Where do states come from?
Learning task representations

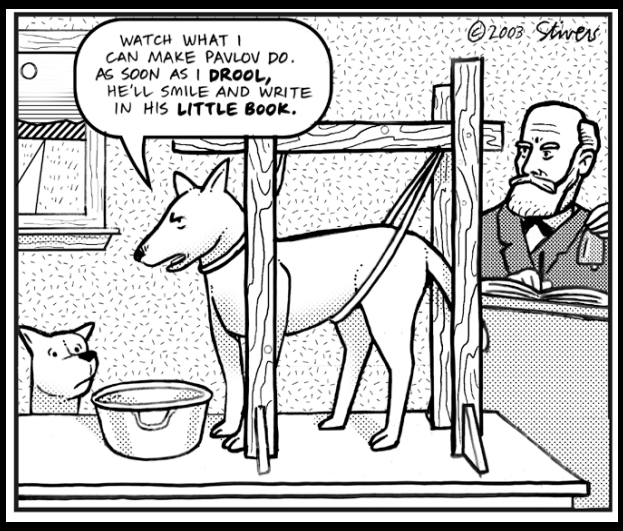

PSY/NEU338 From Animal Learning to Changing People's Minds

No two events are ever the same…

And still, we learn.

all of learning is generalization

the question is: where to set the boundaries?

This is not straightforward even in simple conditioning

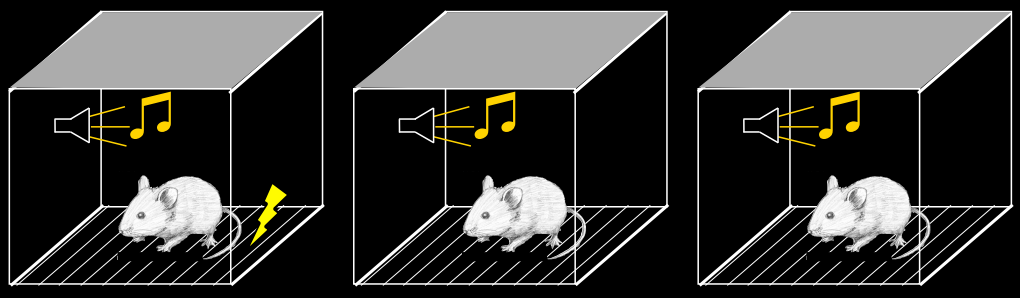

acquisition extinction test

Extinction ≠ Unlearning (rather, new learning, but why?)

Storsve, McNally & Richardson, 2012

Extinction is fragile due to…
spontaneous recovery (over time),
renewal (due to different context), or
reinstatement (due to reminder US) of fear

Learning causal structure

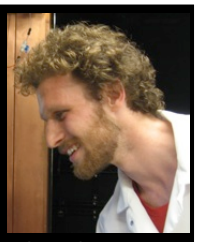

Sam Gershman

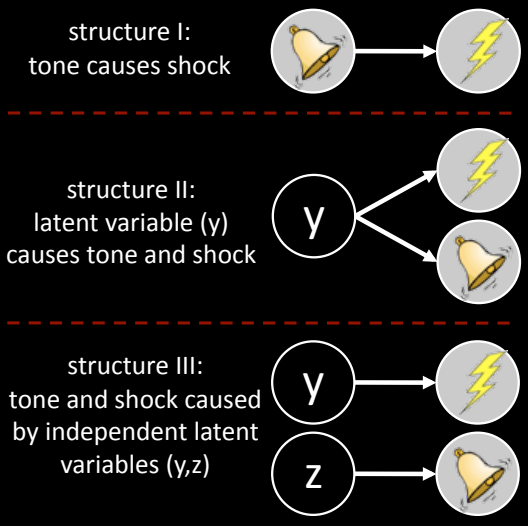

Bayesian inference with an infinite capacity prior
(also called Chinese Restaurant Process Prior or Infinite Mixture Model)

