Building models of the world in replay

'SUPERSTITION' IN THE PIGEON

B. F. Skinner
Indiana University

First published in *Journal of Experimental Psychology*, *38*, 168-172.

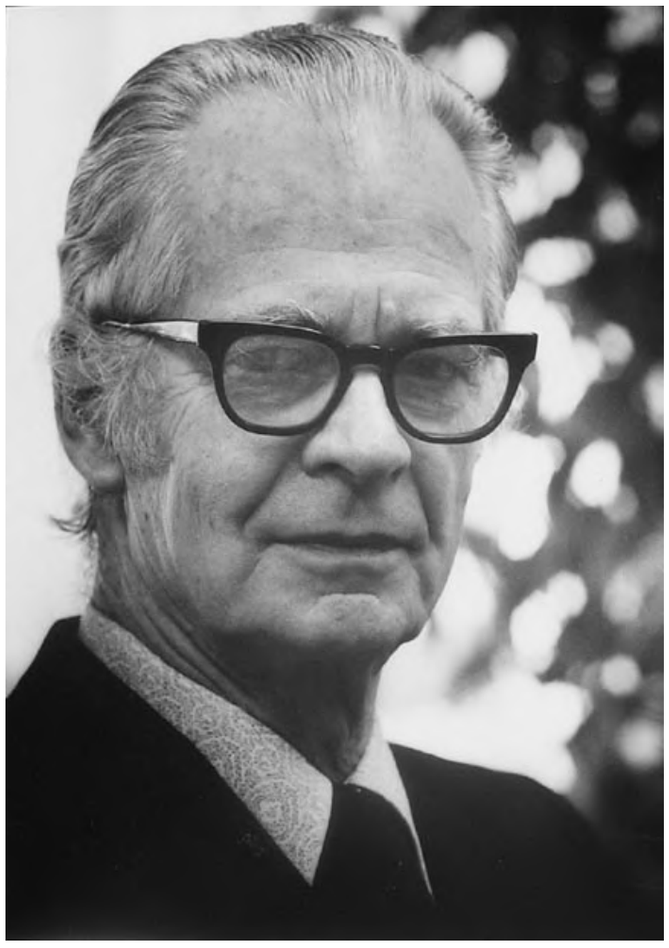

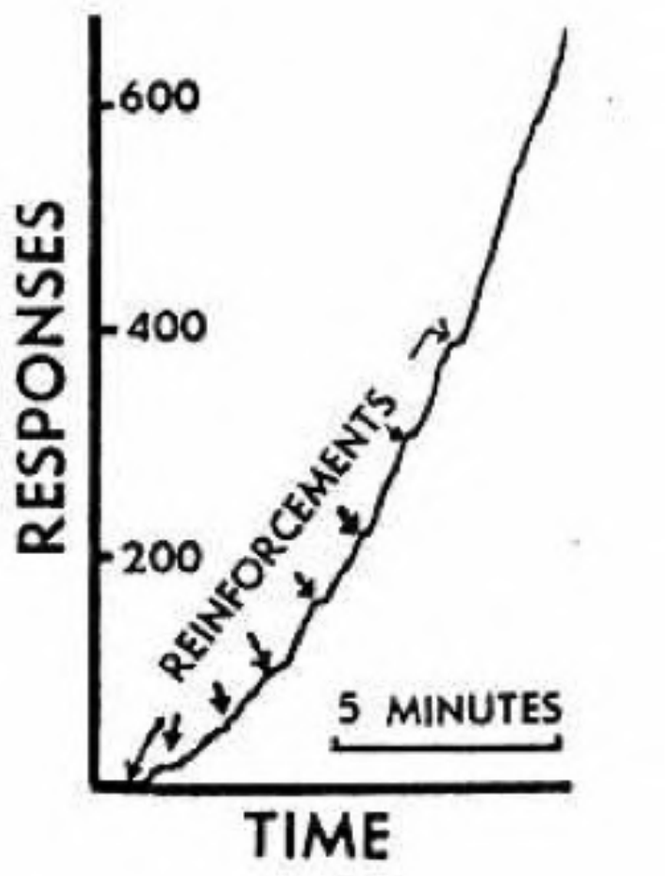

"The real problem is not whether machines think but whether men do."

"Give me a child and I'll shape him into anything."

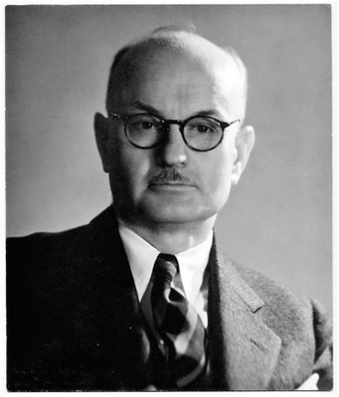

Tolman and the Cognitive Map

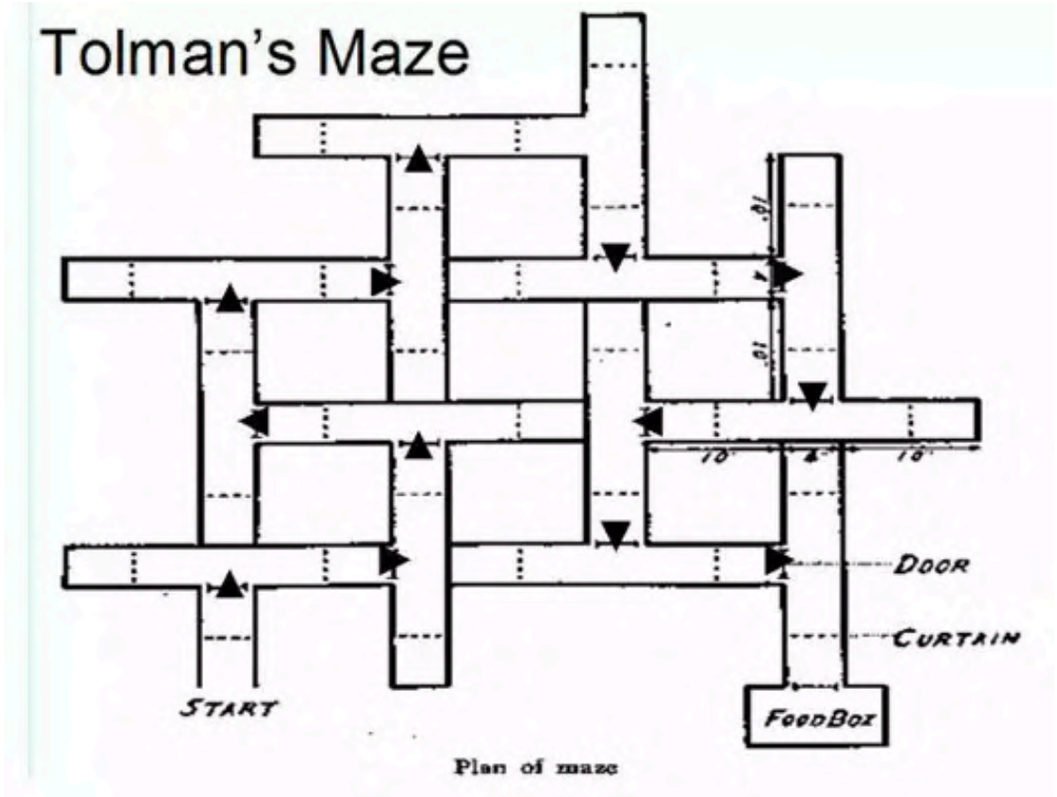

We assert that the central office itself is far more like a map control room than it is like an old fashioned telephone exchange. (Tolman 1948, p.192).

Building models of the world

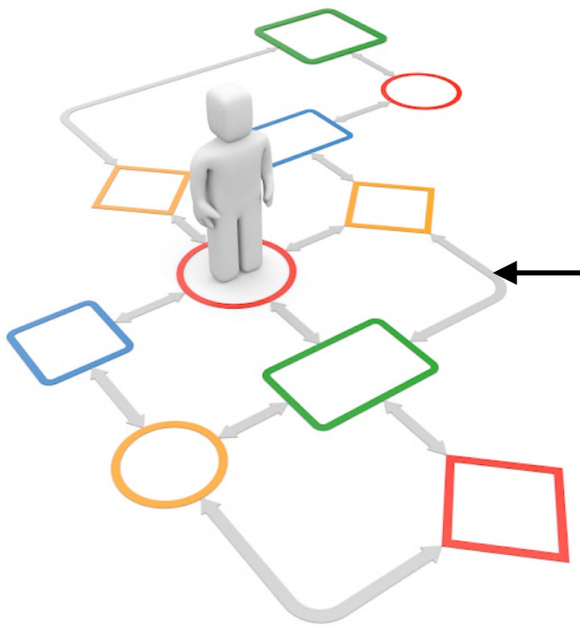

How do I build these links?

What kind of thing are we looking for?

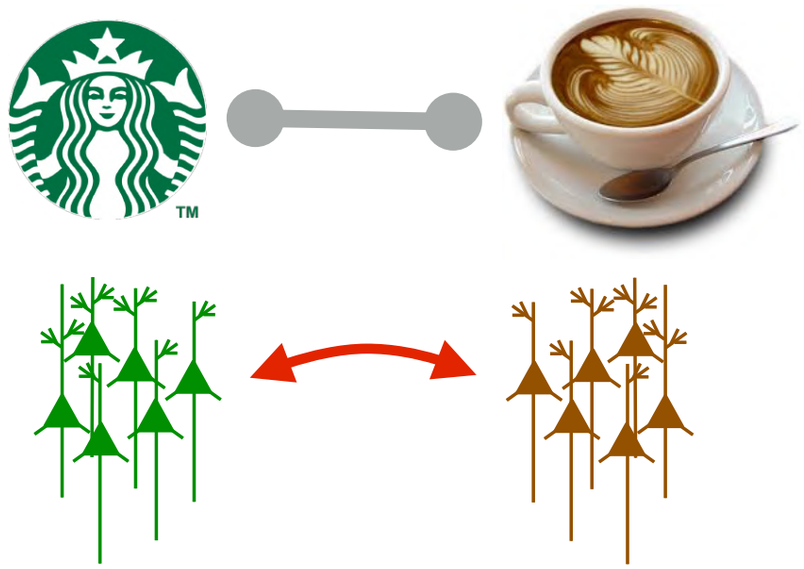

In synapses

In neurons

What kind of thing are we looking for?

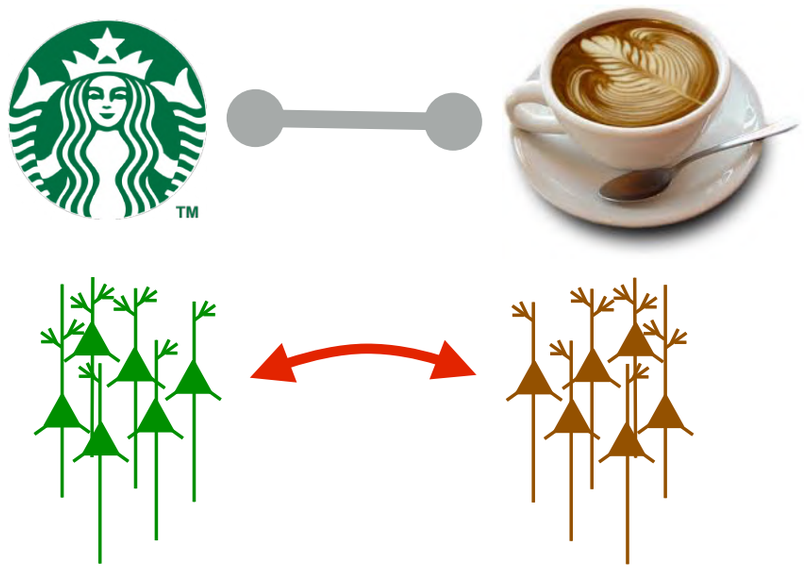

In synapses

In neurons

Using structural knowledge for inference

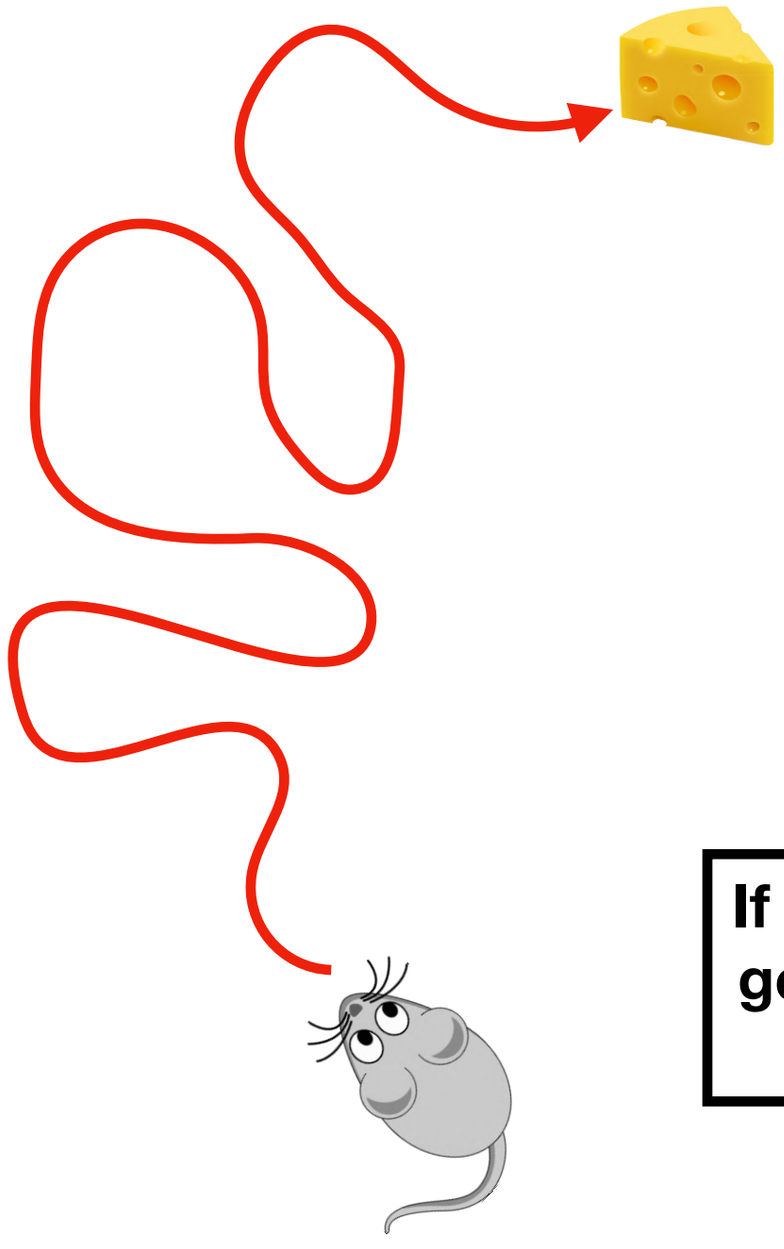

If I am just repeating good actions, I have to do this again

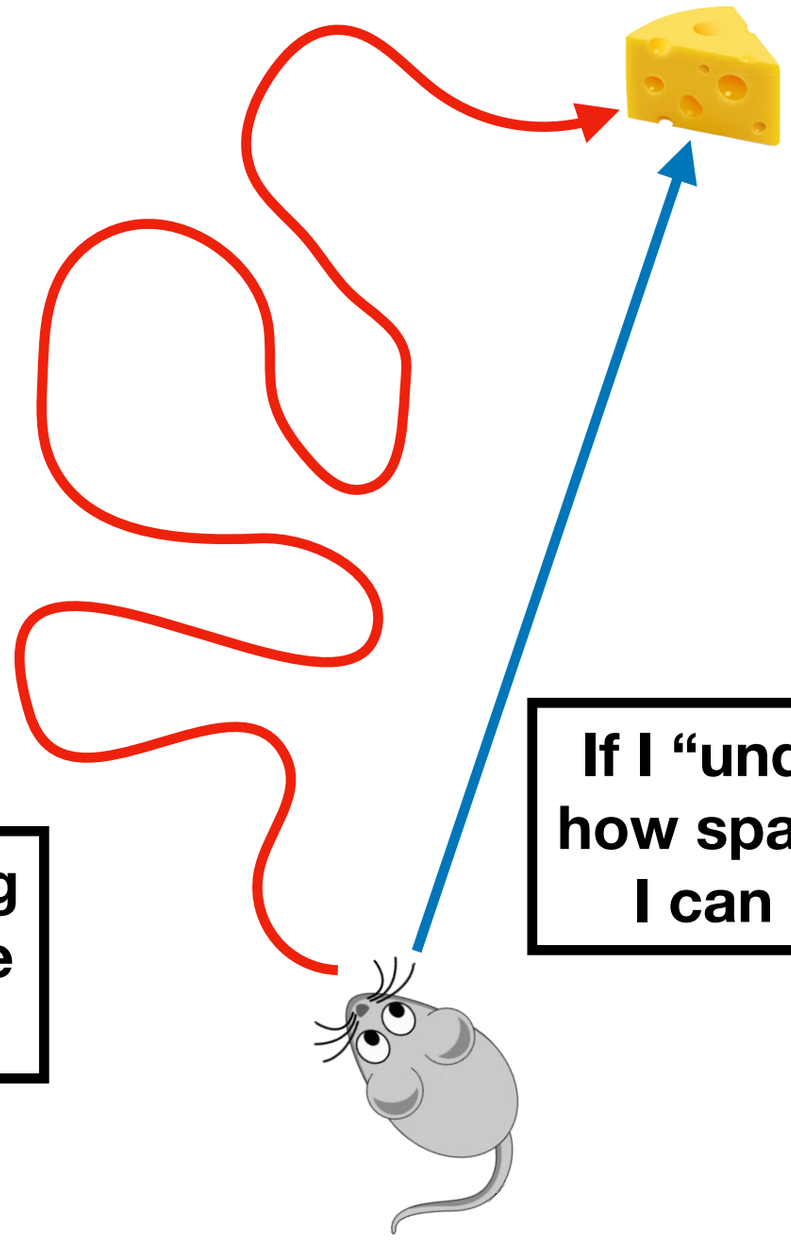

If I “understand” how space works, I can do this.

Next time I want cheese.

Possible because spatial relationships have a known structure

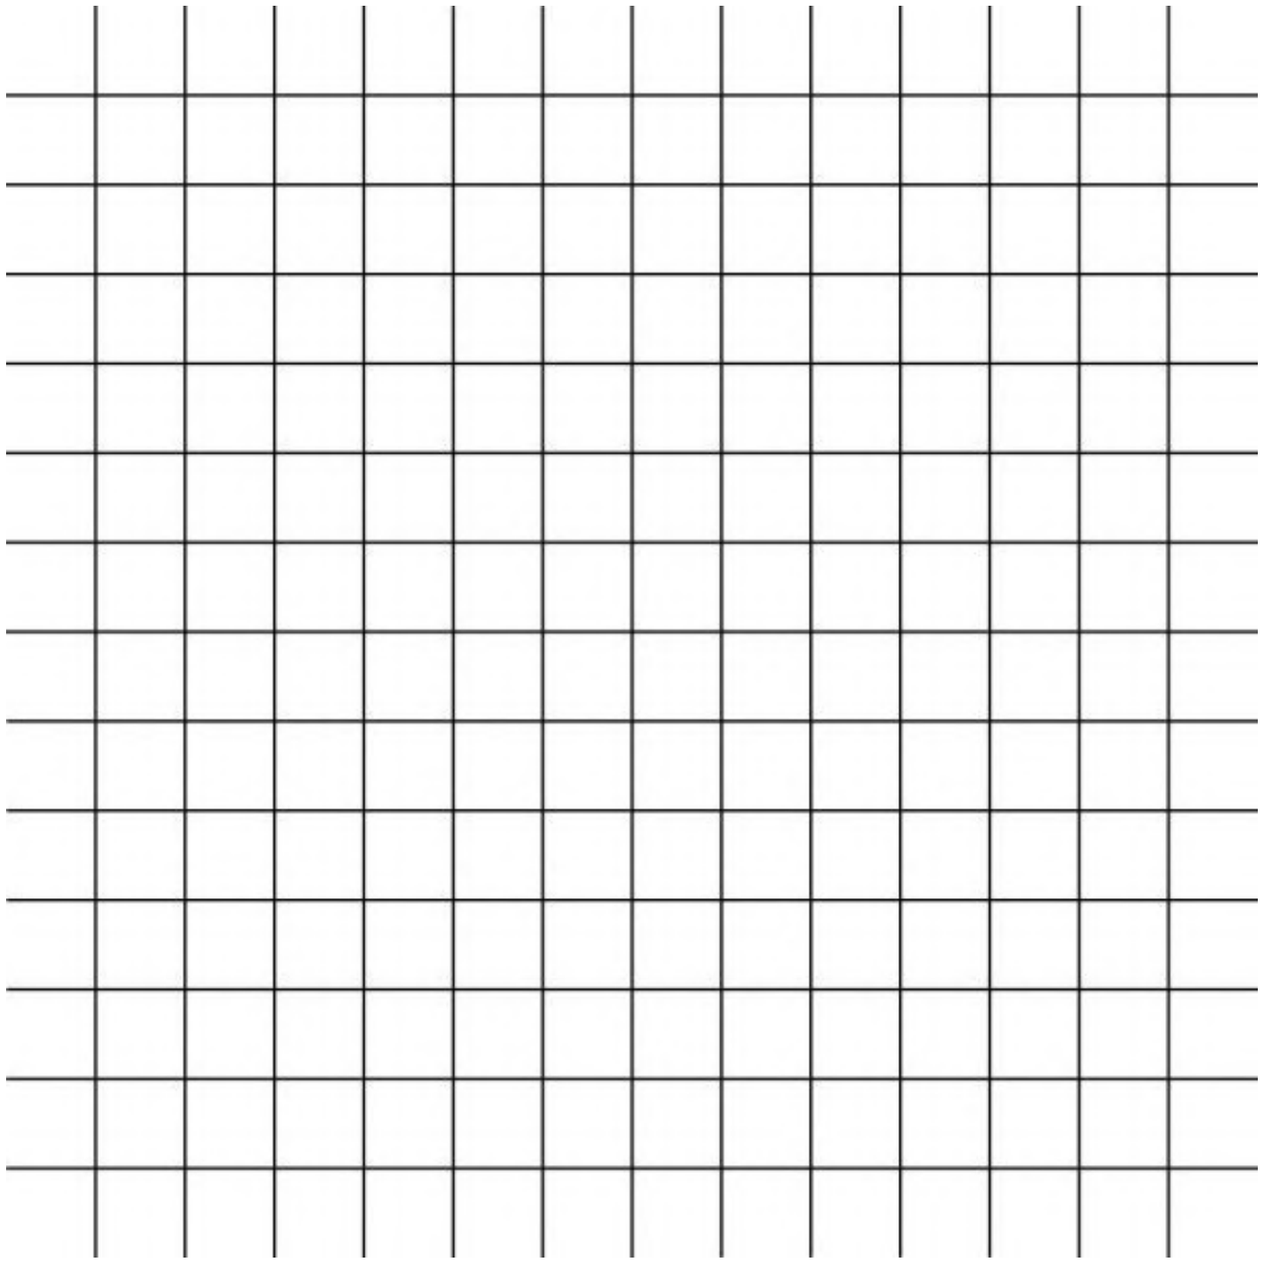

There are many forms of relational structure

Concept hierarchies

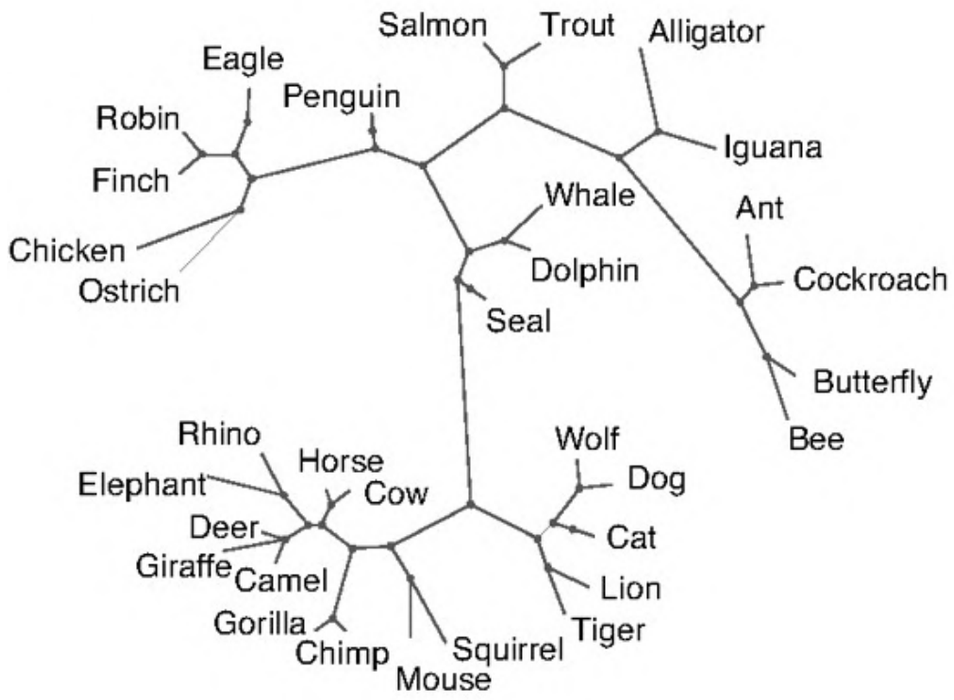

Linear structures

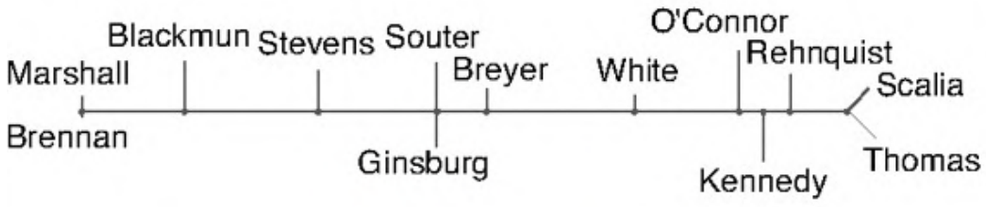

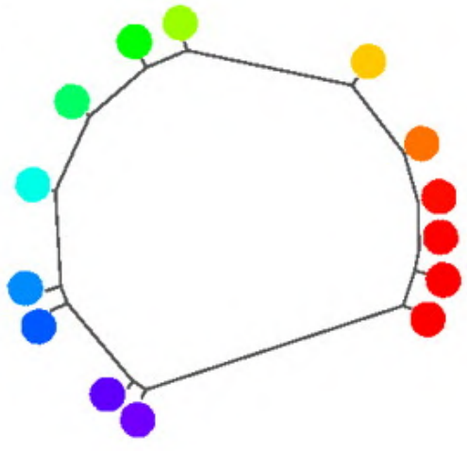

Sequence hierarchies

etc. etc. etc.

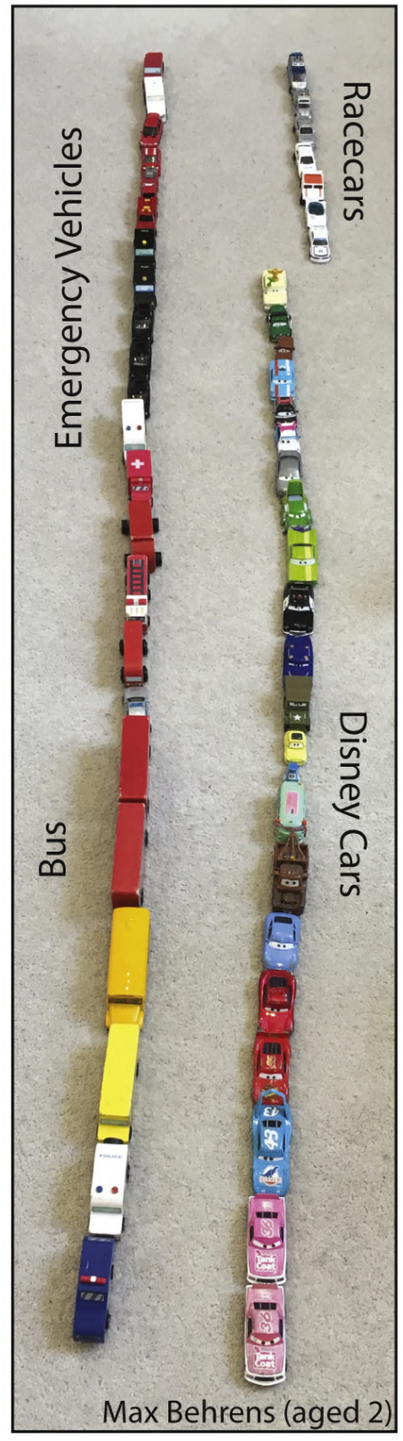

Kemp and Tenembaum 2008

Using structural knowledge for inference

www.noahdanielsmith.com

CHAPTER THREE THE ORIGIN OF O-REN

CHAPTER EIGHT THE CRUEL TUTELAGE OF PAI MEI

CHAPTER SIX MASSACRE AT TWO PINES

CHAPTER TWO THE BLOOD-SPLATTERED BRIDE

CHAPTER FOUR THE MAN FROM OKINAWA

CHAPTER FIVE SHOWDOWN AT HOUSE OF BLUE LEAVES

CHAPTER ONE (2)

