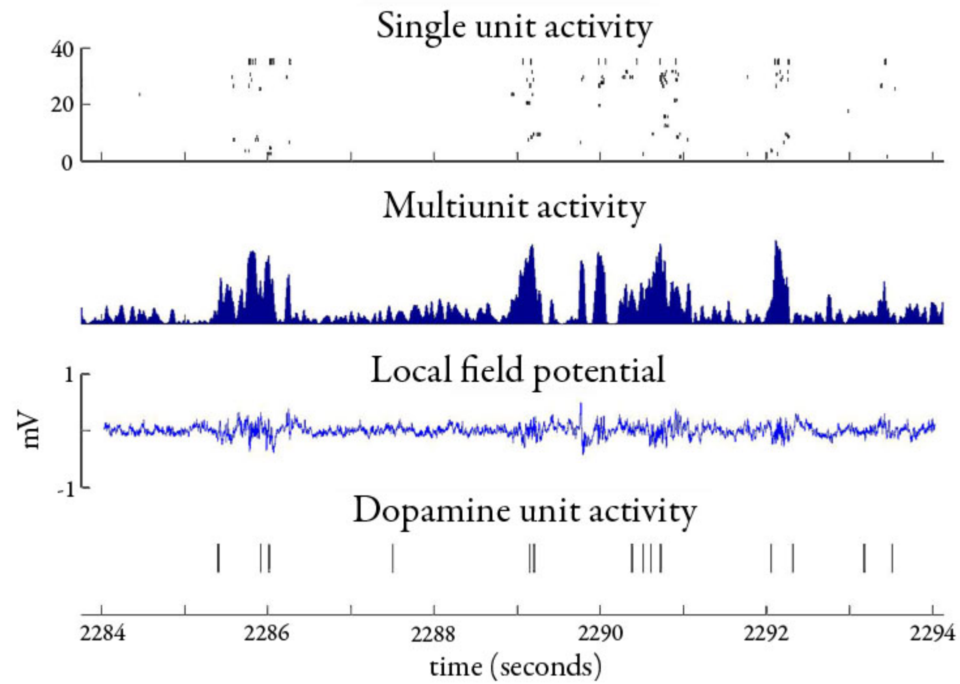

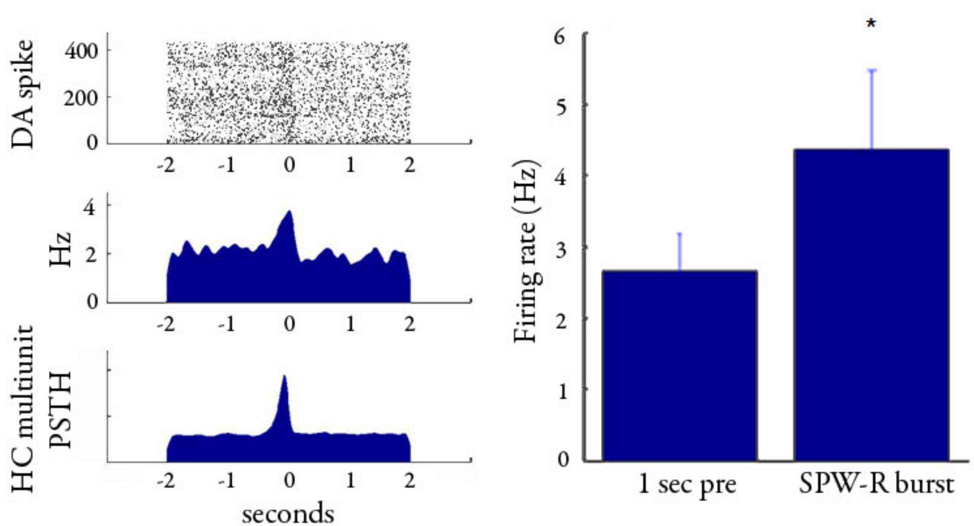

Gomperts, unpublished data

Decoding hippocampal SPW-R associated multiunit bursts with spatial sequence reactivation