Goal: Embue animals with a *generative model* that is flexible enough to capture the complex structure of the environment, but constrained enough to allow learning

1. Each trial is caused by a single latent cause
2. All else equal, a (recently) prolific latent cause is more likely to cause another trial
3. The number of possible latent causes is unbounded.
4. Likelihood: each latent cause tends to produce similar trials

$$P(C_t = k) = \begin{cases} \frac{N_k}{t+\alpha} & \text{if } k \text{ is old cluster} \\ \frac{\alpha}{t+\alpha} & \text{if } k \text{ is new cluster} \end{cases}$$

Gershman, Blei & Niv 2010

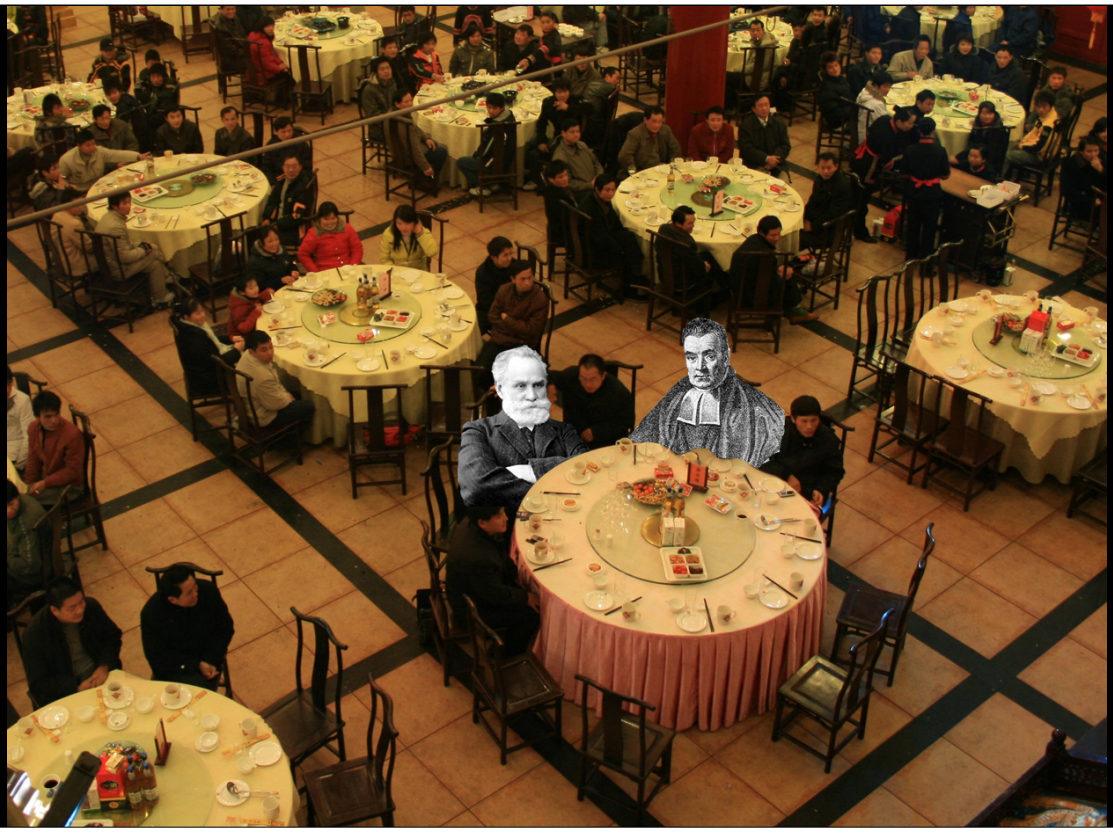

The causal structure of an extinction task

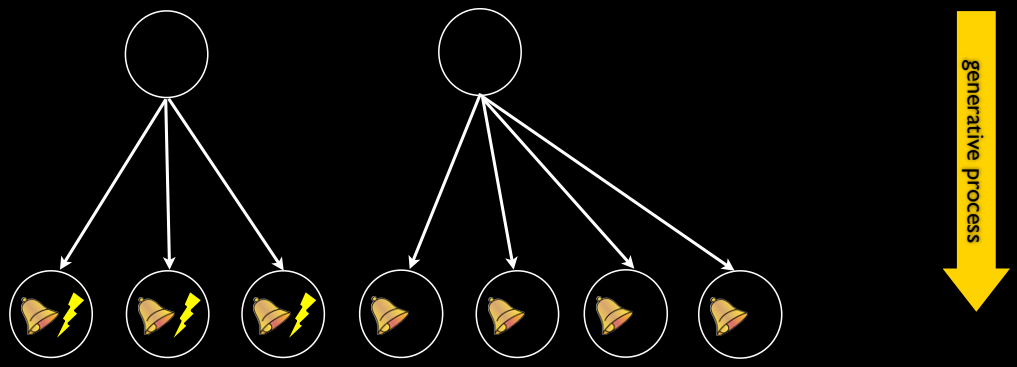

Inferring the causal structure given observations

Inferring the causal structure given observations

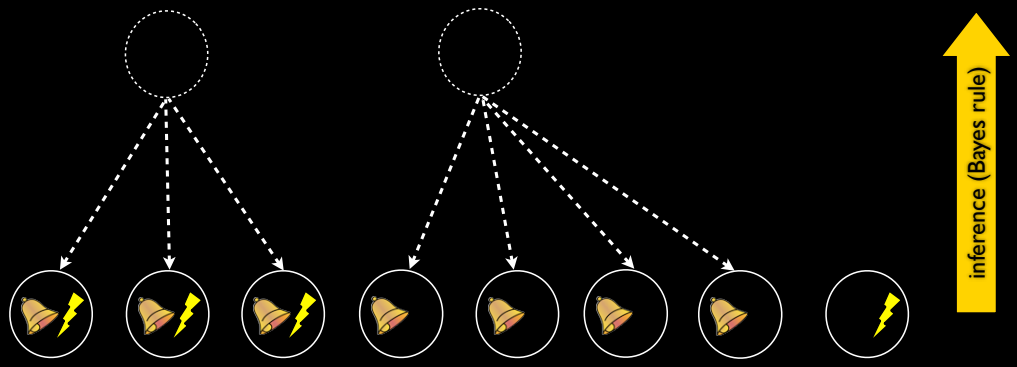

$$p(\text{cause}|\text{data}) \propto p(\text{data}|\text{cause})p(\text{cause})$$

Inferring the causal structure given observations

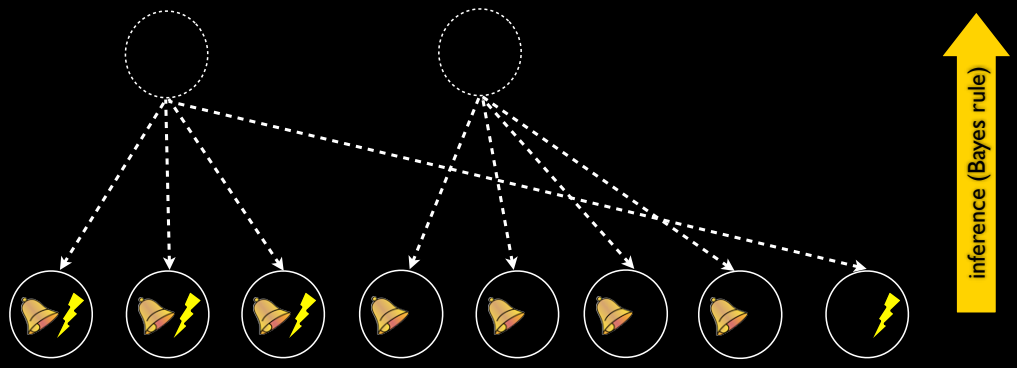

$$p(\text{cause}|\text{data}) \propto p(\text{data}|\text{cause})p(\text{cause})$$