CHAPTER SEVEN THE LONELY GRAVE OF PAULA SCHULTZ

CHAPTER NINE ELLE AND I

LAST CHAPTER FACE TO FACE

www.noahdanielsmith.com

We know something about the neural basis of spatial relationships

Position in space

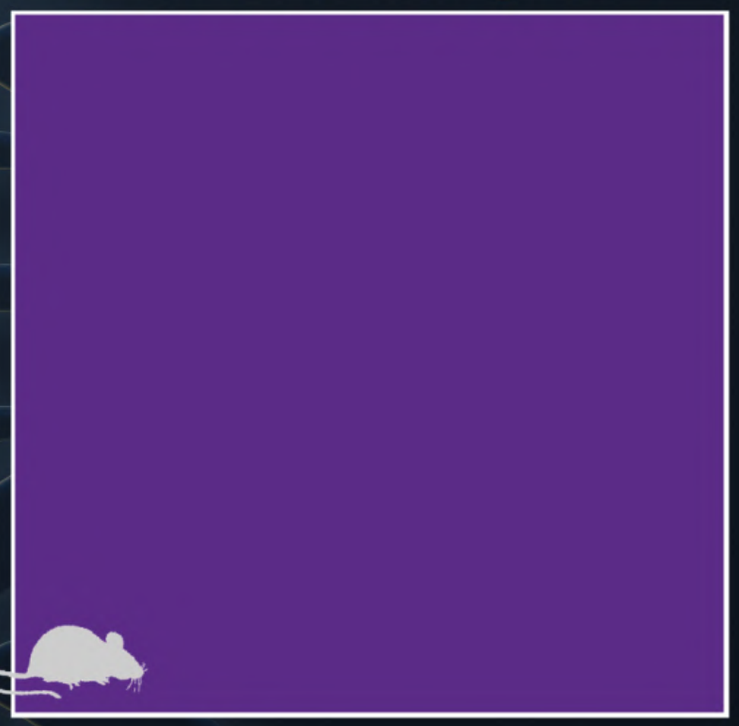

Place Cell

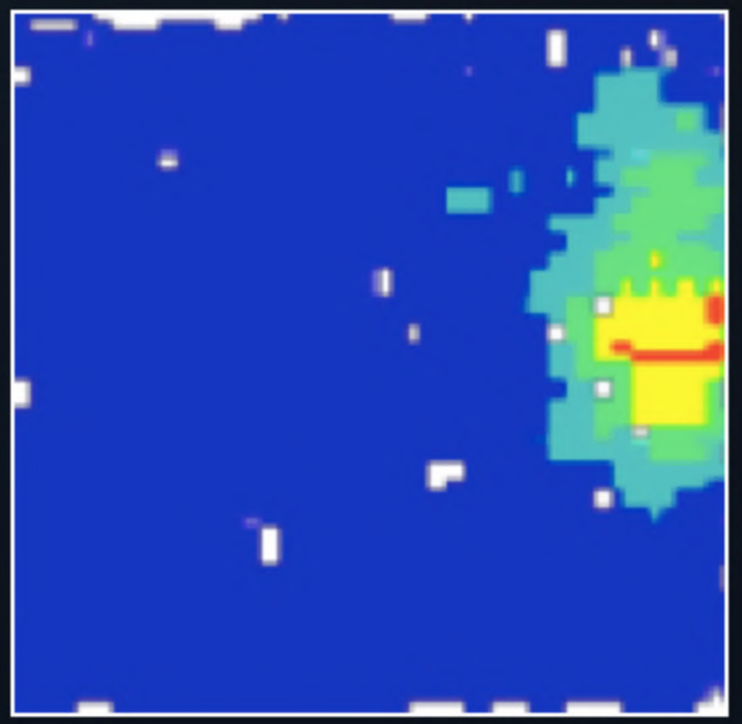

Grid Cell

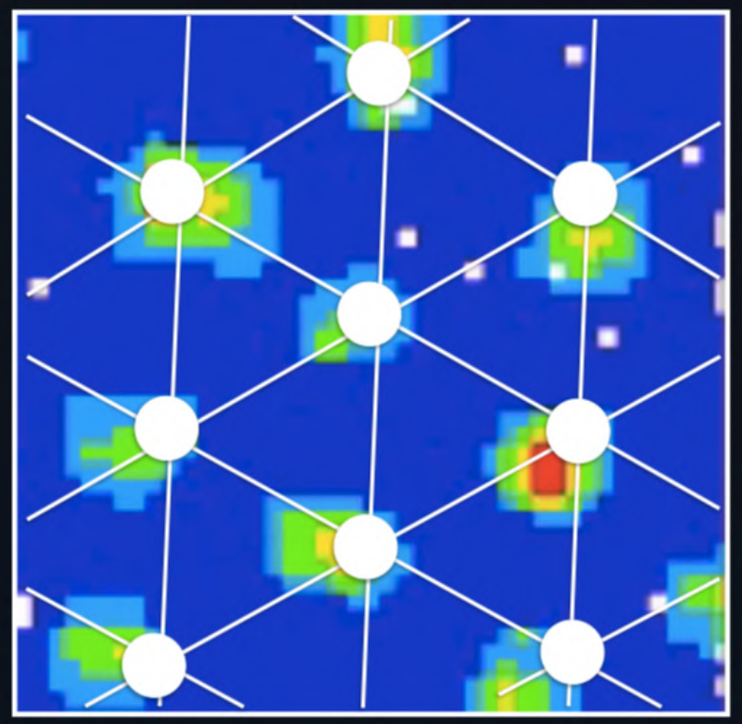

Less Activity

More Activity

We know something about the neural basis of spatial relationships

Human

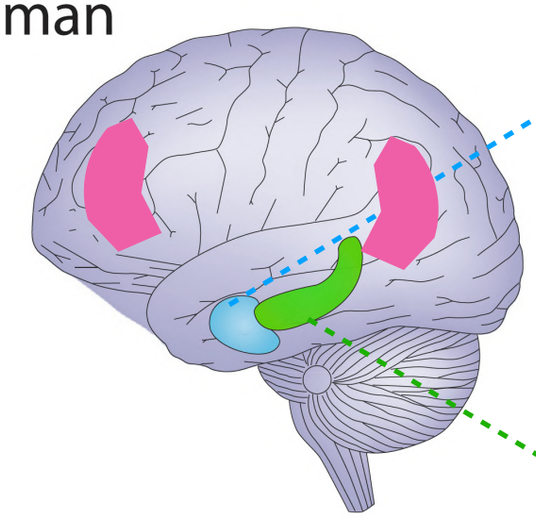

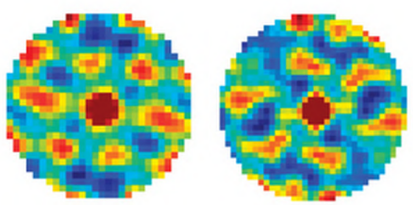

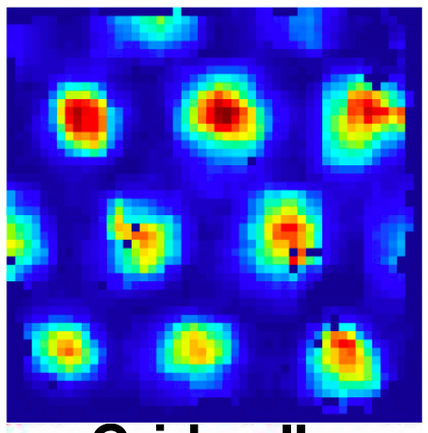

Grid cells
Hafting et al., 2005

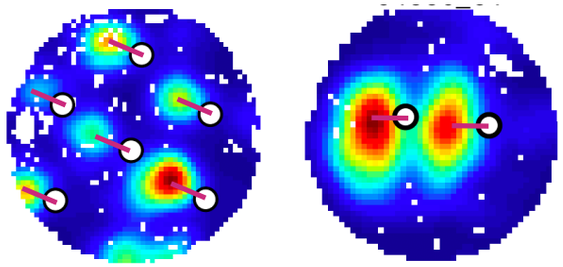

Object vector cells
Hoydal et al., 2018

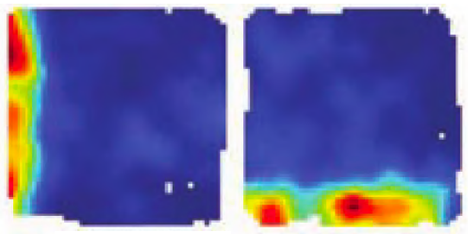

Border cells
Solstad et al., 2008

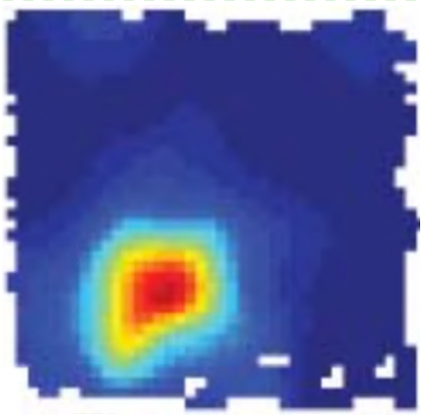

Place cells
O'Keefe & Dostrovsky, 1971

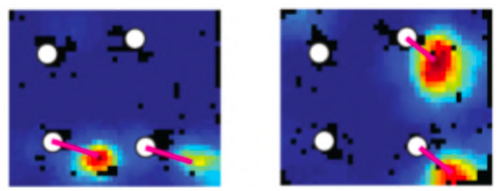

Landmark cells
Deshmukh & Knierim, 2013

Human grid cells also present in mPFC/PCC

Doeller et al. 2010, Jacobs et al. 2013

Similar representations in non-spatial problems

Space

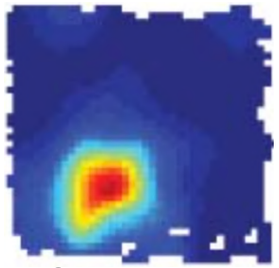

O'Keefe & Dostrovsky, 1971

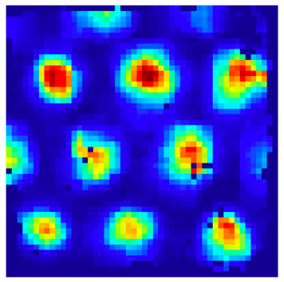

Hafting et al., 2005

Non-Space

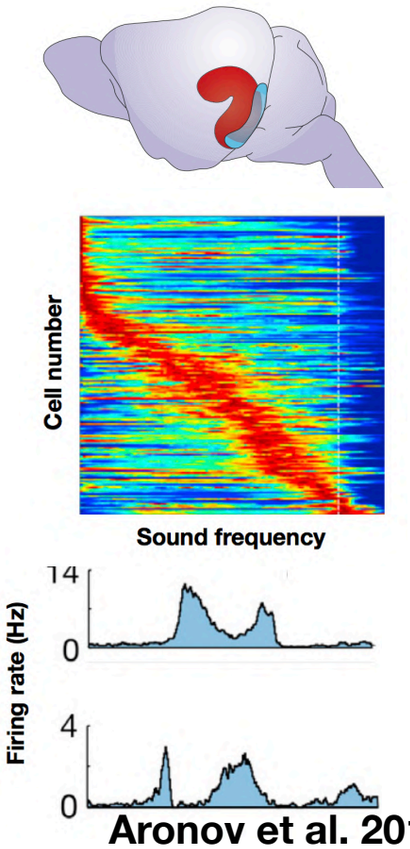

Aronov et al. 2017

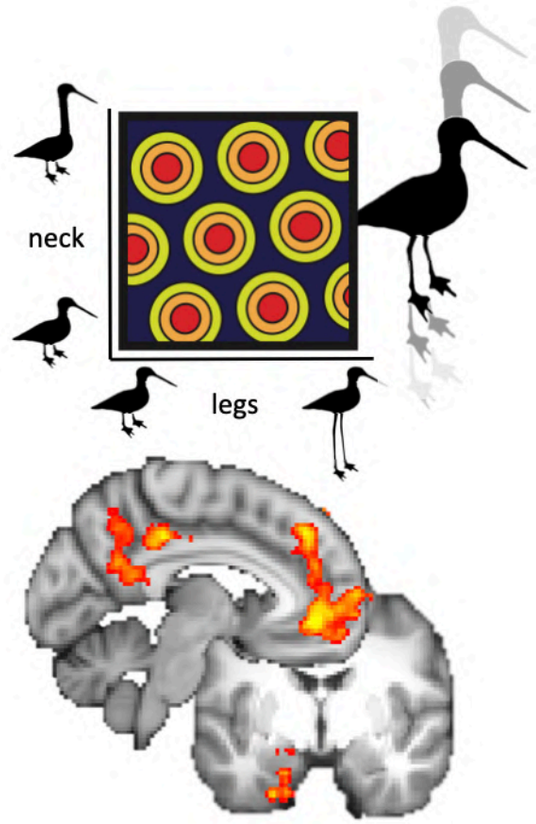

Constantinescu et al. 2016

A smoke signal for grid cells in fMRI

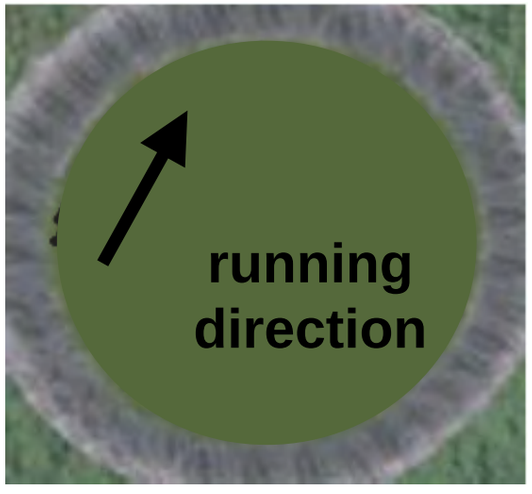

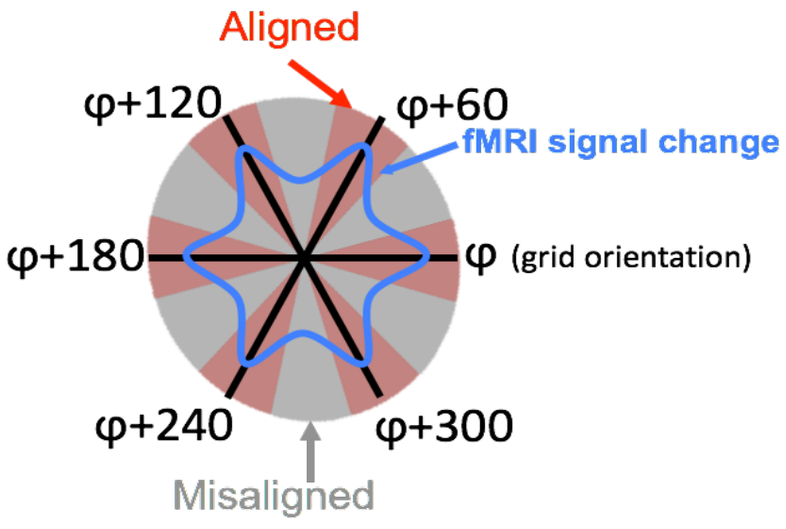

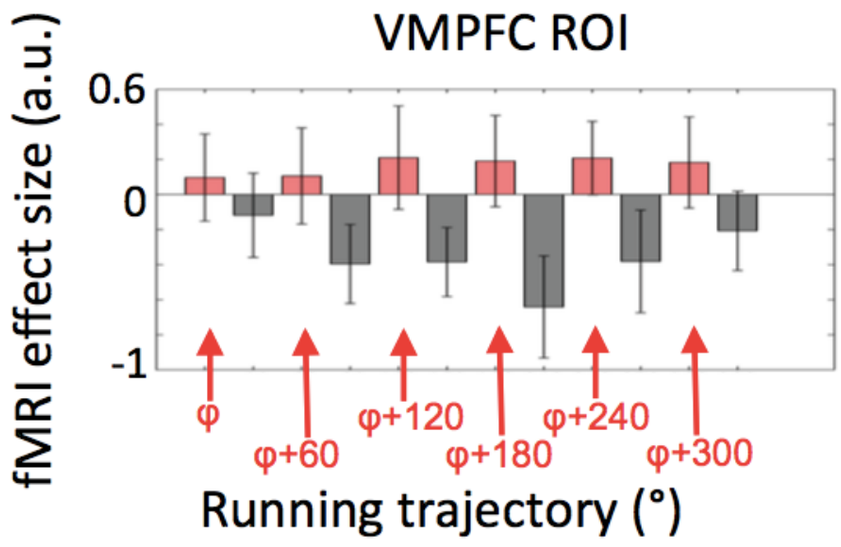

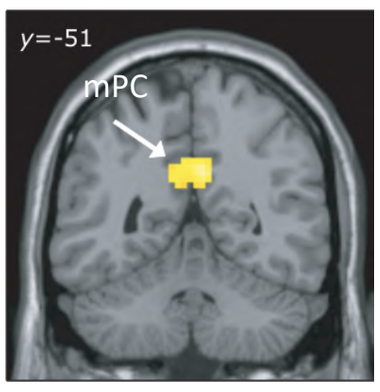

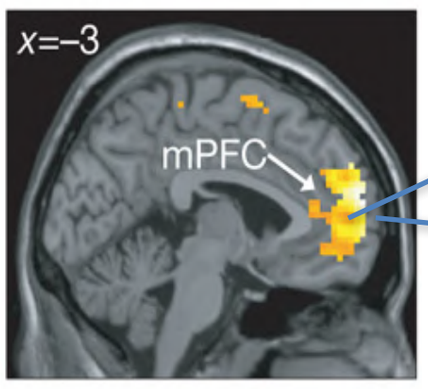

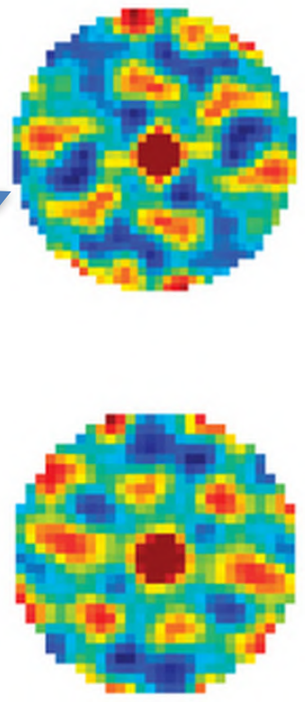

(Doeller et al,, 2010)

(Jacobs et al, 2013)

Coordinate systems for general structures

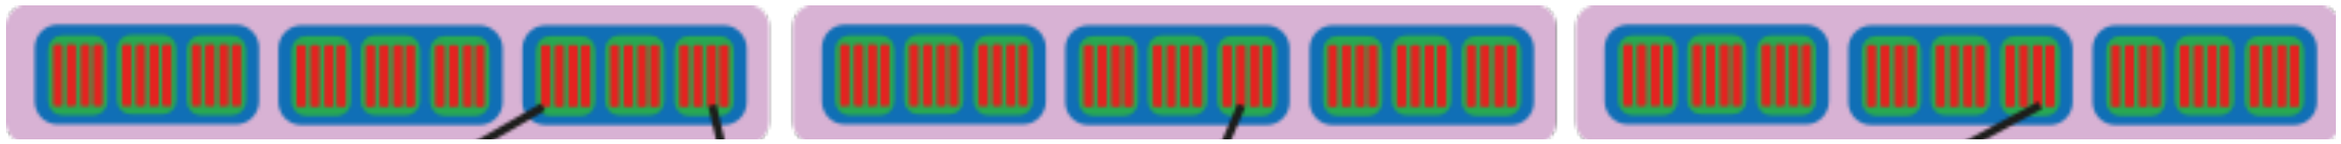

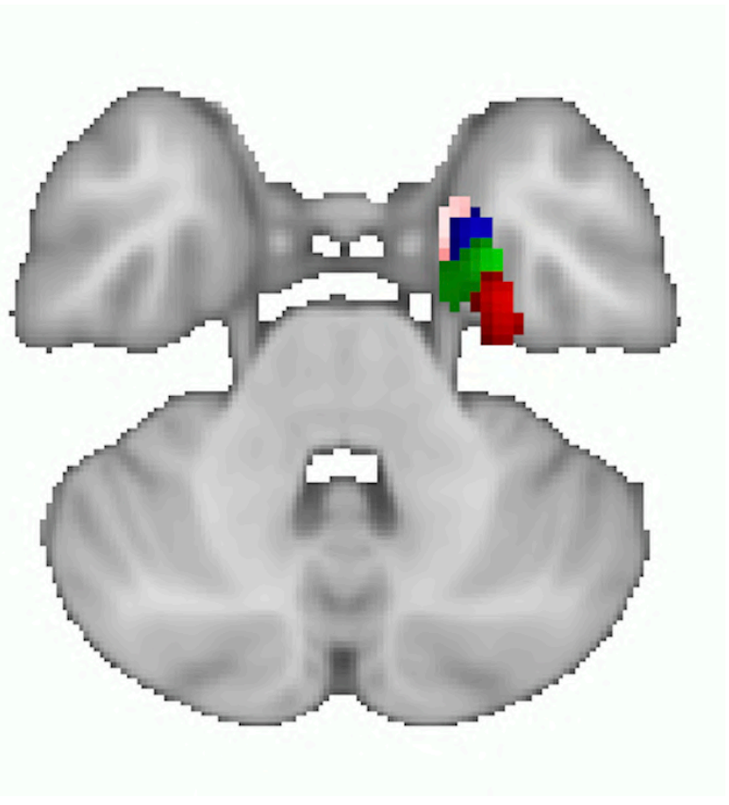

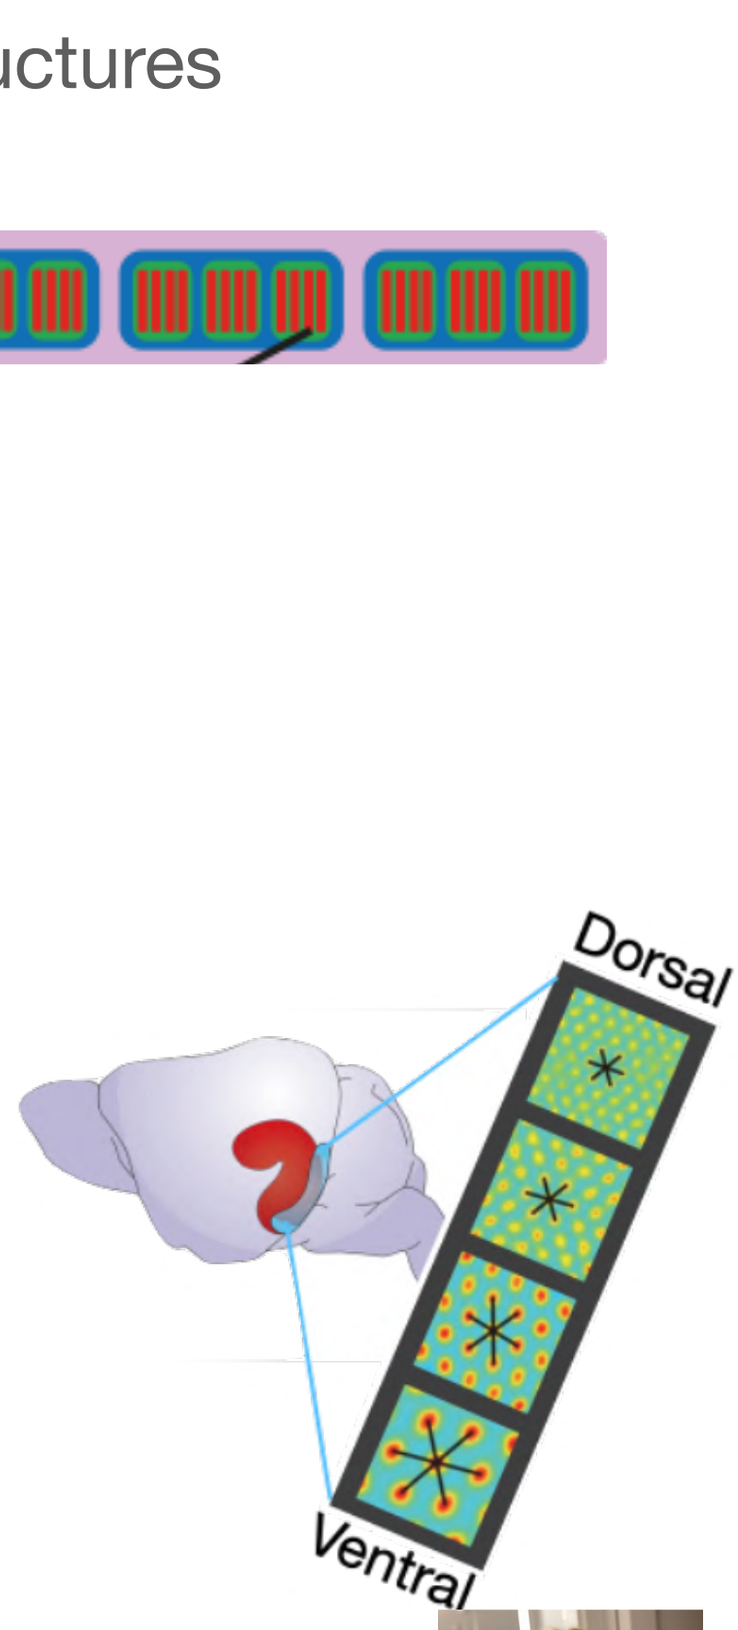

Shpektor et al. biorxiv 2024

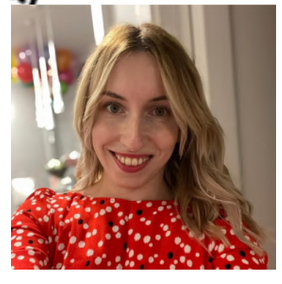

Replay


Resting
Running
Resting
Position Based on Place Cell Activity
Position (cm)
Time (seconds)


Position (cm)
240 milliseconds
150 ms
170 ms
210 ms

Replay

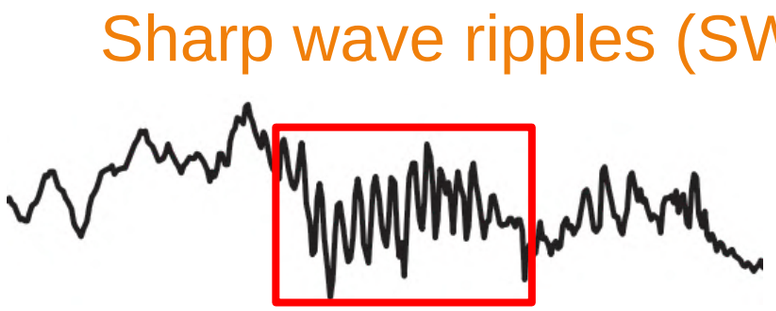

Replay

Forward replay

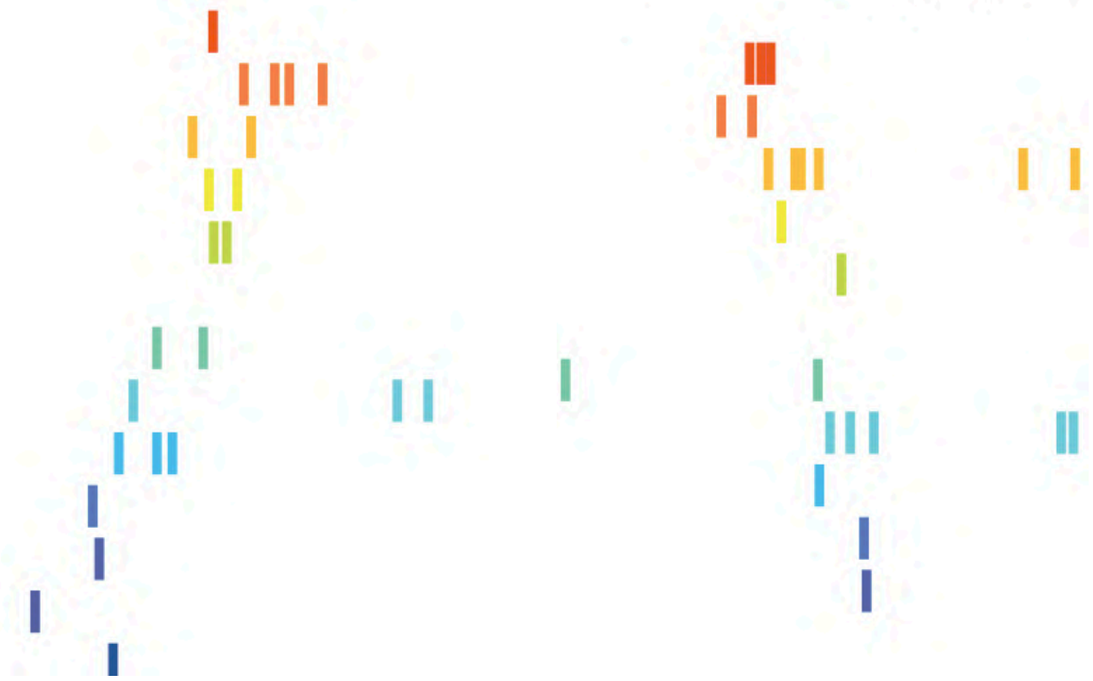

Reverse replay

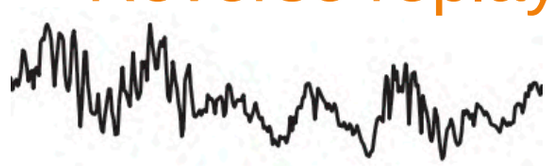

What this talk is going to say

- Replay is not just replaying the past. It is doing inference about possible futures.

- It does so by aligning experience to our structural knowledge

- It is always making these inferences in the background, making you faster at thinking when the time comes.

- If there is time, some data about replay building new structural knowledge.

What this talk is going to say

- Replay is not just replaying the past. It is doing inference about possible futures.

- It does so by aligning experience to our structural knowledge

- It is always making these inferences in the background, making you faster at thinking when the time comes.

- If there is time, some data about replay building new structural knowledge.

Can replay do inference or is it just replaying the past?

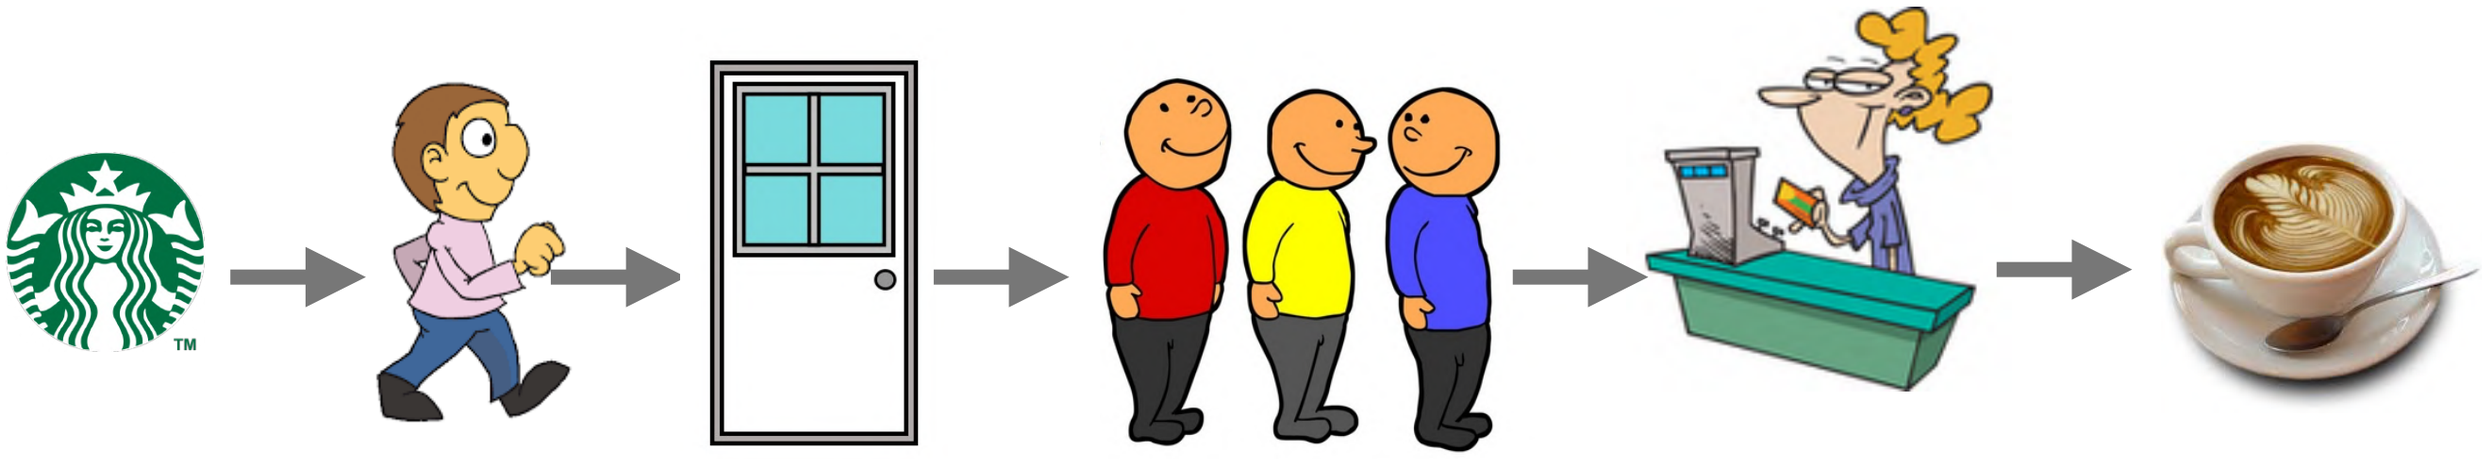

Yunzhe Liu

Zeb Kurth Nelson

Solving Quentin Tarantino in Replay

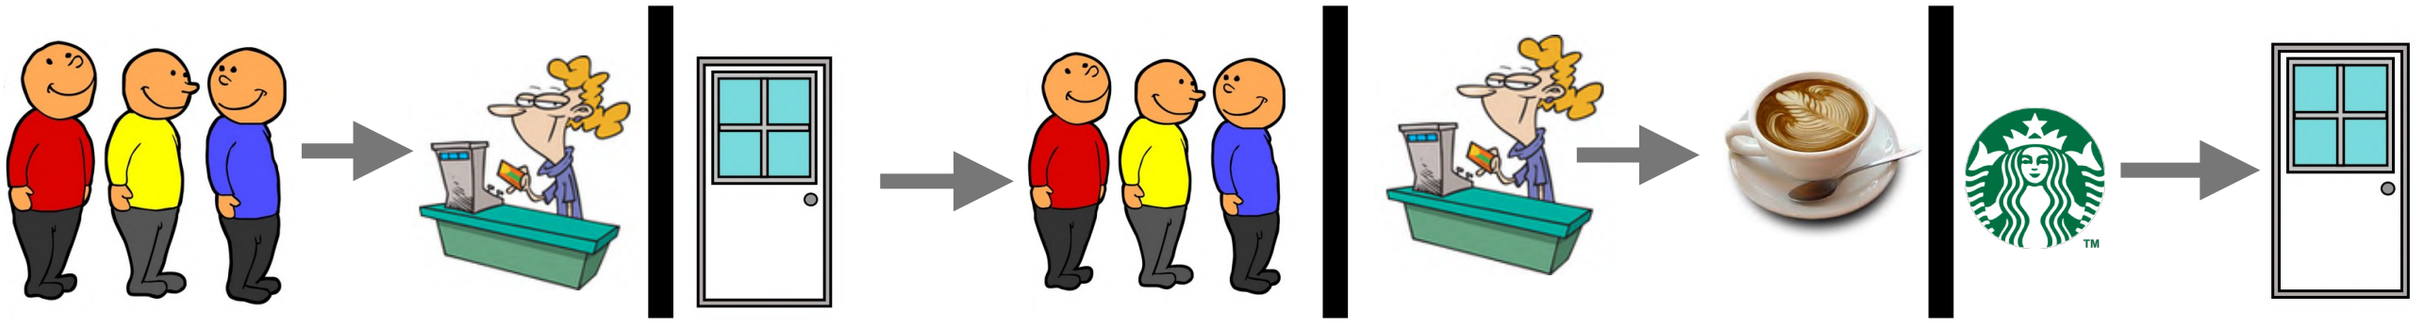

Jumble a sequence up and teach subjects how to unjumble it

Liu et al 2019

Can structural knowledge impact on this replay?

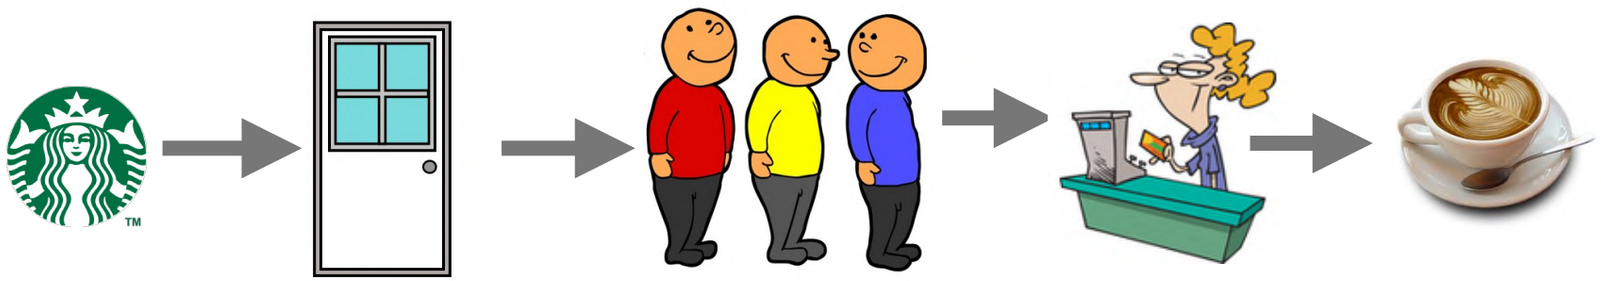

Jumble a sequence up and teach subjects how to unjumble it

Give them **new stimuli** with the same jumble structure

Measure representational sequences (**replay**) in their brain

Which ones will we find. **The jumbled ones, or the unjumbled ones?**

Liu et al 2019

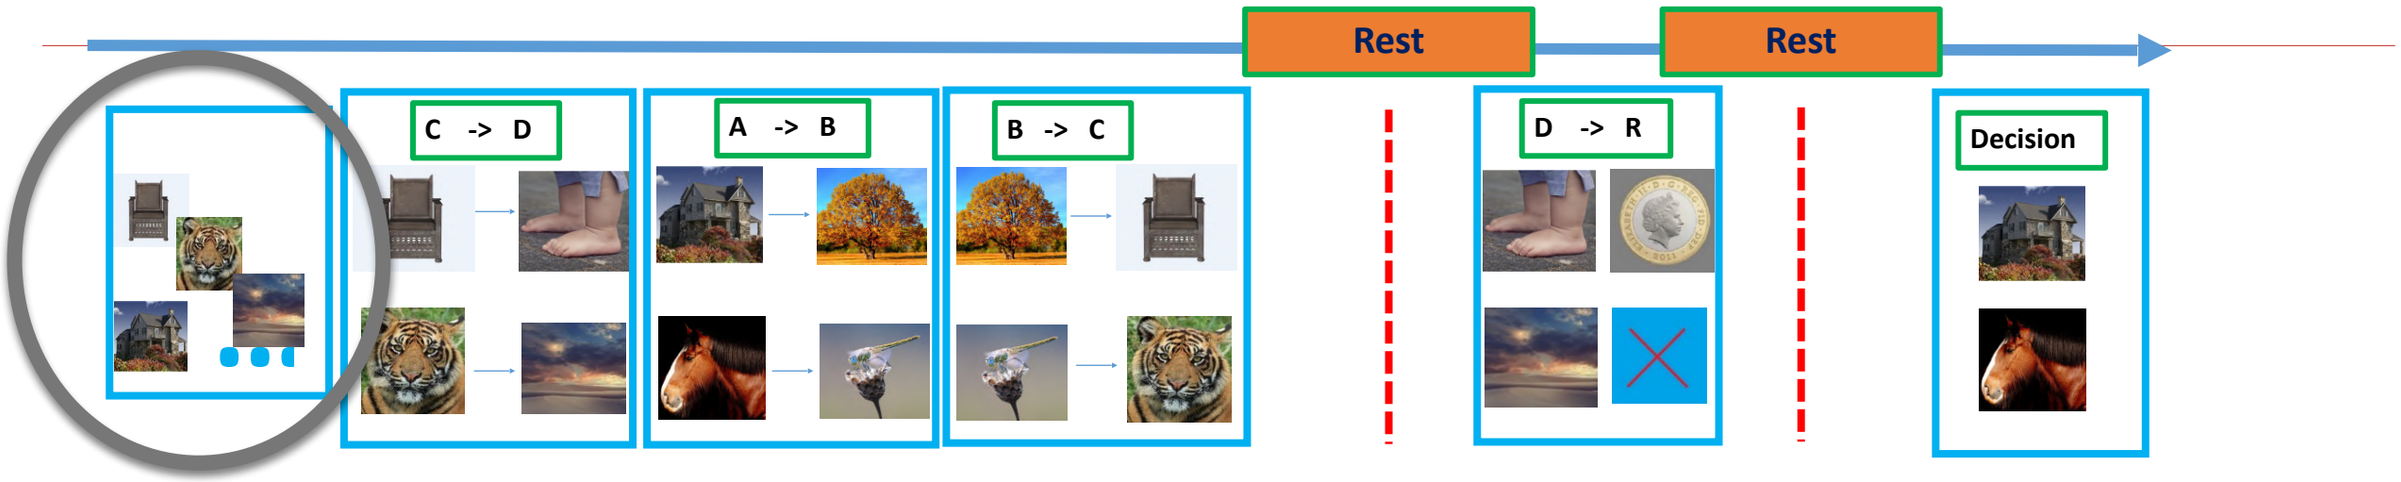

Figure out the **MEG sensor pattern** for each element.

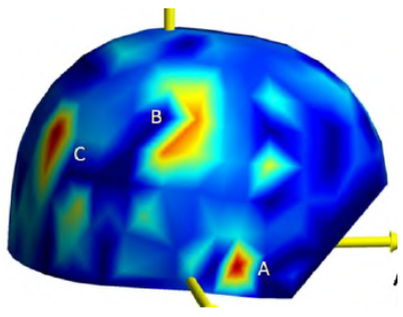

See: Kurth Nelson .. Dayan, Neuron 2016 for similar techniques

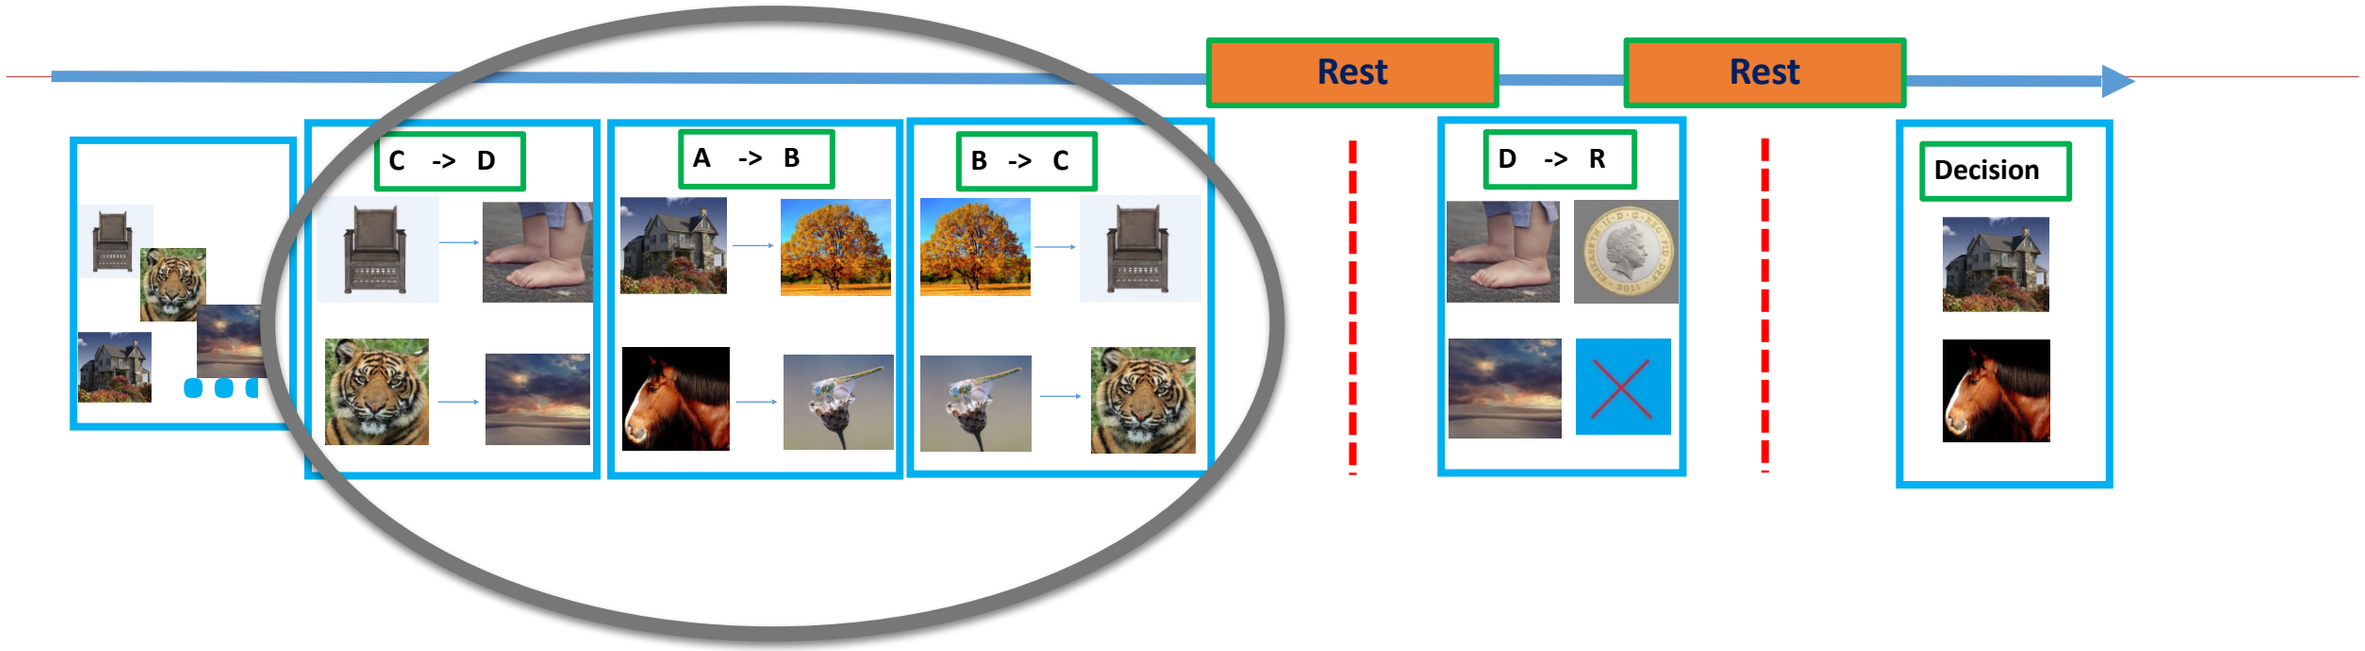

Figure out the MEG sensor pattern for each element.