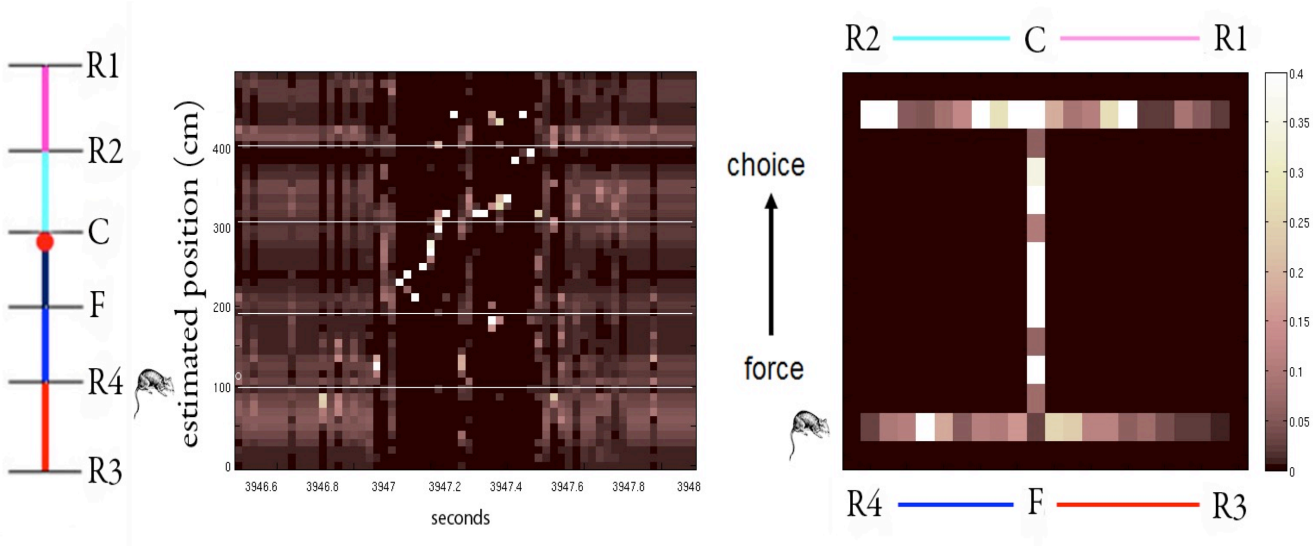

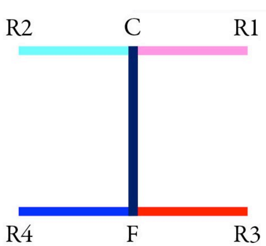

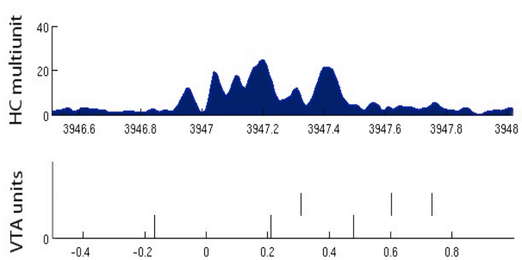

Replay and nonreplay SPW-R events

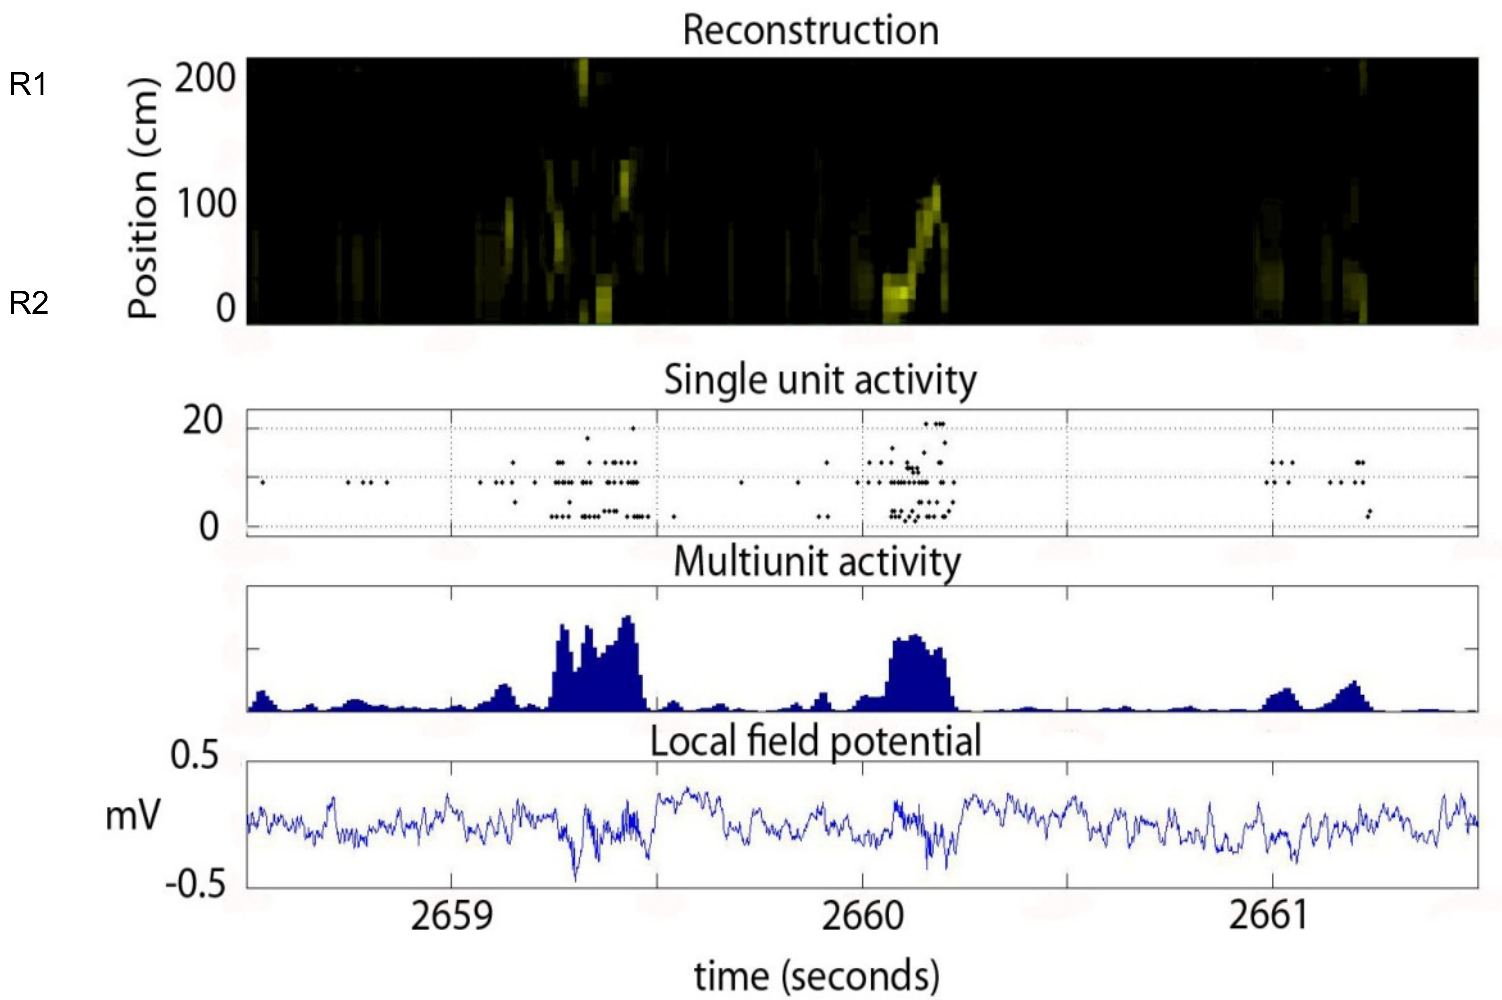

Dopamine unit modulation at hippocampal SPW-R bursts depends on replay content

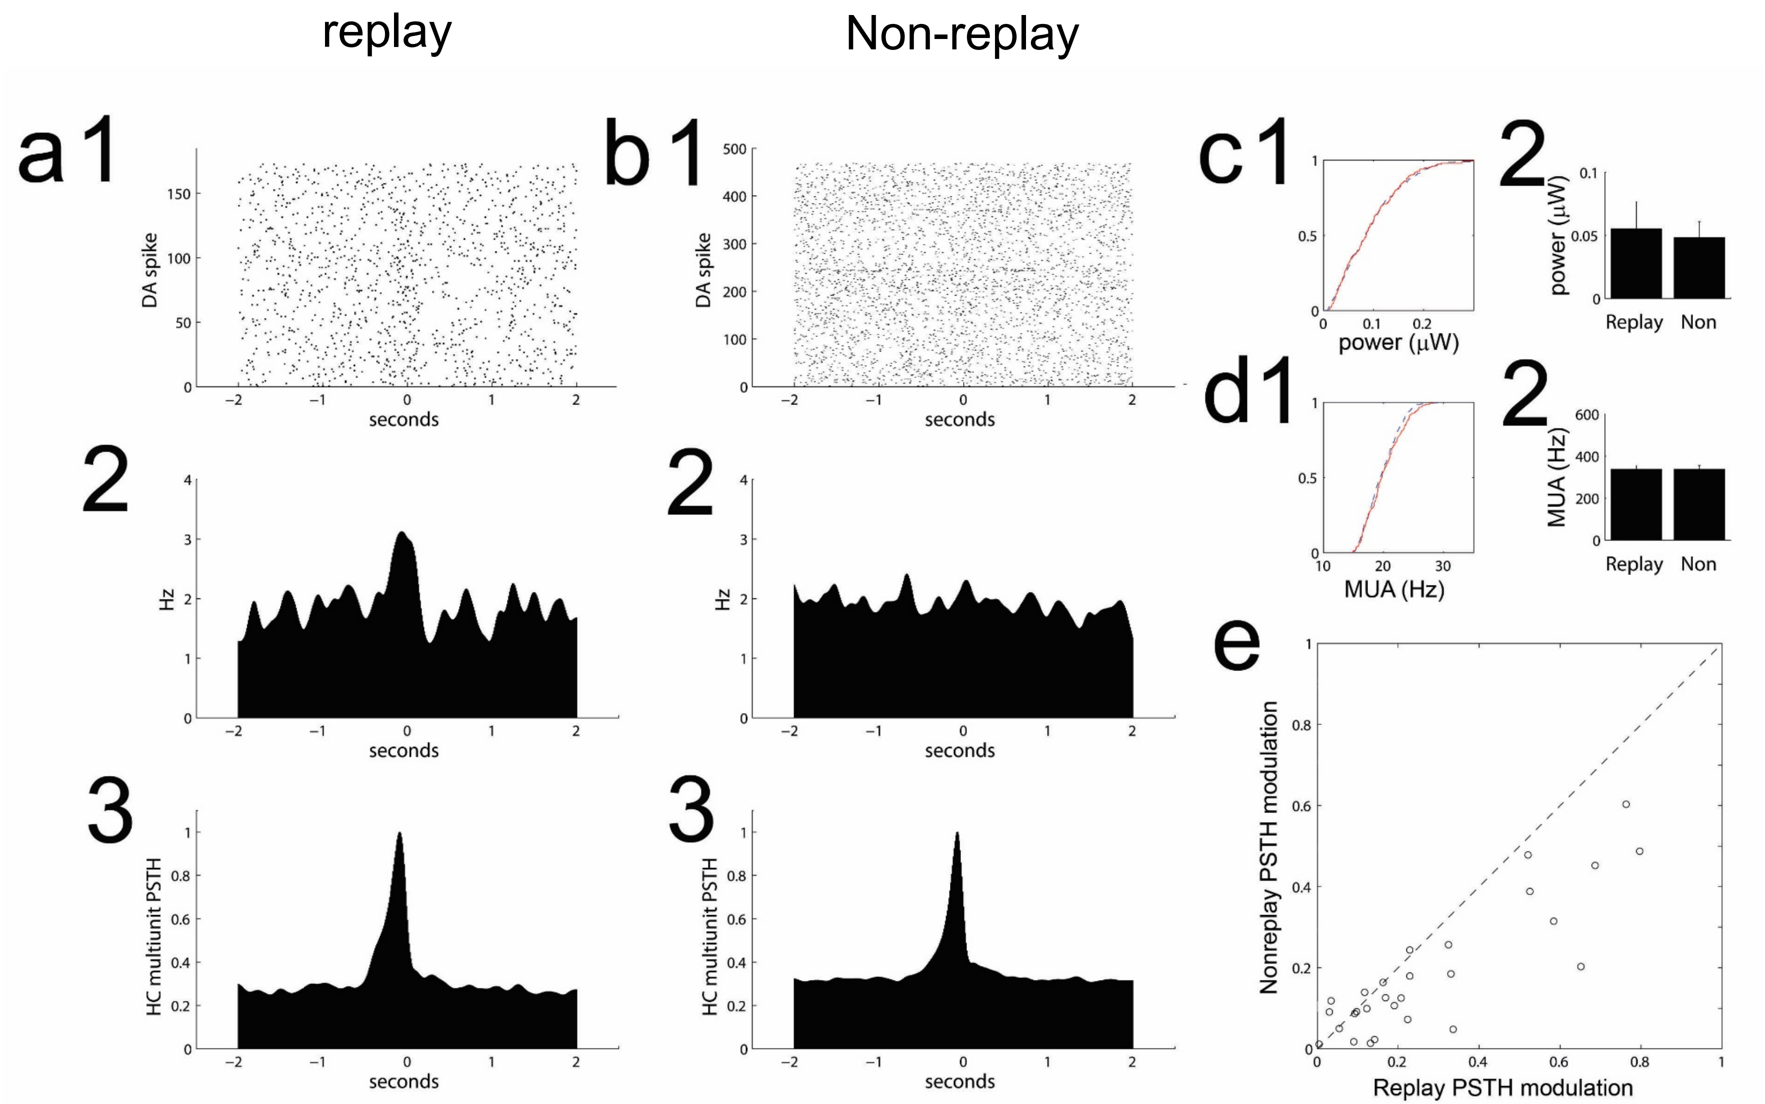

Dopamine units preferentially coordinate with replay of reward site locations on the spatial working memory task and phase lock to hippocampal theta