Conditioning as clustering

Bits of experience impinge on us all the time

We cluster experiences together into memory traces (latent causes, states) based on similarity

Learning happens within a cluster, not across cluster boundaries

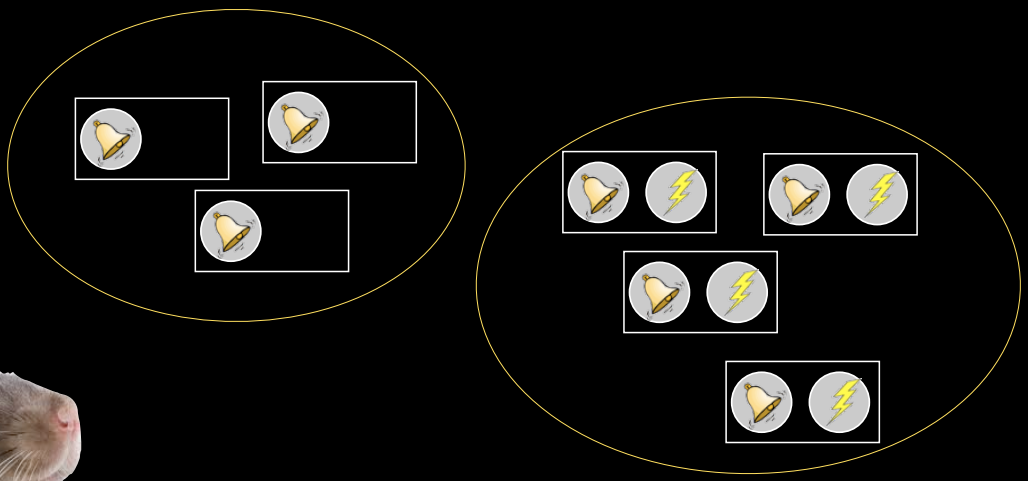

Equivalent to multiple associations

reinforcement learning
(RW/TD) model

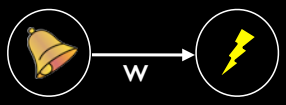

latent-cause modulated
RL model

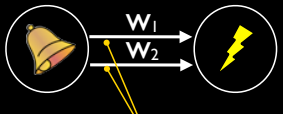

cf. Redish et al. (2007)

Equivalent to "compound cues"

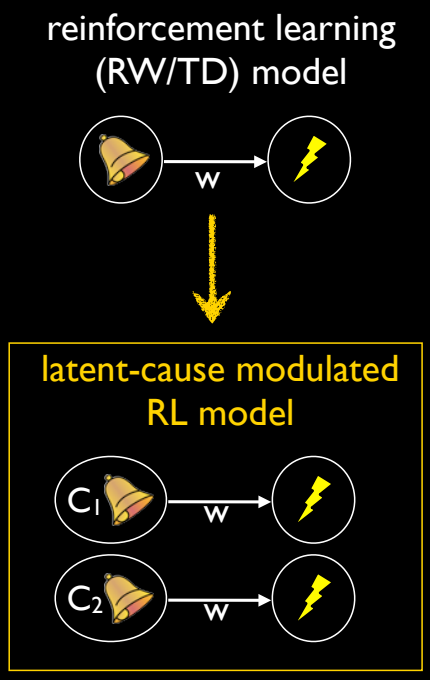

Latent causes as "states" in reinforcement learning

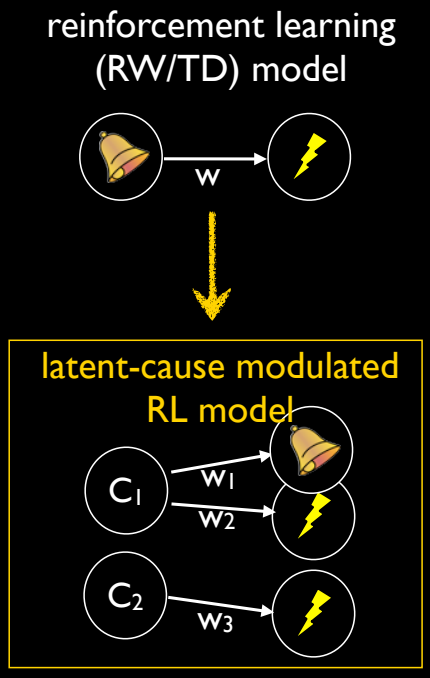

A note about the concept of *causes:*

- Actual causal structure (cause)
- Grouping through time (context)
- Clustering according to similarity (cluster)

all of learning is generalization

Latent cause inference (inferring the current state) sets the boundaries of generalization