Show them 2 sequences **jumbled up** and in the **wrong order**

See: Kurth Nelson .. Dayan, Neuron 2016 for similar techniques

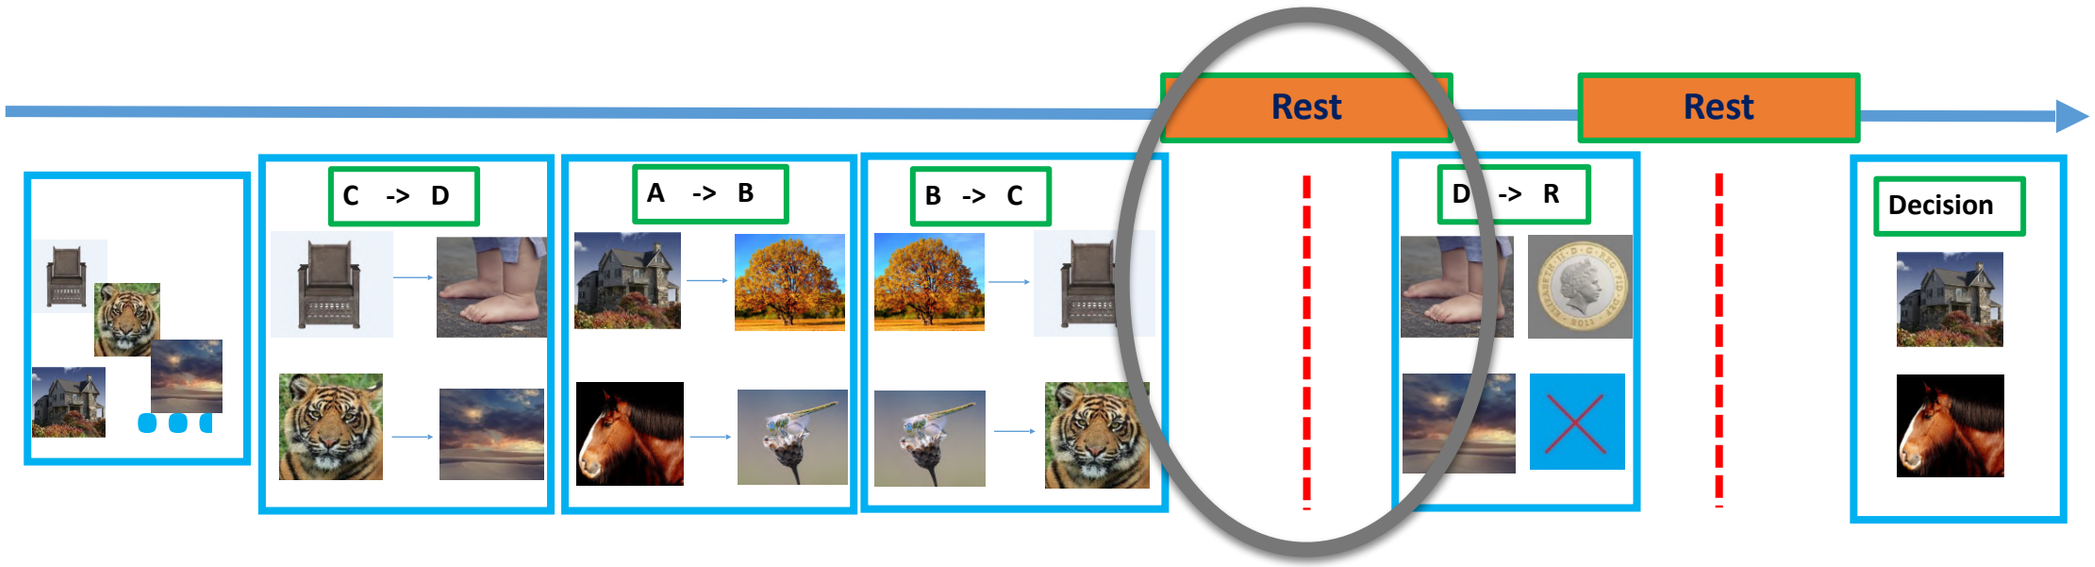

Figure out the MEG sensor pattern for each element.

Show them 2 sequences jumbled up and in the wrong order

Do the sensor patterns **replay** during rest? If so, In **what order**?

Liu et al 2019

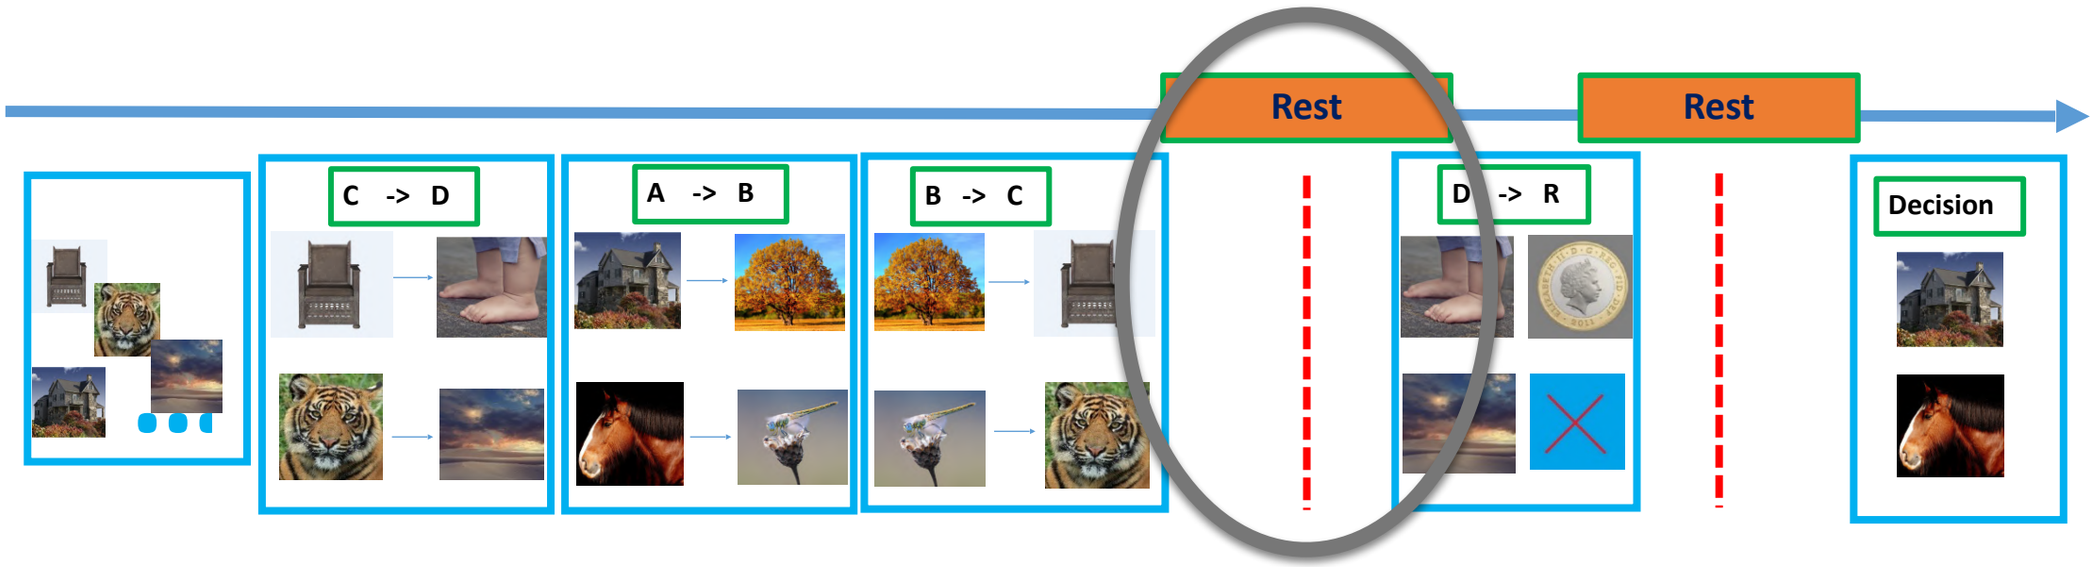

Figure out the MEG sensor pattern for each element.

Show them 2 sequences jumbled up and in the wrong order

Do the sensor patterns **replay** during rest? If so, In **what order**?

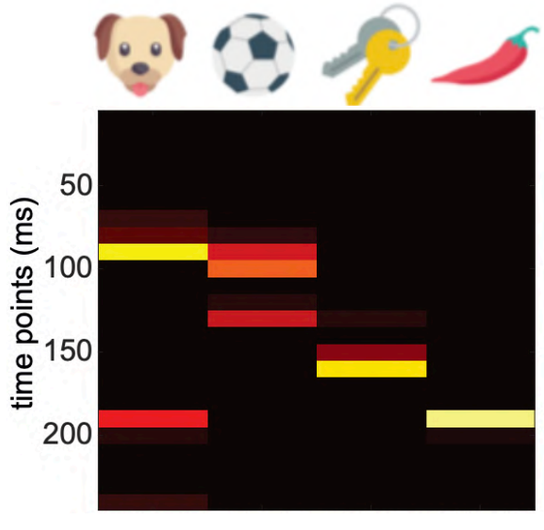

Liu et al 2019

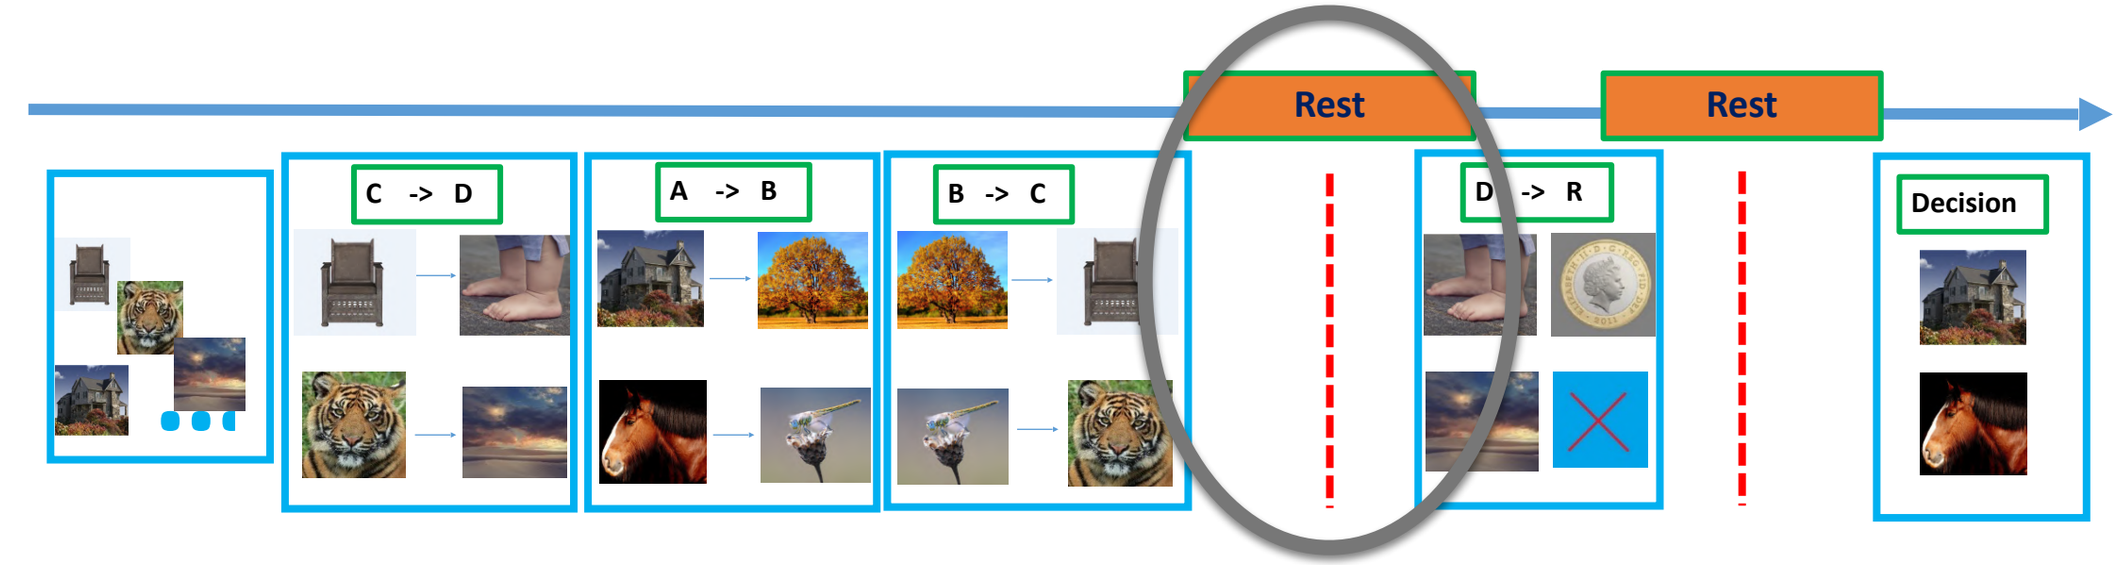

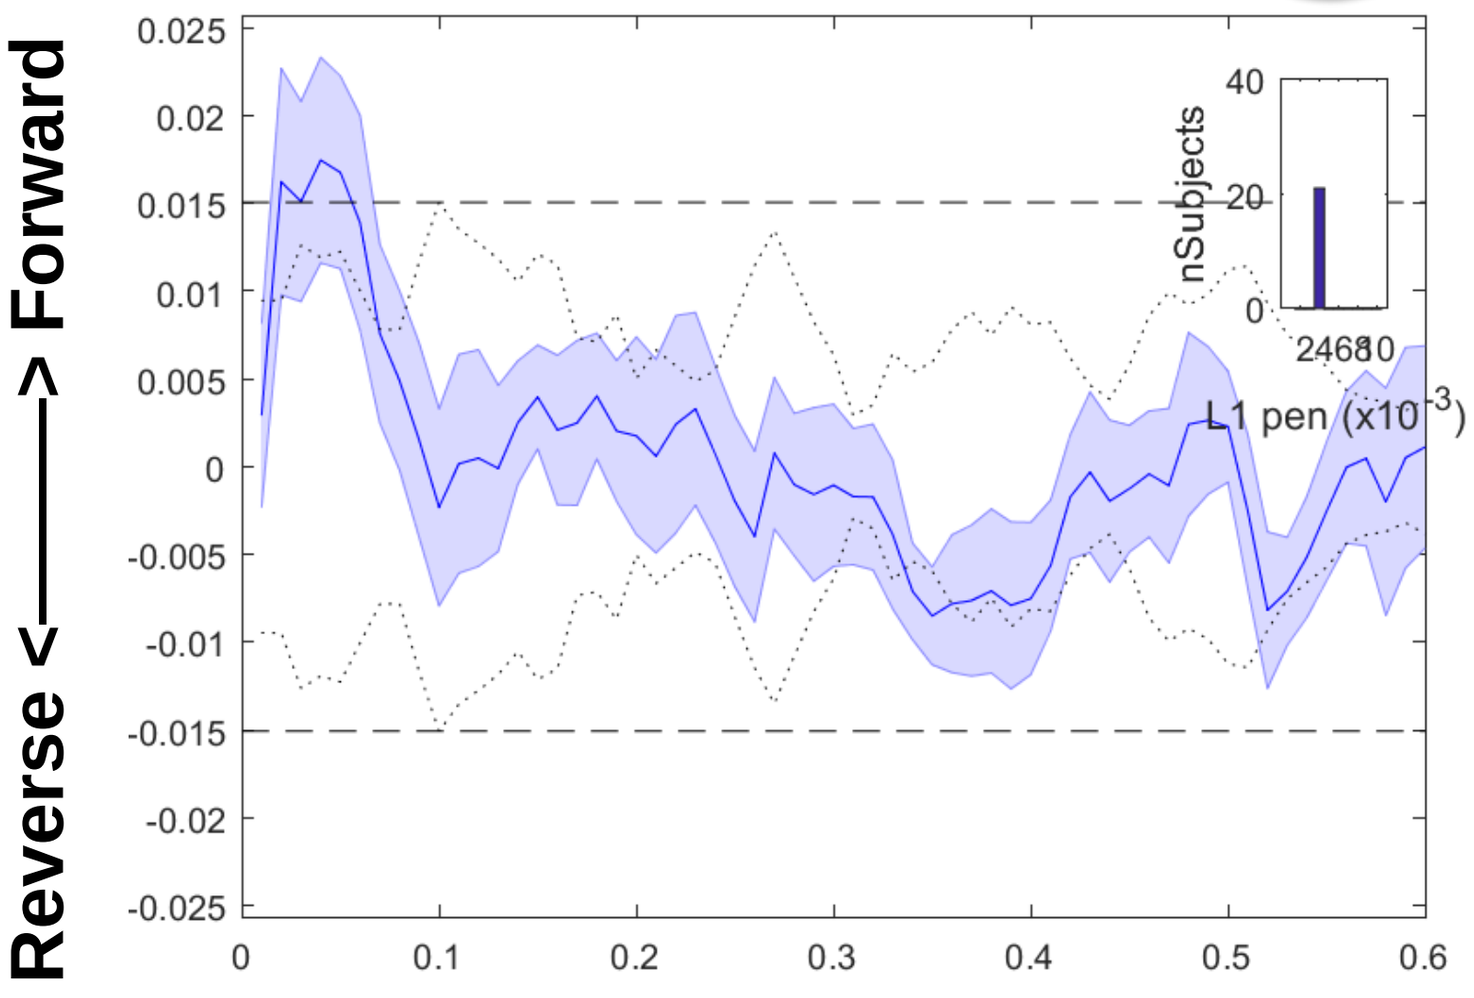

Time lag between elements (seconds)

Liu et al 2019

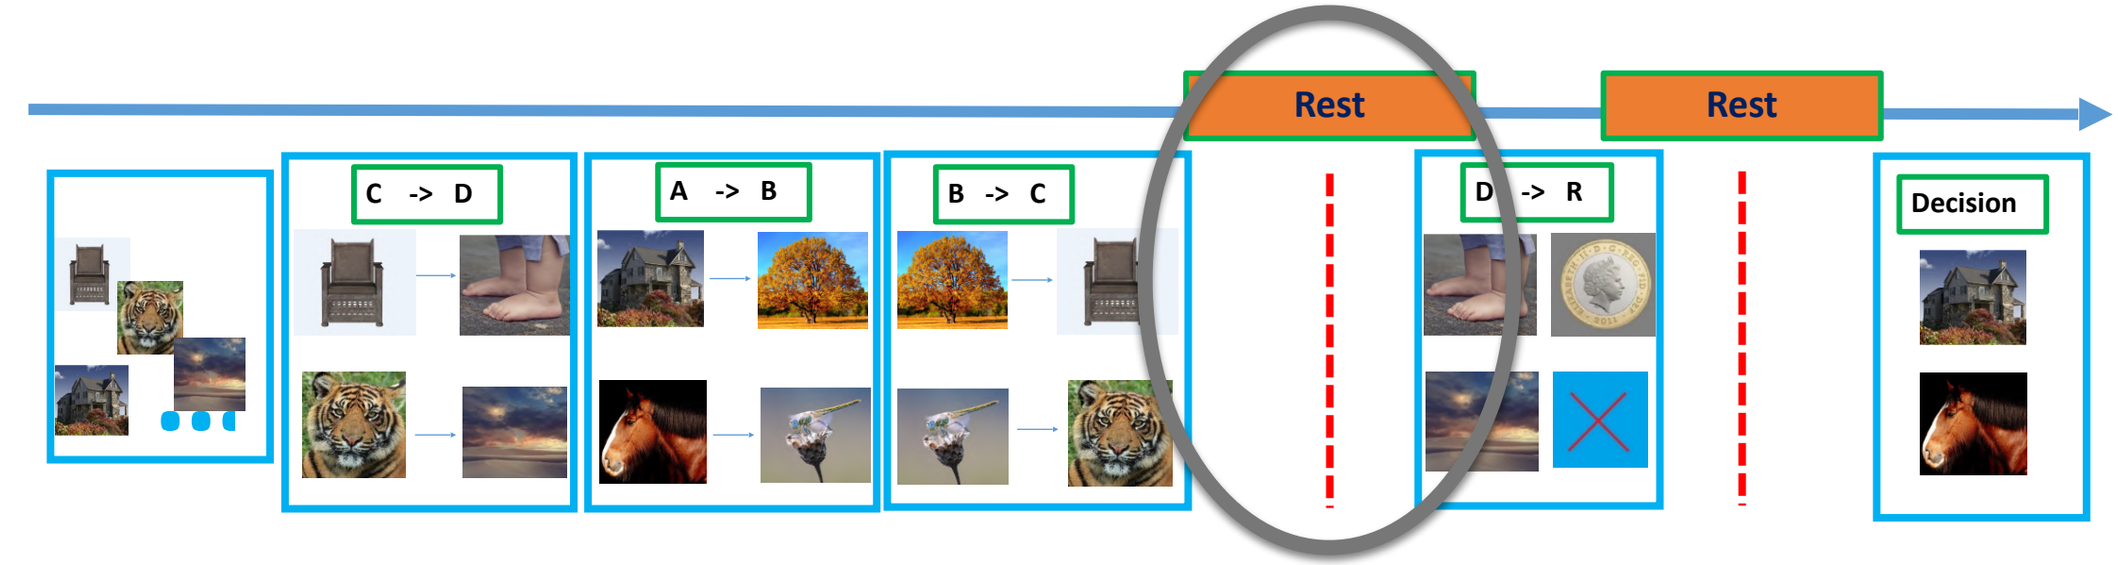

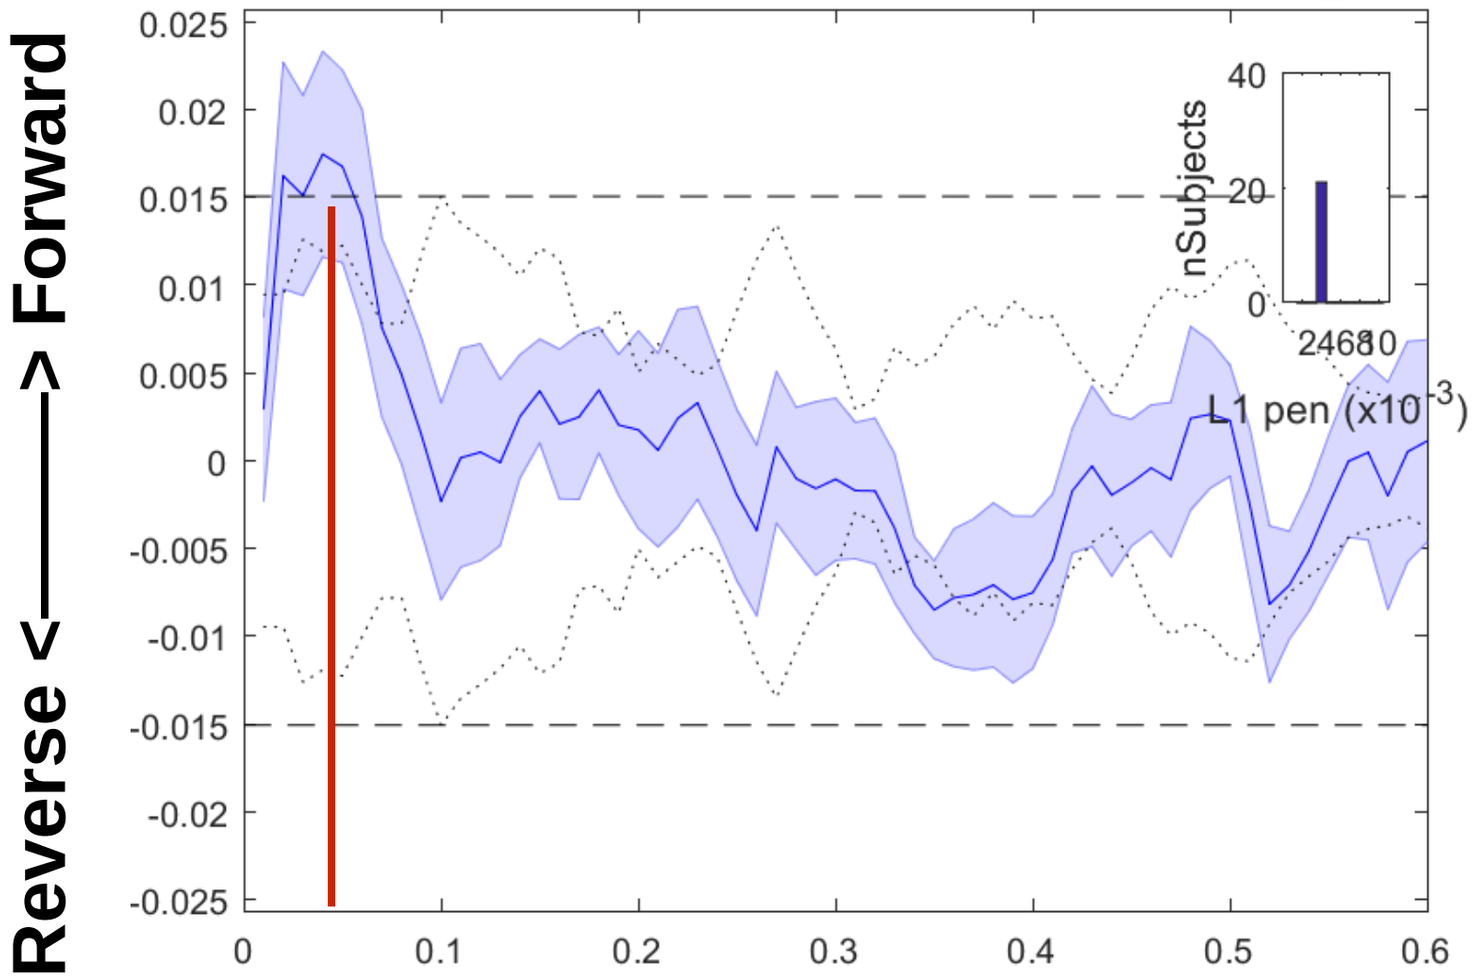

Time lag between elements (seconds)

The replay is **really quick**
(40ms)

Liu et al 2019

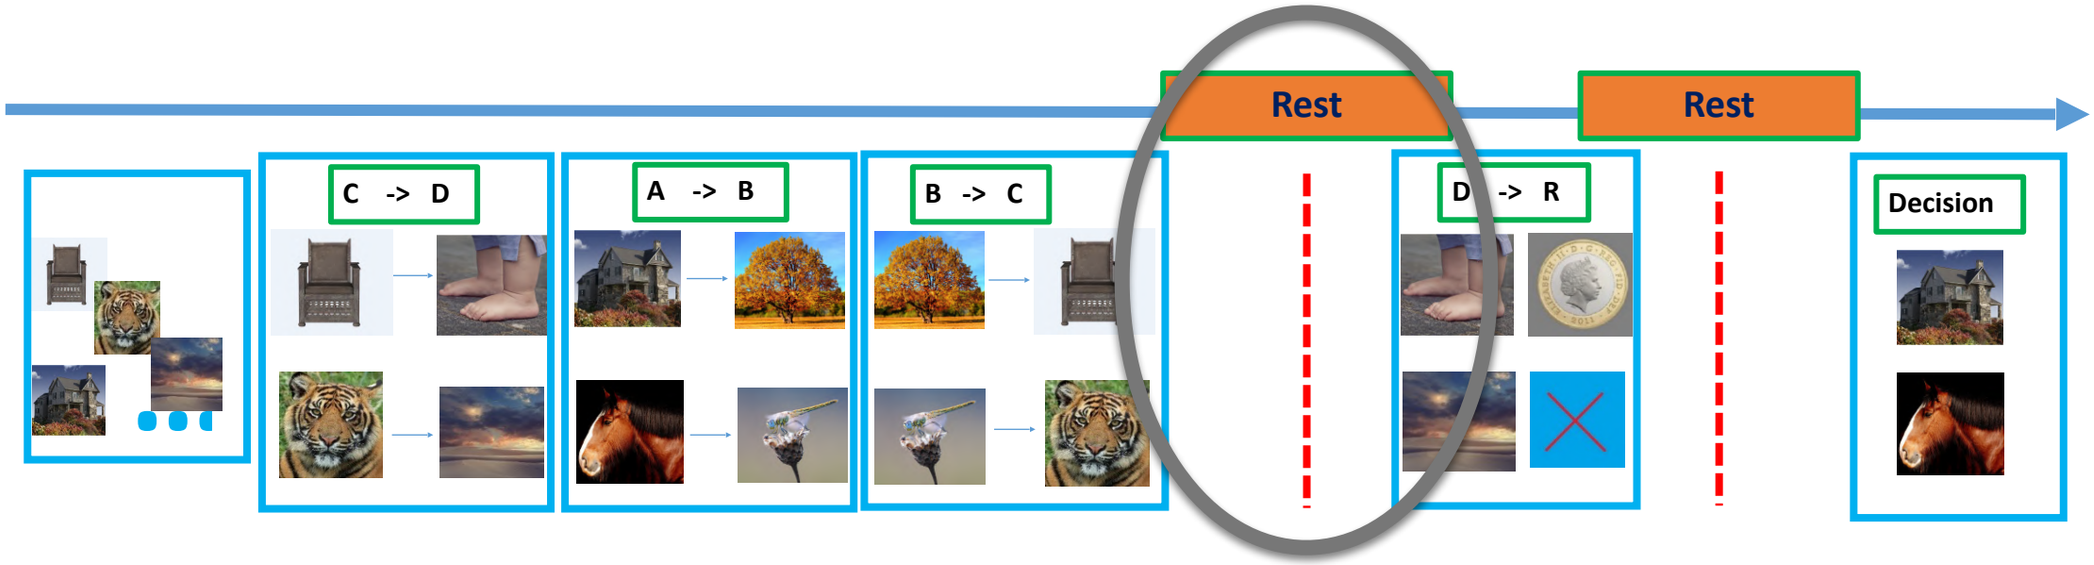

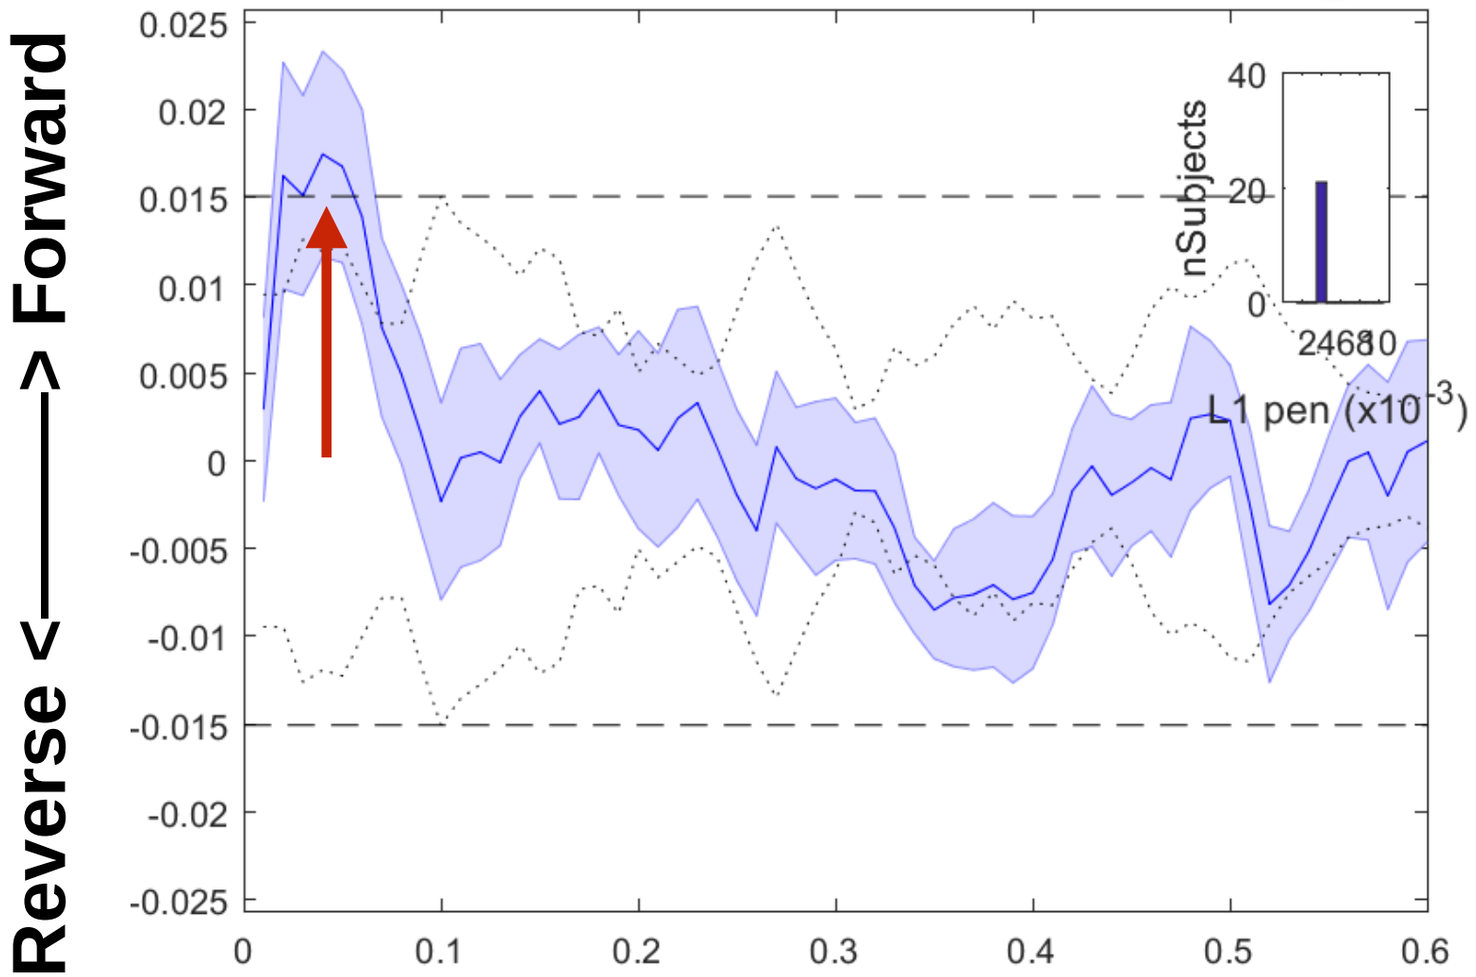

Time lag between elements (seconds)

The replay is **really quick** (50ms)