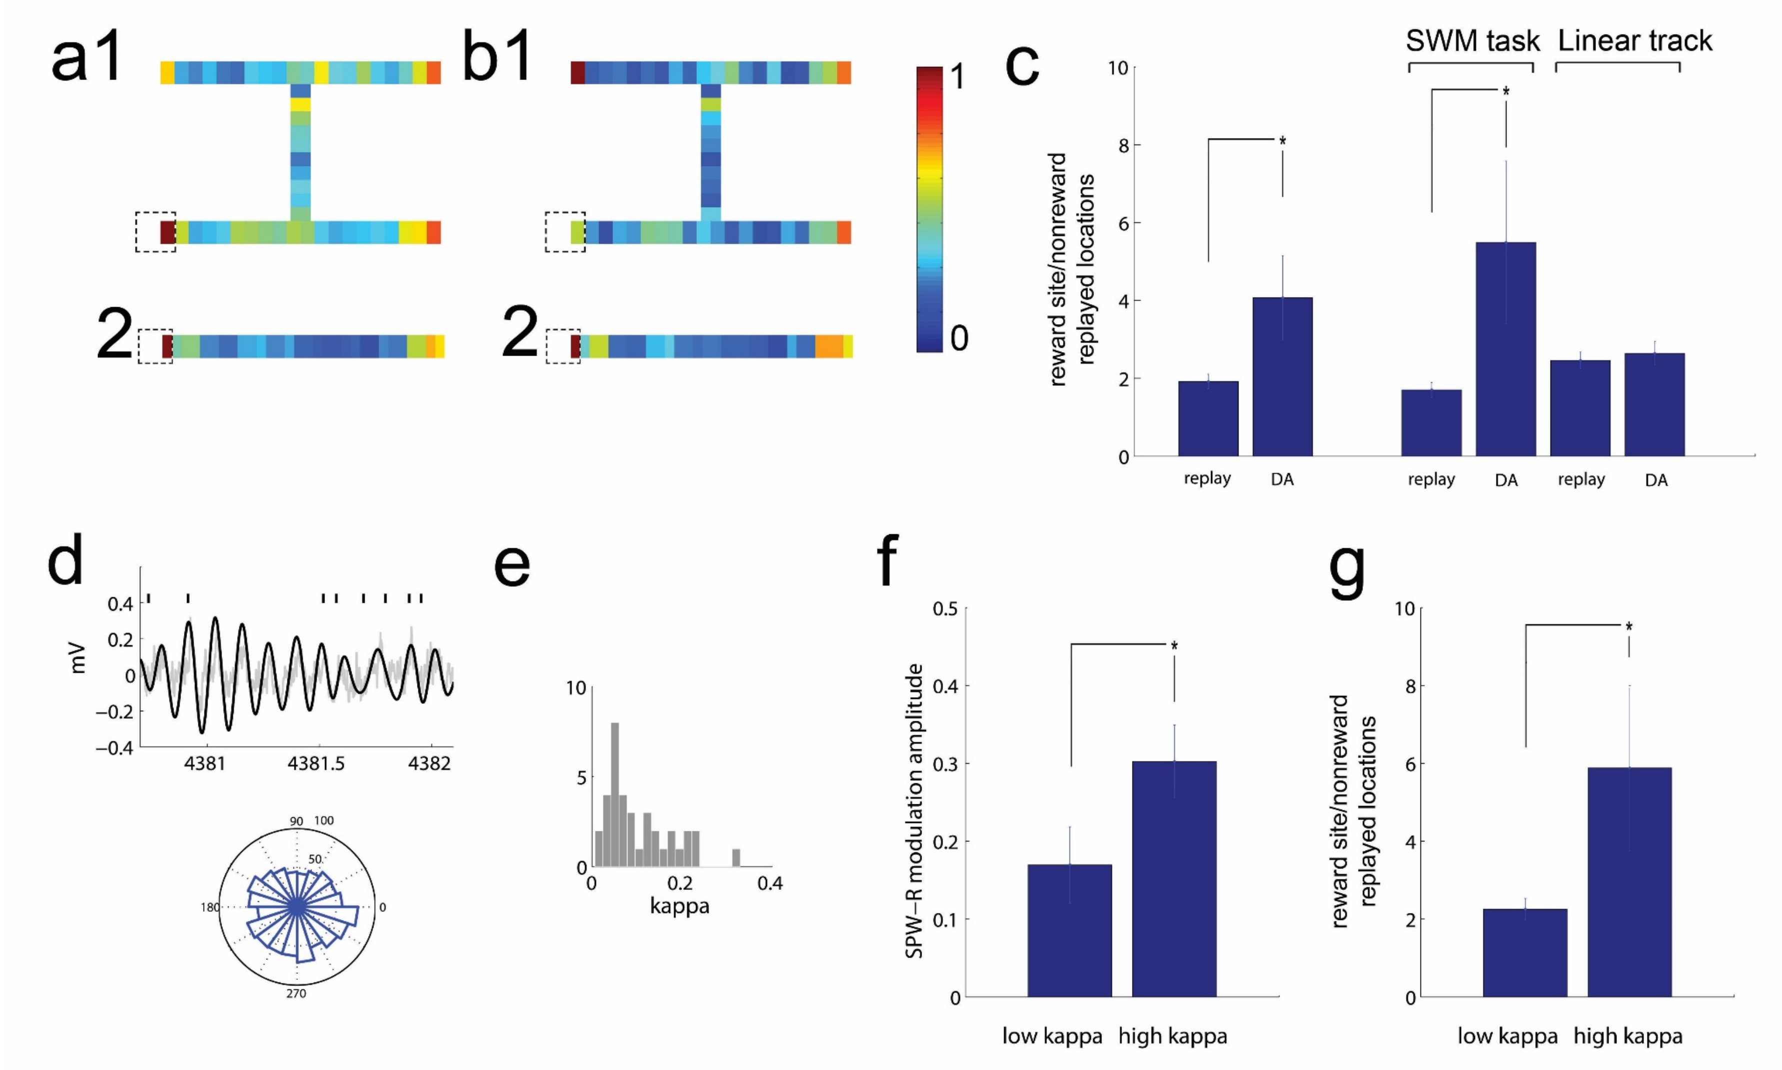

Modulation of DA activity at SPW-Rs is reduced in slow wave sleep

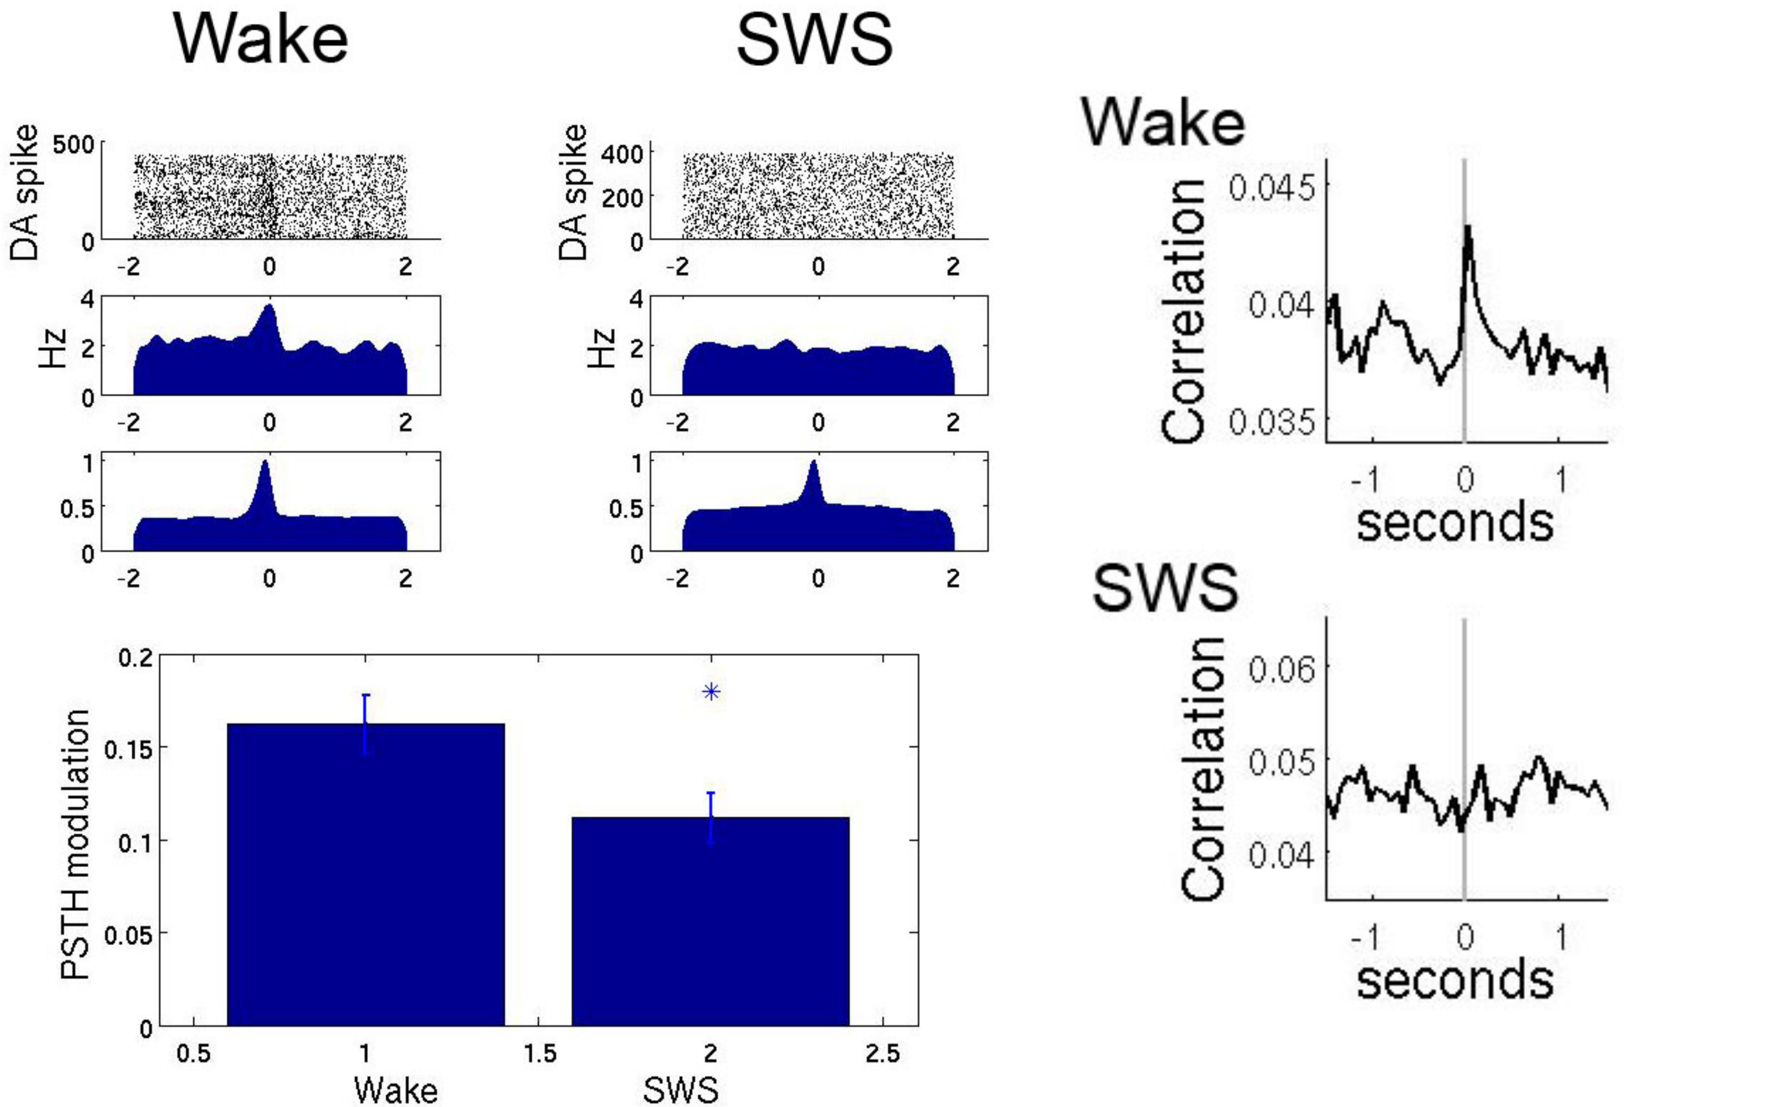

Pairing reward (MFB stim) during sleep can influence learning

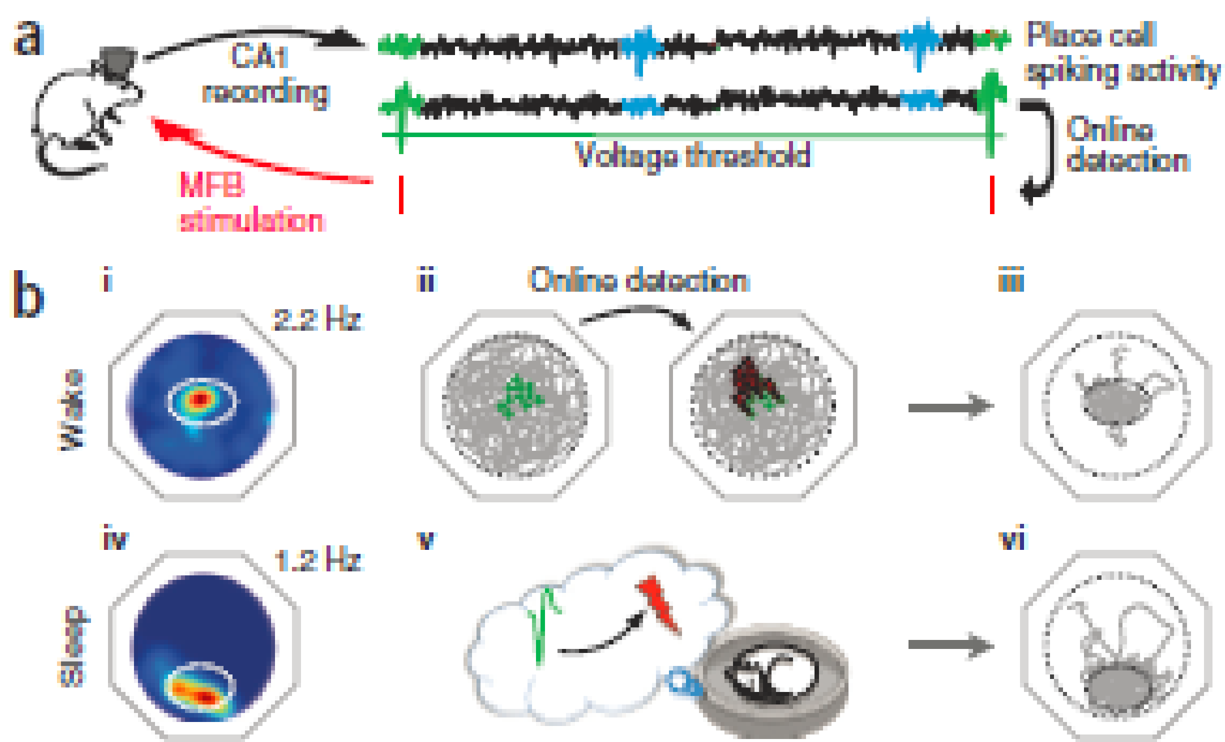

Lavilleon et al, Nature Neuroscience, 2015

Summary

- DA unit activity increases during trajectories to rewards, differentially represents correct over error trials, and correlates with Q-TD prediction error in a spatial task.
- Hippocampal SPW-R bursts are associated with the modulation of DA units.
- Hippocampal theta phase-locking of DA unit activity correlates with the degree of SPW-R associated modulation.
- DA coordination with SPW-R bursts depends on replay content:
    - Replay of spatial sequences is associated with greater modulation.
    - DA units preferentially relate to replay of reward locations.
- SPW-R modulation of DA units is reduced in slow wave sleep.

Predicting the future for sequential decision-making

Rodent models of predictive evaluation – What happens next?

Sam Gershman, Matt Wilson, Hector Penagos

Gershman dual-system model

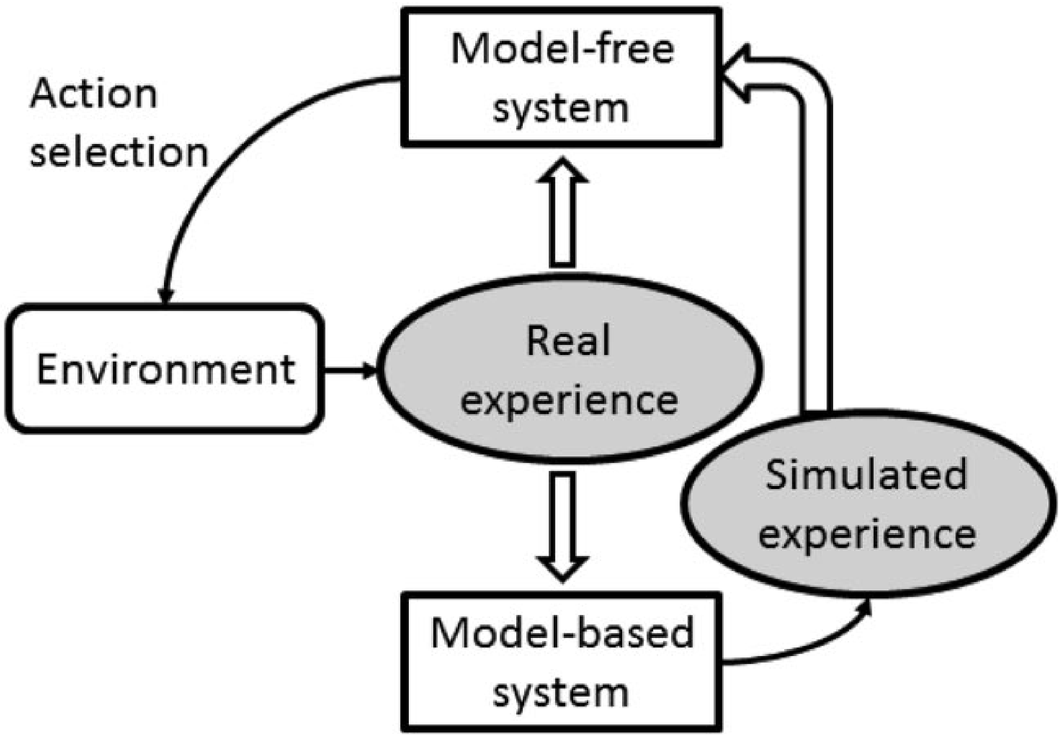

Figure 1. The DYNA architecture (from Sutton, 1990). The environment furnishes the agent with real experience (transitions and rewards), whereas the model-based system furnishes the agent with simulated experience by replaying state–action pairs from memory and then using a learned model of the environment to generate transitions and rewards. The model-free system learns in the same way from both real and simulated experience, and uses its learned values to control action selection. (Gershman,Markman,Otto 2012)

Gershman human expt design

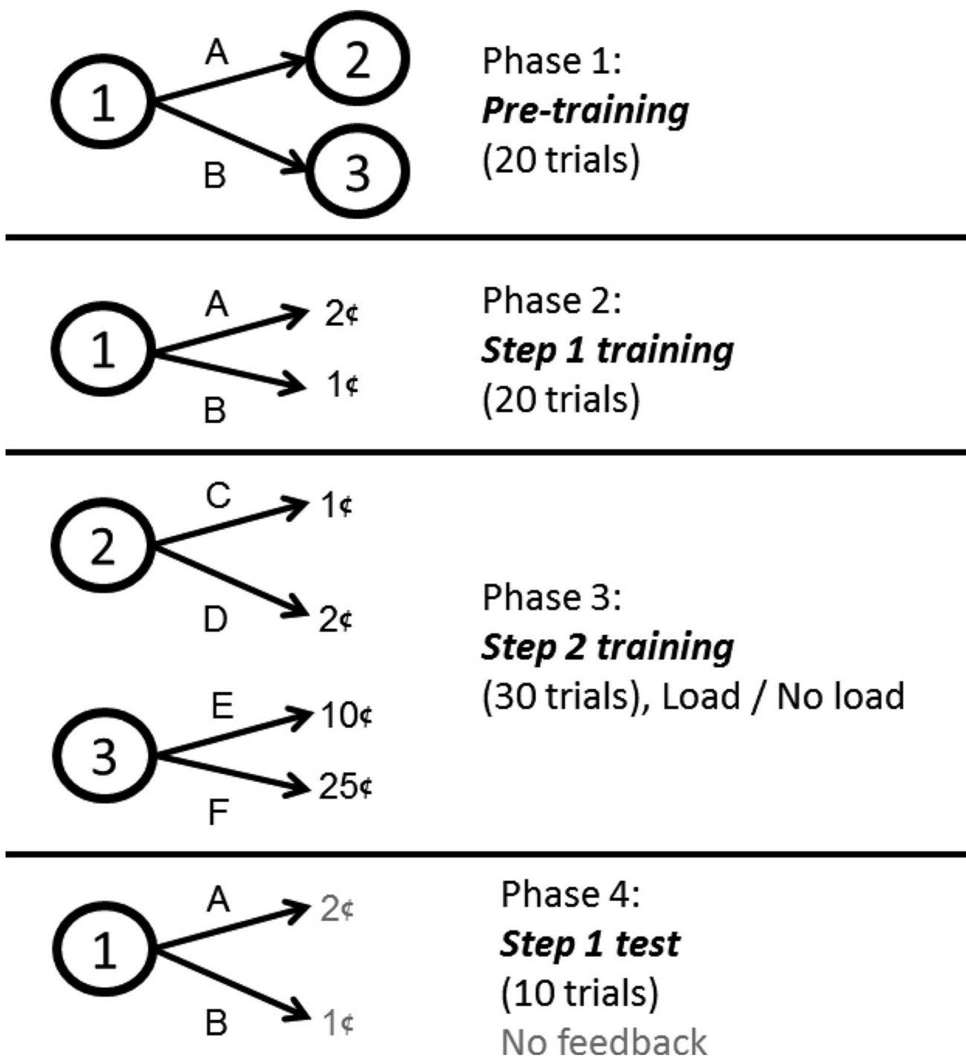

Figure 2. Experimental design. The sequential decision problem consists of three states (indicated by numbered circles) and two mutually exclusive actions in each state (indicated by letters). Deterministic transitions between states conditional upon the chosen action are indicated by arrows. Rewards for each state–action pair are indicated by amounts (in cents). In Phase 4, reward feedback is delayed until the end of the phase. In the task interface, states are signaled by background colors, and actions are signaled by fractal images. (Gershman,Markman, Otto 2012)