Testing the model I :"Circles Task"

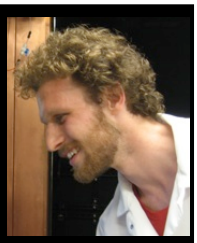

Sam Gershman

hypothesis: similarity determines the boundaries of generalization

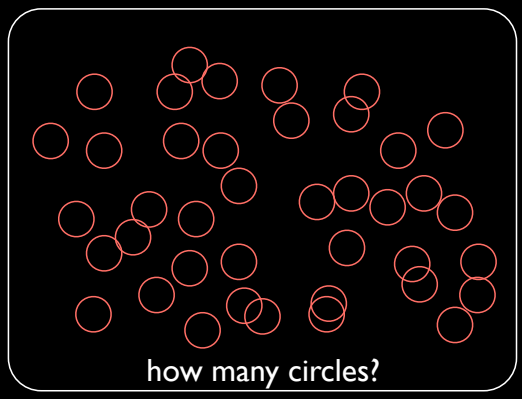

Gershman & Niv 2013

Testing the model I :"Circles Task"

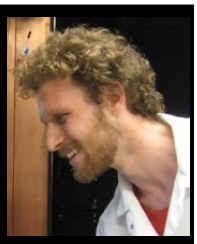

Sam Gershman

there were 39 circles

Gershman & Niv 2013

Testing the model 1 : "Circles Task"

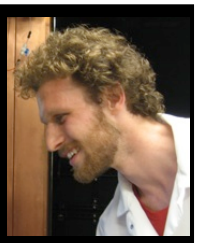

Sam Gershman

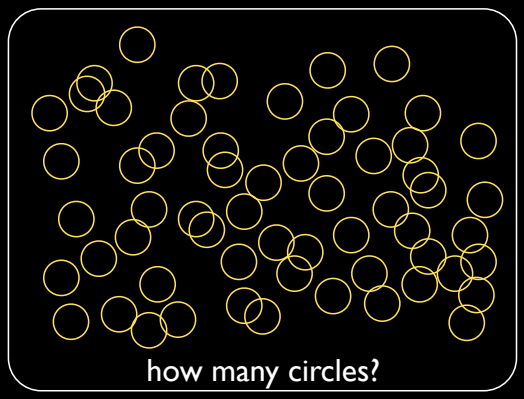

Gershman & Niv 2013

Testing the model 1 : "Circles Task"

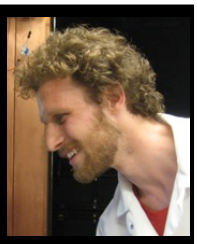

Sam Gershman

there were 61 circles

Gershman & Niv 2013

Testing the model I :"Circles Task"