The replay goes **forwards**

Liu et al 2019

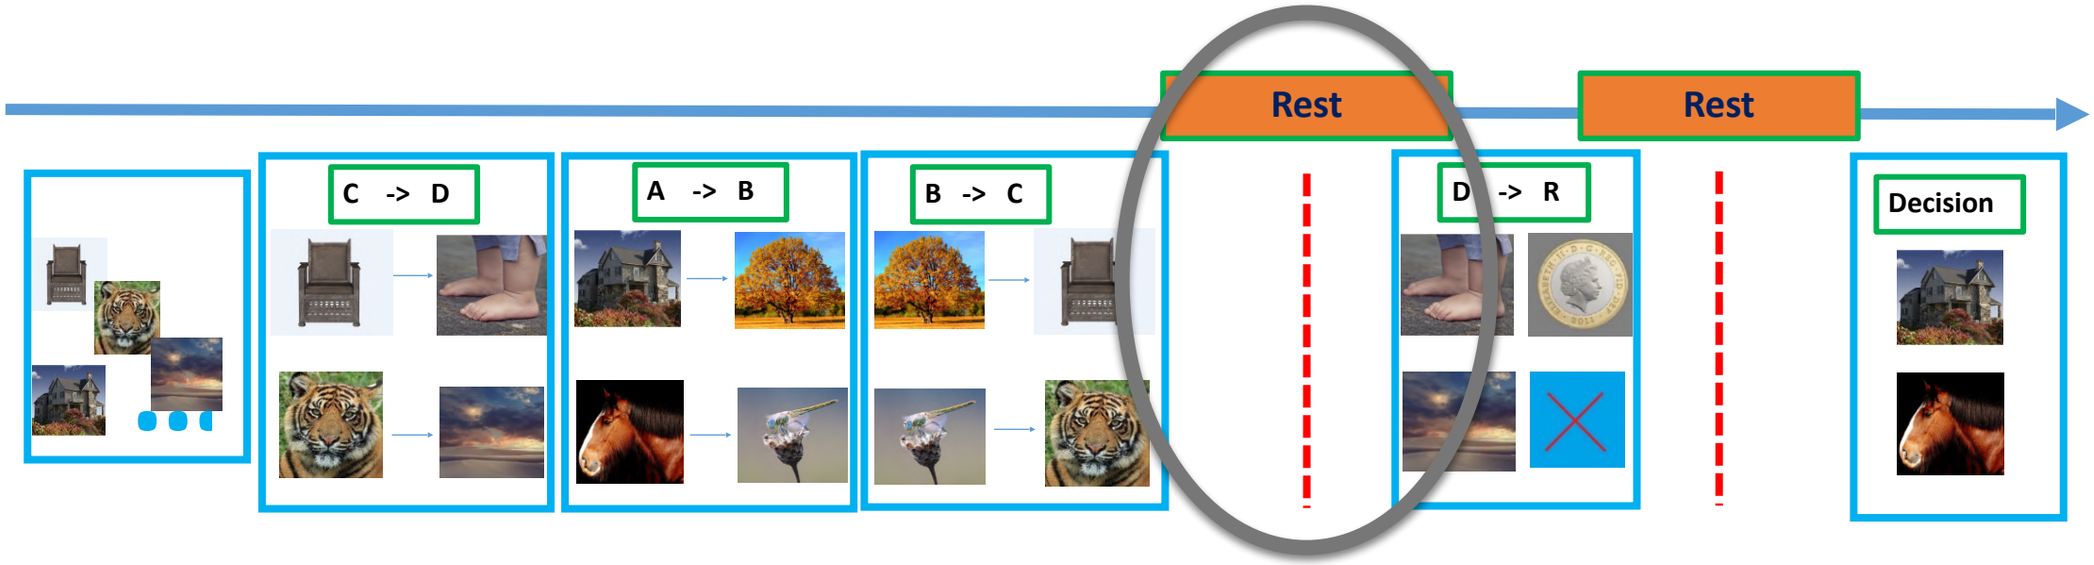

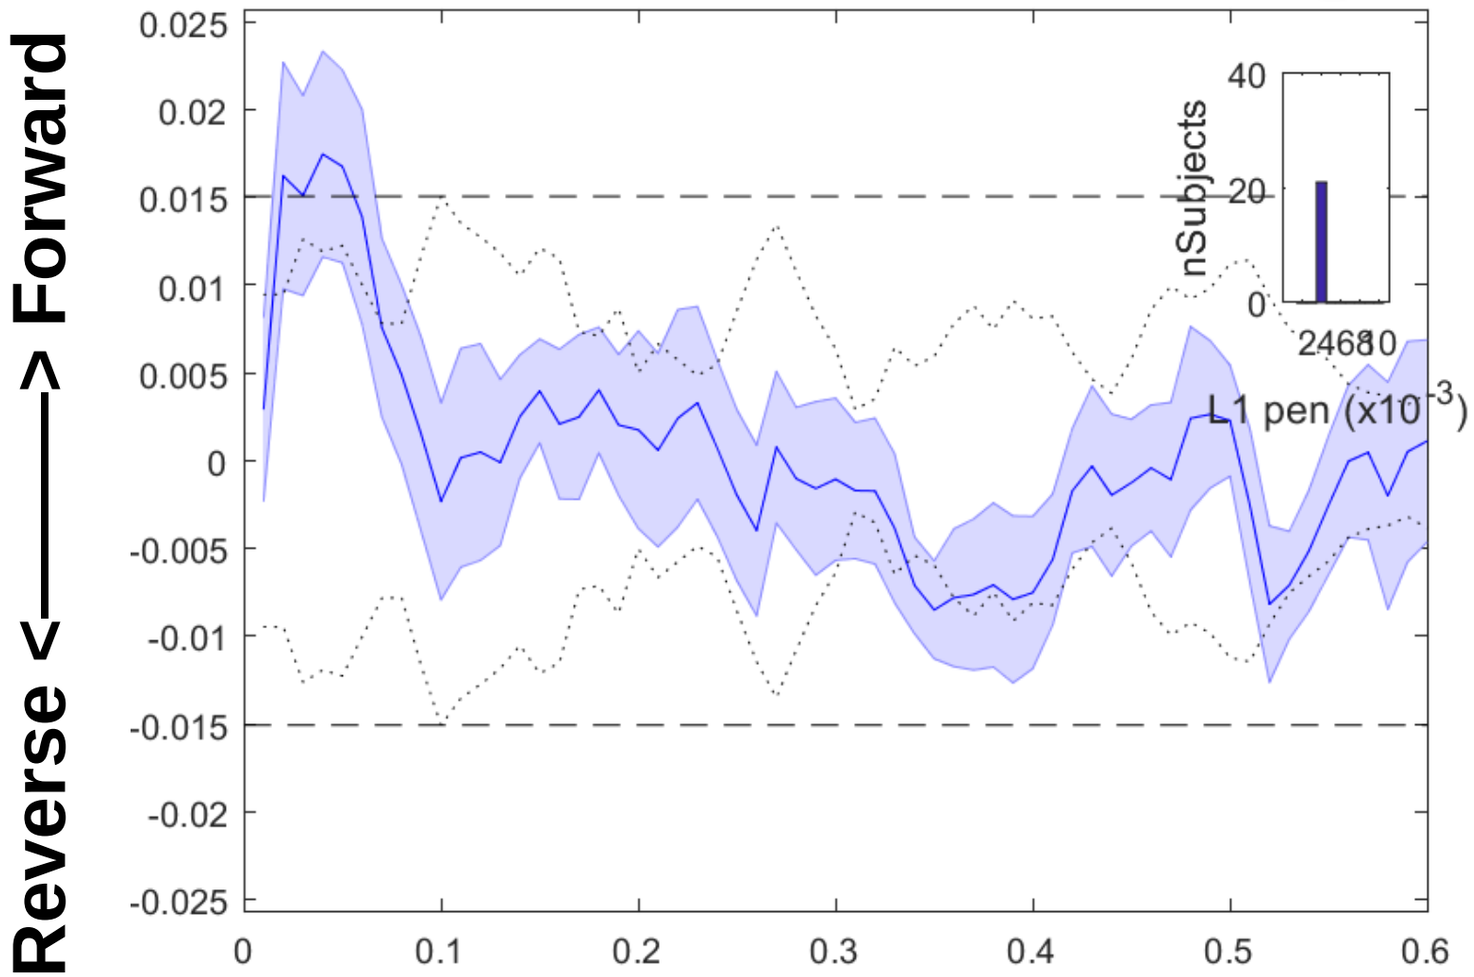

Time lag between elements (seconds)

The replay is **really quick** (50ms)

The replay goes **forwards**

These are the **unjumbled** sequences not the jumbled ones

Liu et al 2019

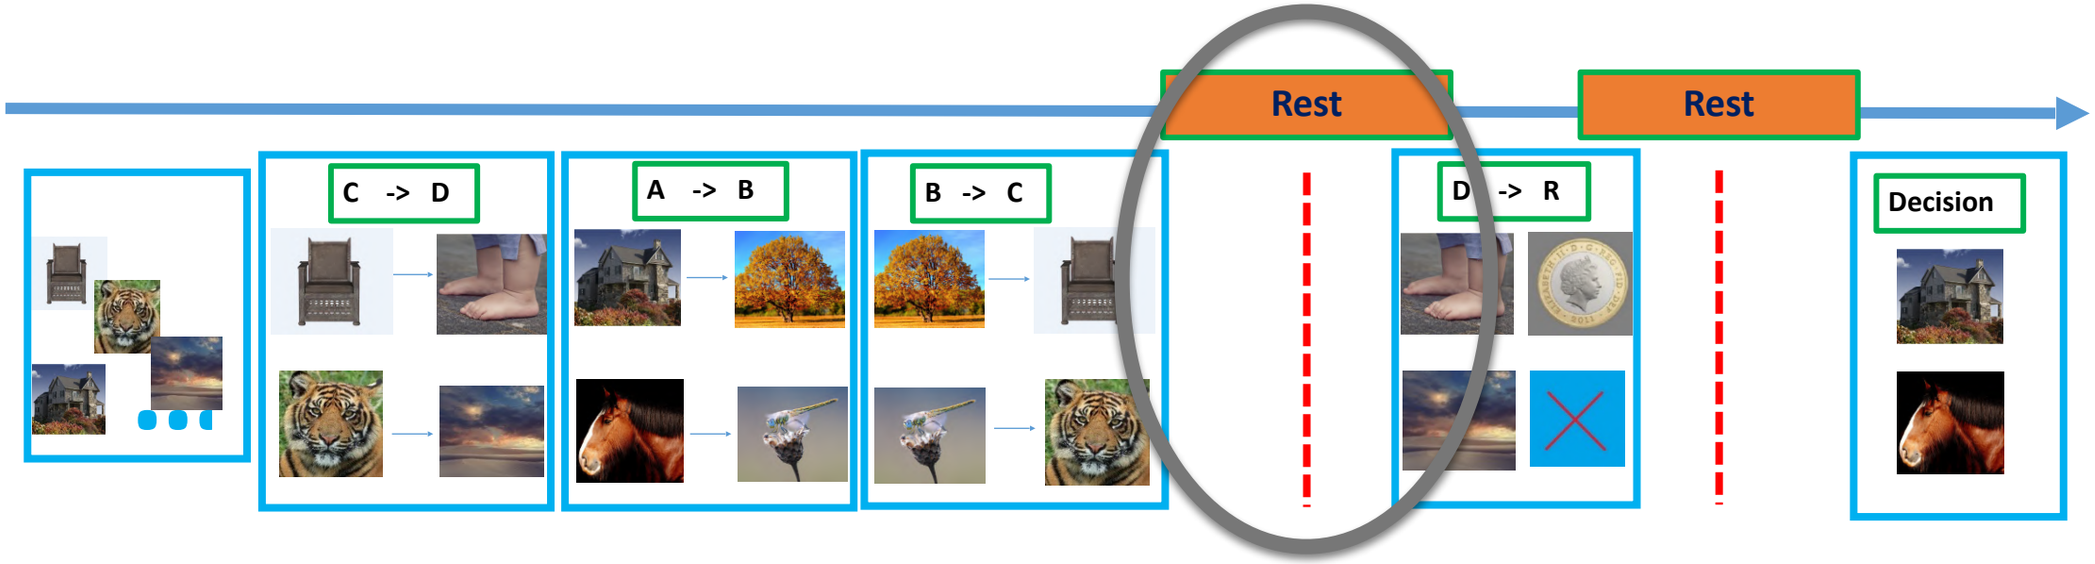

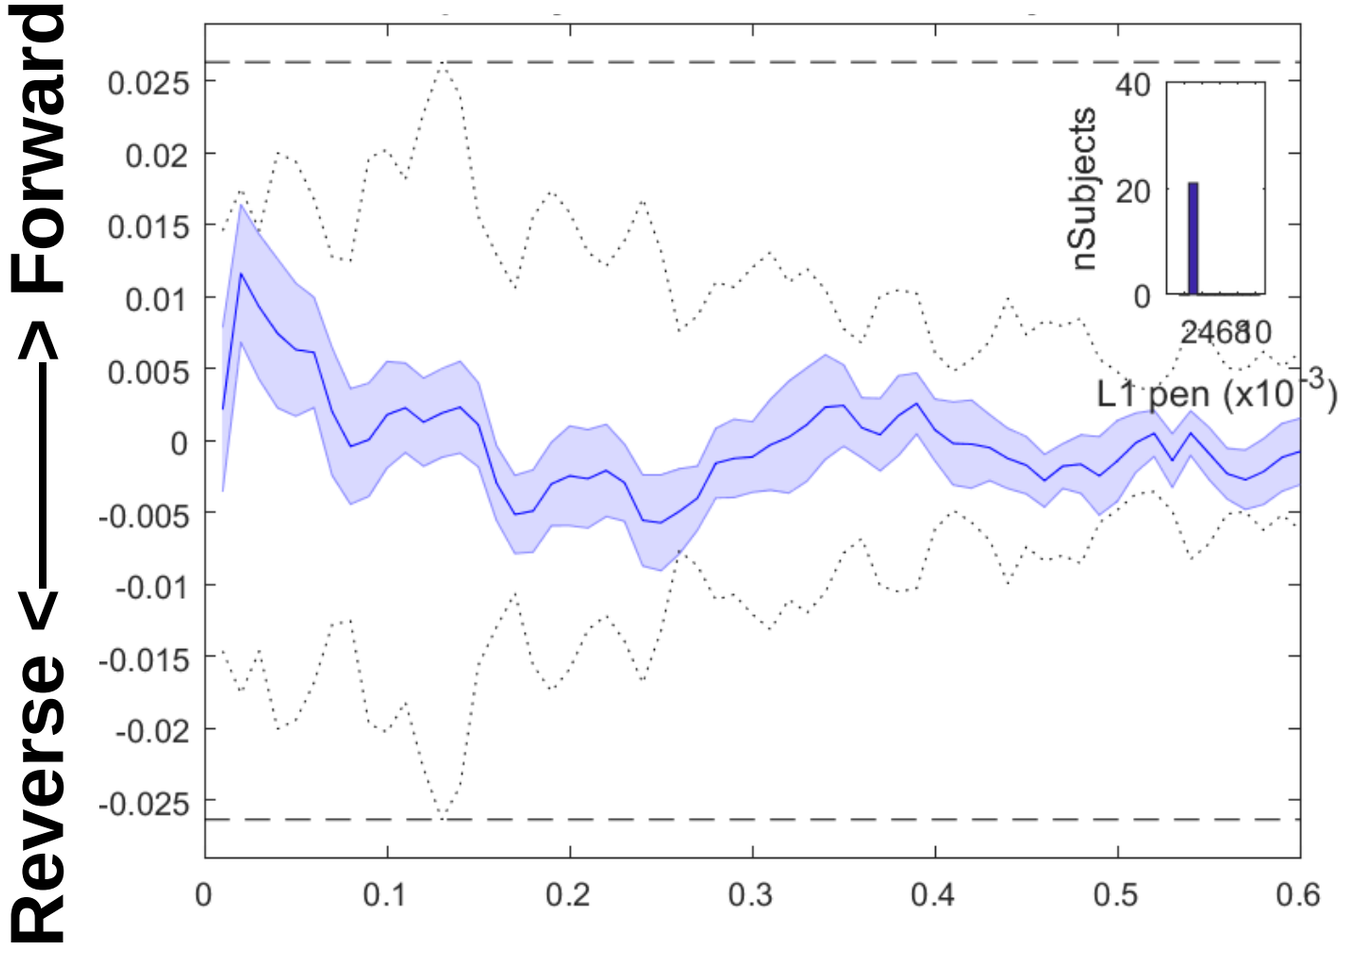

Nothing for the jumbled ones that they actually saw.

Time lag between elements (seconds)

Liu et al 2019

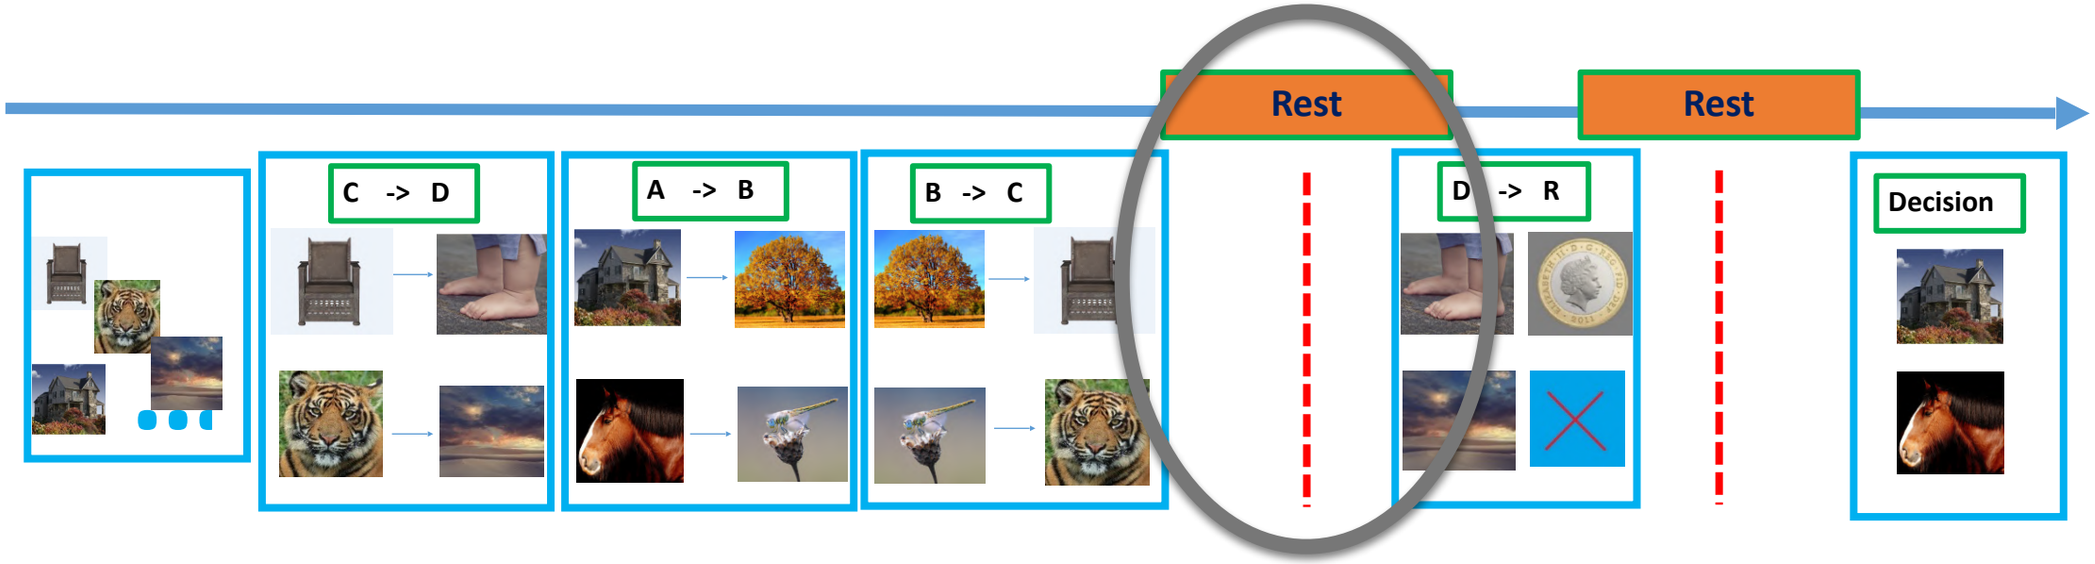

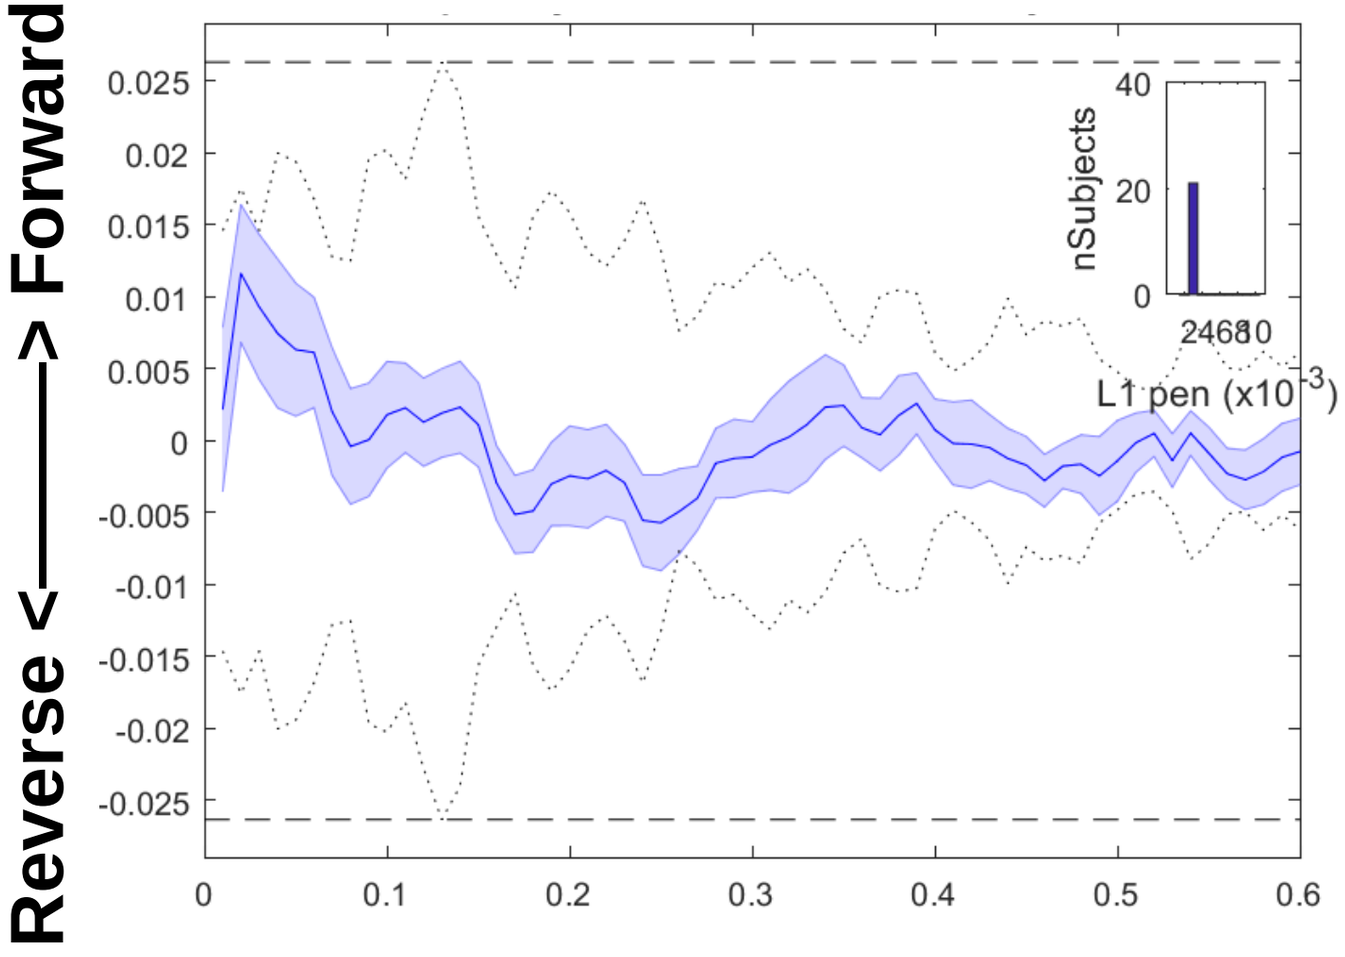

Nothing for the jumbled ones that they actually saw.

The learnt structure is unjumbling replay for new stimuli.

Time lag between elements (seconds)

Liu et al 2019

Liu et al 2019
Hahamy et al 2023

Watching the Film

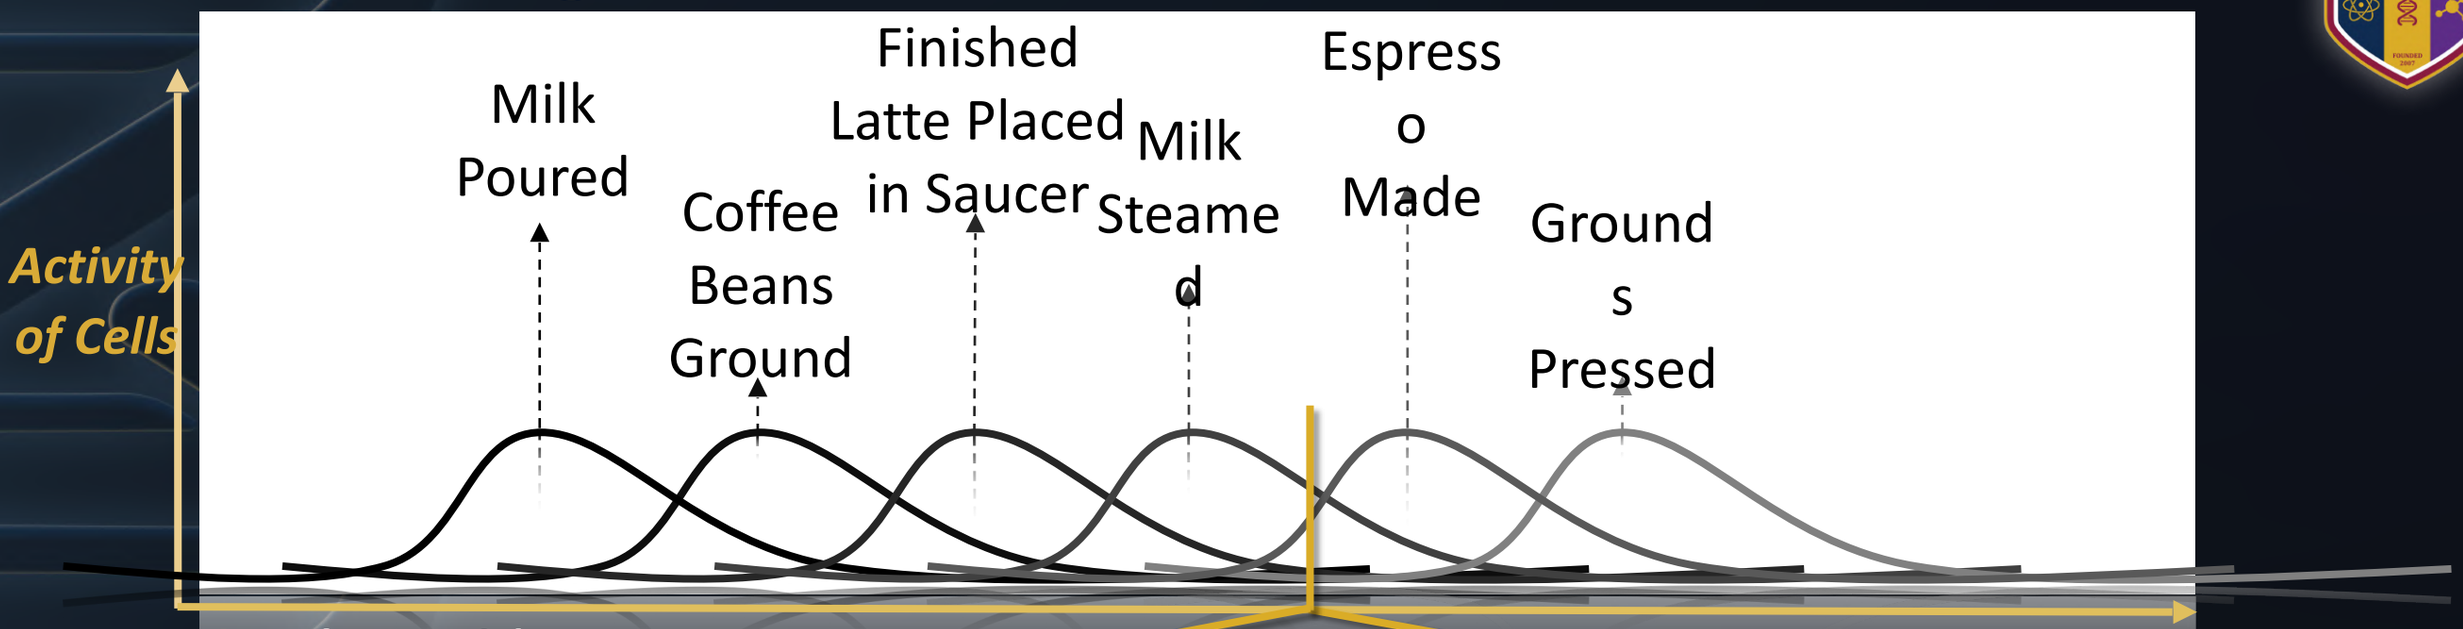

Replay Event

Time (milliseconds)

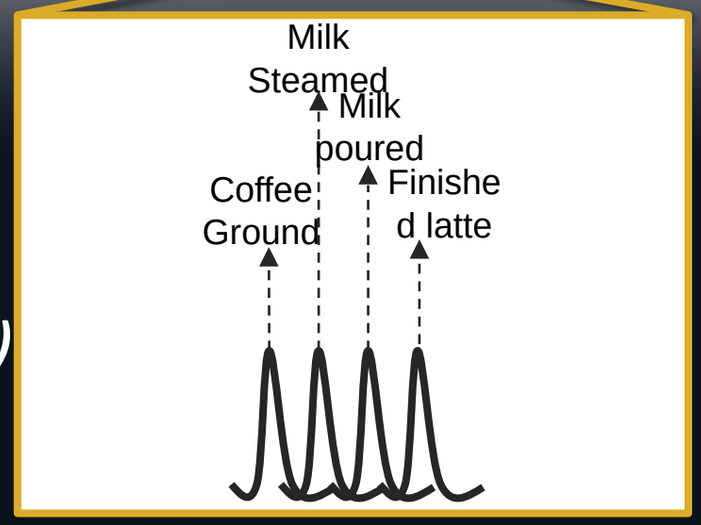

These replays happen at scene boundaries in actual movies.
Pause when you are speaking if you want people to understand!

Hahamy et al 2023

Liu et al 2019

What this talk is going to say

- Replay is not just replaying the past. It is doing inference about possible futures.

- It does so by aligning experience to our structural knowledge

- It is always making these inferences in the background, making you faster at thinking when the time comes.

- If there is time, some data about replay building new structural knowledge.

If replay is doing inference, how did this inference actually work?

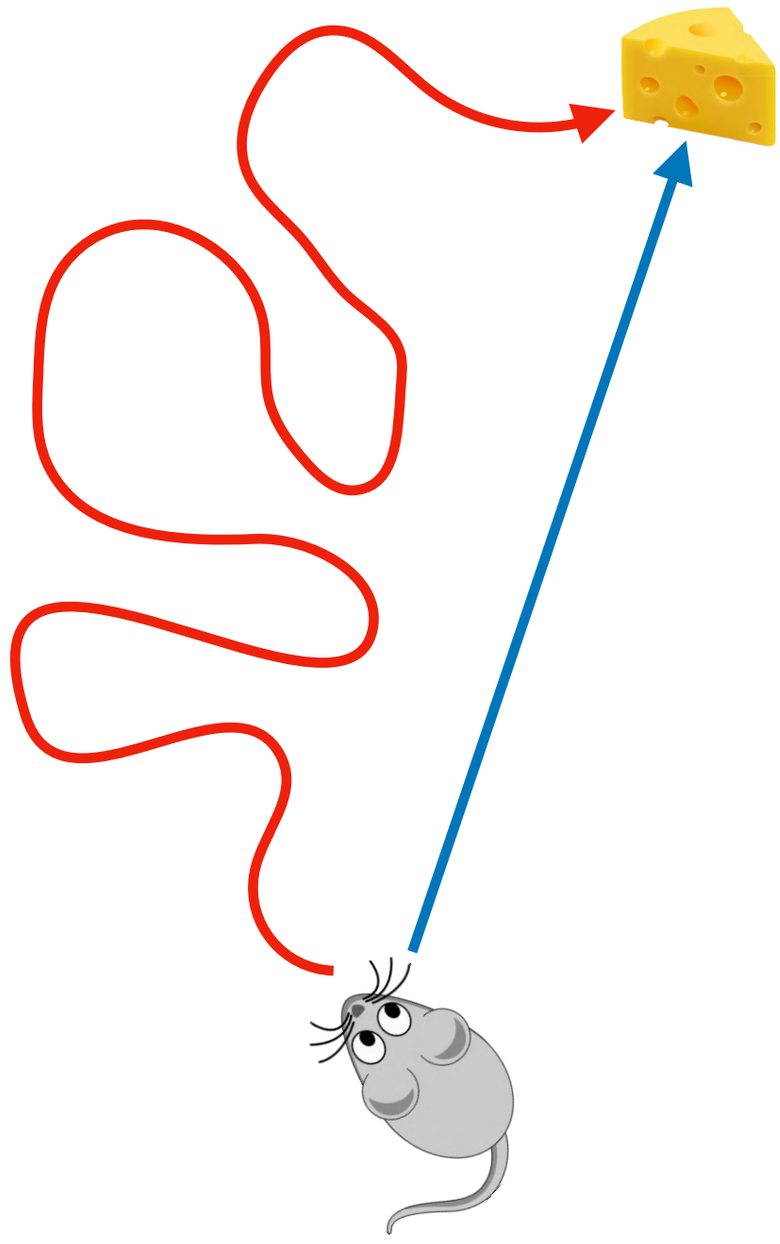

We need to see the cells so let’s go back to this spatial problem in rats

Jacob Bakermans, James Whittington Bakermans et al. 2025

If I map first, I don’t have to think later

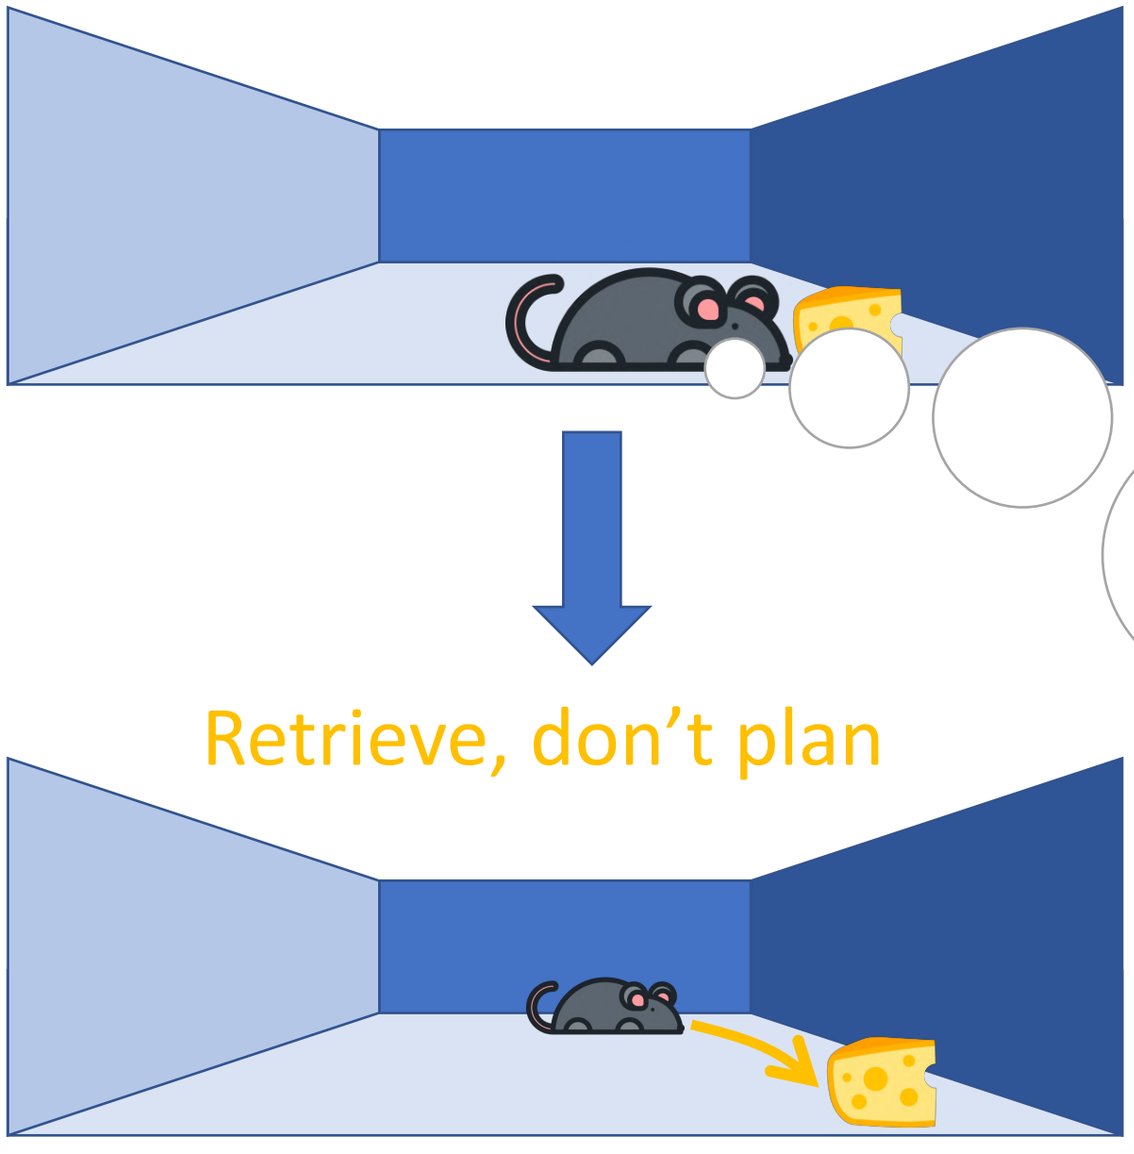

Retrieve, don’t plan

Replay builds cheese map

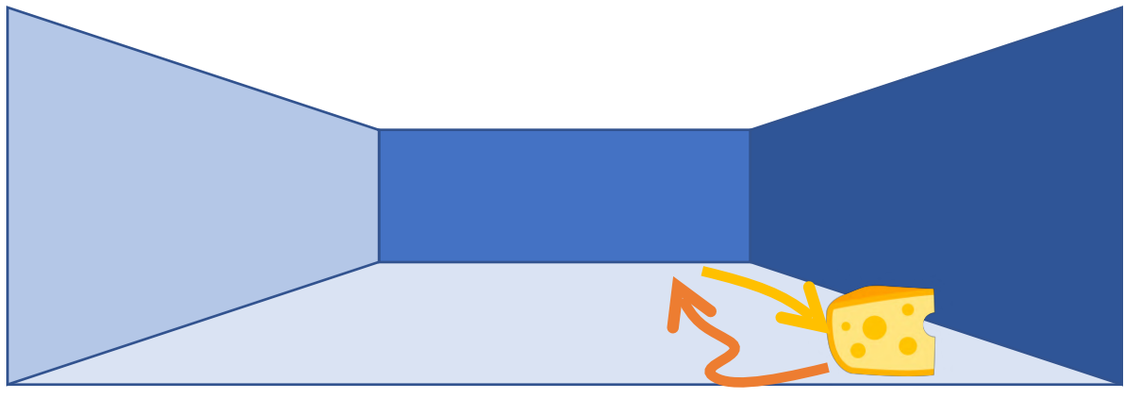

Cheese over there!