Rest period revaluation dependency

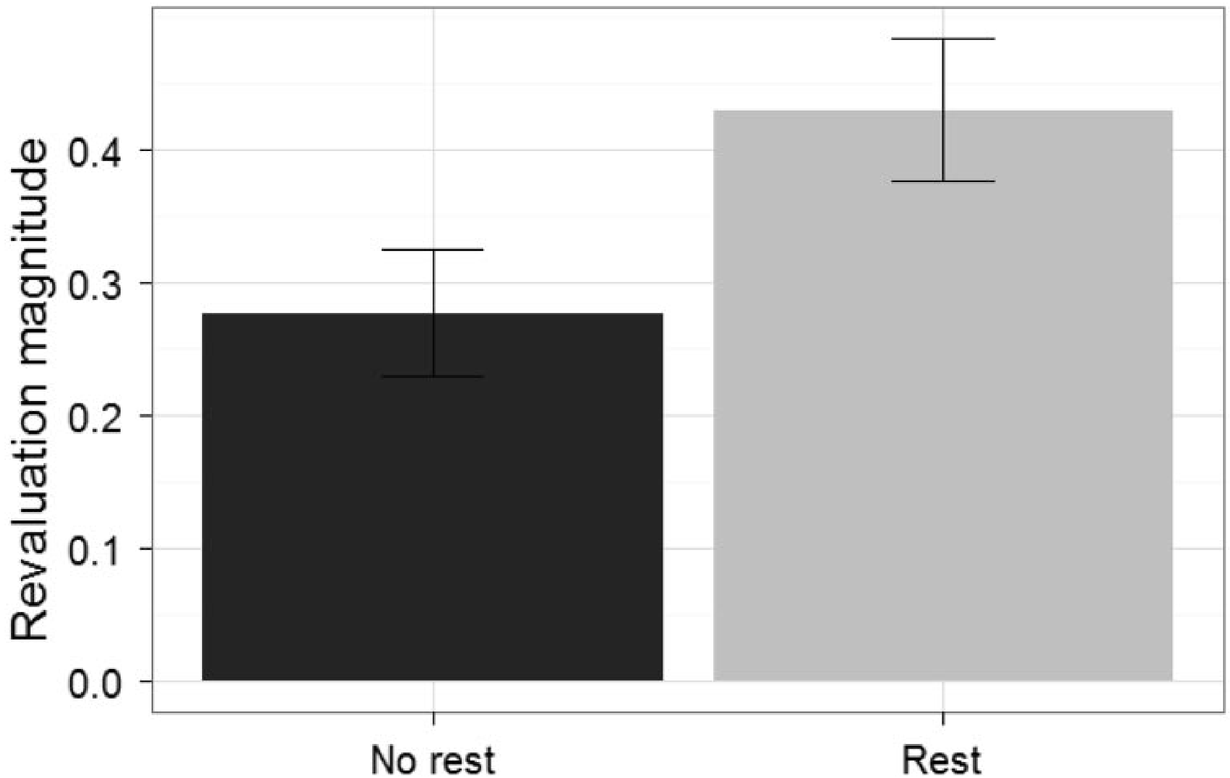

Figure 5. Experiment 3 results. Revaluation magnitude in Phase 4 as a function of rest condition. Both conditions were performed under working memory load. Error bars indicate standard error of the mean. (Gershman,Markman, Otto 2012)

Rodent Retrospective Revaluation task

Training sessions:

1. Starting at 1, differentially reward A and B. Assess choice preference.

2. Starting at , differentially reward C and D. Assess choice preference.

3. Starting at , differentially reward E and F. Assess choice preference.

Test sessions:

4. Starting at 1, equally reward A and B. Assess choice preference. Compare to preference during initial training session.

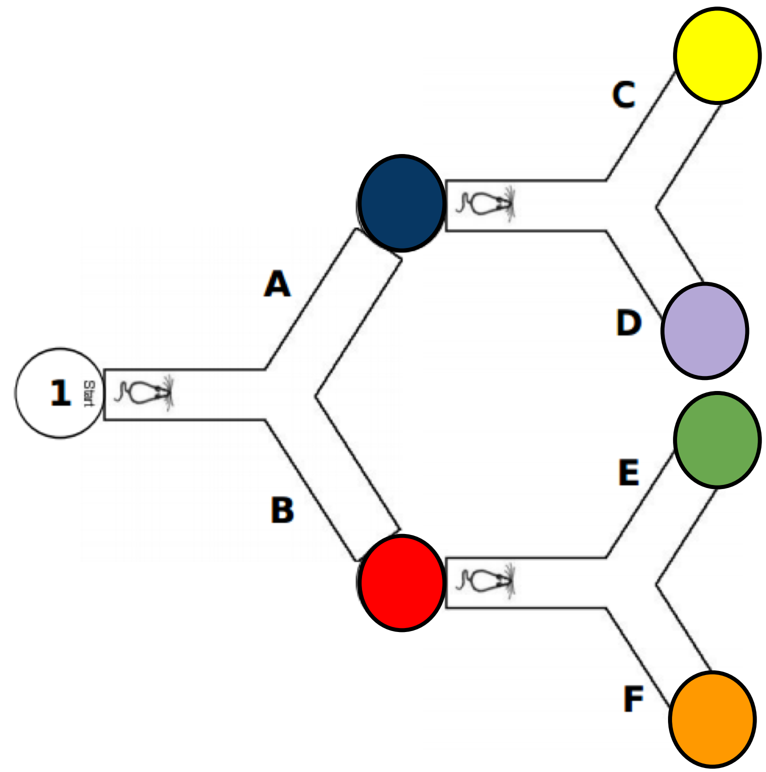

How can offline state contribute to integration across experiences?