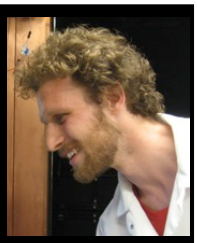

Sam Gershman

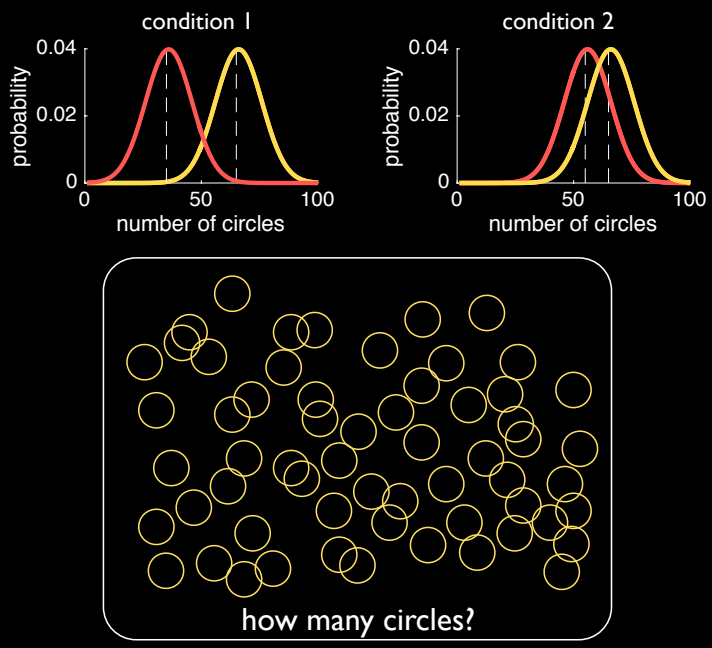

Gershman & Niv 2013

Similarity determines the boundaries of learning

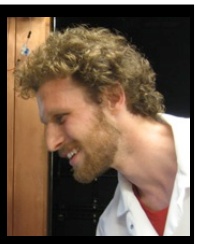

Sam Gershman

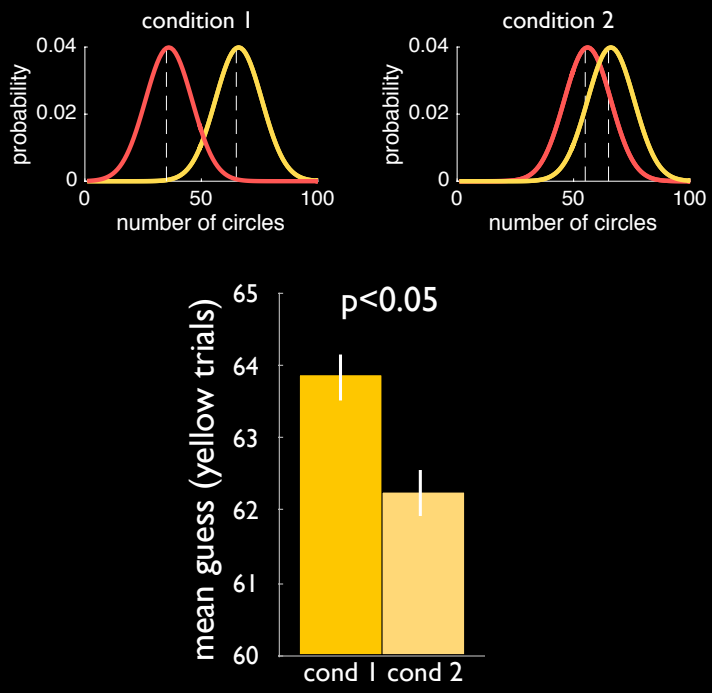

Gershman & Niv 2013

Inference about latent causes determines the boundaries of generalization

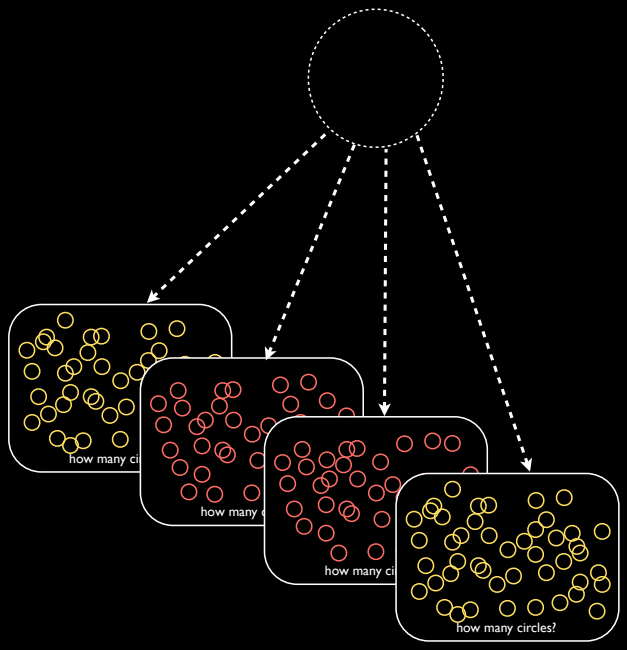

Gershman, Blei & Niv 2010