Bakermans et al. 2025

How to do this with neurons

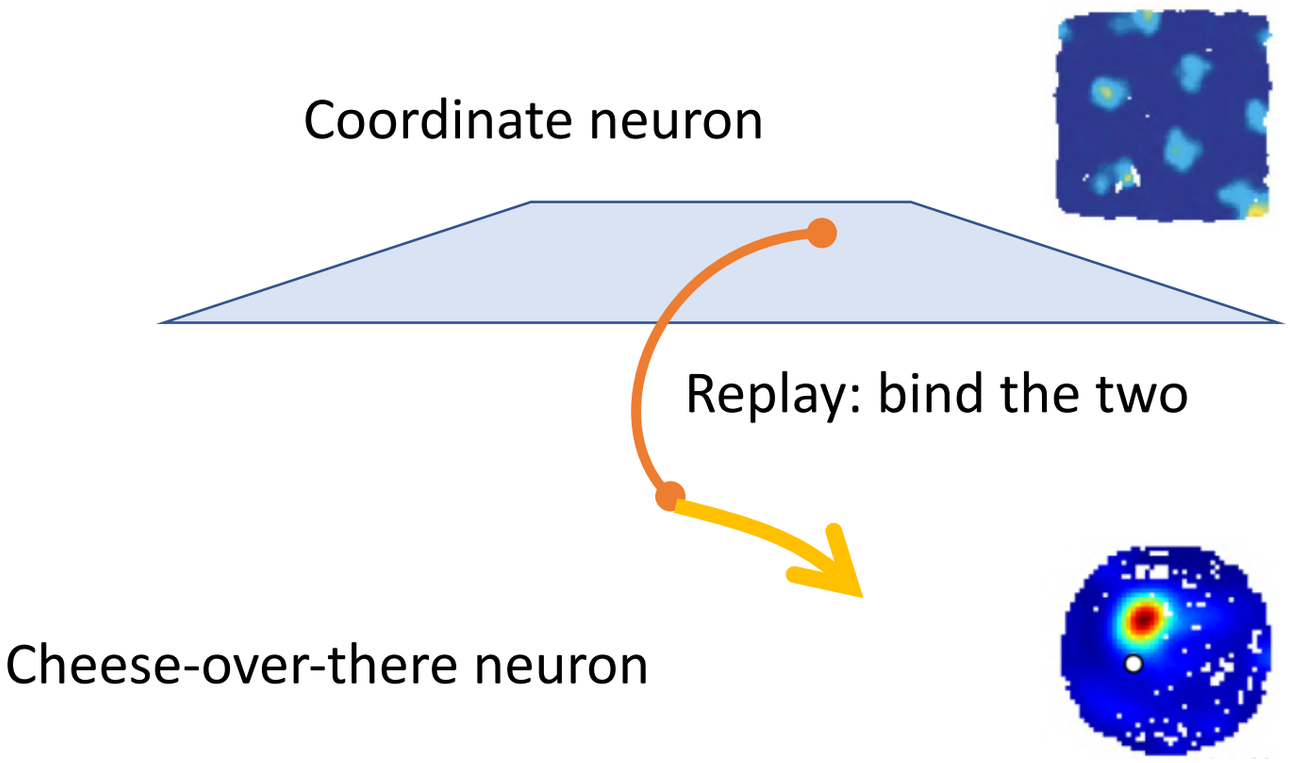

Mapping replay creates new ‘linking’ place cells

Bakermans et al. 2025

OK so let’s get some rats doing this

Find the hidden cheese

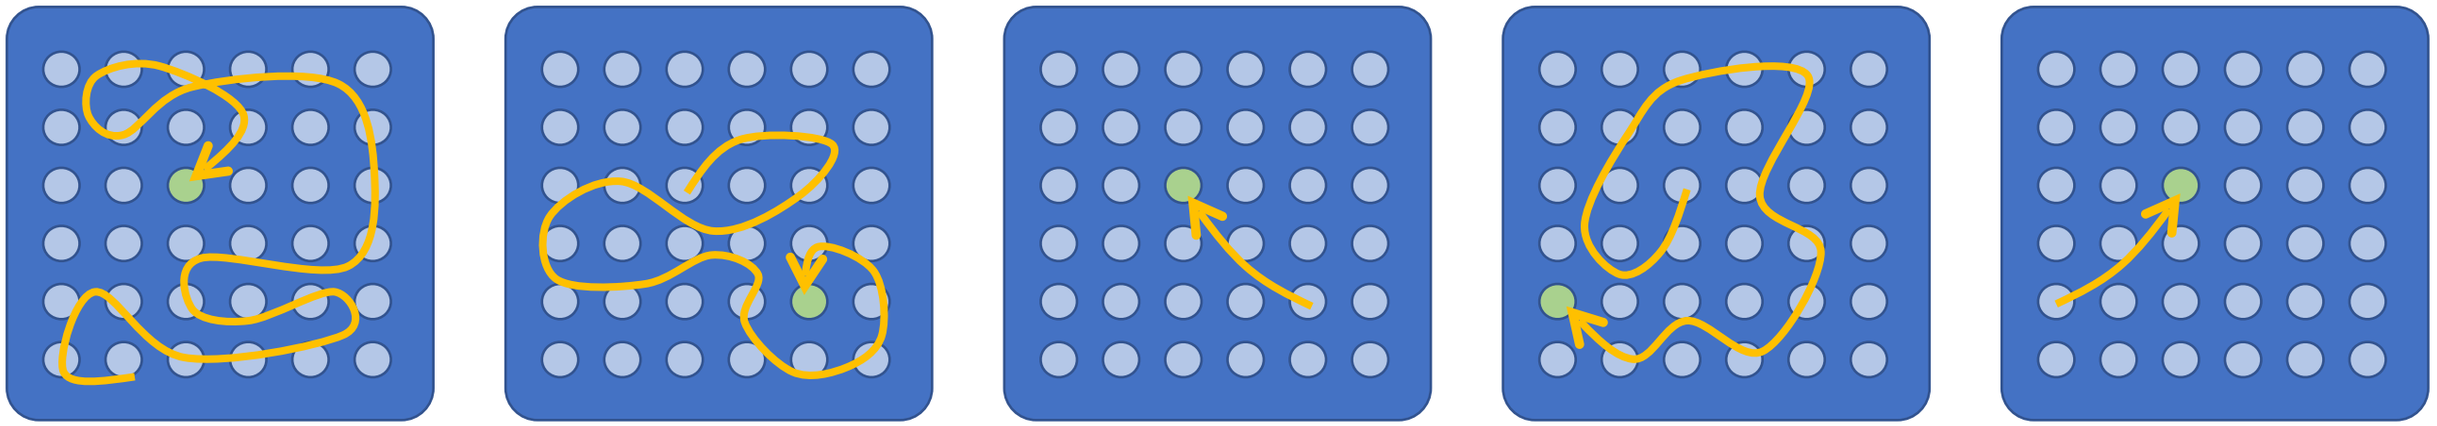

1. Unknown location

2. Random location

3. Back to 1 (home)

4. Random location

5. Back to 1… Etc.

They need to learn to get home from anywhere

Data from Pfeiffer and Foster 2013

What happens at home?

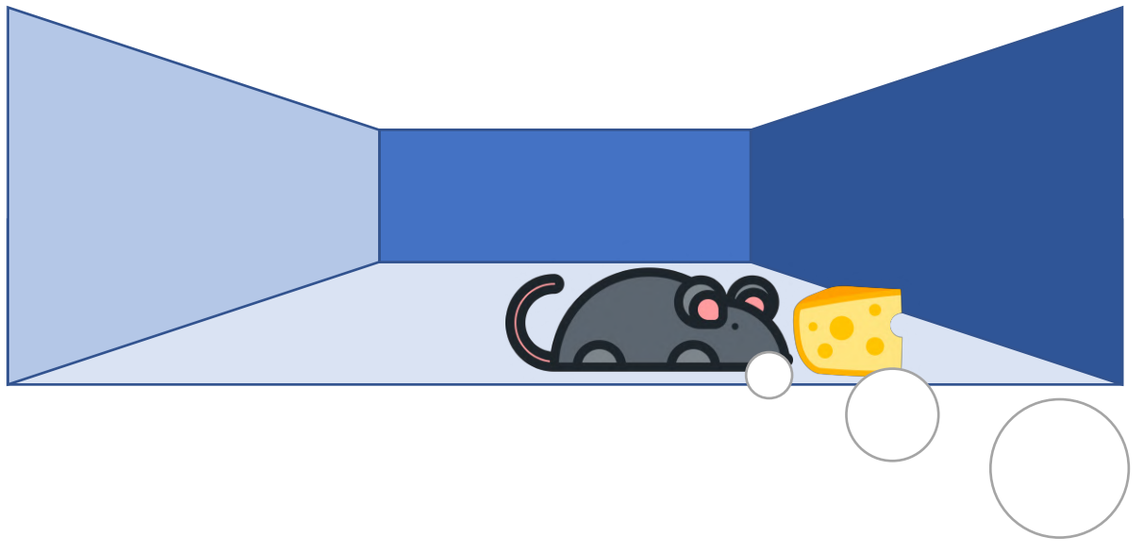

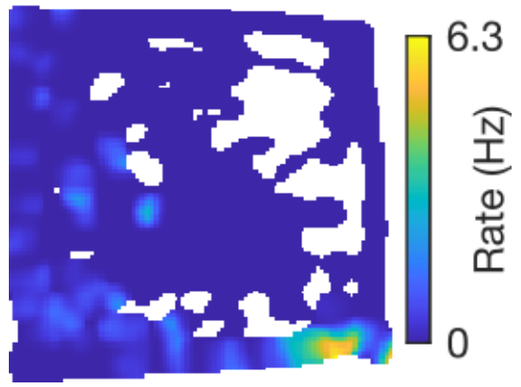

Before replay

Replay trajectory

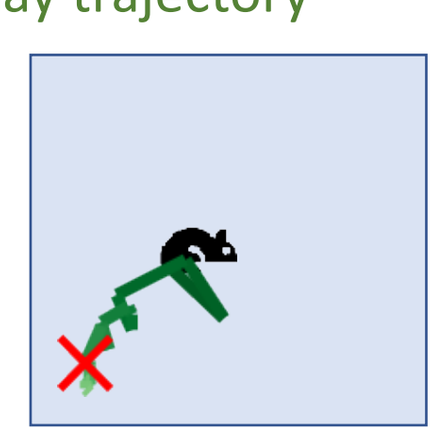

New field at spike

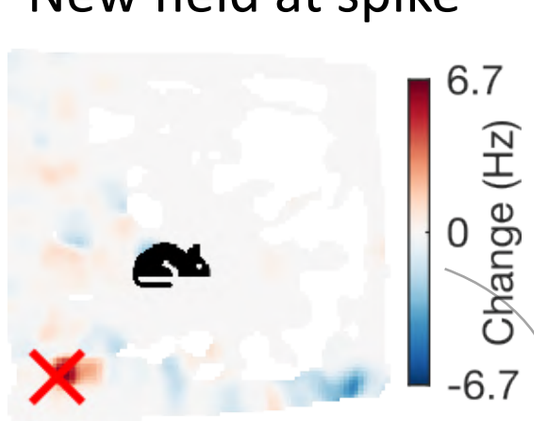

This neuron’s spike

After replay

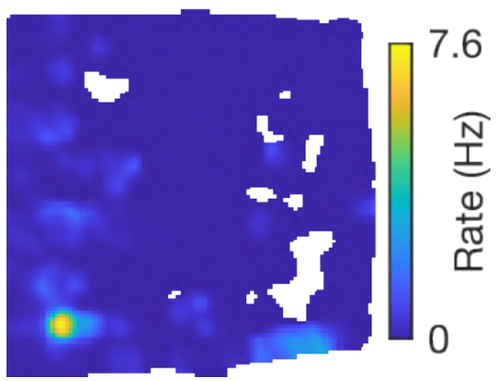

Bakermans et al. 2025

Here’s another example

Before replay

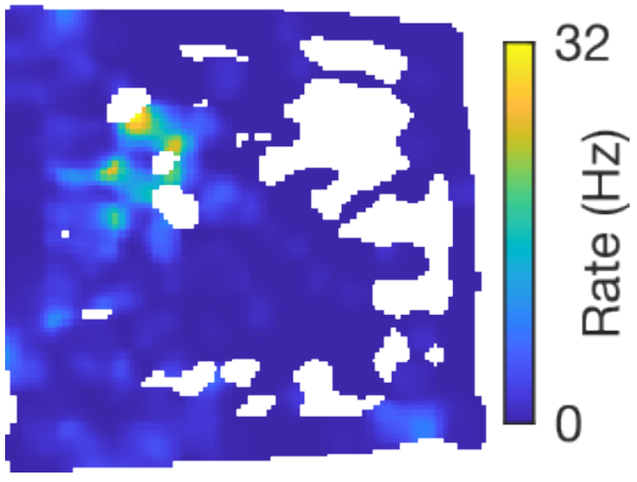

Replay and change

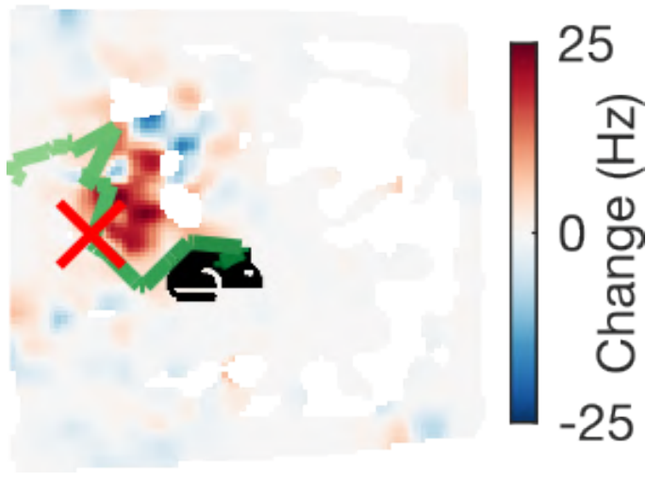

After replay

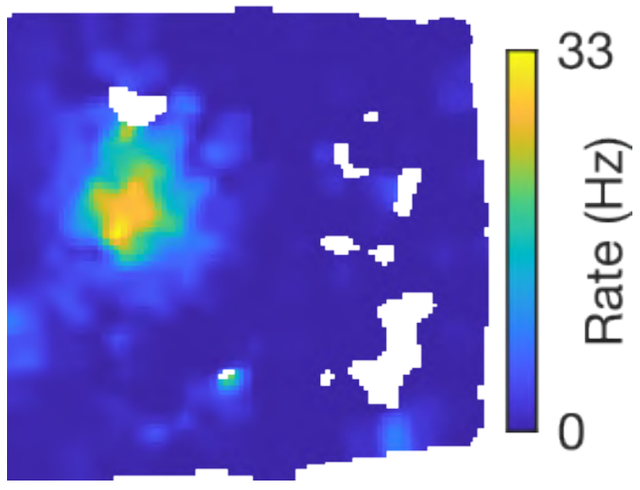

New place field at replayed location

Bakermans et al. 2025

This is true across the population

Take this change map

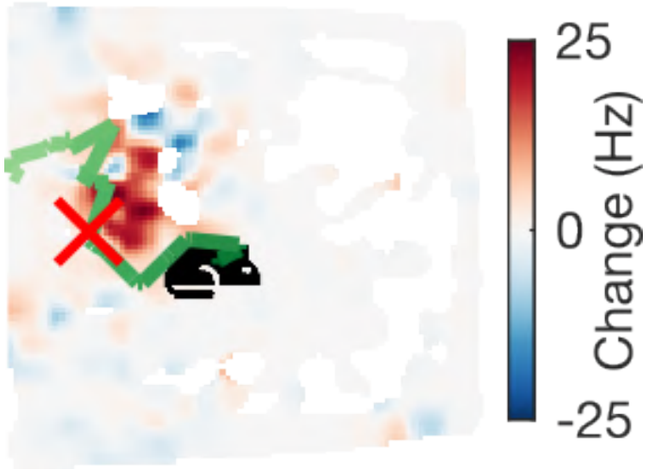

Cut out around X

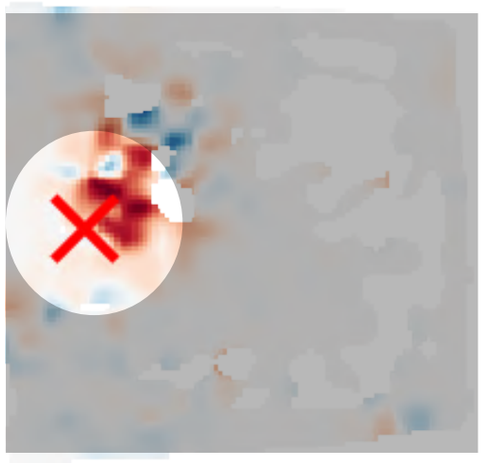

Aligned change on average

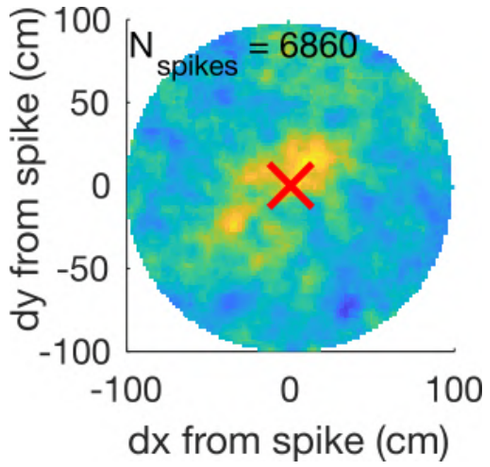

But not for replay in other places

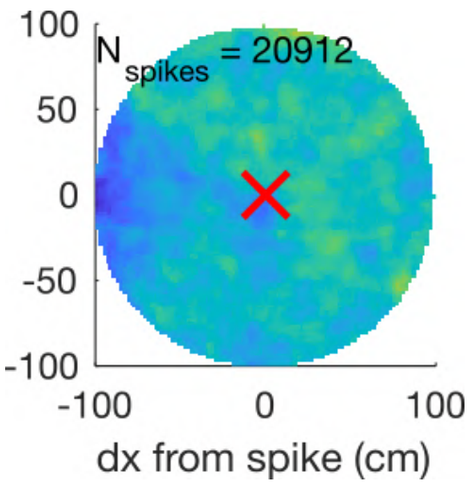

… And average across all replays across all cells

Replay creates new place fields

Bakermans et al. 2025

What if the cheese moves?

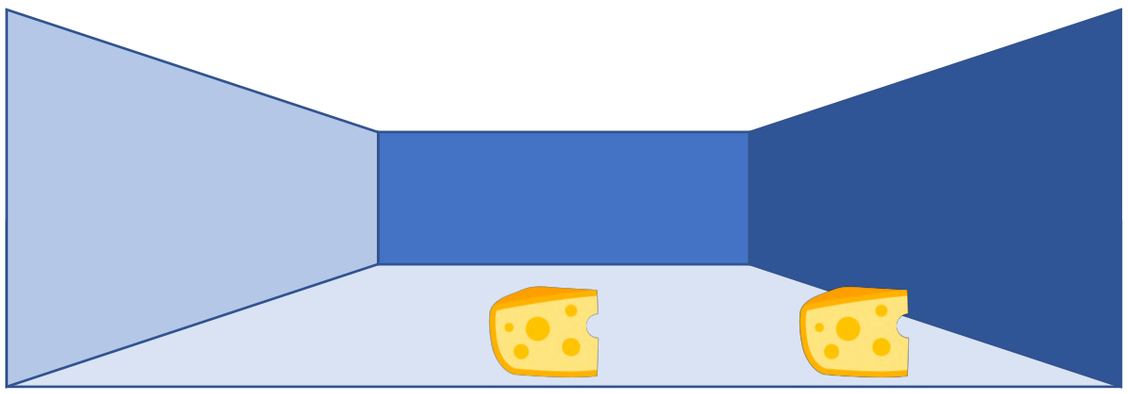

Remove link with respect to old location
Create new place field at same relation

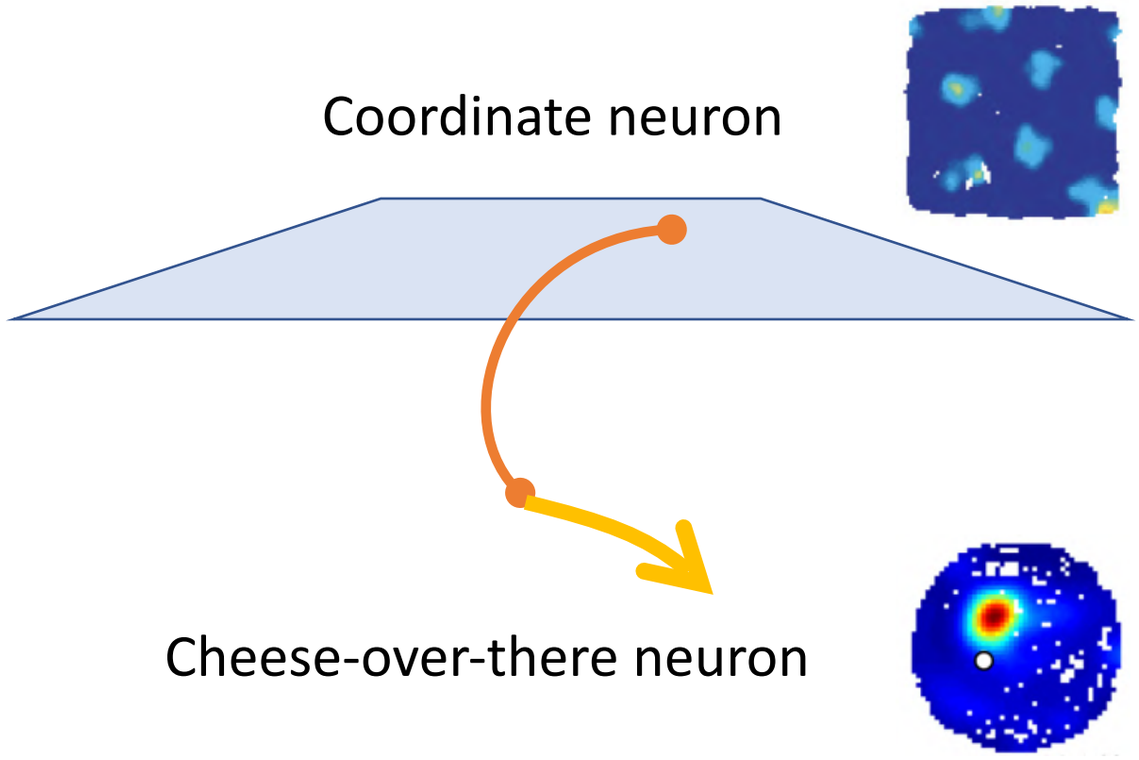

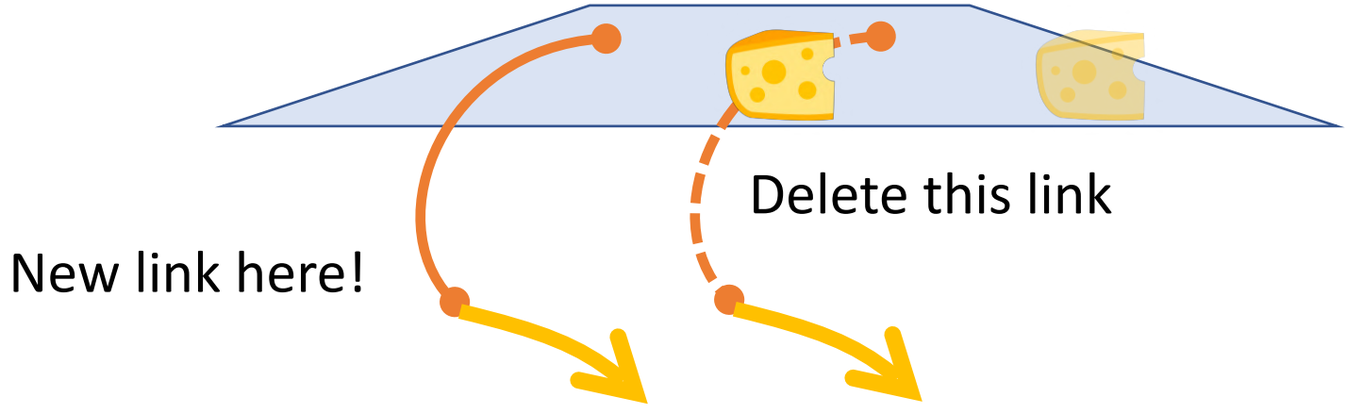

Bakermans et al. 2025

We also find neurons like that

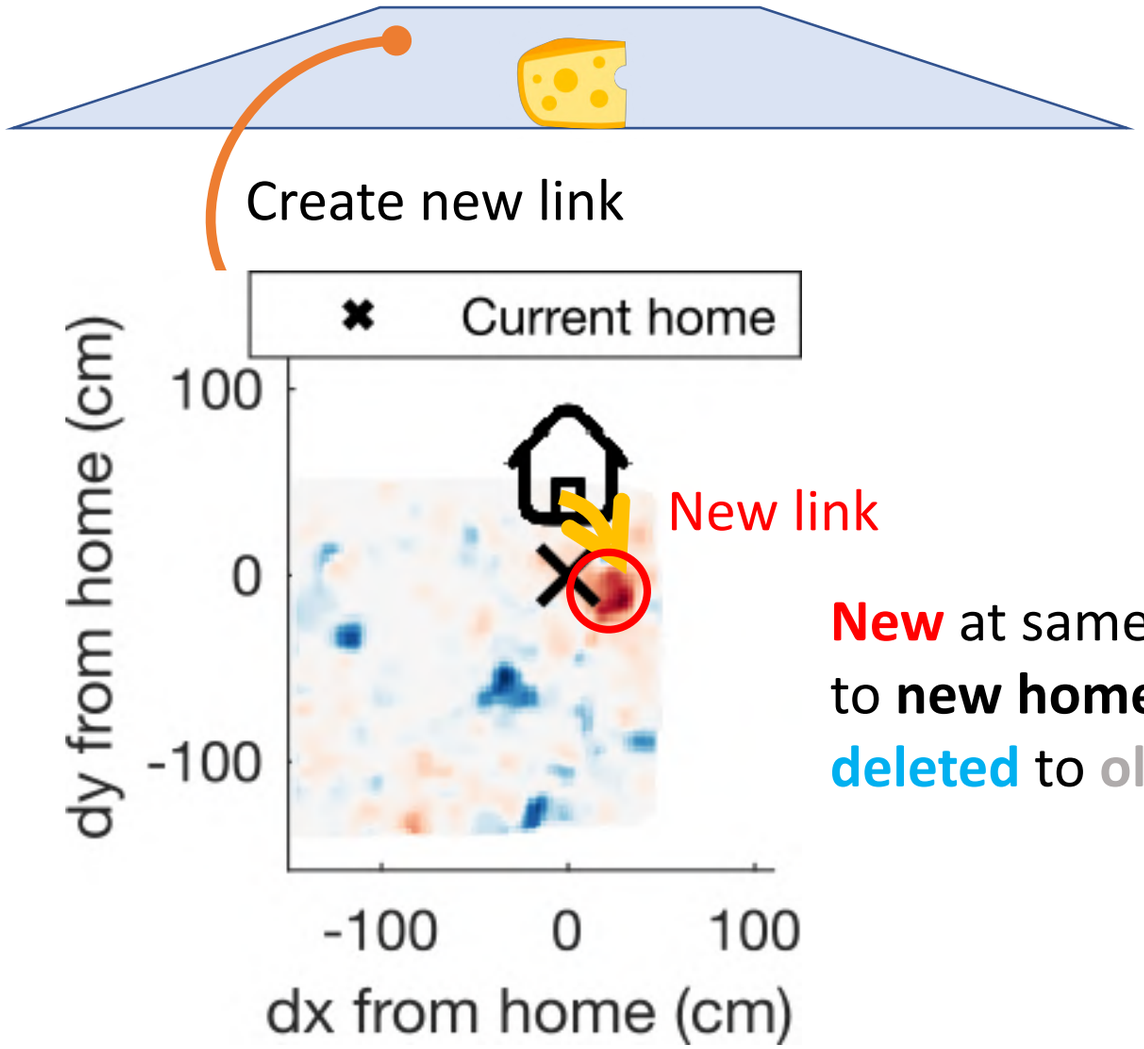

New at same relation to **new home** as **deleted** to **old home**

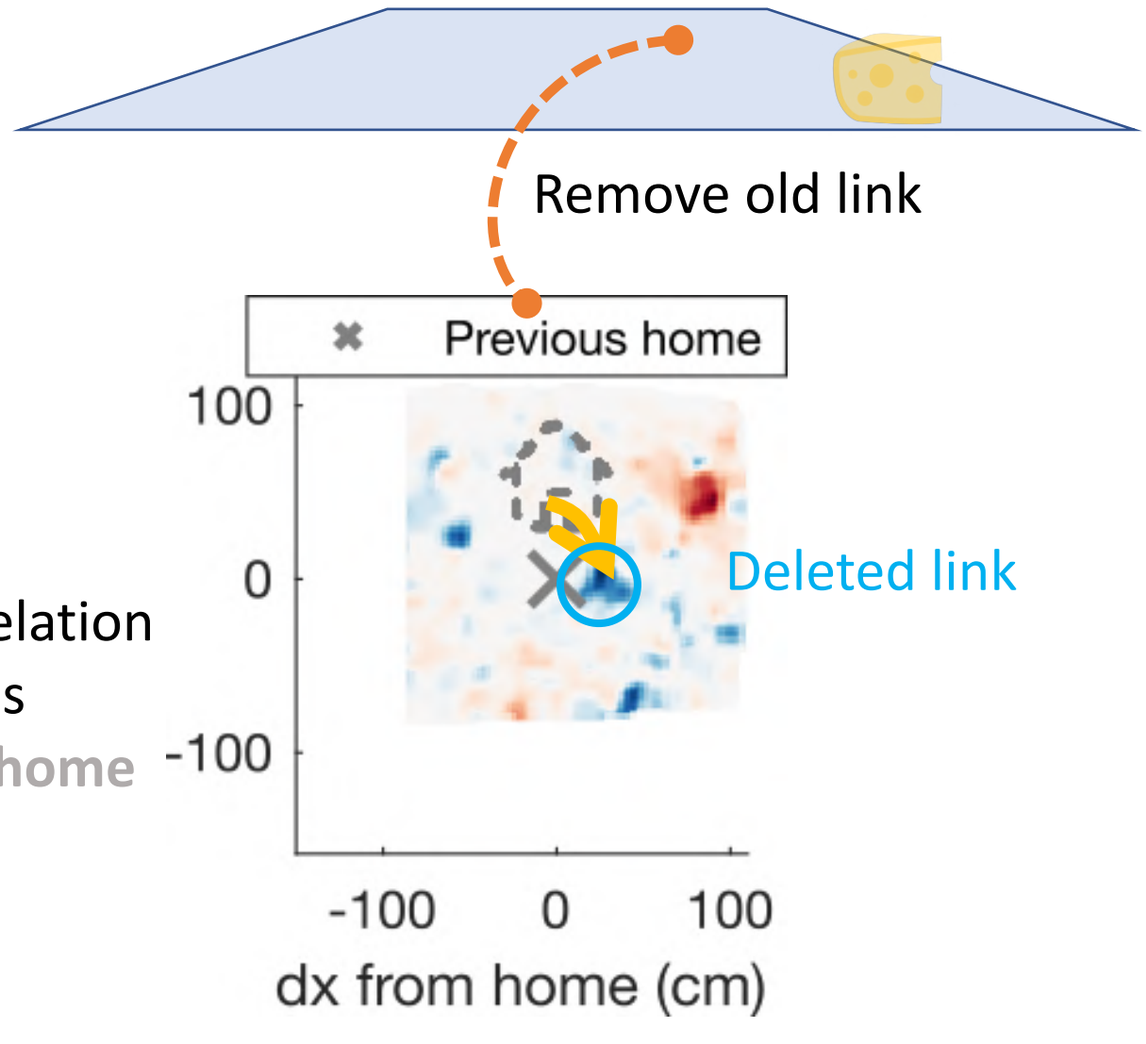

Neuron that generalises home relation

Bakermans et al. 2025

And these are the same neurons that replay

Change wrt new home

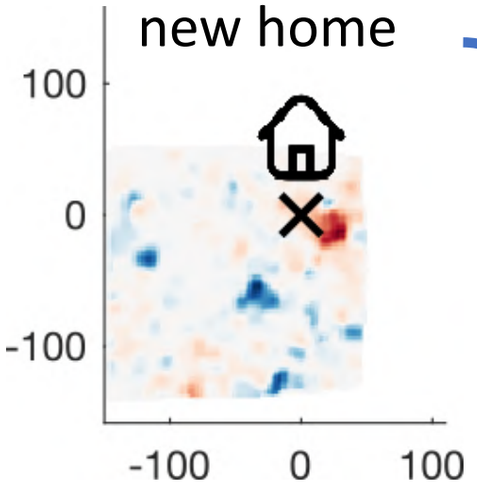

Change wrt old home

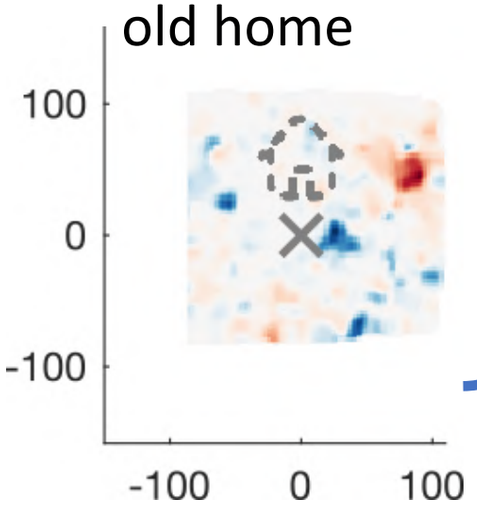

Correlate these maps: red on top of blue, so anticorrelation

These cells shift their response

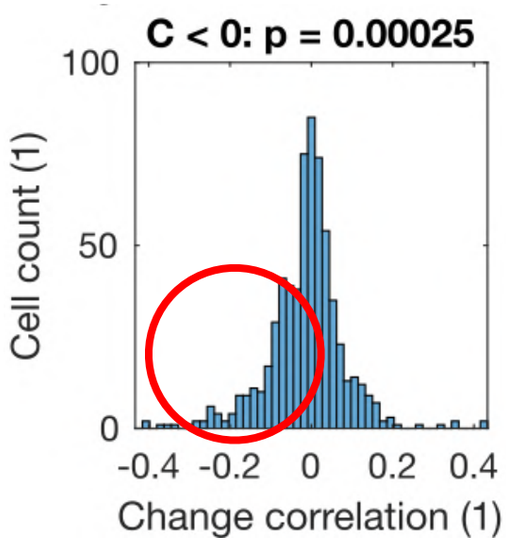

There is a subpopulation of these generalising neurons

Same cells replay more

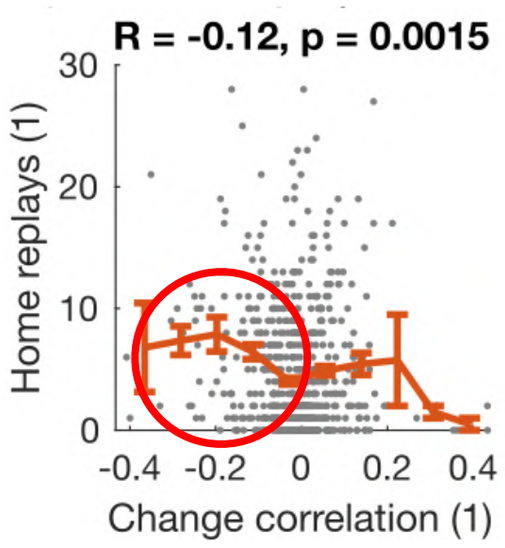

And these are the same neurons that replay!

Replay is shifting these vector cells to new coordinates in the map structure

What this talk is going to say

- Replay is not just replaying the past. It is doing inference about possible futures.

- It does so by aligning experience to our structural knowledge

- It is always making these inferences in the background, making you faster at thinking when the time comes.

- If there is time, some data about replay building new structural knowledge.

- Does this same mechanism work in more complicated human tasks?

- If you bind to a coordinate system offline, does it really mean you don’t have to think online.

- Let’s go back to humans to find out!