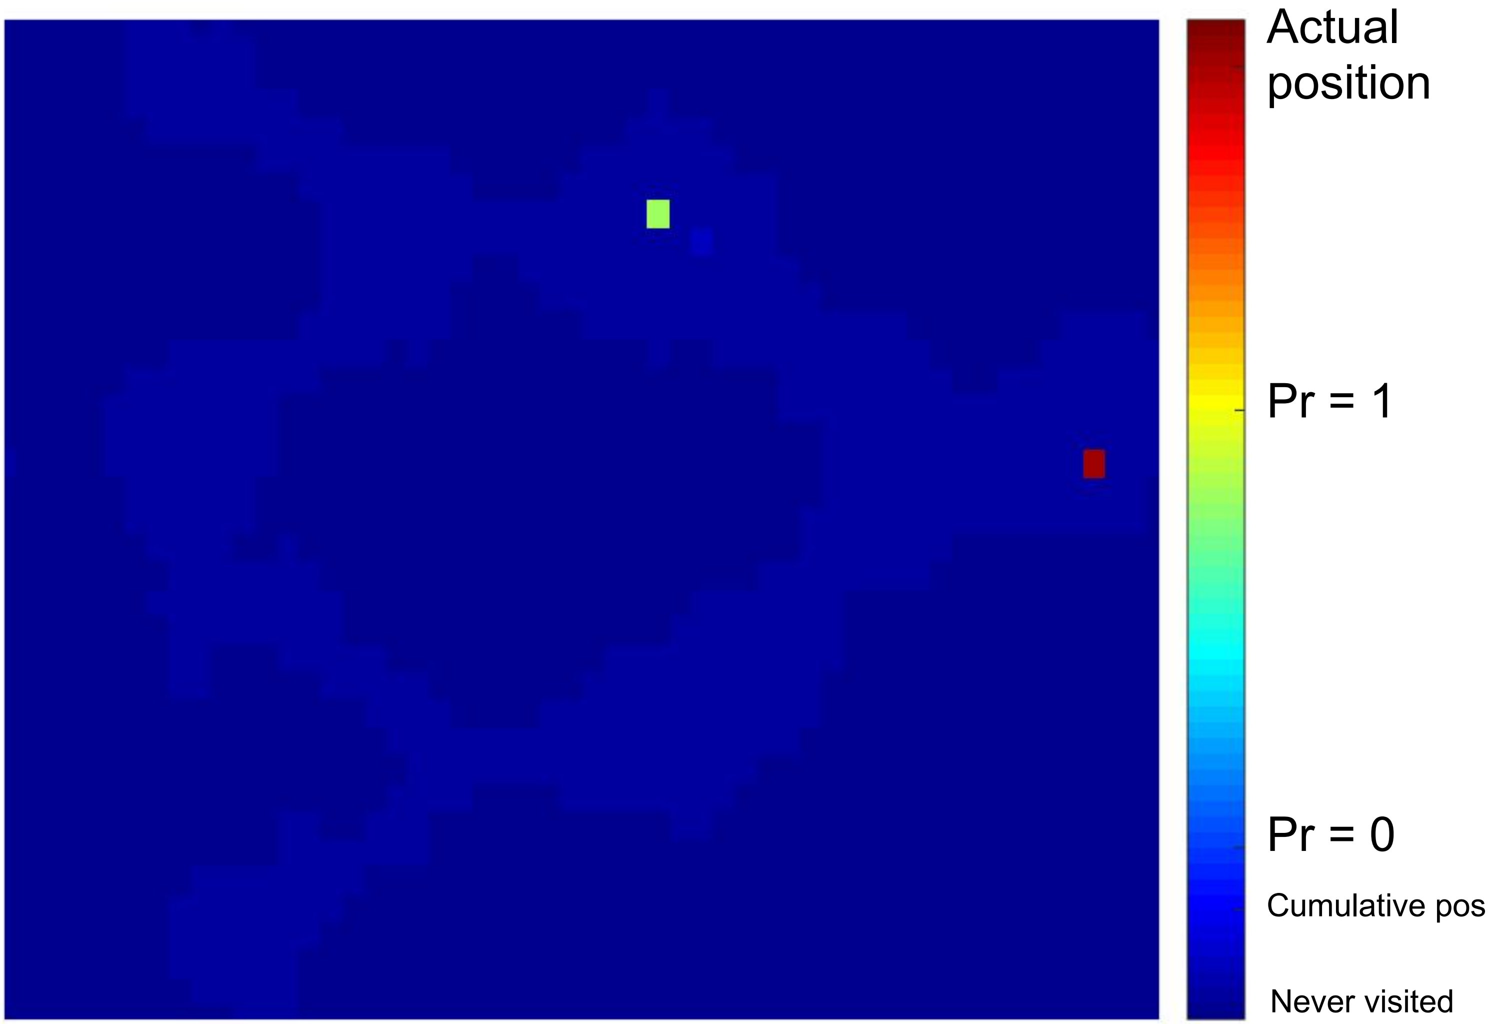

Unsorted spikes above 65 μV, decoding bin = 250 ms

Track 1

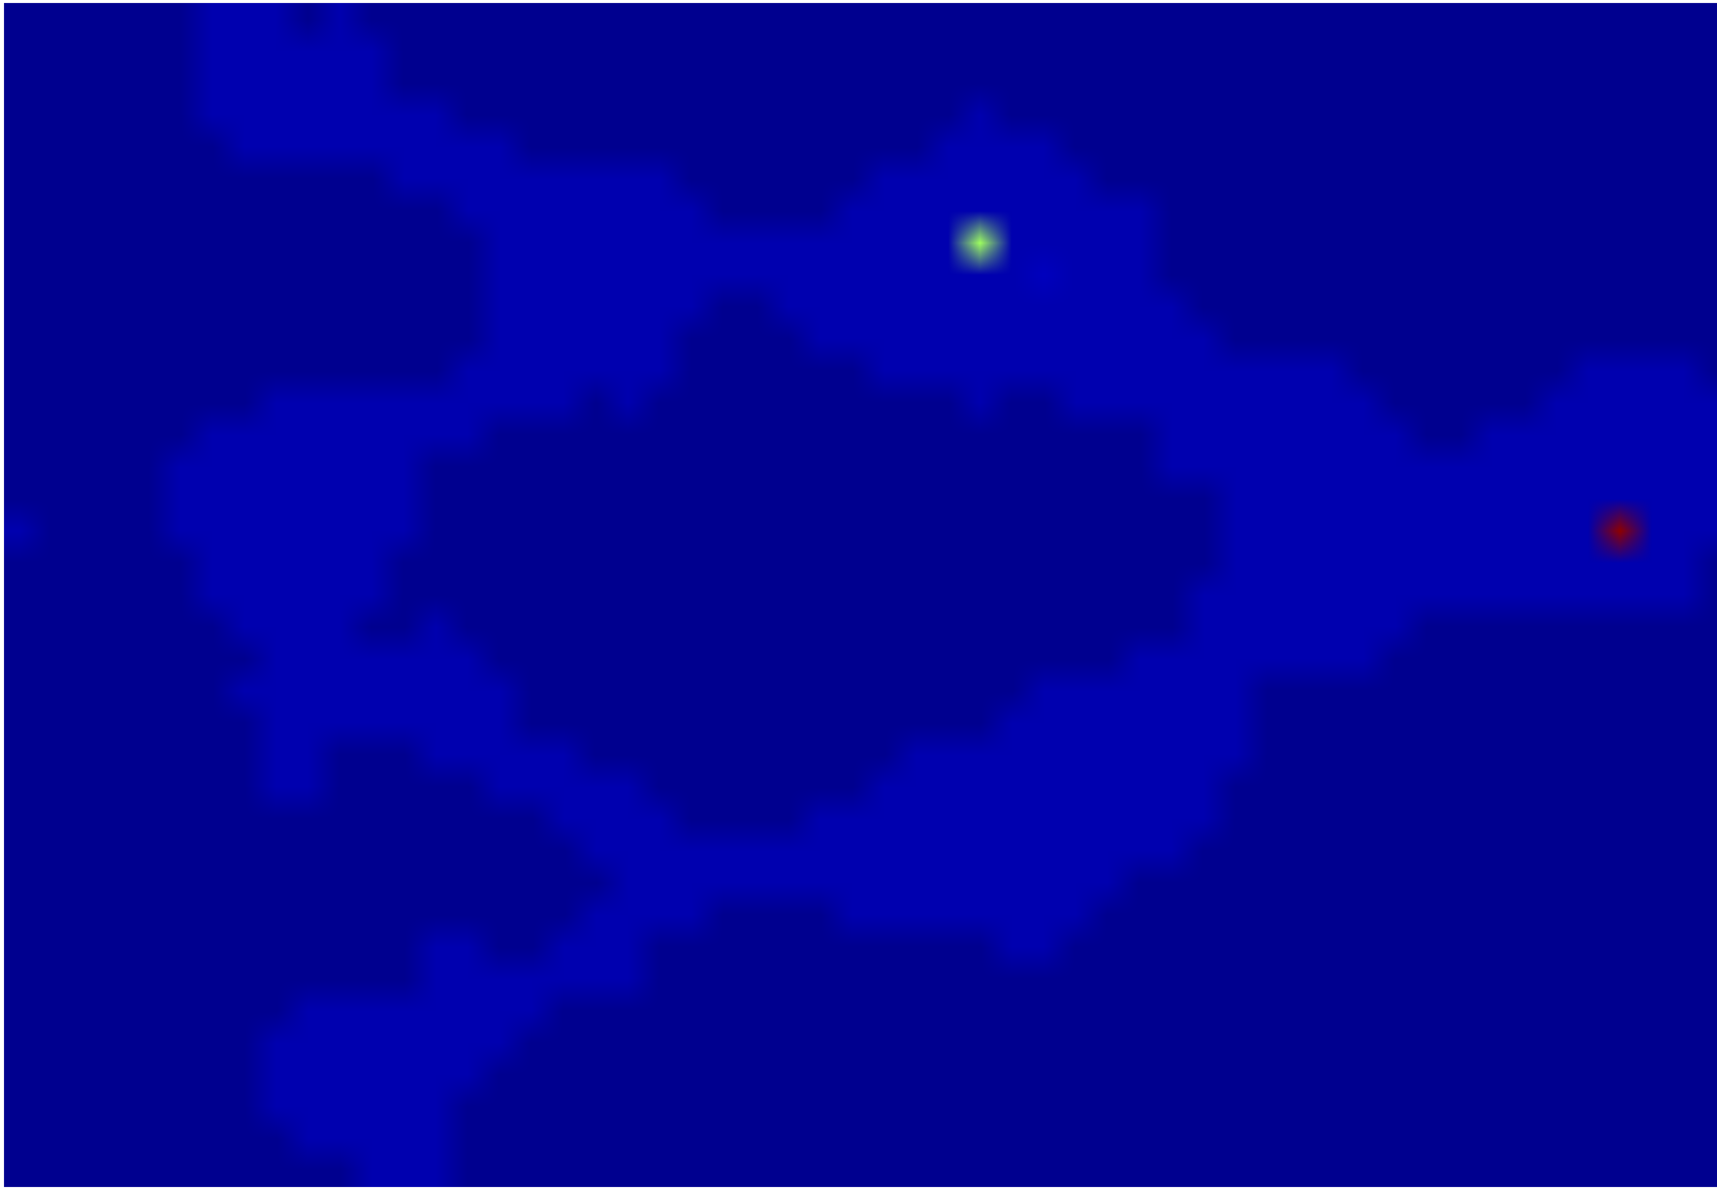

Track 2

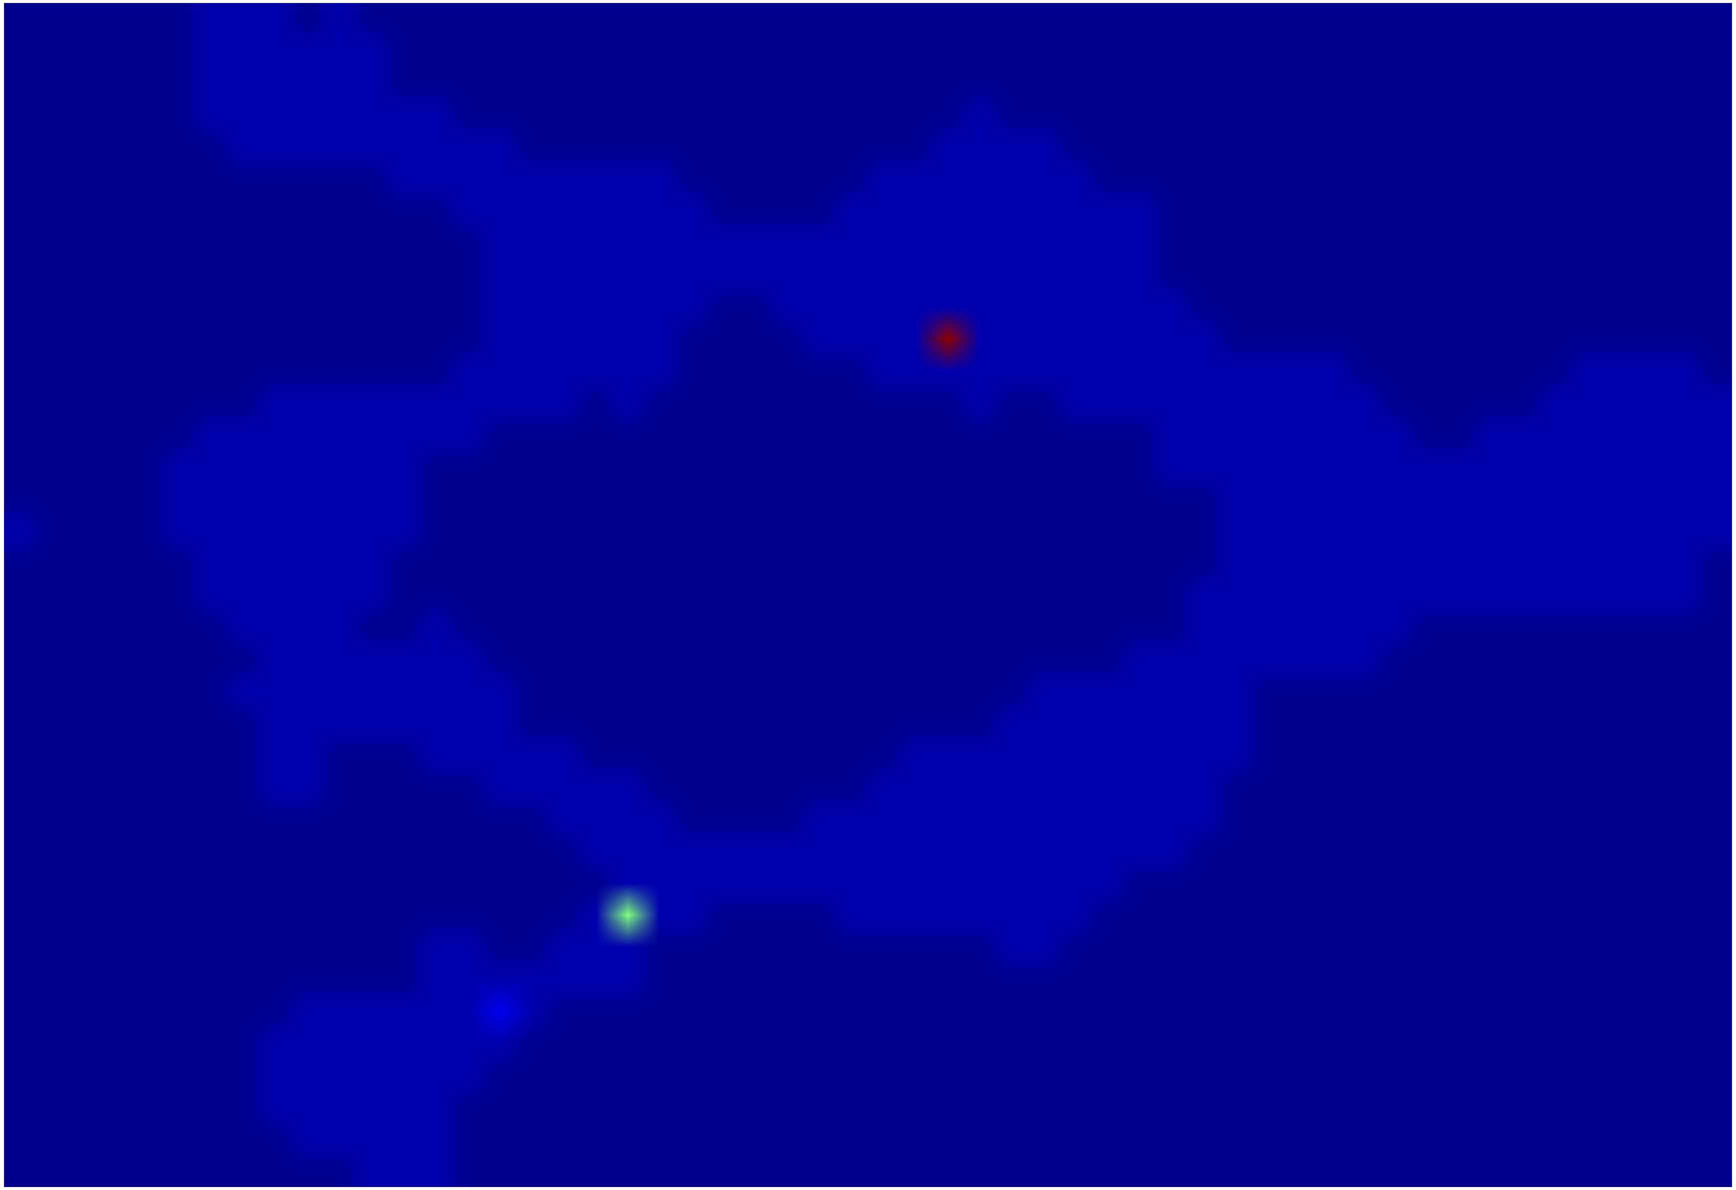

Track 3

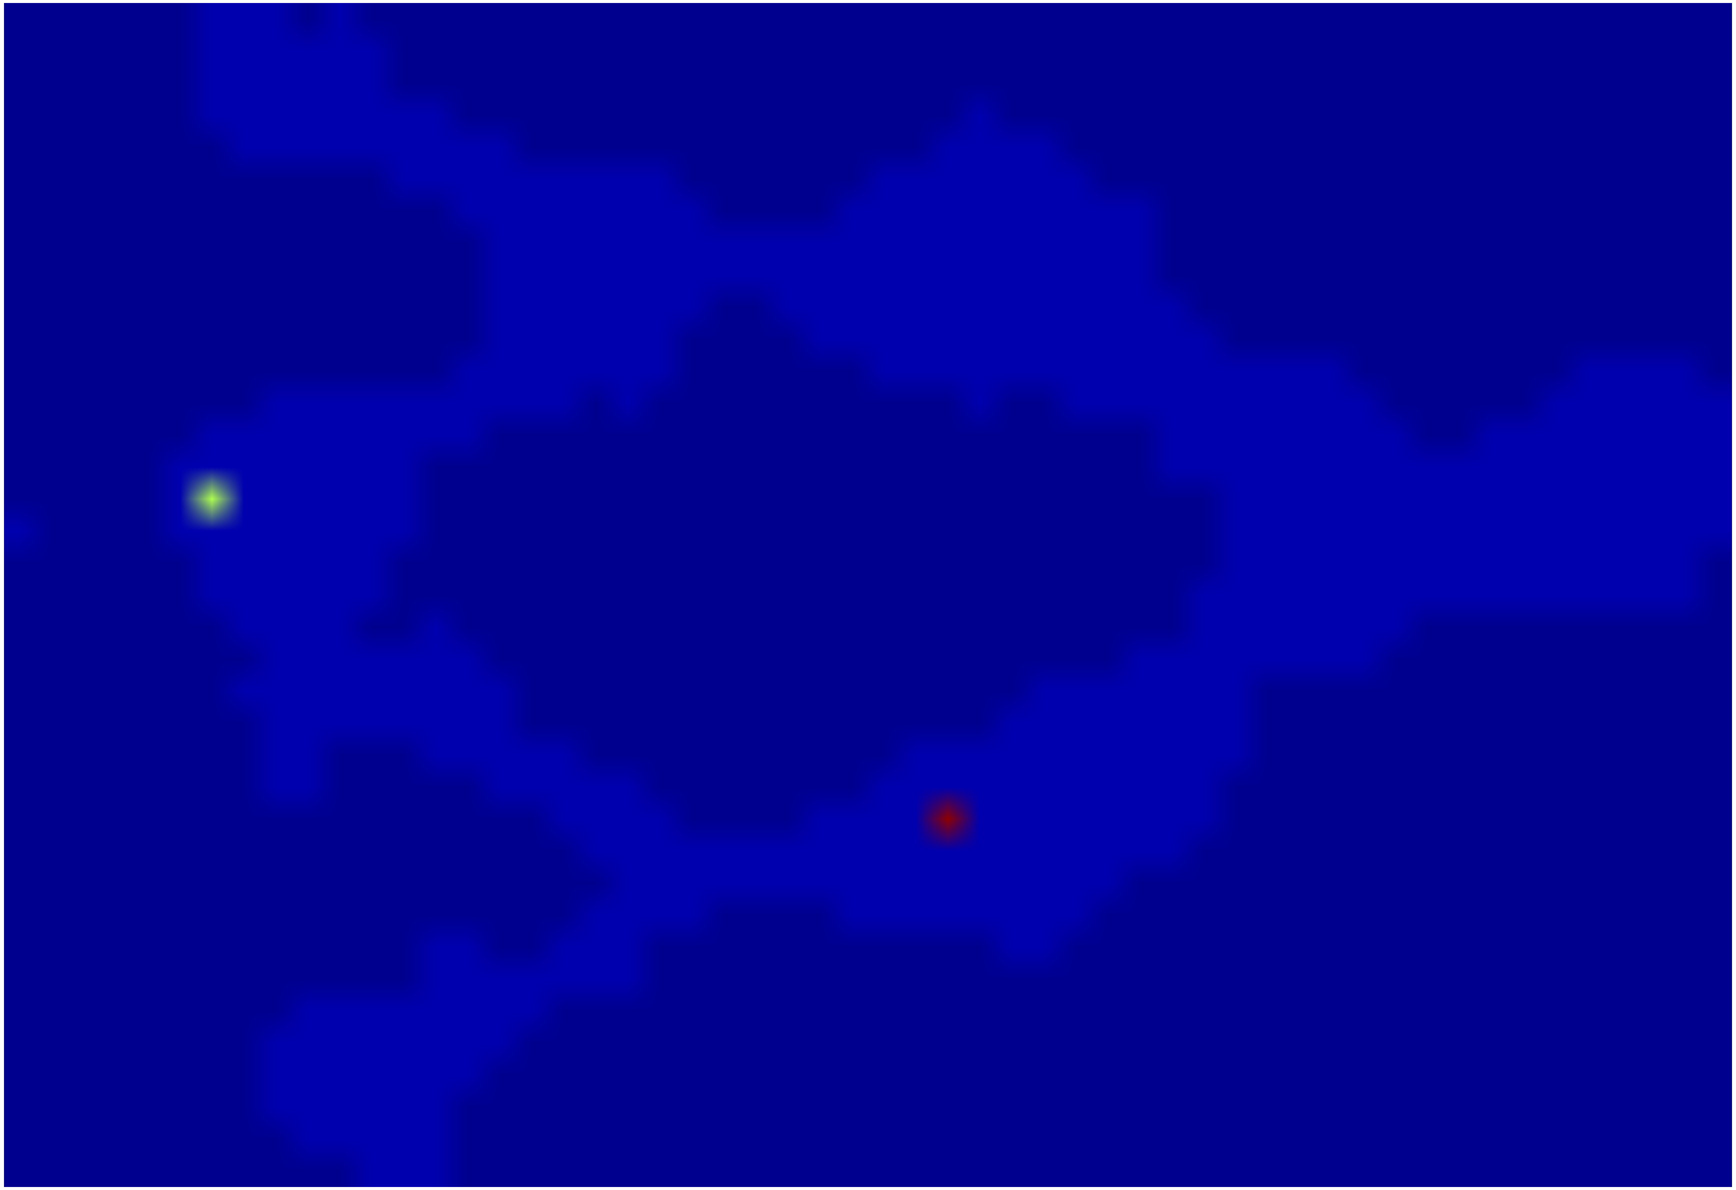

Test

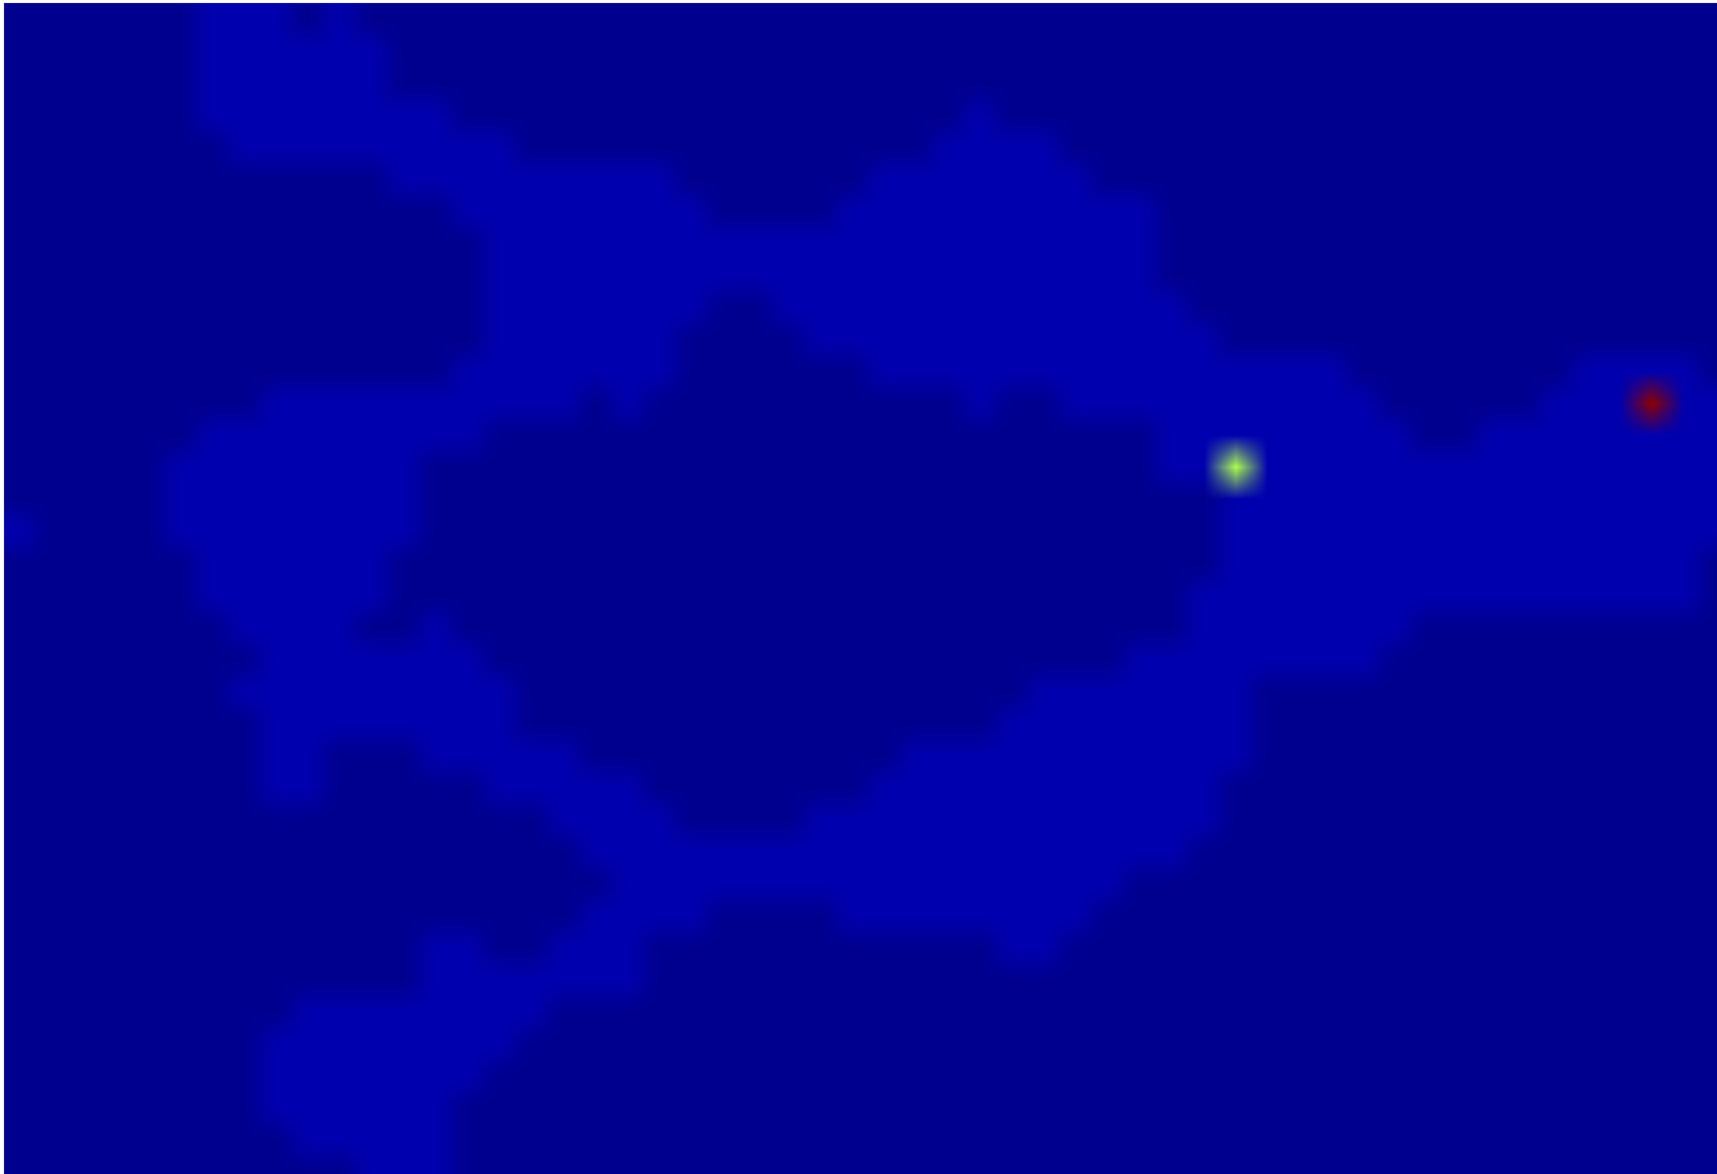

Hippocampus may incorporate maze representations to cortical targets during expression of offline state

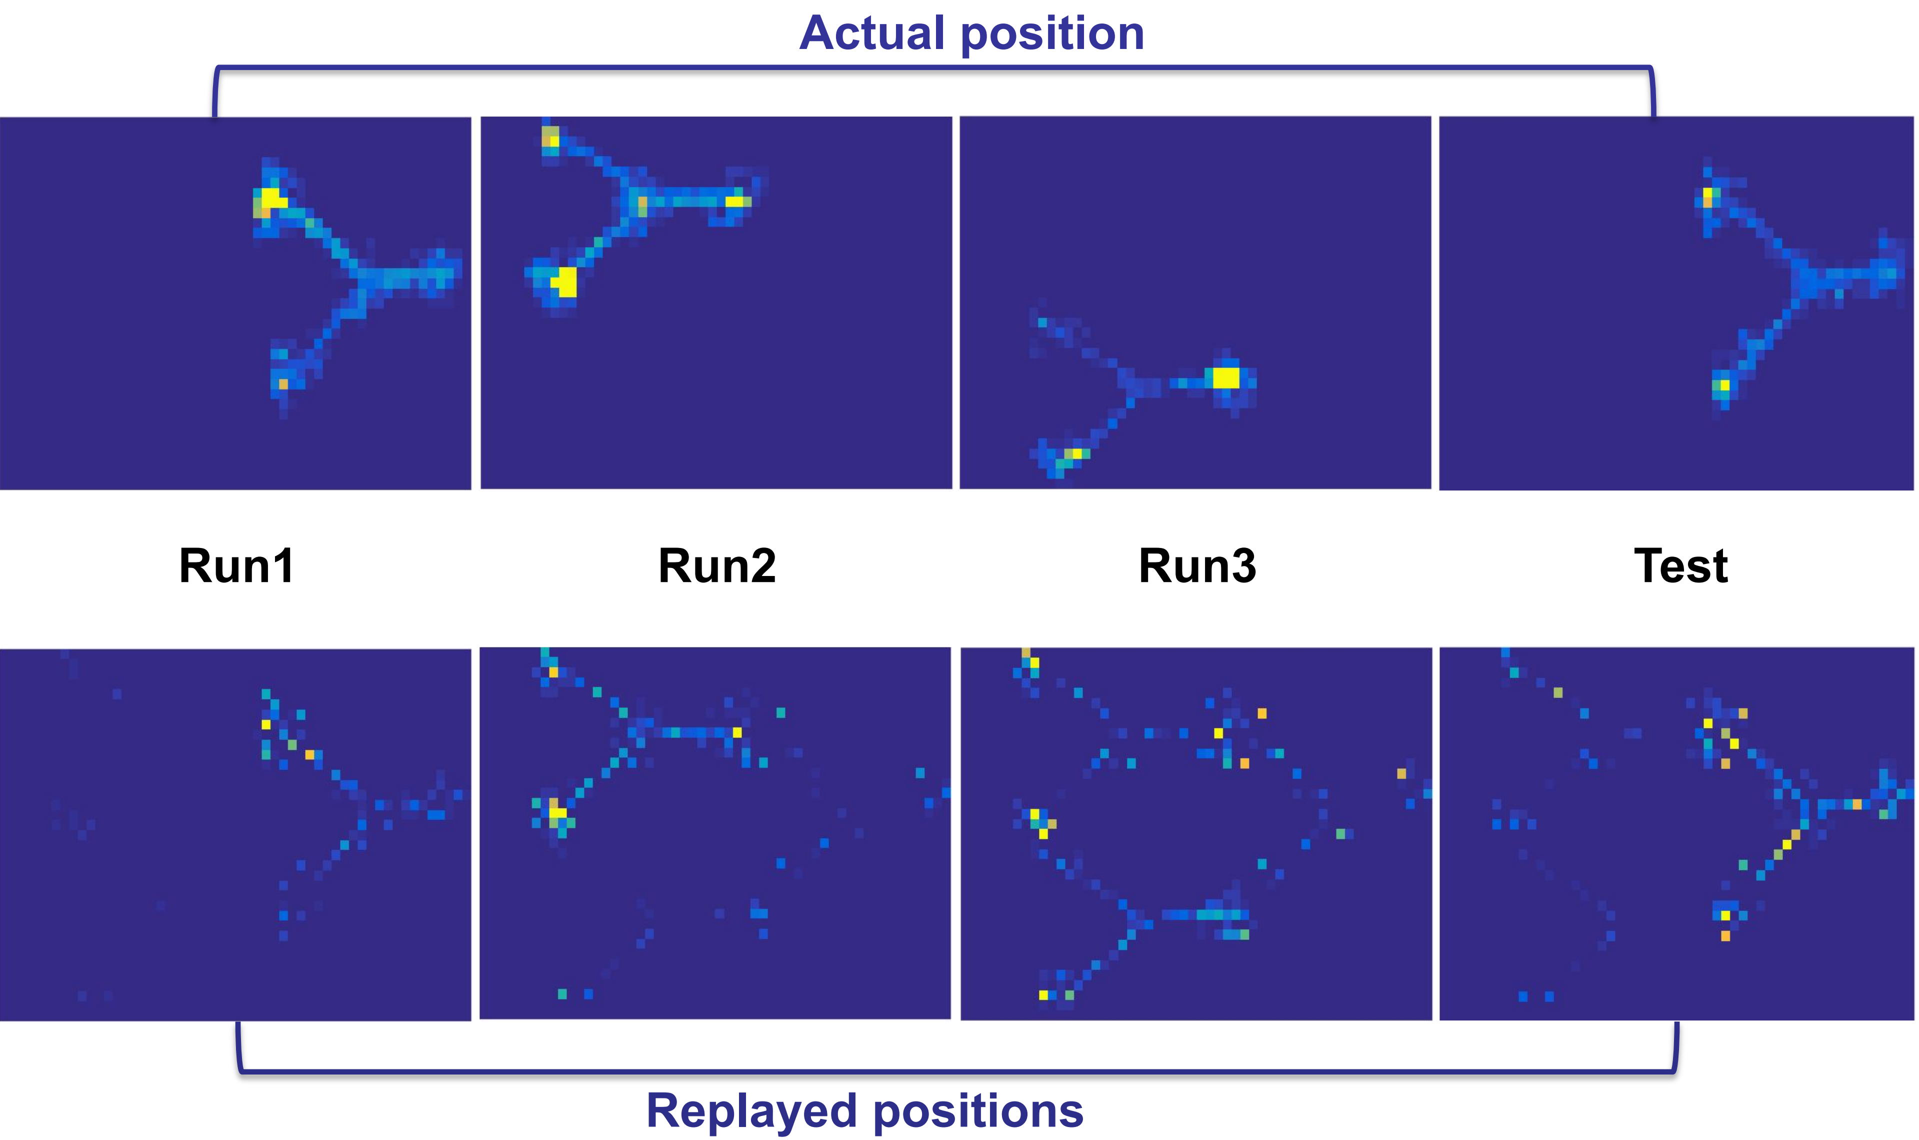

Hippocampal representations during sleep include three mazes consistent with action/state evaluations