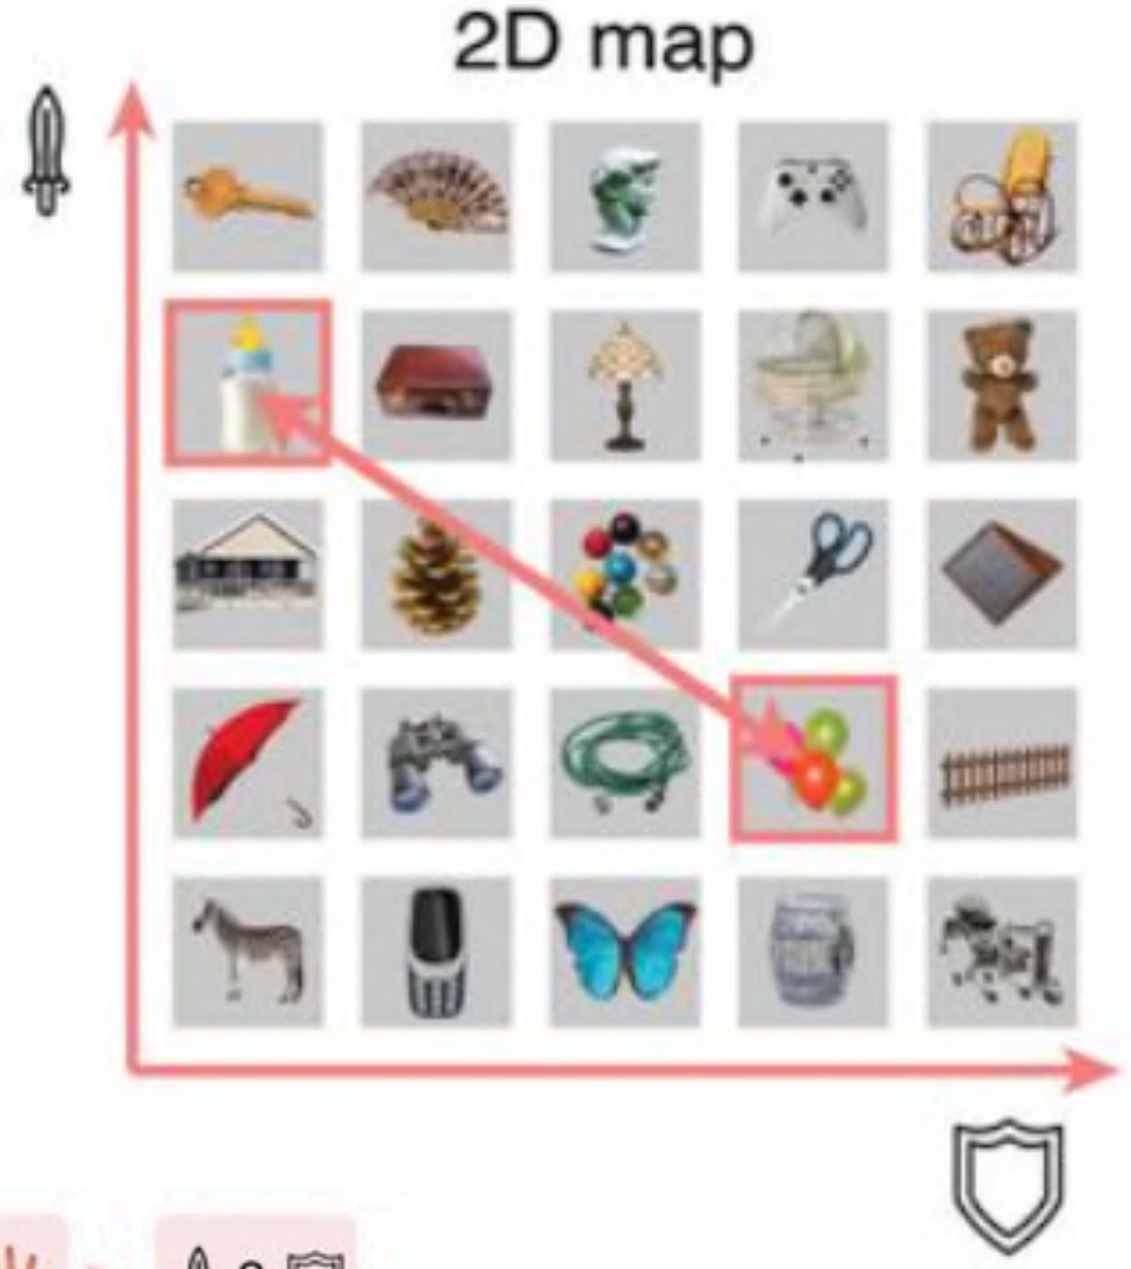

We’re going to play this same game, except you have to remember the strengths

Replay during the task

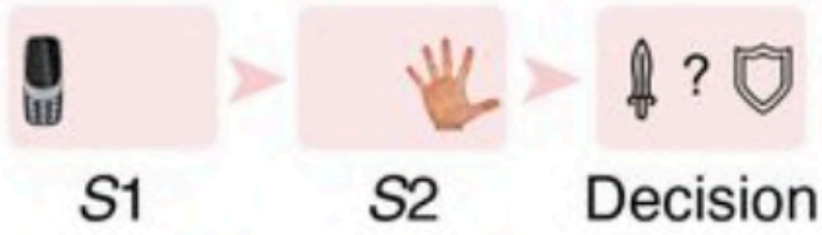

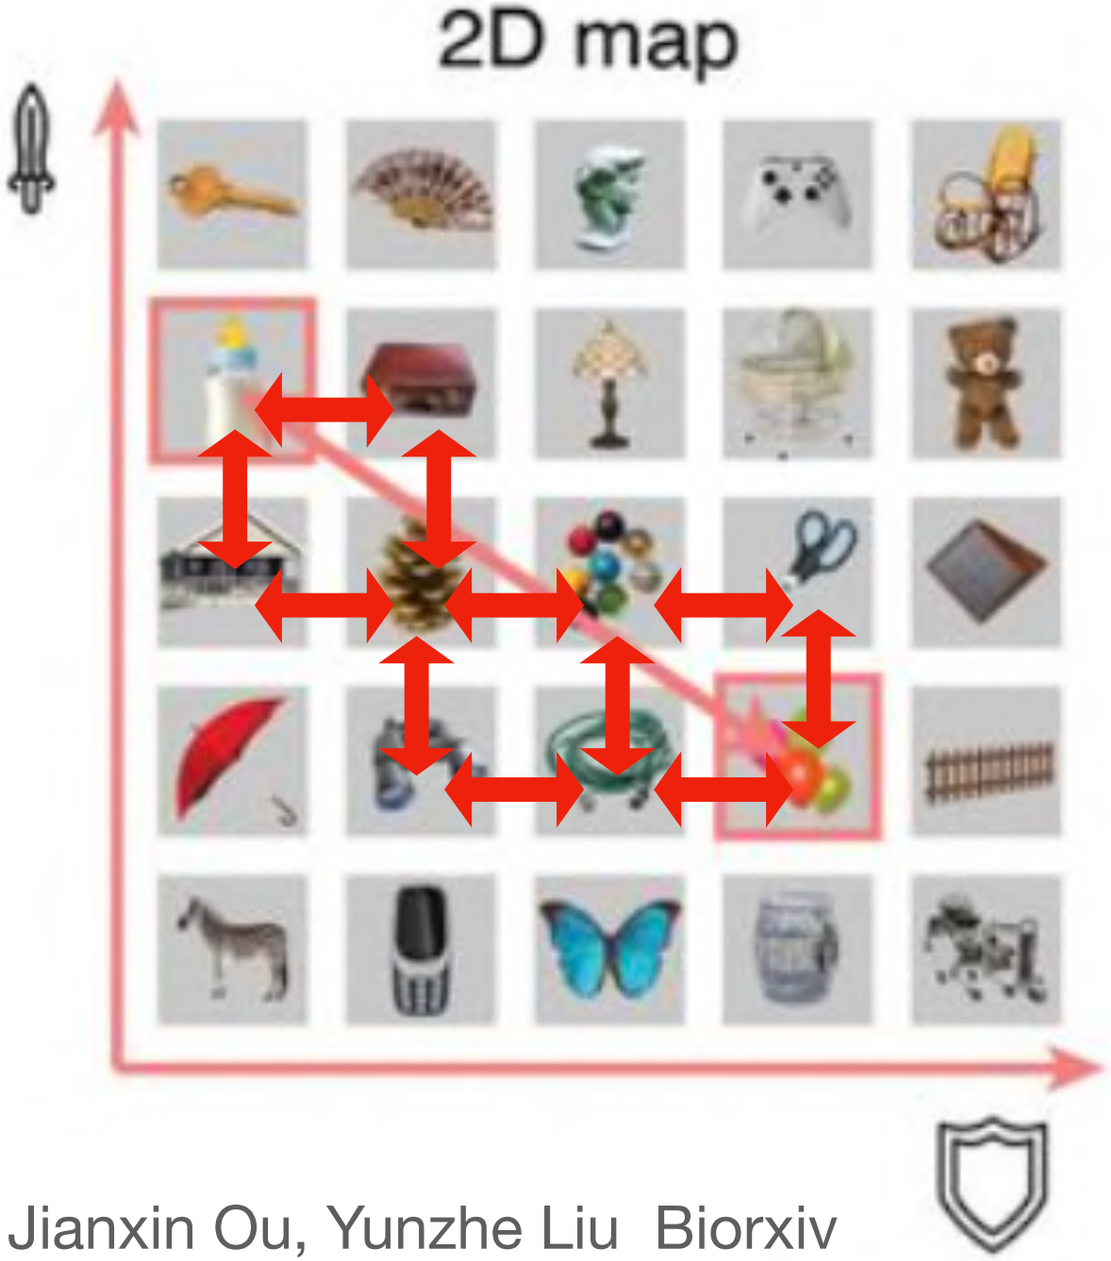

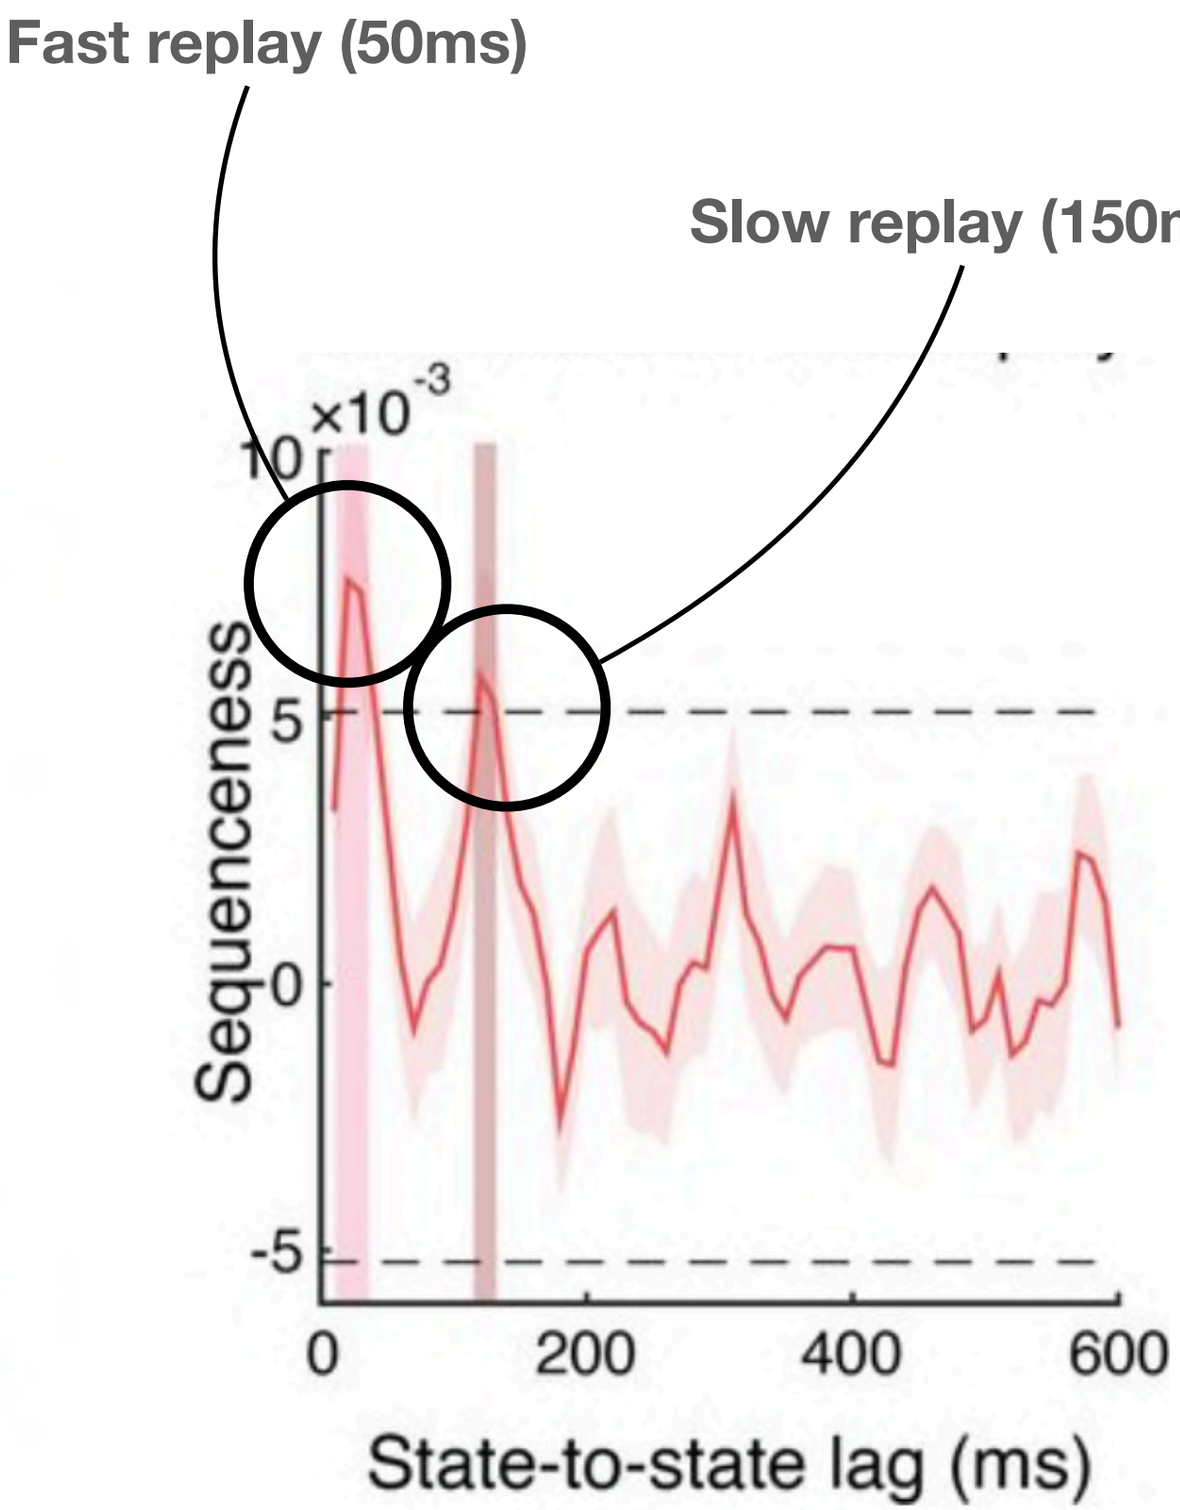

Jianxin Ou, Yunzhe Liu Biorxiv

Replay during the task

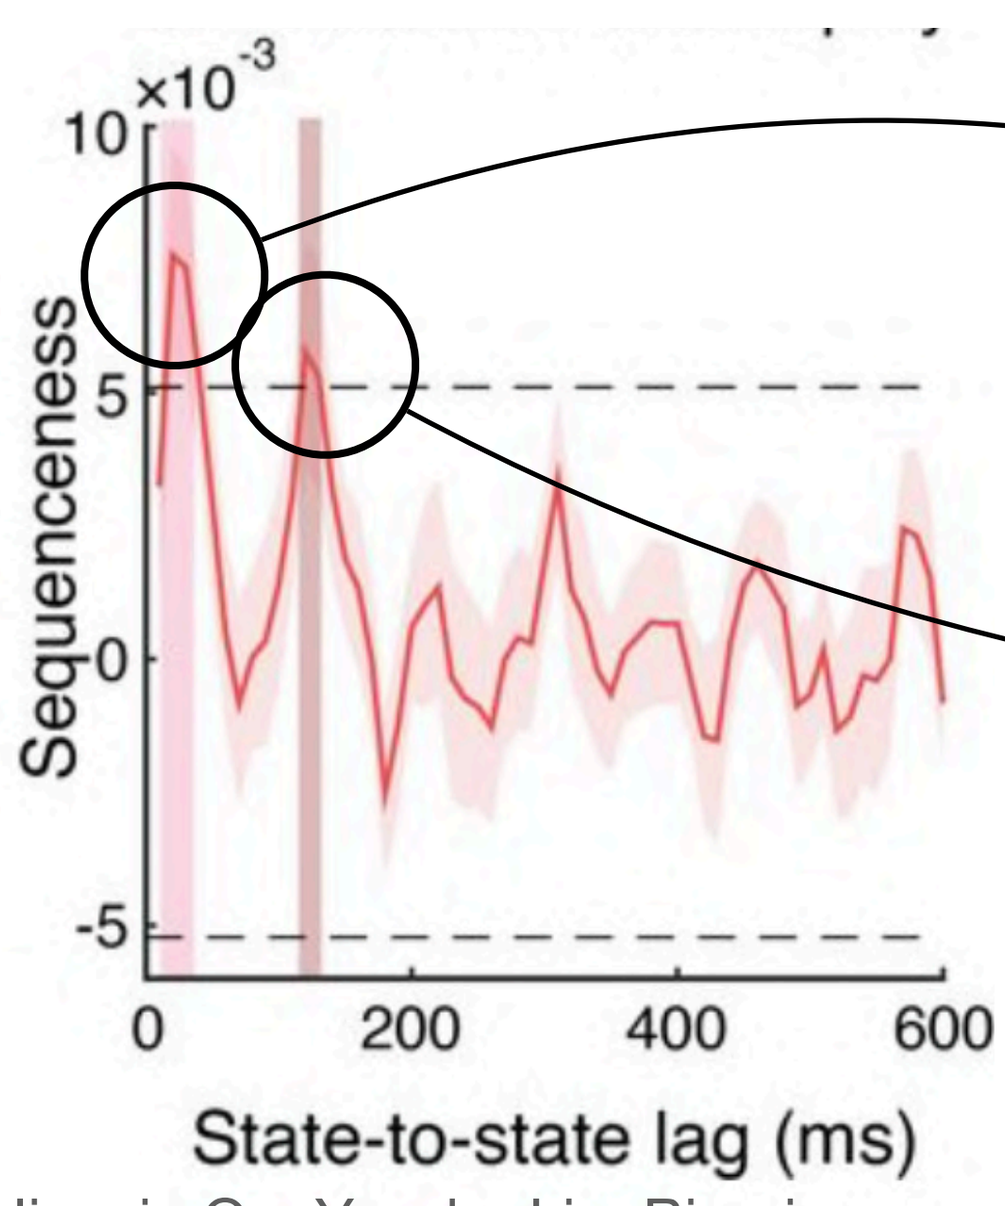

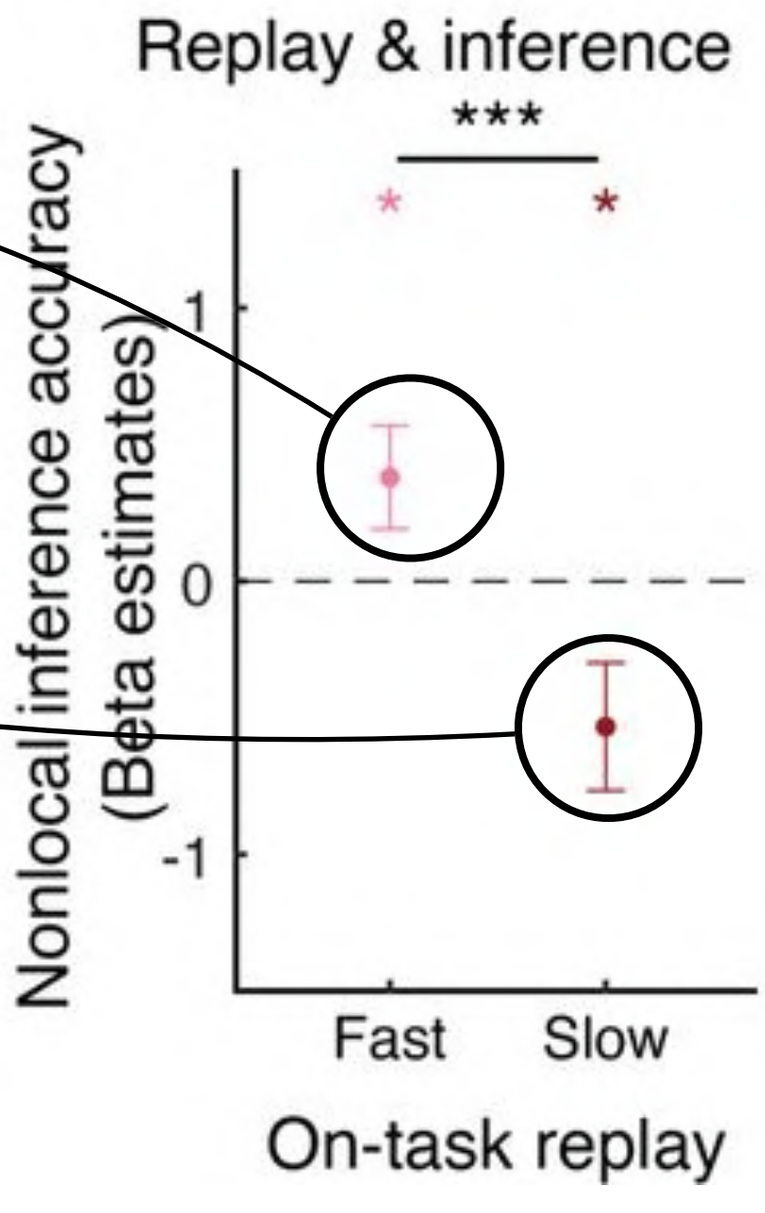

Jianxin Ou, Yunzhe Liu Biorxiv

Replay during the task

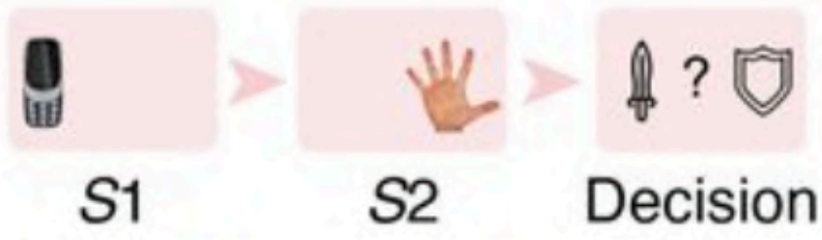

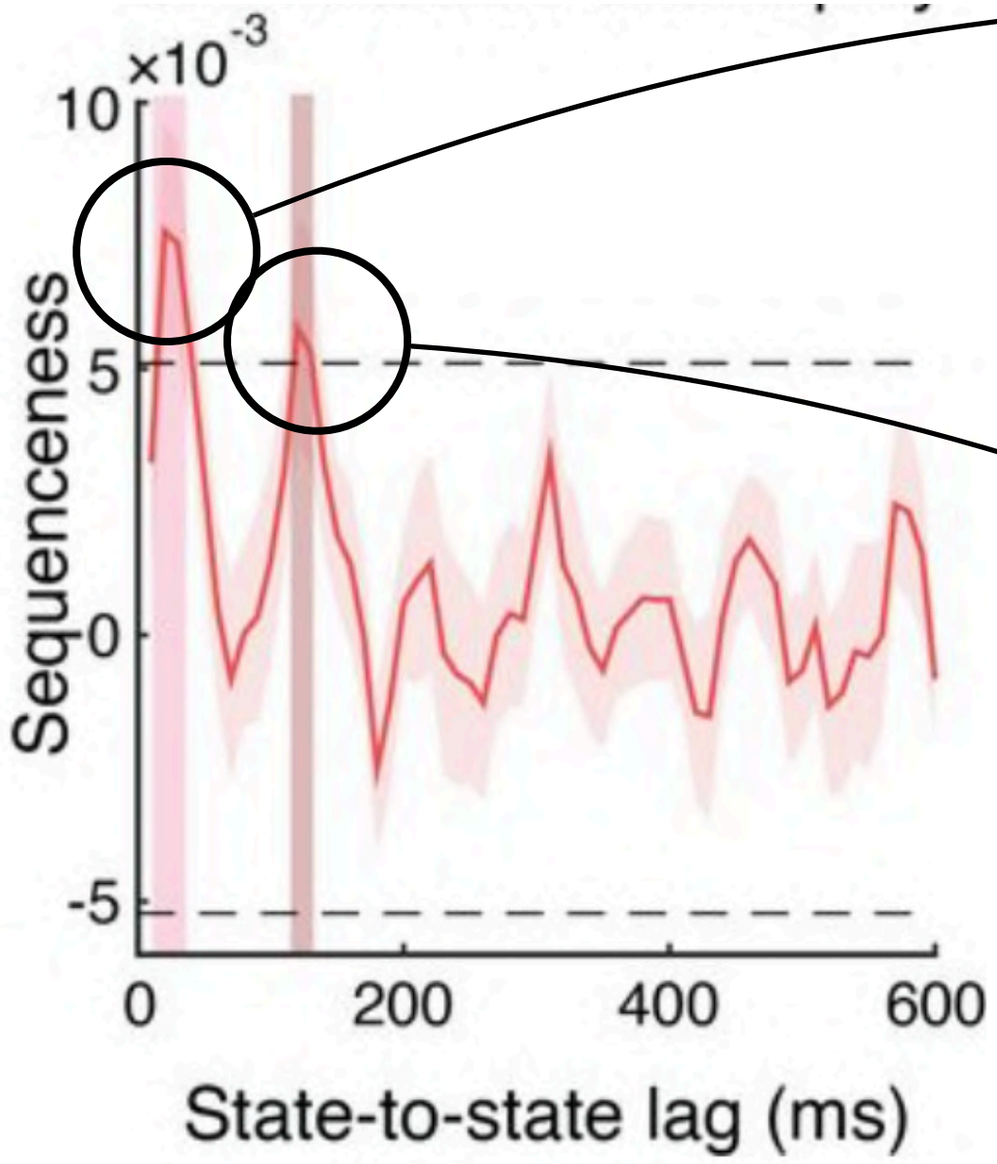

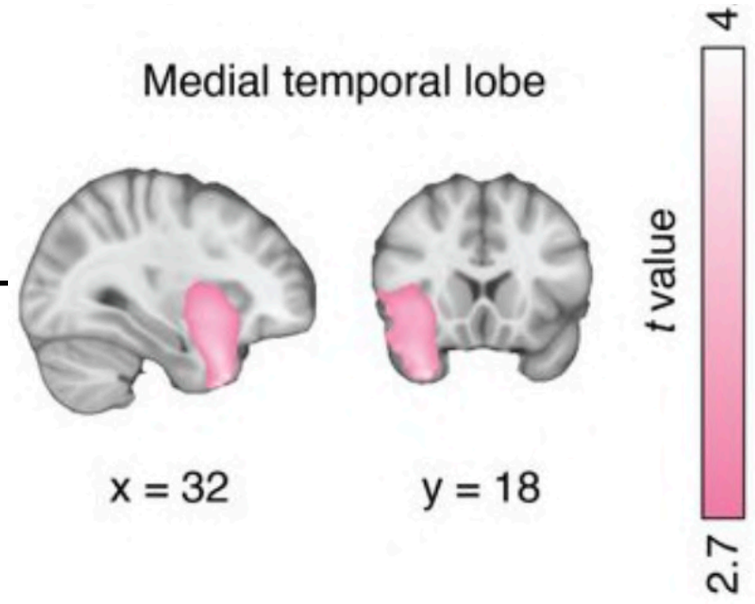

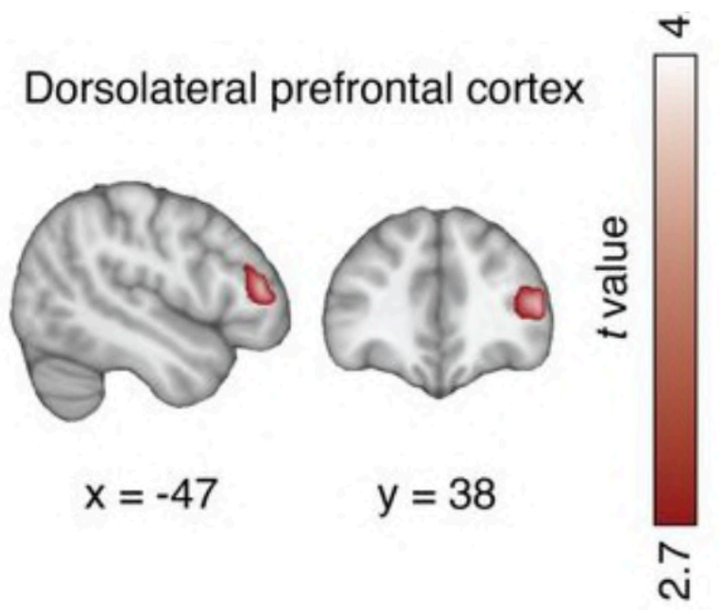

Jianxin Ou, Yunzhe Liu Biorxiv

But wait - if you know the coordinates, you don't need to sequence.

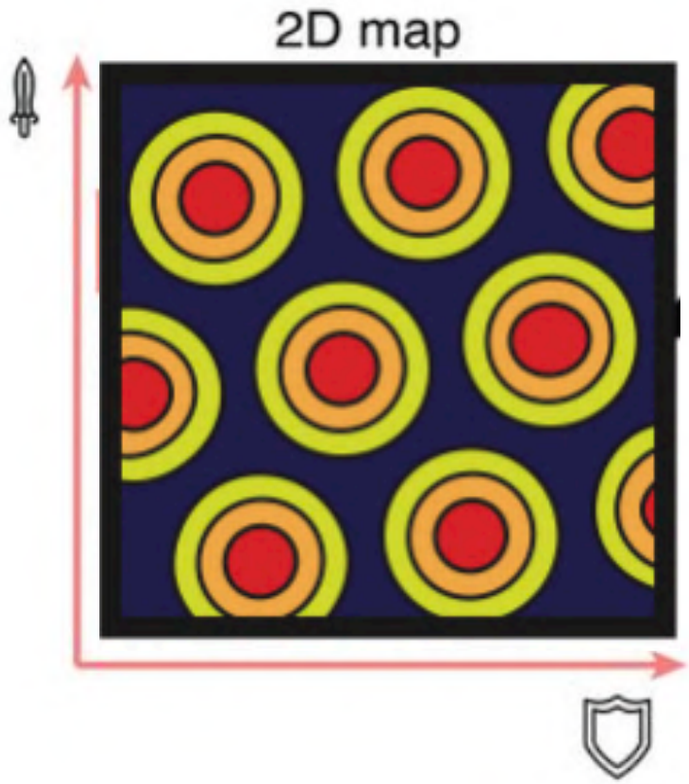

Grid cell smoke signal!

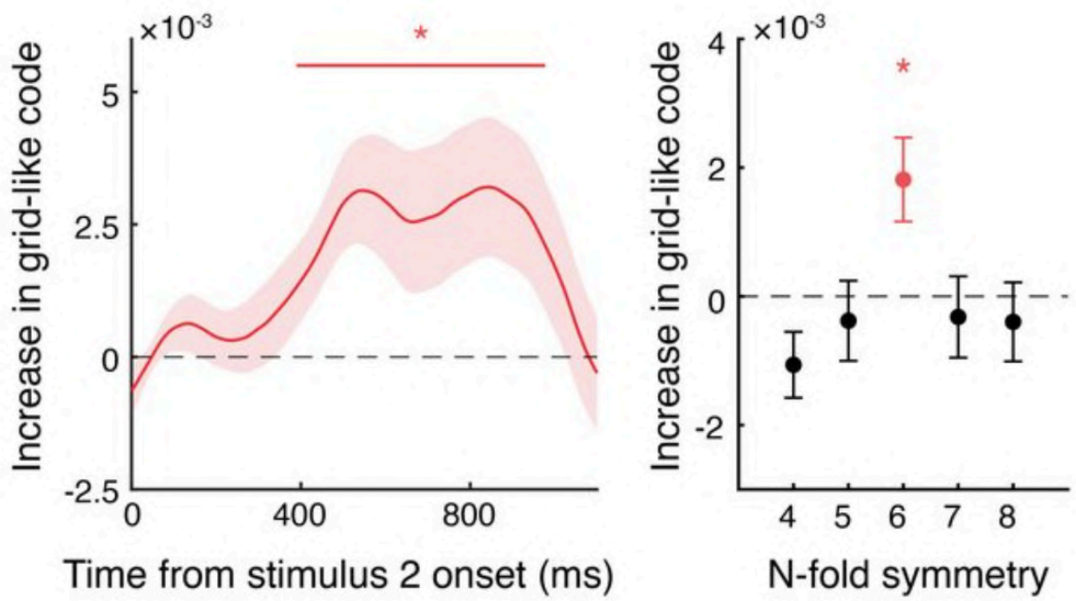

If you have the map, you don't have to think!

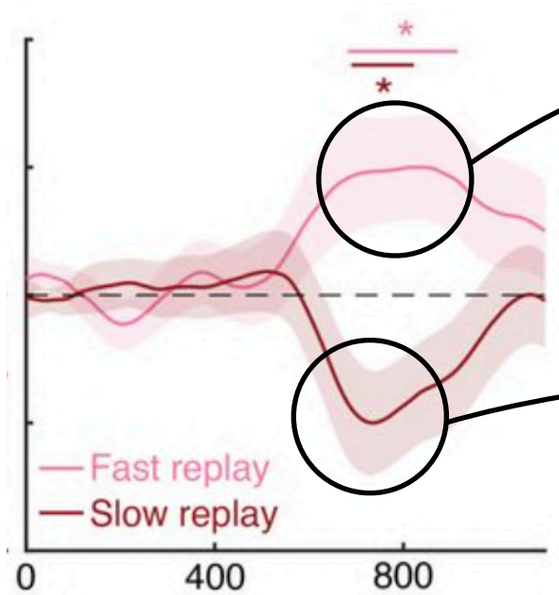

Grid signal correlates
+ve with fast replay

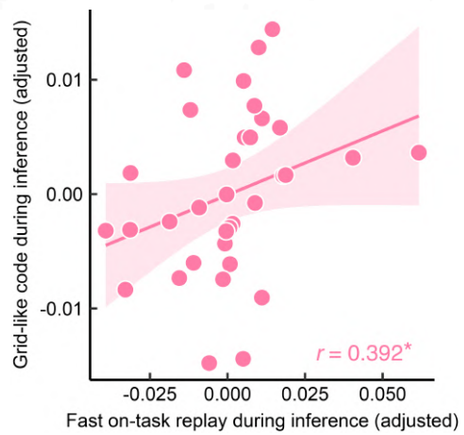

Grid signal correlates
-ve with slow replay

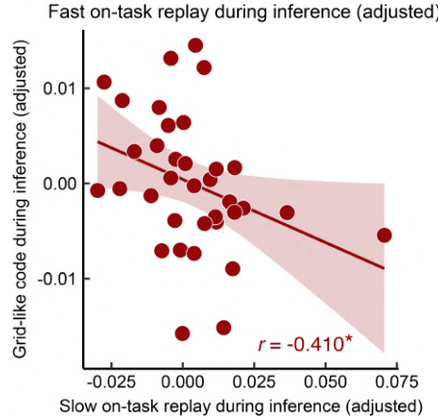

Jianxin Ou, Yunzhe Liu Biorxiv

But how did they get the map?

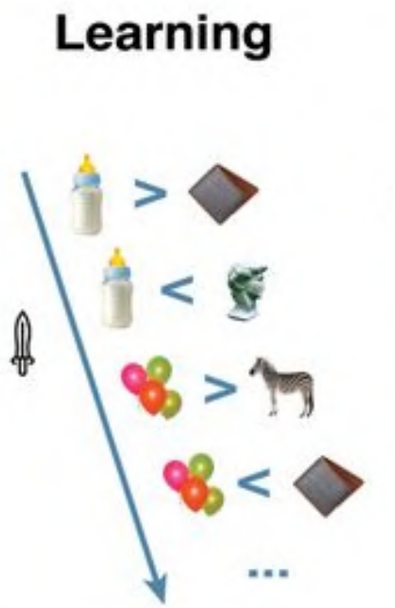

Rest

?

Fast replay builds map offline!

They were only ever shown the 1D hierarchies. They had to build the map

Jianxin Ou, Yunzhe Liu Biorxiv

Replay during Rest

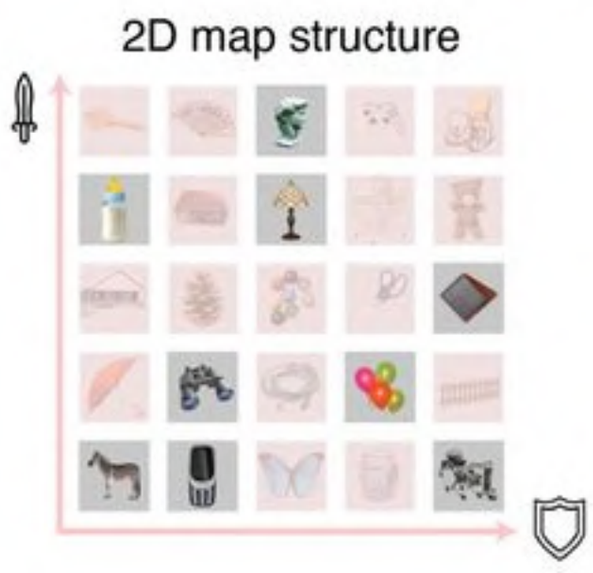

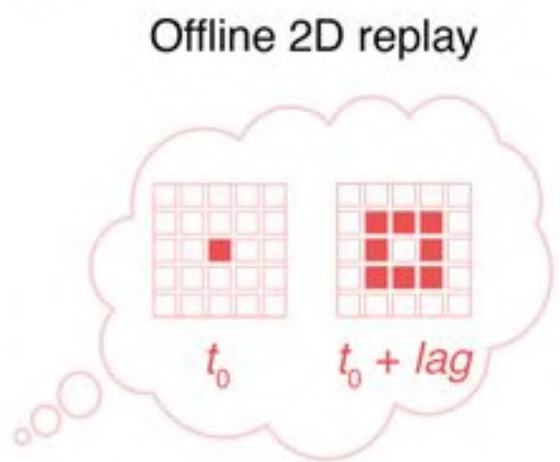

Like the mouse shortcuts!