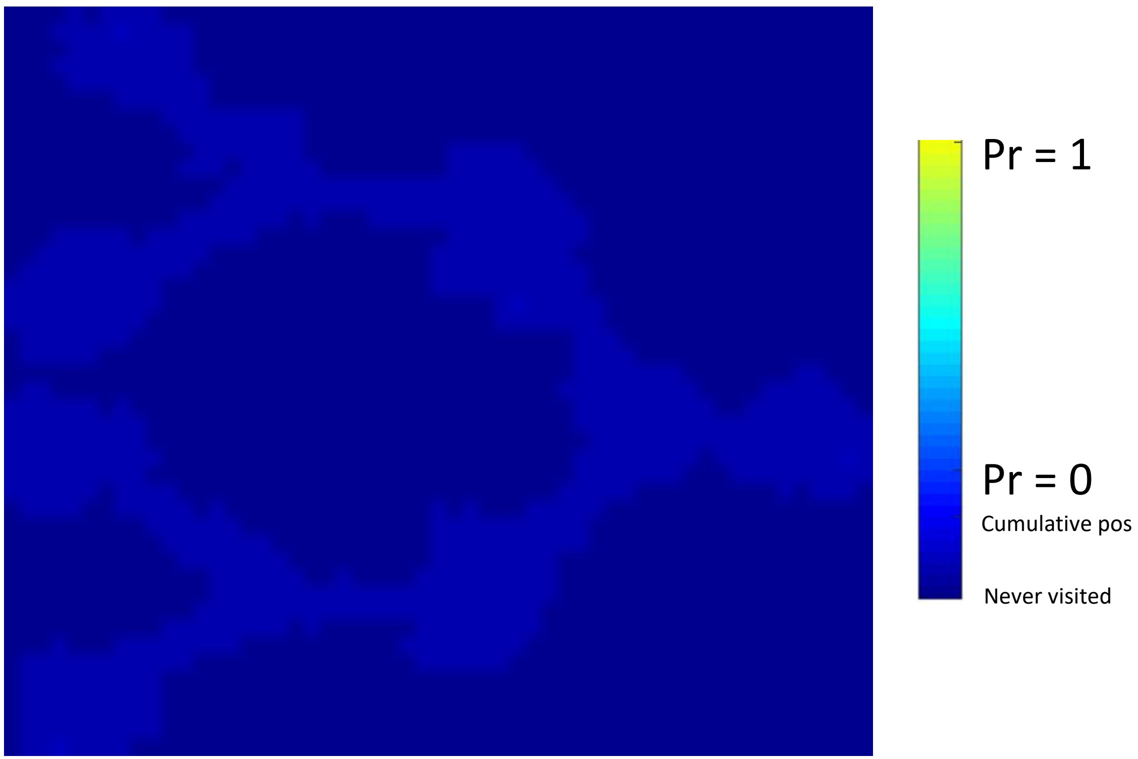

Chaining of sequences

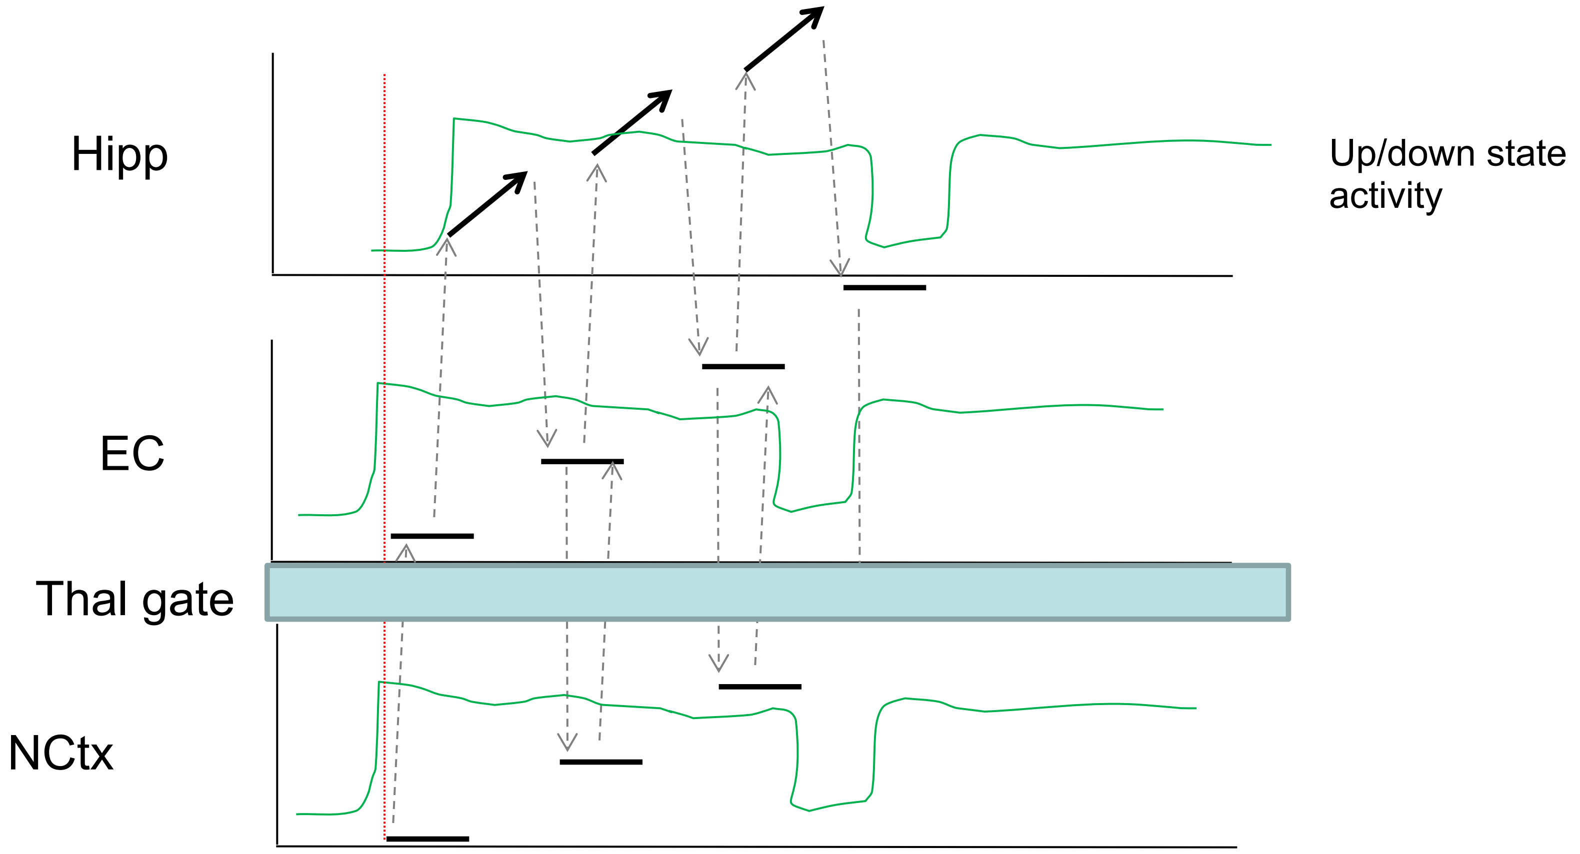

Hinton's Recurrent Temporal Restricted Boltzmann Machine (RTBM) architecture

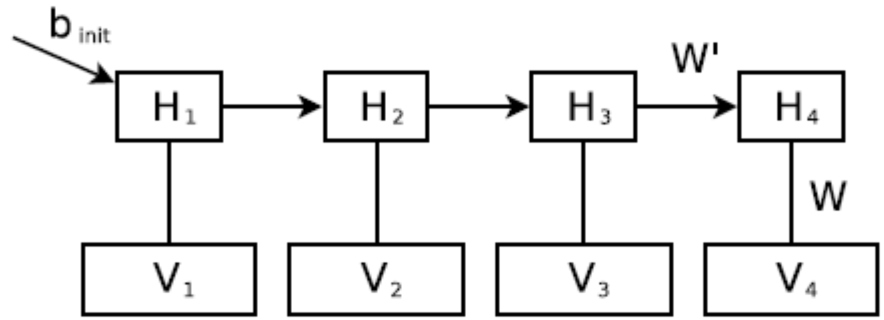

Figure 1: The graphical structure of a TRBM: a directed sequence of RBMs.

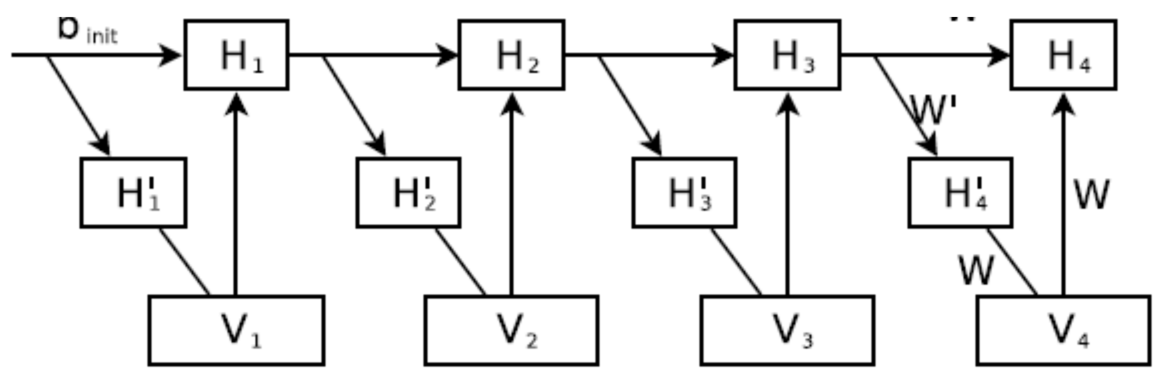

Figure 2: The graphical structure of the RTRBM, Q. The variables H_t are real valued while the variables H'_t are binary. The conditional distribution $Q(V_t, H'_t|h_{t-1})$ is given by the equation $Q(v_t, h'_t|h_{t-1}) = \exp\left(v_t^\top W h'_t + v_t^\top b_V + h'_t(b_H + W' h_{t-1})\right)/Z(h_{t-1})$, which is essentially the same as the TRBM's conditional distribution P from equation 5. We will always integrate out H'_t and will work directly with the distribution $Q(V_t|h_{t-1})$. Notice that when V_1 is observed, H'_1 cannot affect H_1.

Sleep and Learning

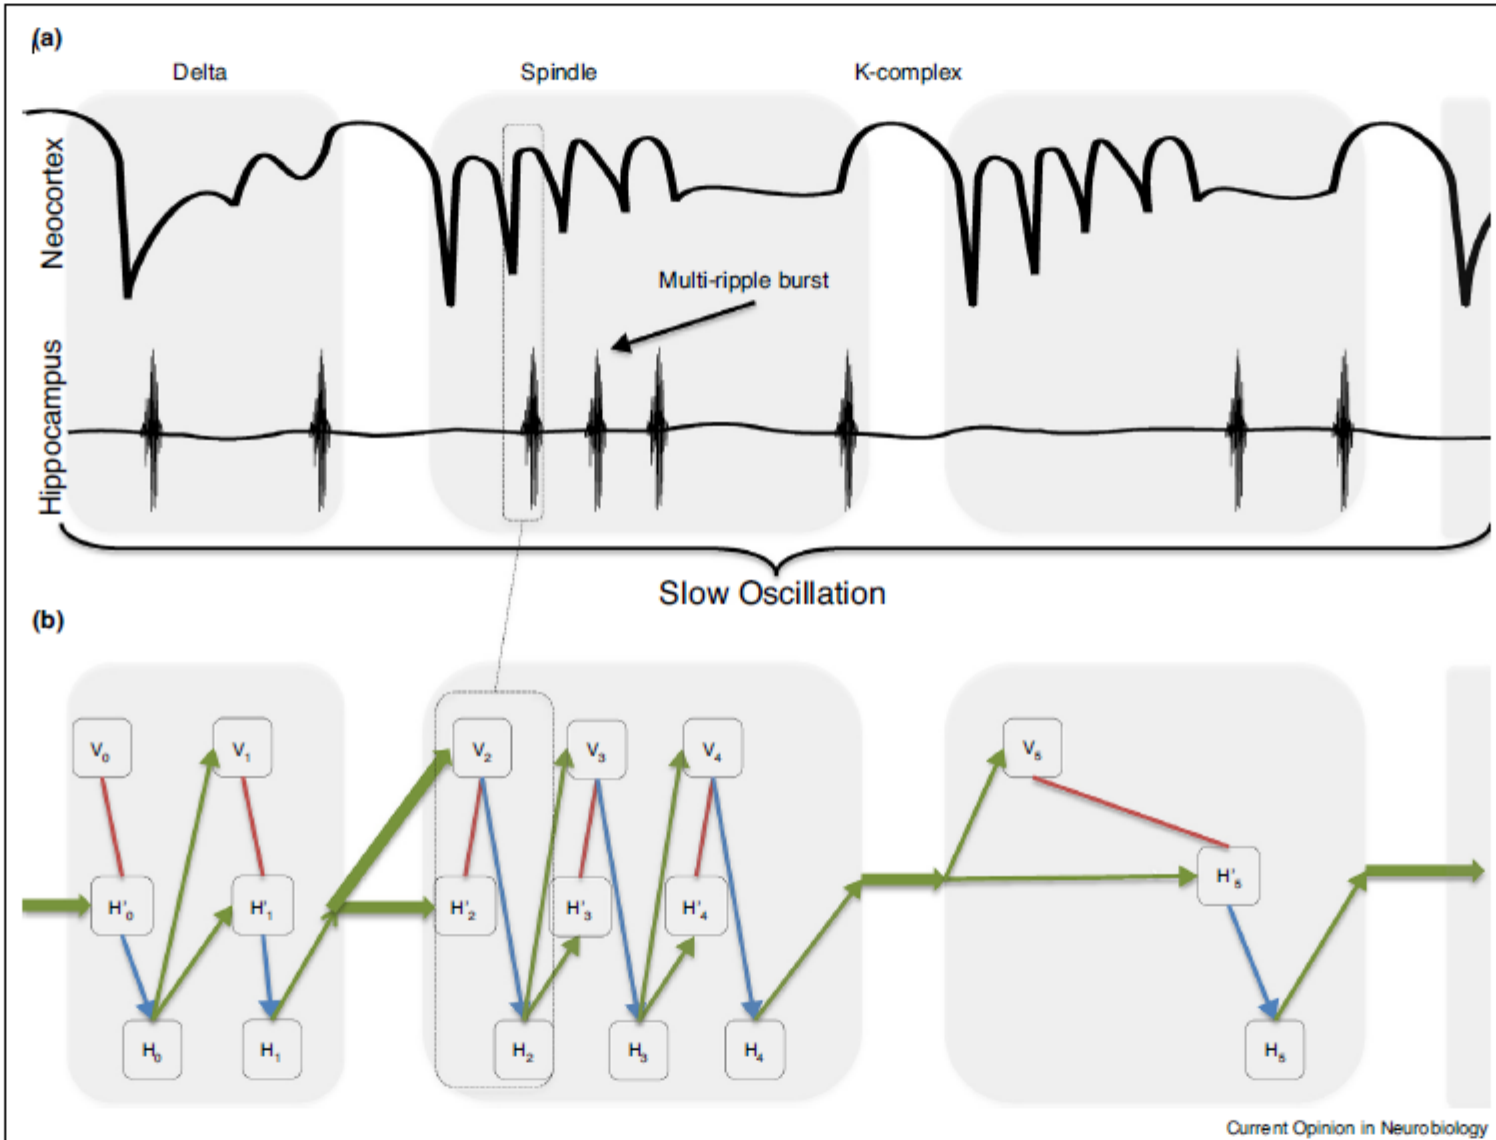

Penagos, Varela, Wilson, *Curr. Opin. Neuro.* 2017

Mastering the game of Go with deep neural networks and tree search

David Silver[1*], Aja Huang[1*], Chris J. Maddison[1], Arthur Guez[1], Laurent Sifre[1], George van den Driessche[1], Julian Schrittwieser[1], Ioannis Antonoglou[1], Veda Panneershelvam[1], Marc Lanctot[1], Sander Dieleman[1], Dominik Grewe[1], John Nham[2], Nal Kalchbrenner[1], Ilya Sutskever[2], Timothy Lillicrap[1], Madeleine Leach[1], Koray Kavukcuoglu[1], Thore Graepel[1] & Demis Hassabis[1]

The game of Go has long been viewed as the most challenging of classic games for artificial intelligence owing to its enormous search space and the difficulty of evaluating board positions and moves. Here we introduce a new approach to computer Go that uses 'value networks' to evaluate board positions and 'policy networks' to select moves. These deep neural networks are trained by a novel combination of supervised learning from human expert games, and reinforcement learning from games of self-play. Without any lookahead search, the neural networks play Go at the level of state-of-the-art Monte Carlo tree search programs that simulate thousands of random games of self-play. We also introduce a new search algorithm that combines Monte Carlo simulation with value and policy networks. Using this search algorithm, our program AlphaGo achieved a 99.8% winning rate against other Go programs, and defeated the human European Go champion by 5 games to 0. This is the first time that a computer program has defeated a human professional player in the full-sized game of Go, a feat previously thought to be at least a decade away.

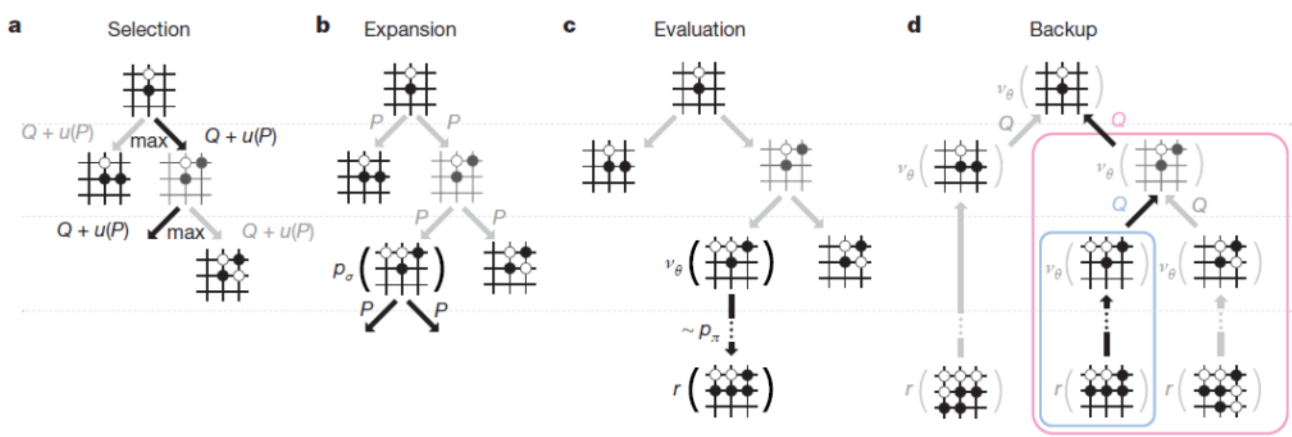

Figure 3 | Monte Carlo tree search in AlphaGo. a, Each simulation traverses the tree by selecting the edge with maximum action value Q, plus a bonus $u(P)$ that depends on a stored prior probability P for that edge. **b,** The leaf node may be expanded; the new node is processed once by the policy network p_σ and the output probabilities are stored as prior probabilities P for each action. **c,** At the end of a simulation, the leaf node is evaluated in two ways: using the value network v_θ; and by running a rollout to the end of the game with the fast rollout policy p_π, then computing the winner with function r. **d,** Action values Q are updated to track the mean value of all evaluations $r(\cdot)$ and $v_\theta(\cdot)$ in the subtree below that action.

Silver et al., *Nature* 2016

Vector-based navigation using grid-like representations in artificial agents

Andrea Banino, Caswell Barry, [...] Dharshan Kumaran

Nature **557**, 429–433 (2018) | Download Citation

Fig. 2: One-shot open field navigation to a hidden goal.

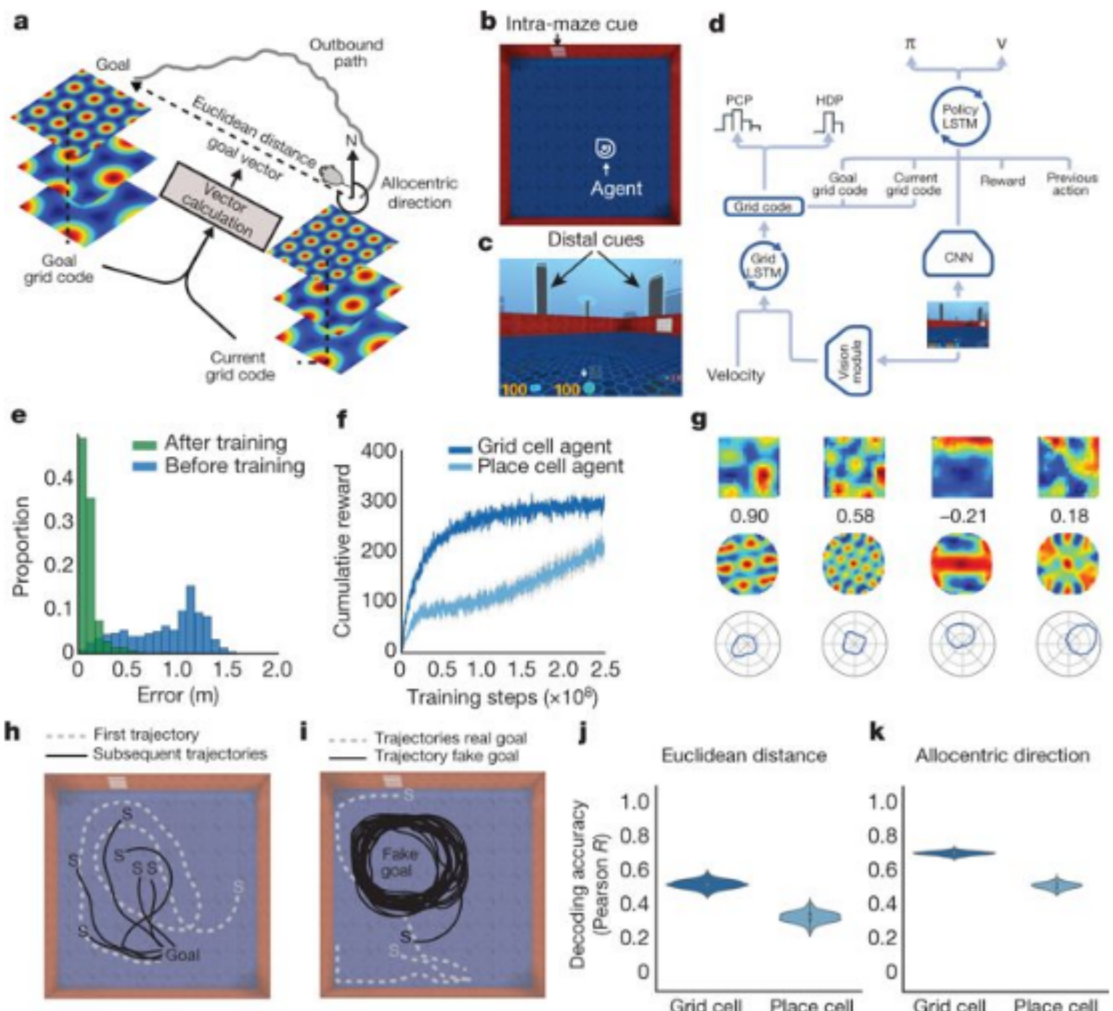

Compositional action planning

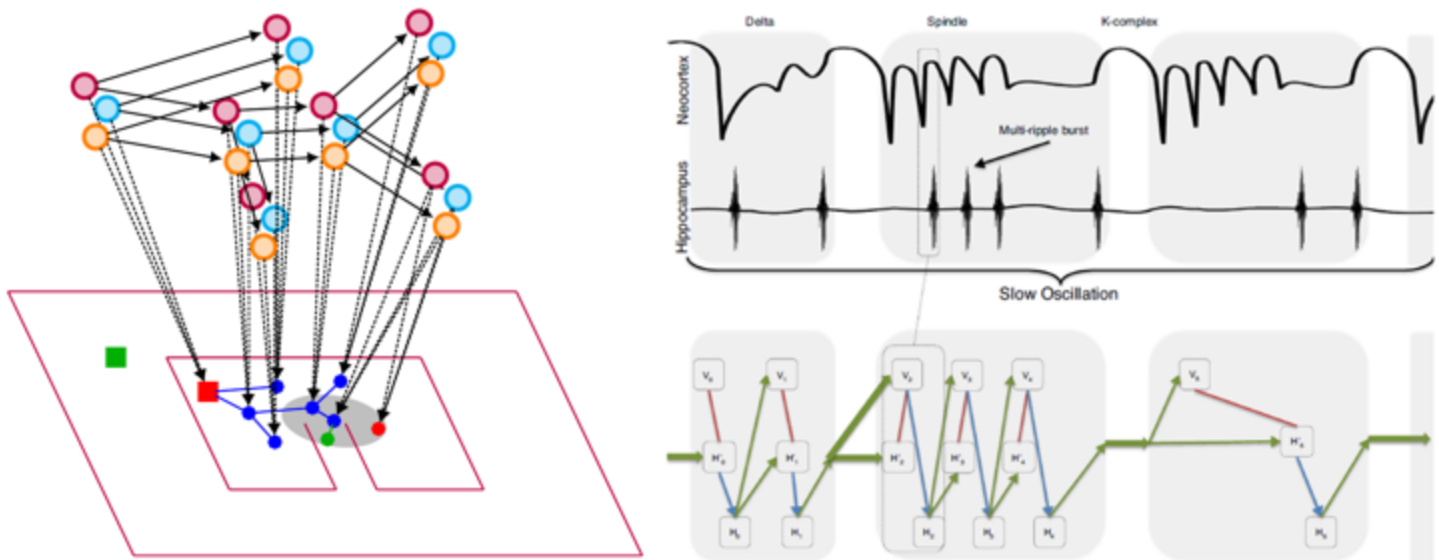

Figure 5. *(left) Schematic of an extended compositional model for action planning[42]. Recurrent networks (circles at the top) are the primitives that combine together (as in* ***Fig. 2****). To reason about actions, the network maintains a tree of decisions, shown in blue, while it explores the map below it from its start position (red) to its goal (green.) It makes decisions by expanding the tree into the future, reasoning about which decisions are more likely to pay off, and then committing to a particular decision. (right) Recordings from rat hippocampus and neocortex (top). Memories are acquired, used, and consolidated through this pathway where information is encoded by slow oscillations in neocortex and ripples in hippocampus[58]. Reactivated memories take the form of spatial state trajectories with recurrent Hidden Markov-like structure (bottom). We hypothesize that this recurrent state sequence processing is reflected in the coordinated interactions between the hippocampus and neocortex, and, when expressed during rodent navigational tasks (left, and* ***Fig. 6****), take the form of compositional spatial trajectories. In analogy to Challenge 1, here we will decode plans as they are recalled using combinations of sequence models.*

Kuo Y, Barbu A, Katz B. Deep sequential models for sampling-based planning. In proceedings of the International Conference on Intelligent Robots (IROS), 2018.

Overall summary

- Sequence memory can be encoded in the hippocampus during active behavior.
- Sequence memory is subsequently replayed during sleep in both the hippocampus and neocortex.
- The content of reactivated memory during sleep can be biased by external manipulation.
- Sequence memory replayed during quiet wakefulness is associated reward information and may serve a different role in learning than replay during sleep.

Wilson Lab contributors

Josh Siegle (Allen Institute)	Theta optogenetics
Matt Jones (Bristol)	Hipp-Prefrontal theta
Albert Lee (Janelia Farm)	Non-REM replay
Daoyun Ji (Baylor)	H-Visual cortex
David Foster (J. Hopkins)	Awake replay
Fabian Kloosterman (Leuven)	Extended awake replay
Tom Davidson (UCSF)	Extended awake replay
Dan Bendor (UCL)	Biased sleep replay
Steve Gomperts (Harvard/MGH)	VTA and reward
Hector Penagos (MIT)	Reward and planning