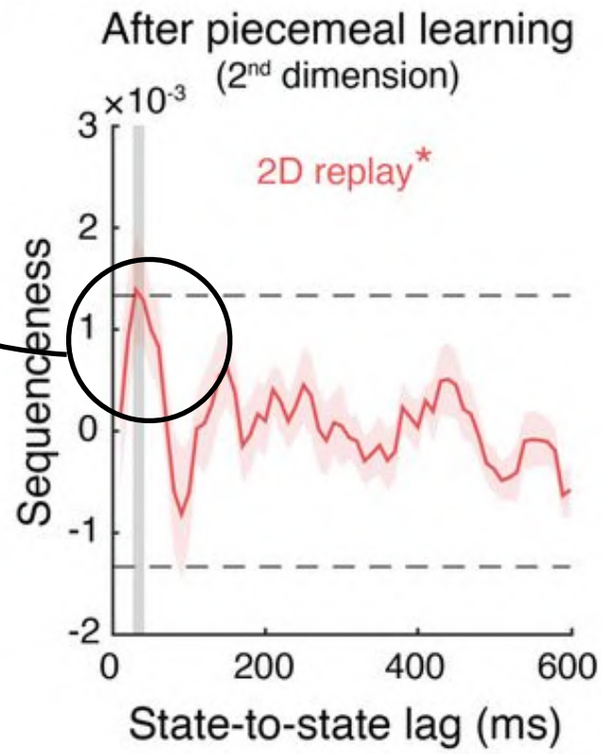

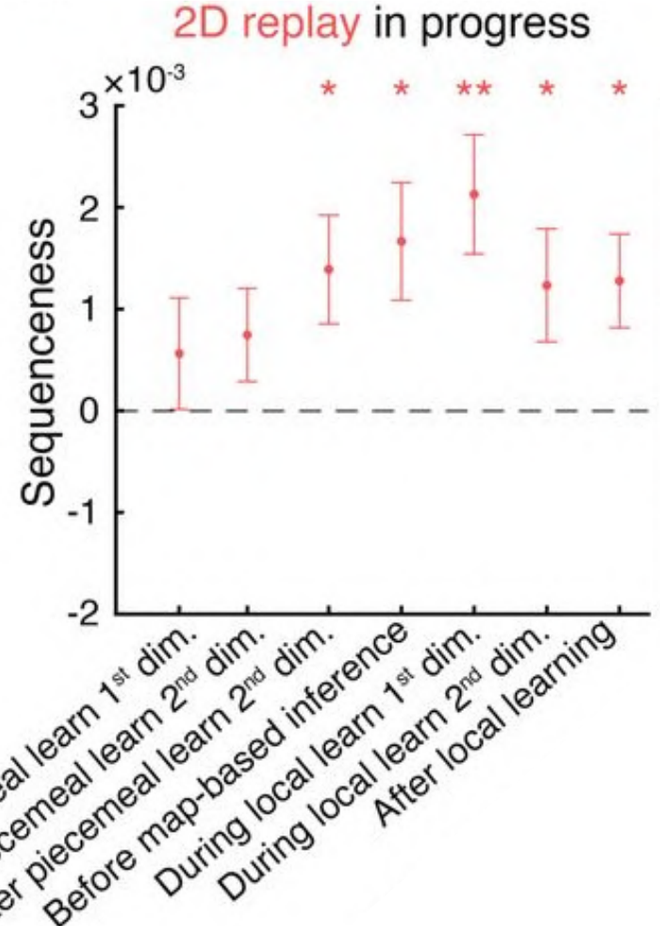

If you build the map offline, you can use it online

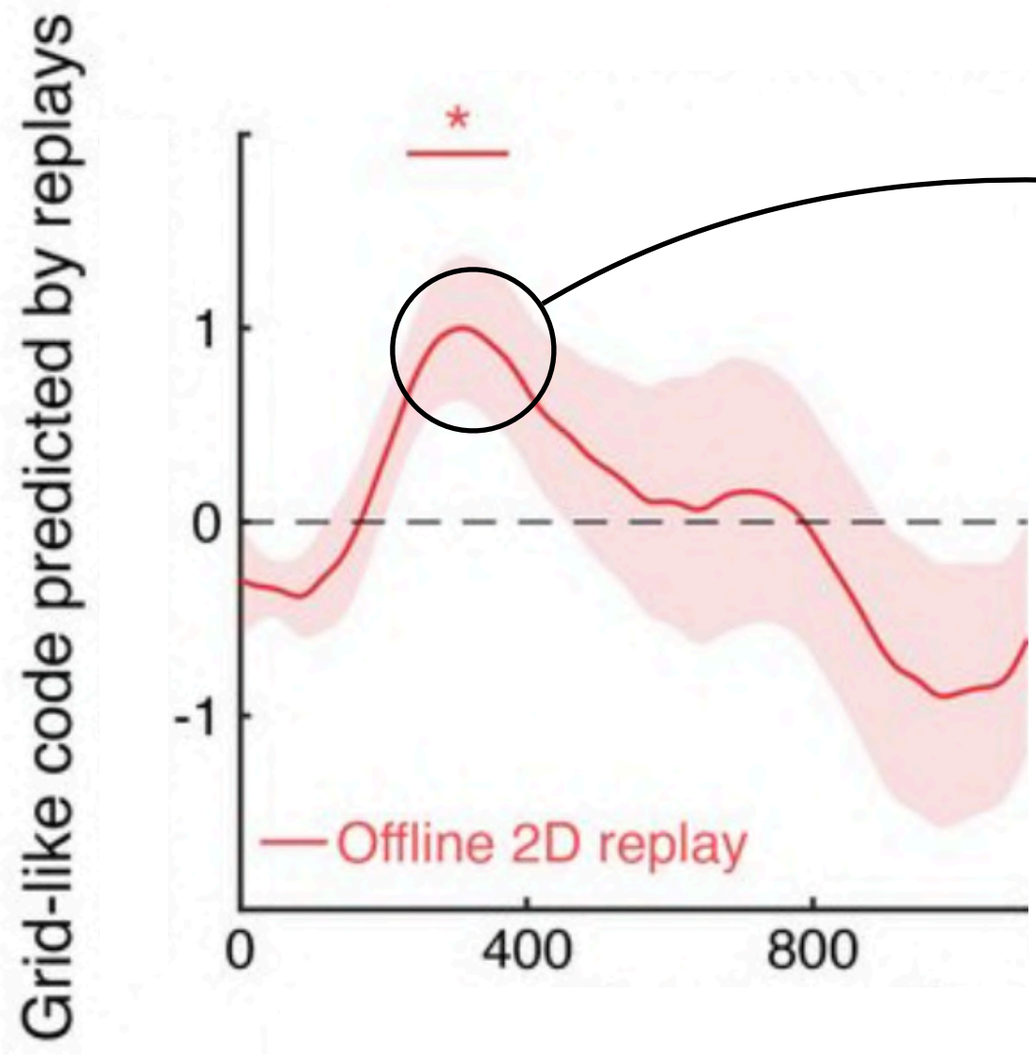

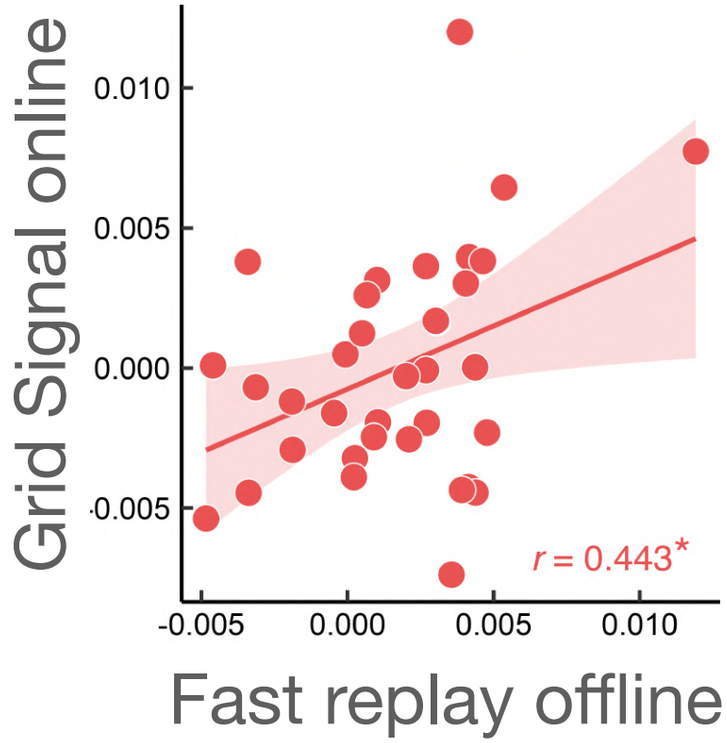

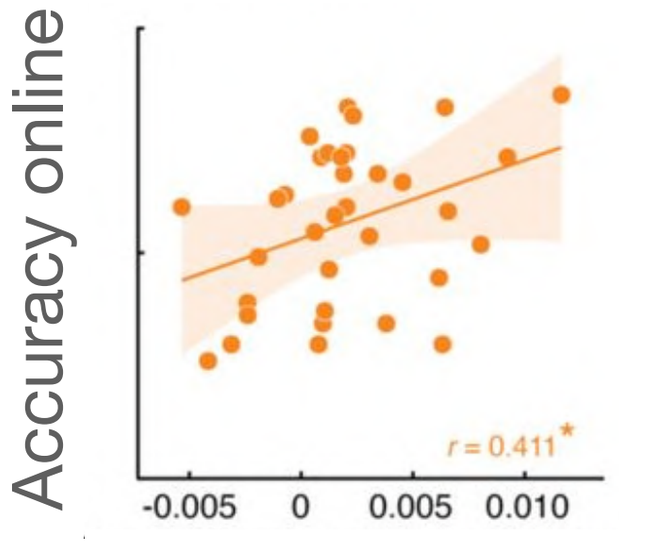

The more offline replay you have the stronger your online grid-code and the better you inferences!

Get some rest!!!

(An aside - grid like code predicts IQ in development - take that skinner!!)

y=-4

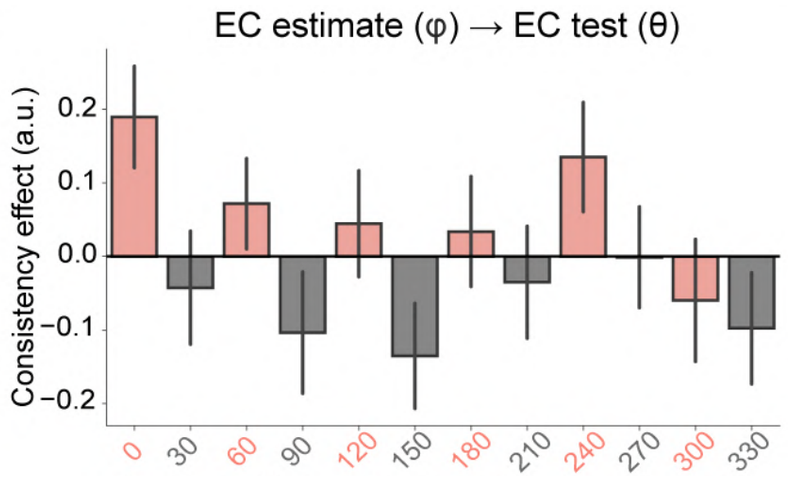

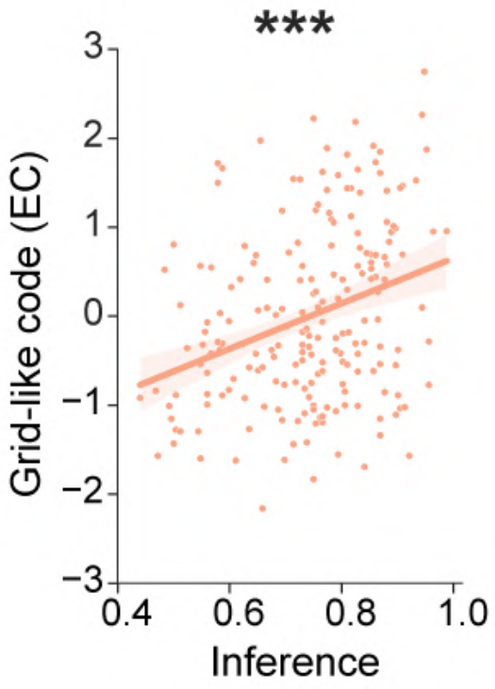

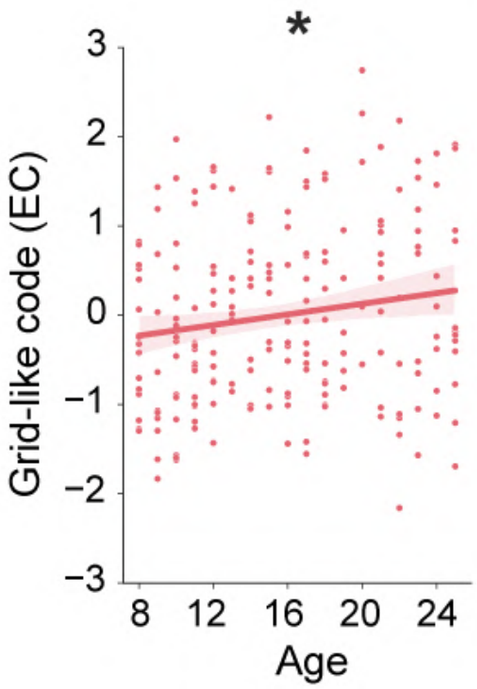

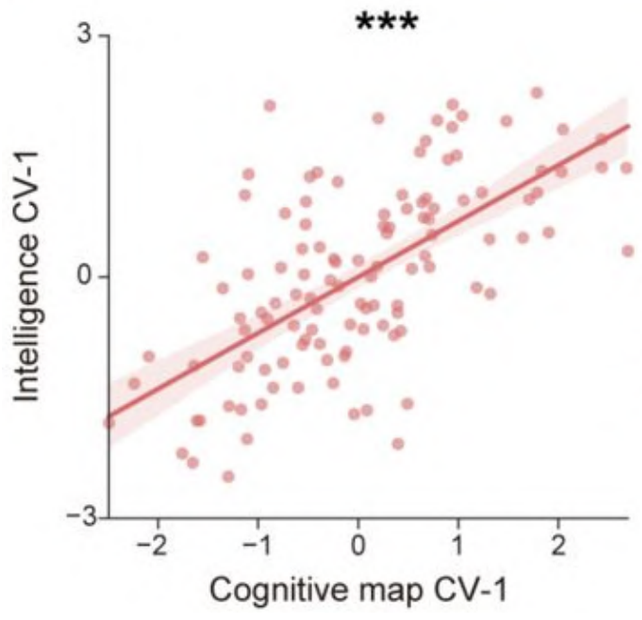

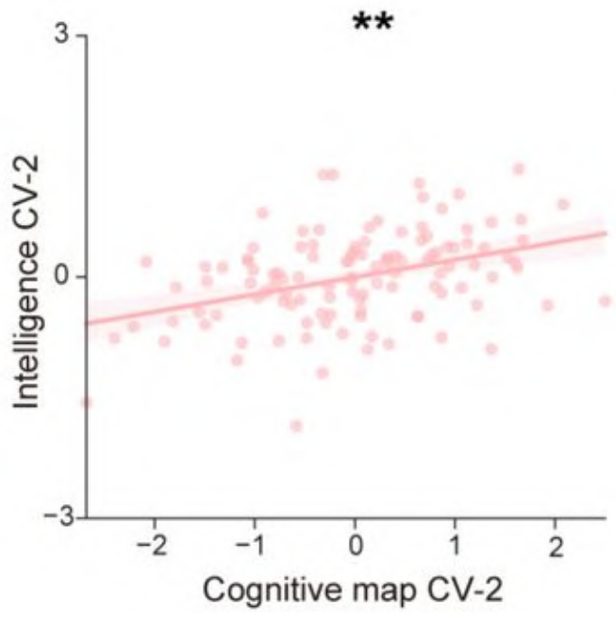

Yukun Qu, Yunzhe Liu biorxiv

What this talk is going to say

- Replay is not just replaying the past. It is doing inference about possible futures.

- It does so by aligning experience to our structural knowledge

- It is always making these inferences in the background, making you faster at thinking when the time comes.

- If there is time, some data about replay building new structural knowledge.

A different kind of world model in frontal cortex

Is there any time left?

El Gaby et al 2024

A different kind of world model in frontal cortex

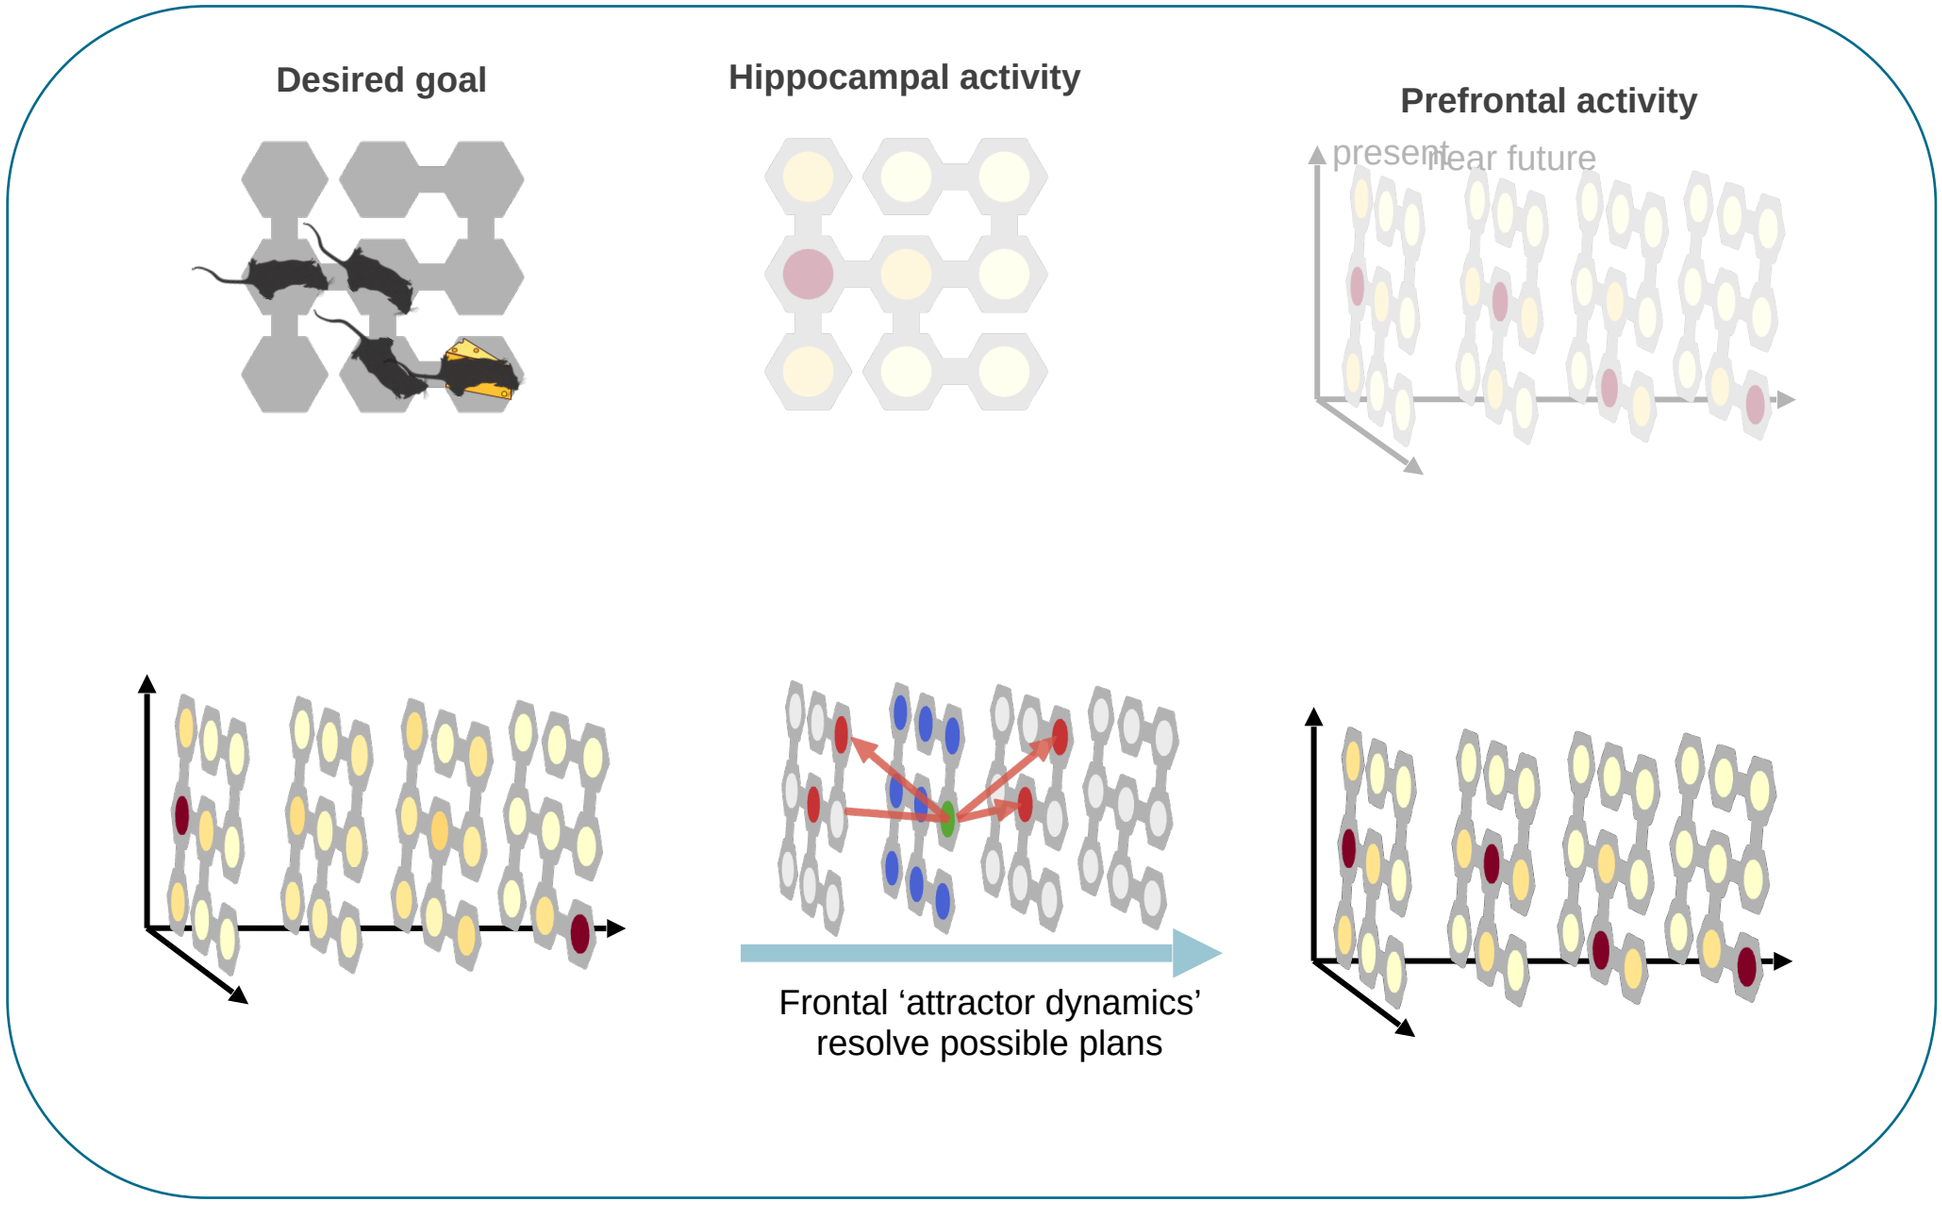

Now an instantaneous representation is an entire plan!

El Gaby et al 2024

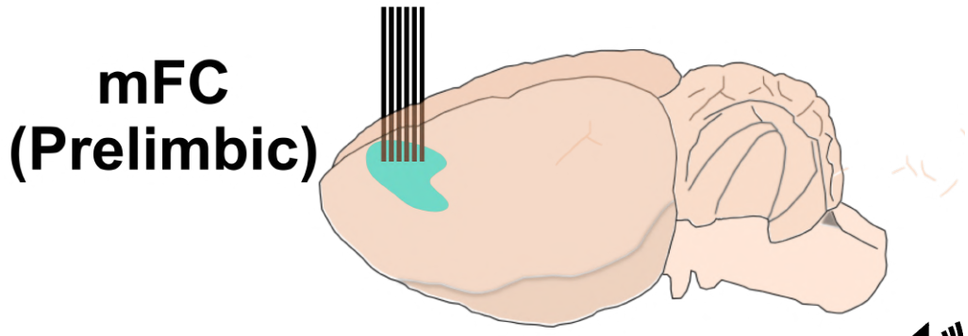

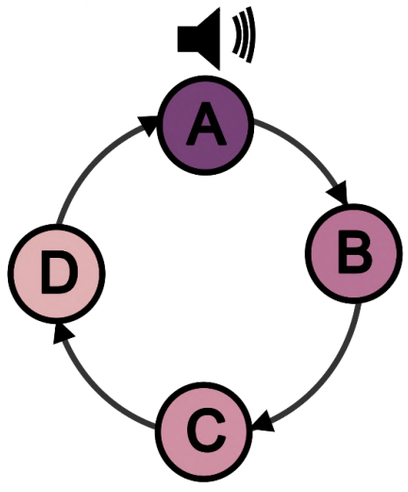

Abstract task structure

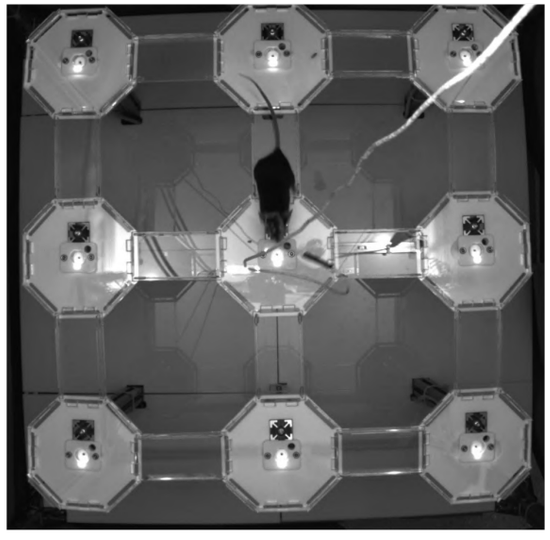

Spatial maze structure

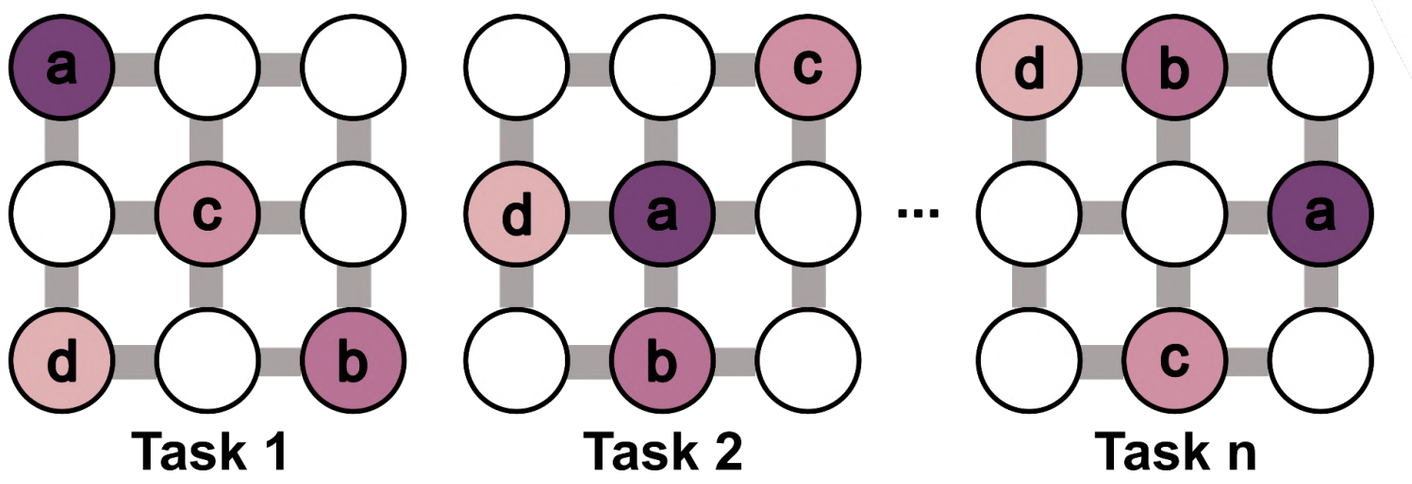

El Gaby et al 2024

Is knowledge abstracted?

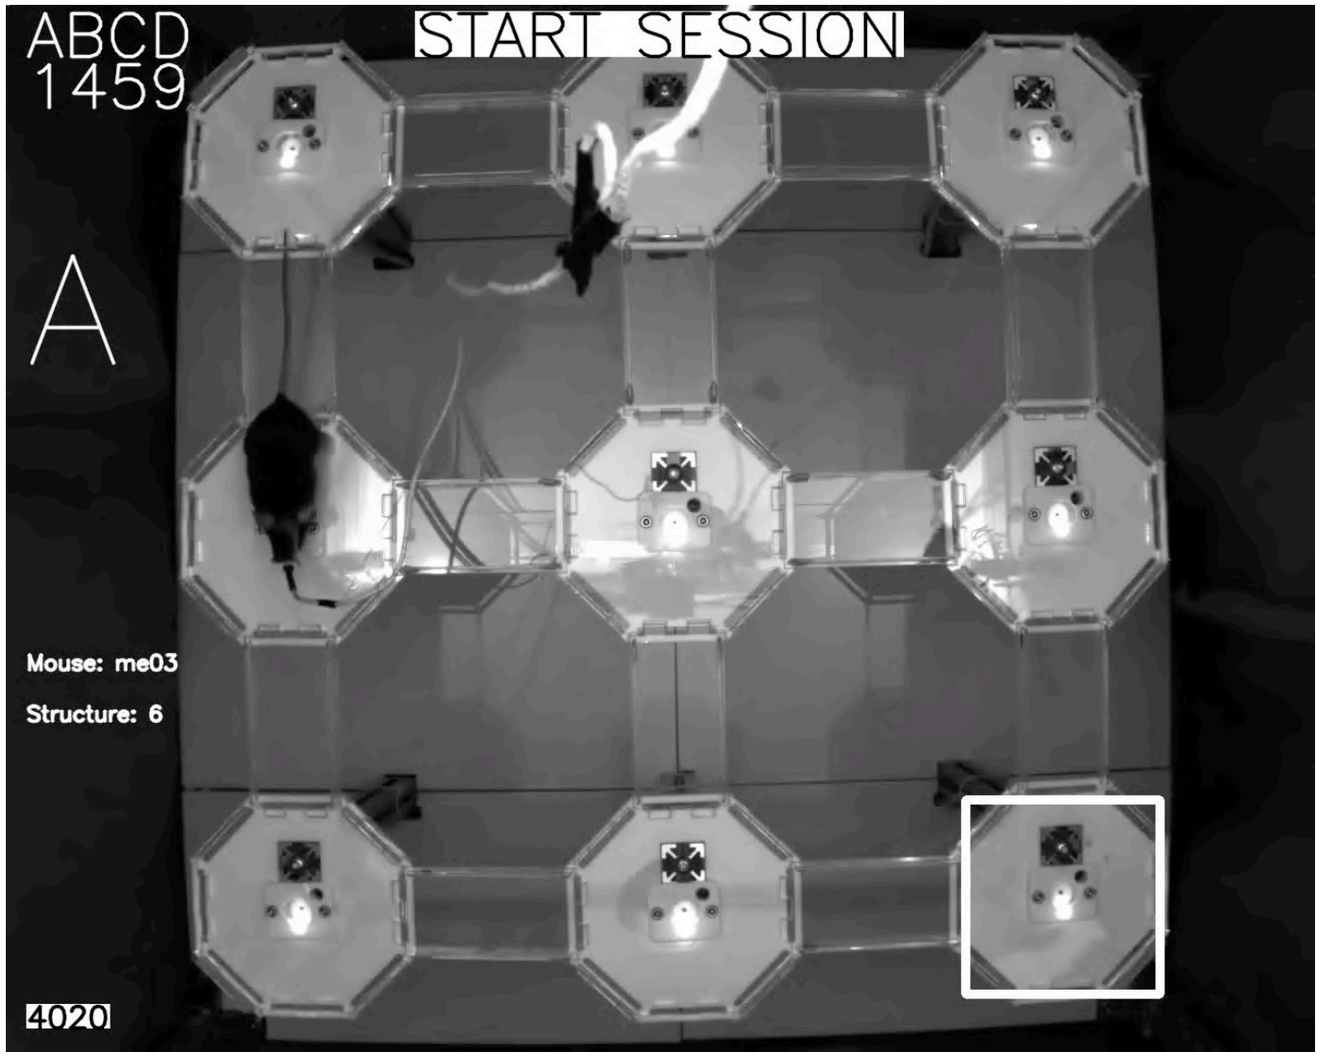

D		
	C	
	B	A

He knows something like this
“Find 4 things then go back to the first”

Consistent place fields for where you **will be**

Now

2 rewards from now

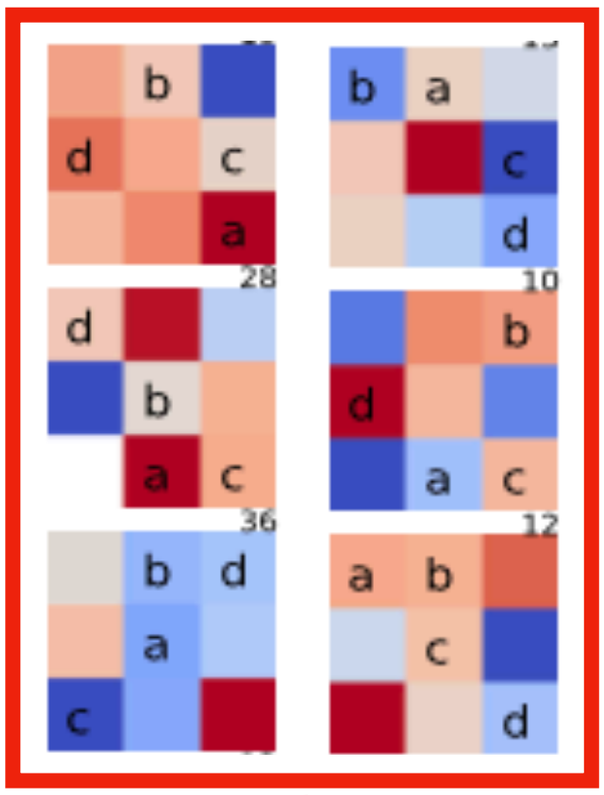

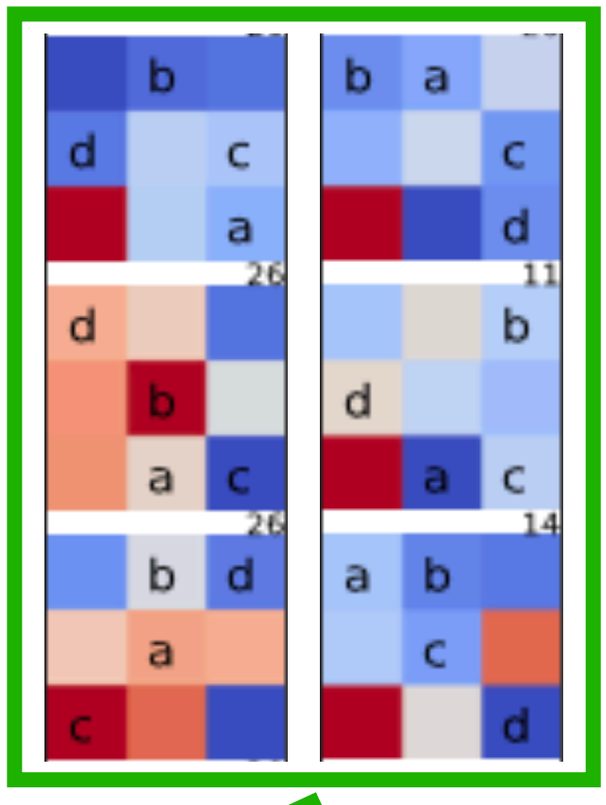

Predict lagged-spatial-map
in held out task.

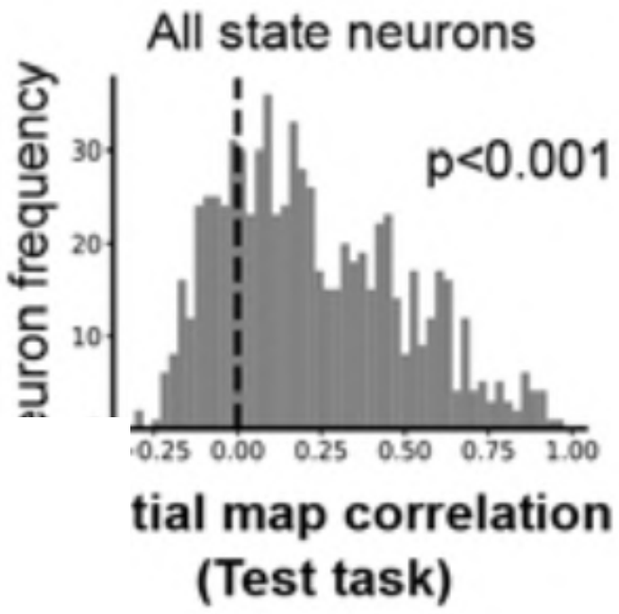

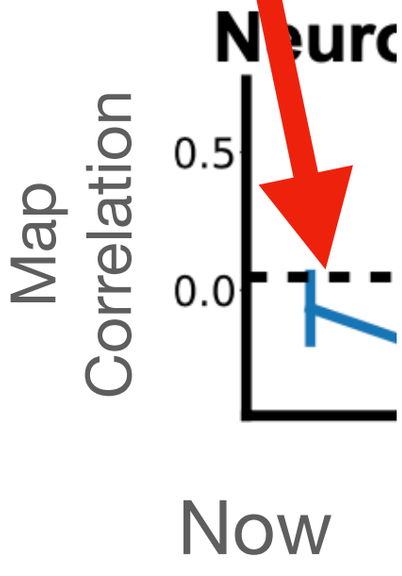

Some local cells but lots of cells prefer long lags

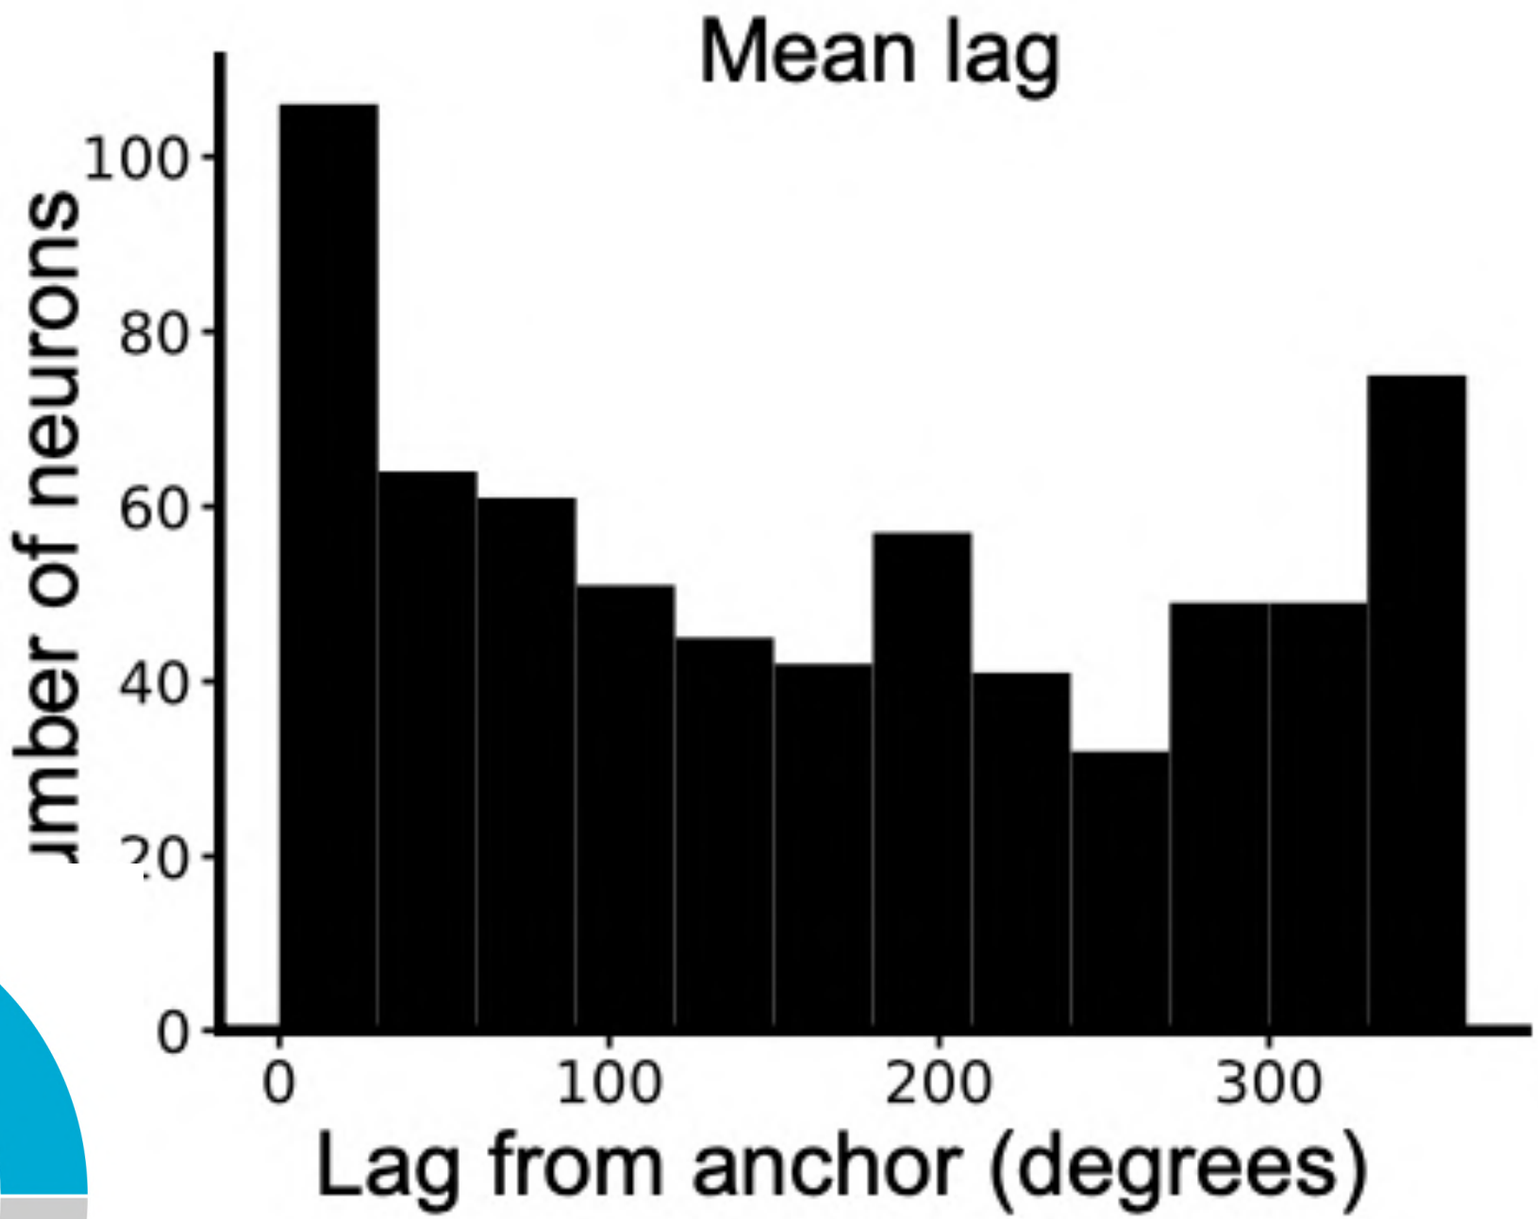

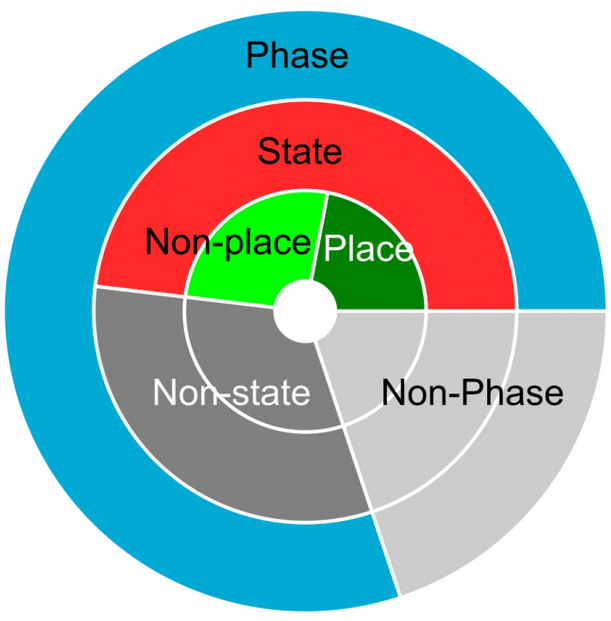

Together these cells have learnt the task structure

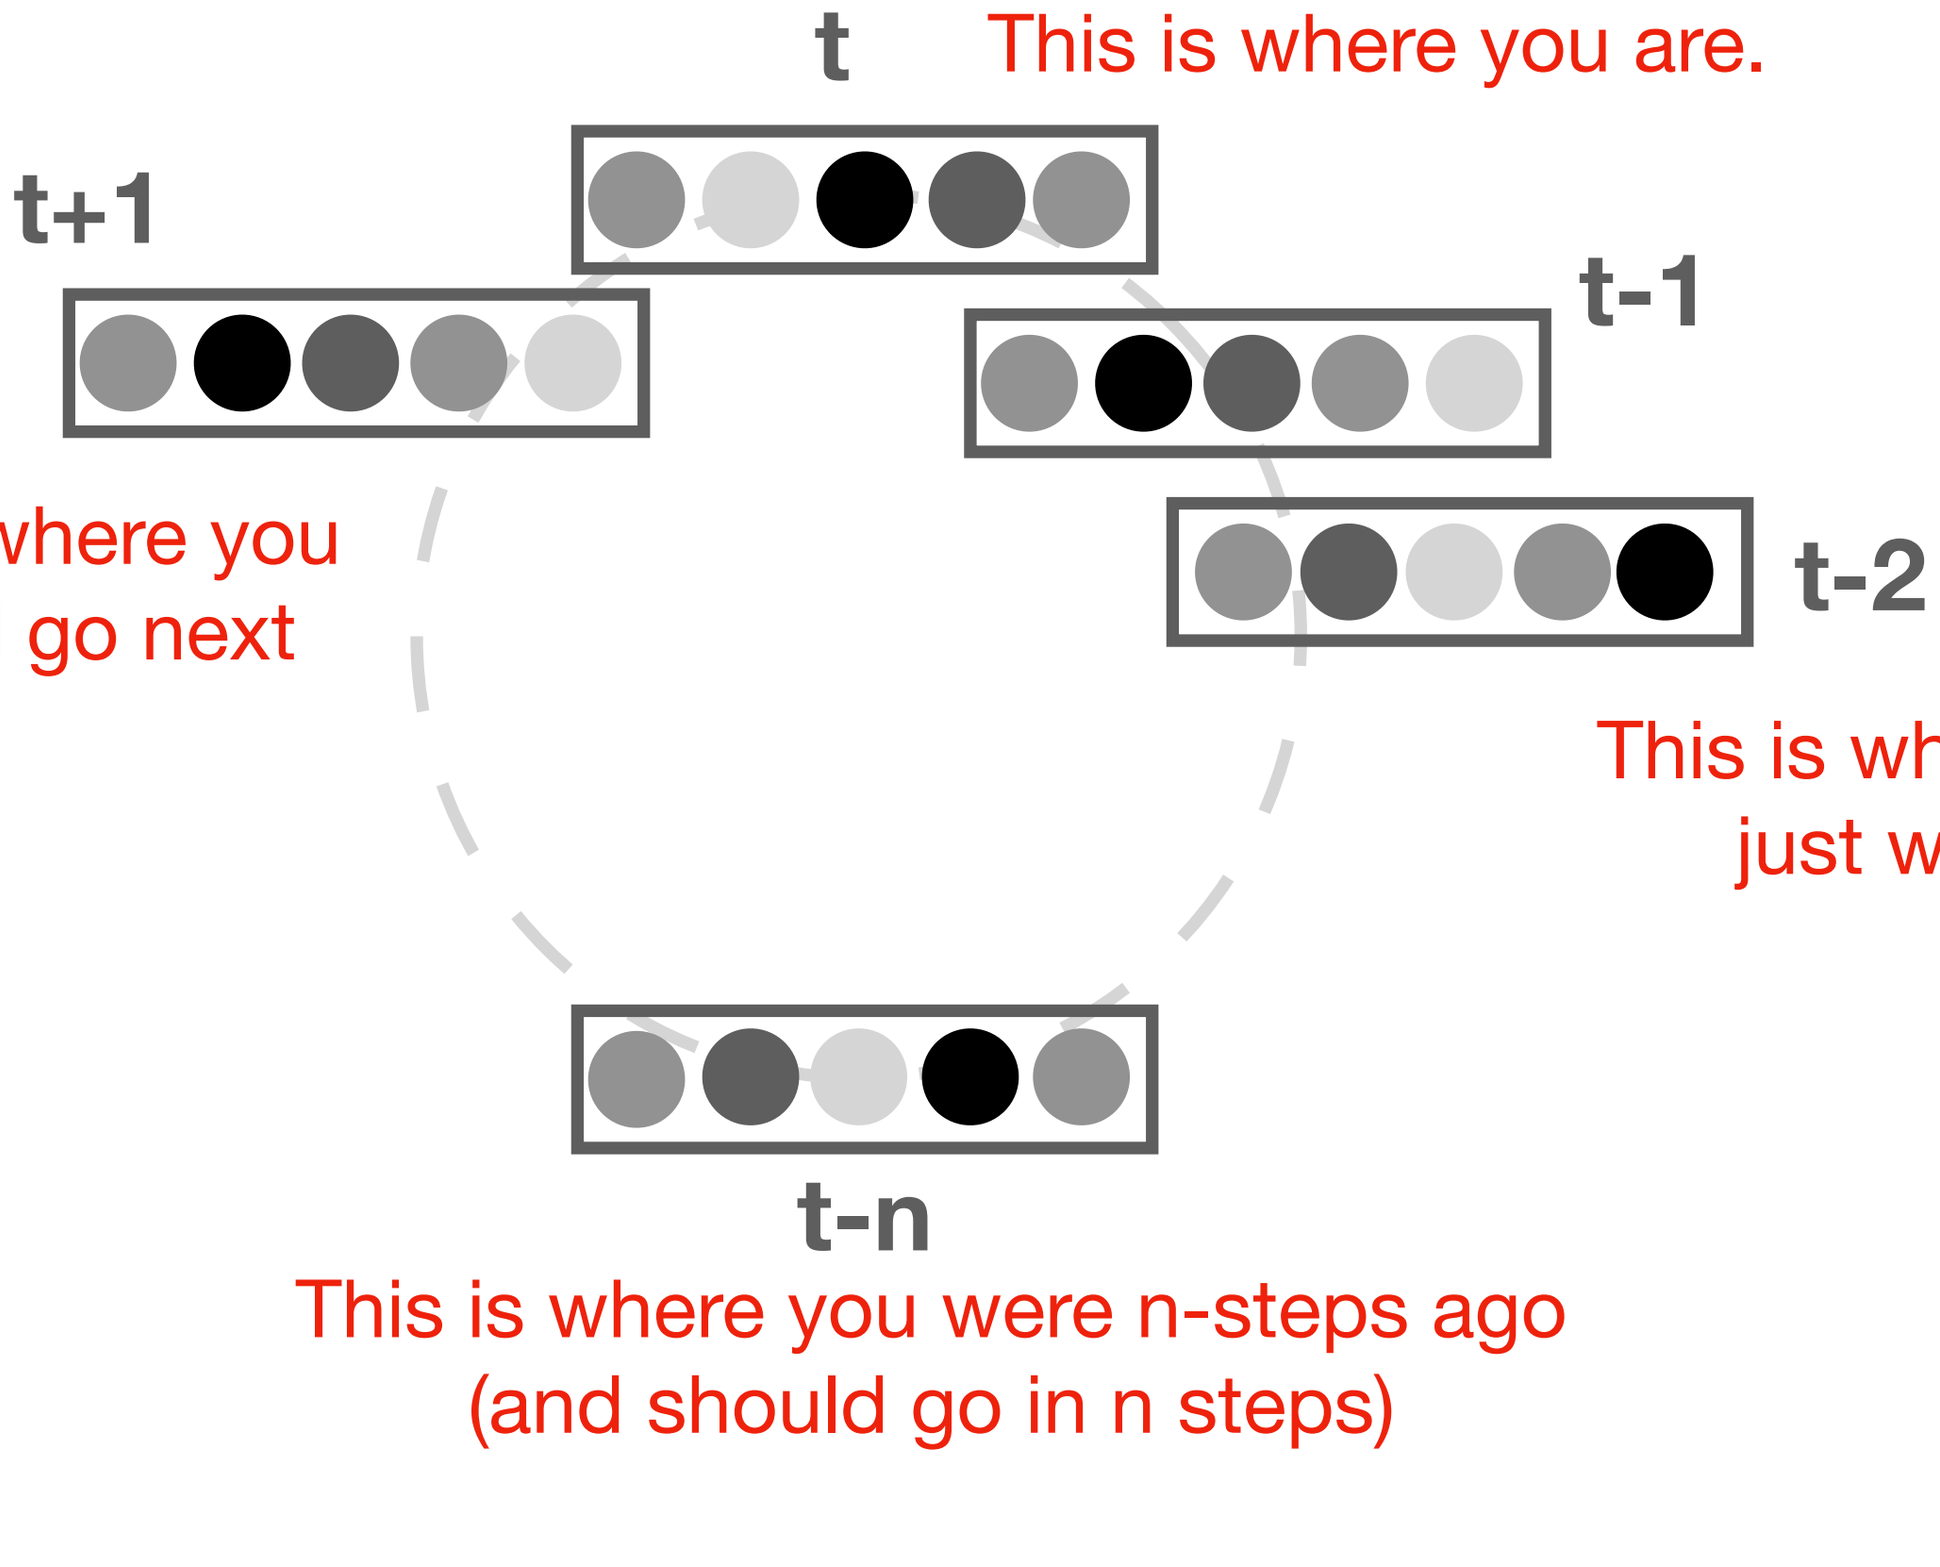

You could just pass the whole representation around the loop.

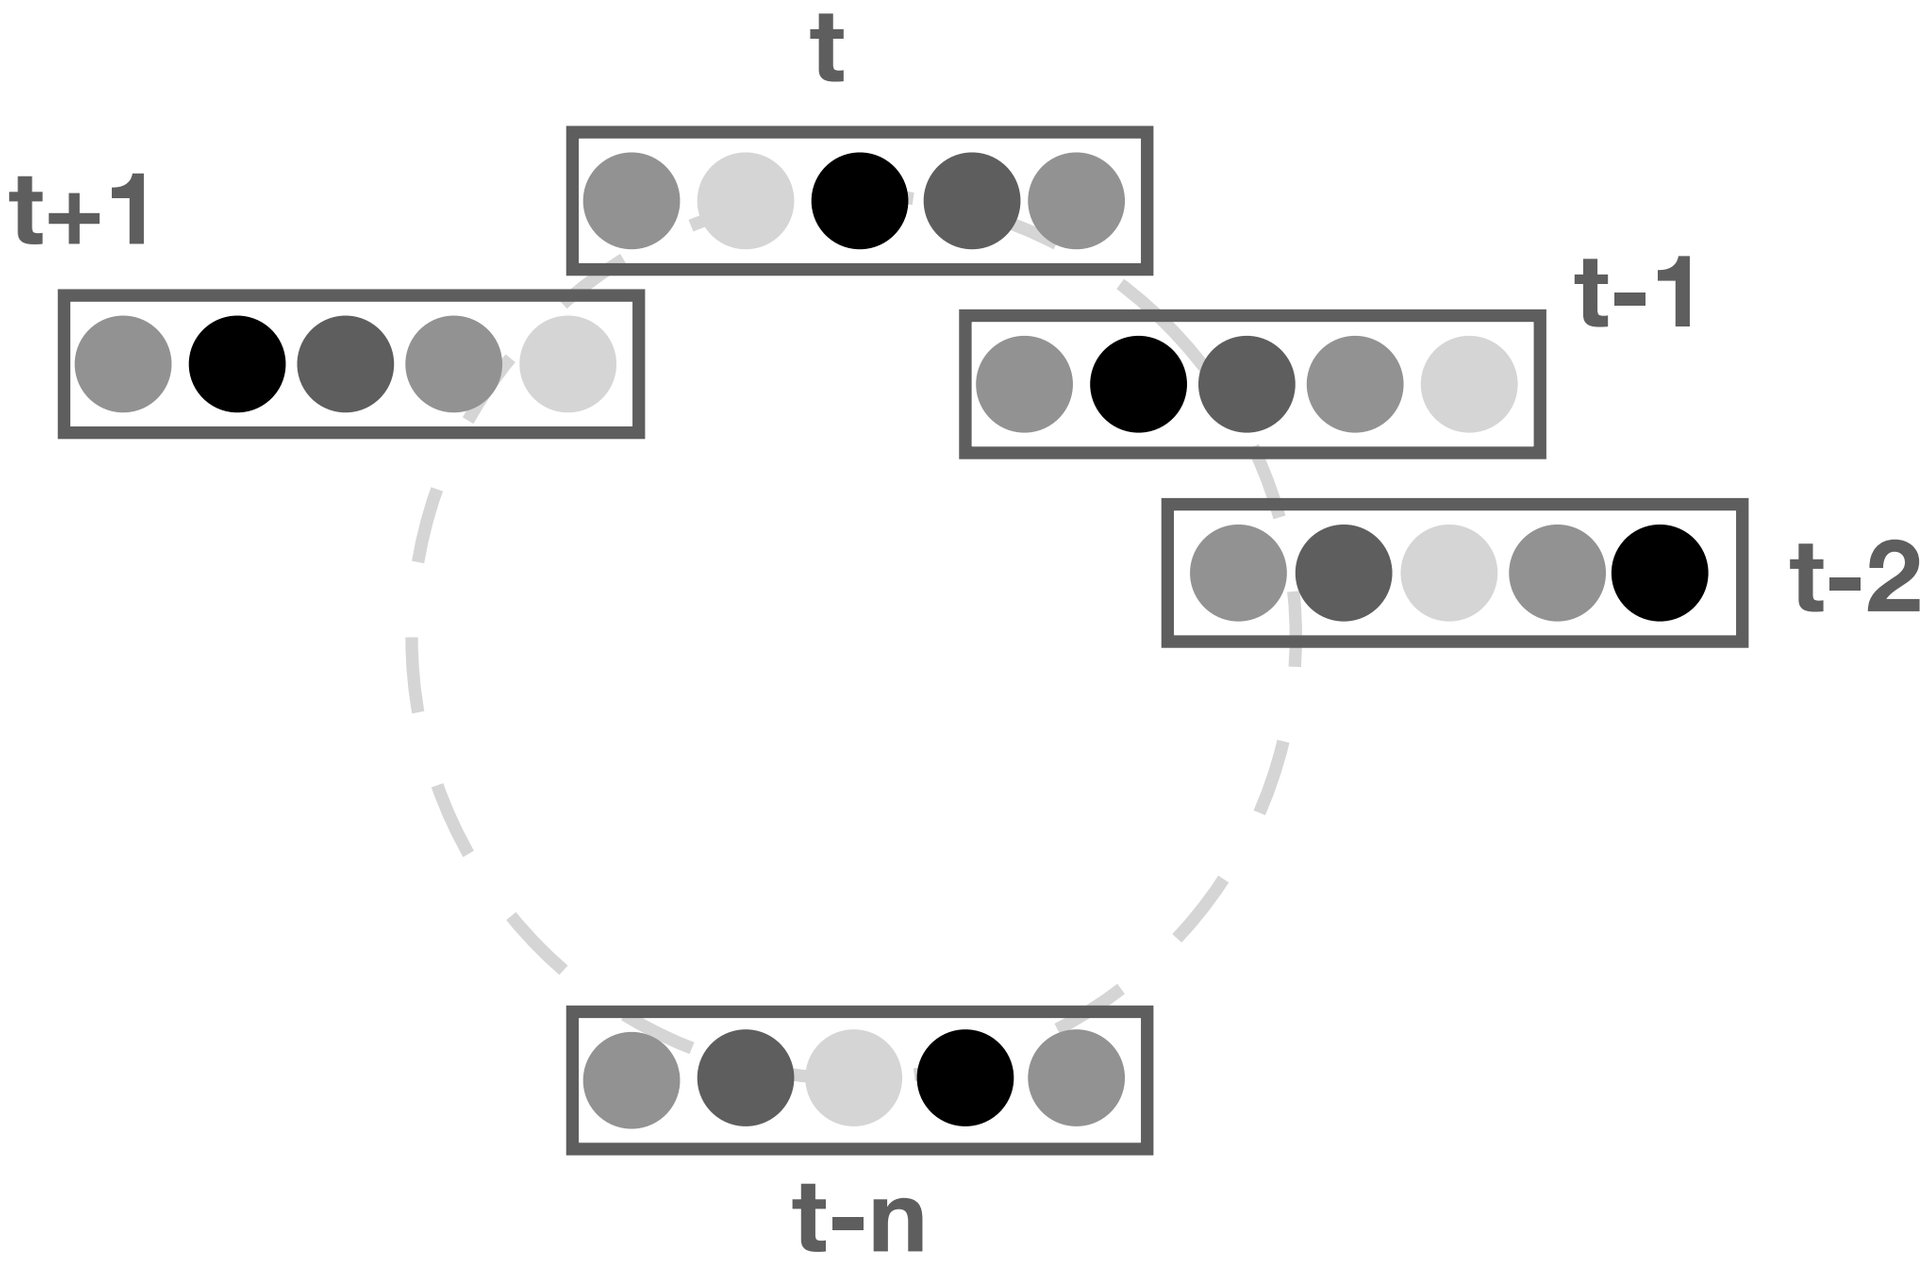

All these representations need to be in the brain at the same time

This is called a Drum Sequencer. Or a music box

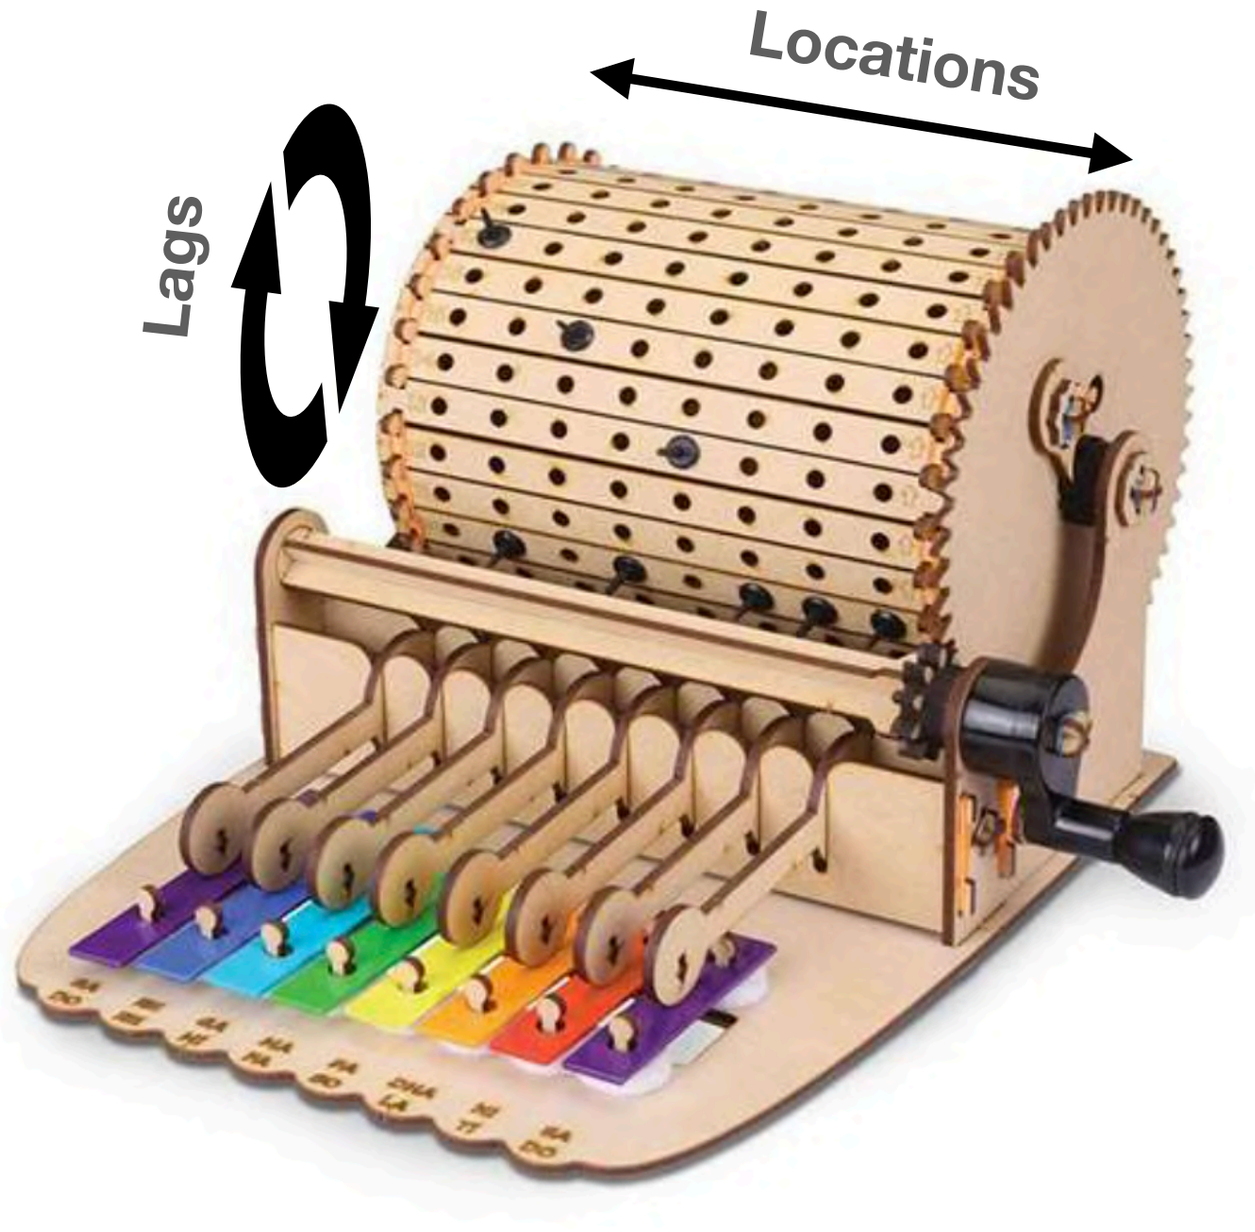

It is programmable to new sequences by just moving the activity bumps

Task transitions can be seen in sleep

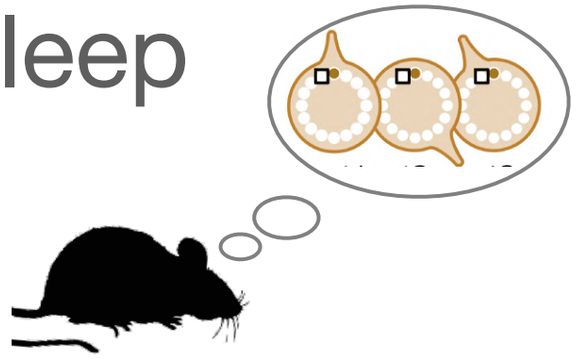

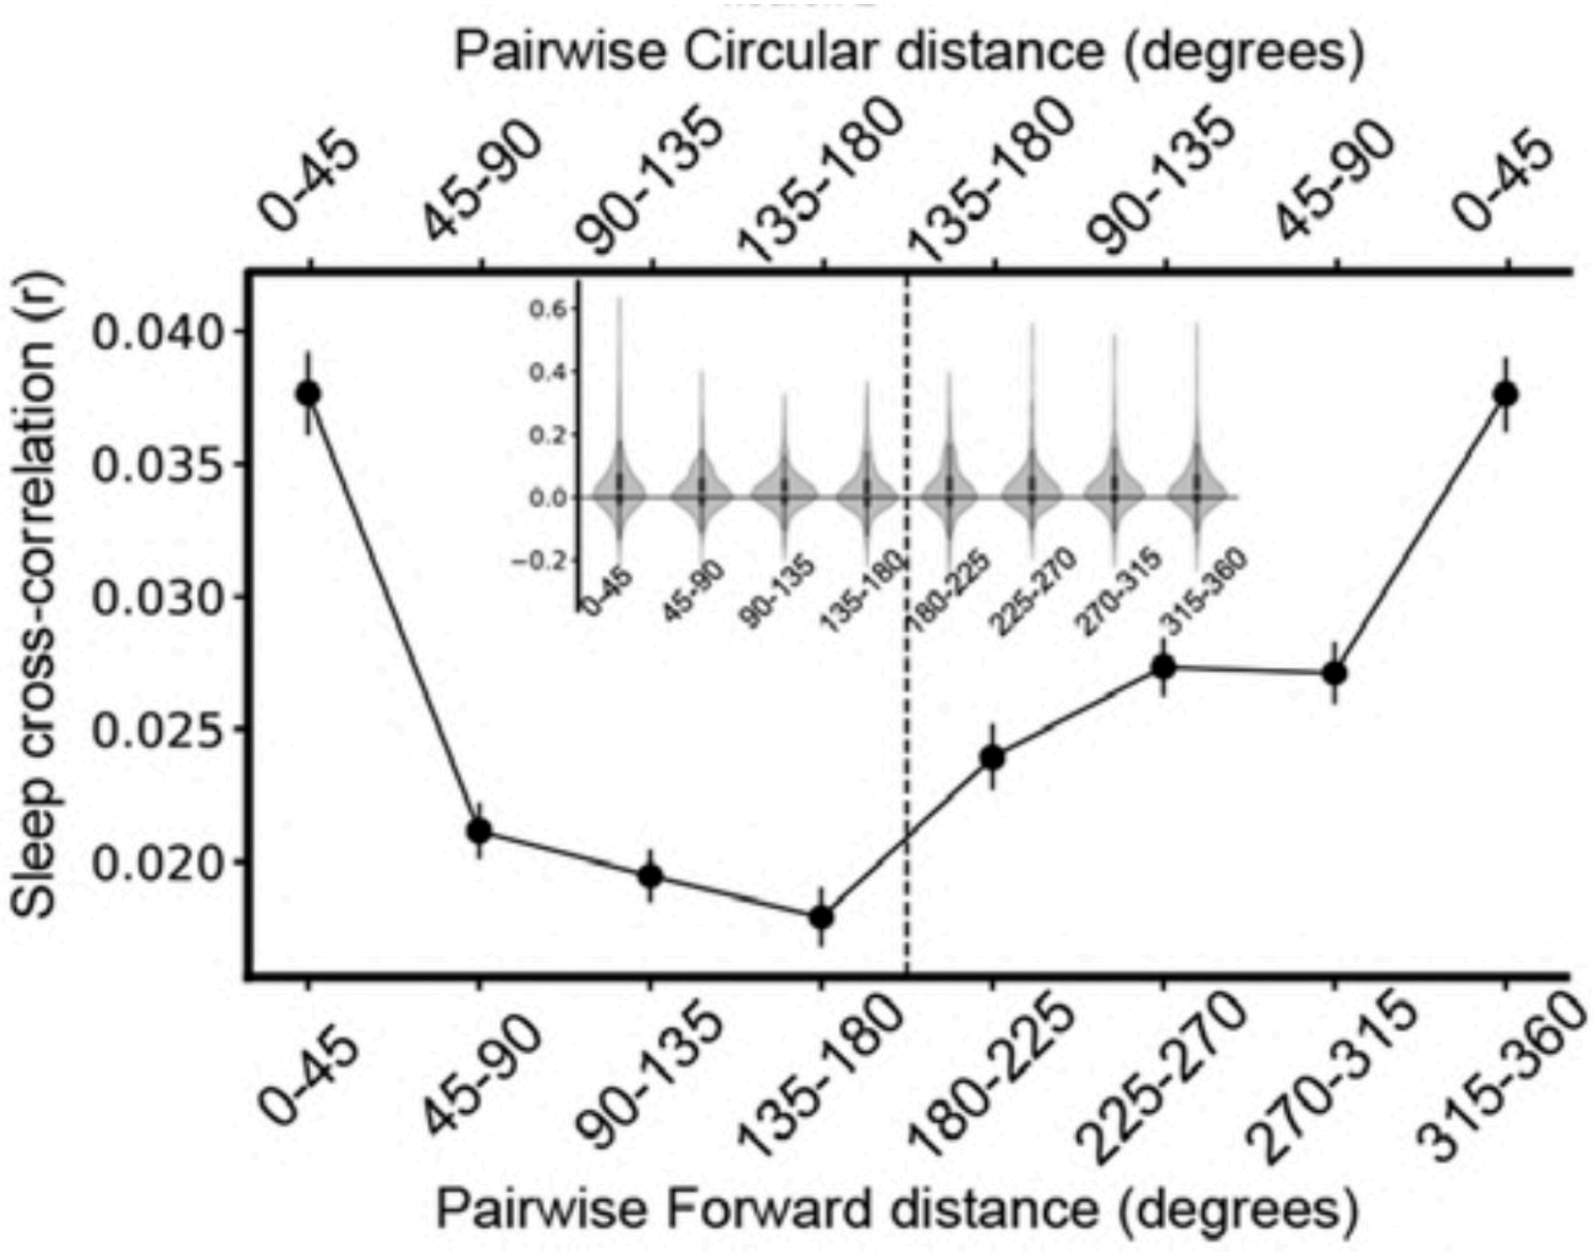

The closer you are on the ring the more you co-fire during sleep.

Conclusions

- Get loads of rest!!!

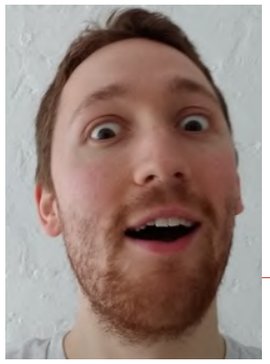
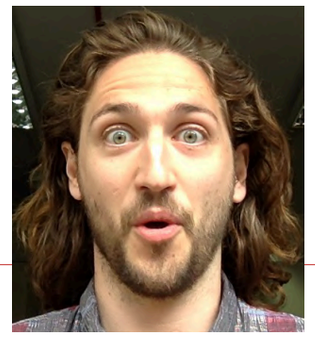




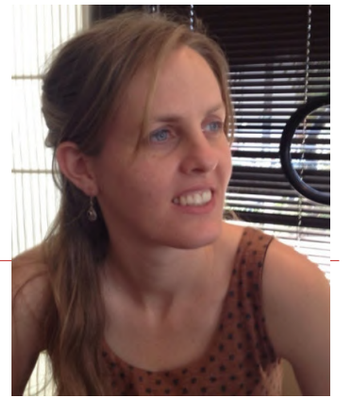

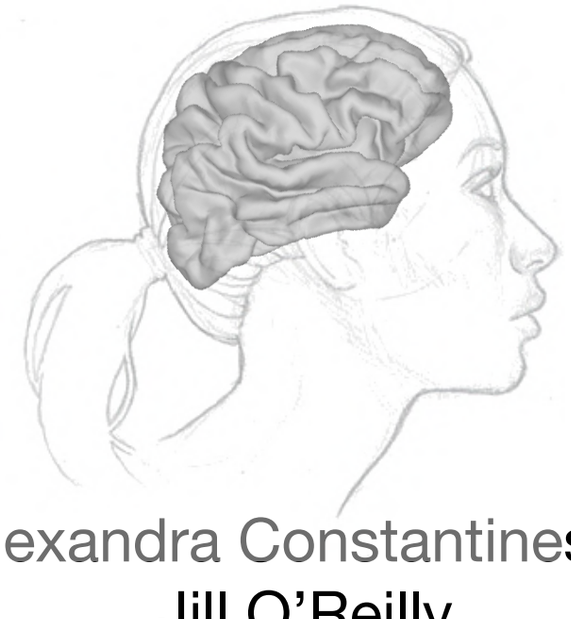


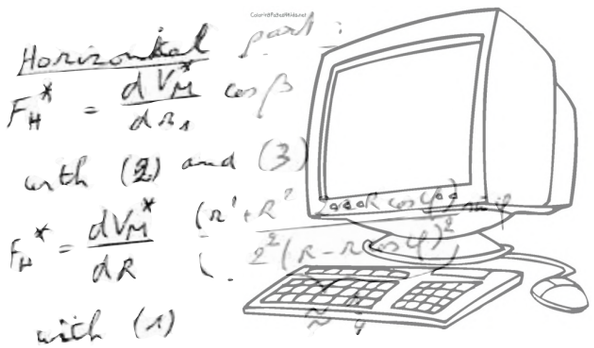

Alexandra Constantinescu
Jill O'Reilly
Mona Garvert
Helen Barron
Shirley Mark
Anna Shpektor
Philipp Schwartenbeck
Alon Baram
Yunzhe Liu

Steve Kennerley
James Butler

Mark Walton
Thomas Akam
Veronika Samborska
Mohamady el Gaby
Yves Weissenberger
Adam Harris

James Whittington
Jacob Bakermans
Tim Muller
Zeb Kurth Nelson
Jo Warren