Inference about latent causes determines the boundaries of generalization

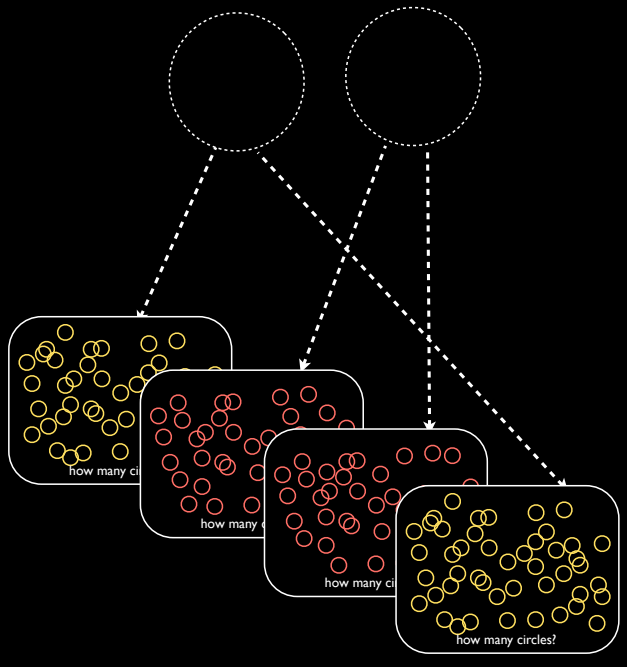

Gershman, Blei & Niv 2010

Conditioning as clustering

Bits of experience impinge on us all the time

We cluster experiences together into memory traces (latent causes, states) based on similarity

Learning happens within a cluster, not across cluster boundaries

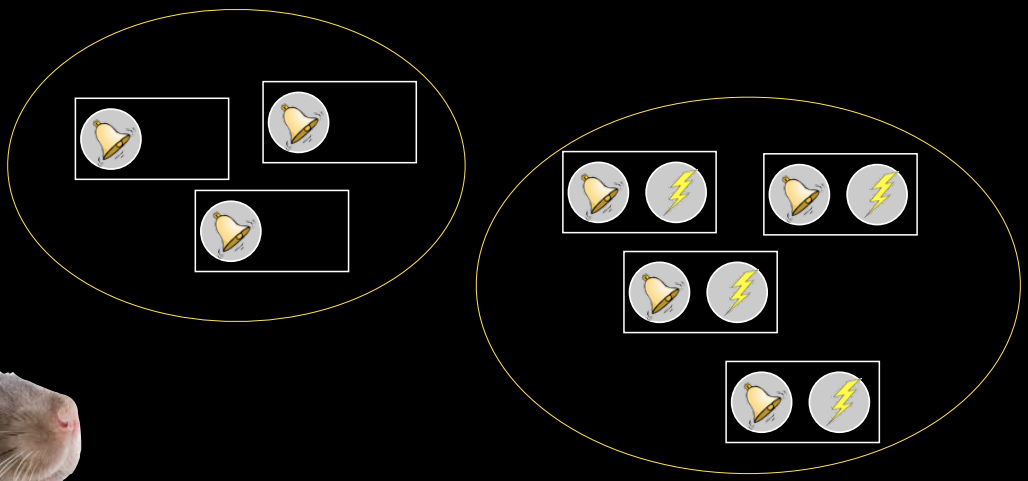

Inferred structure as memory traces

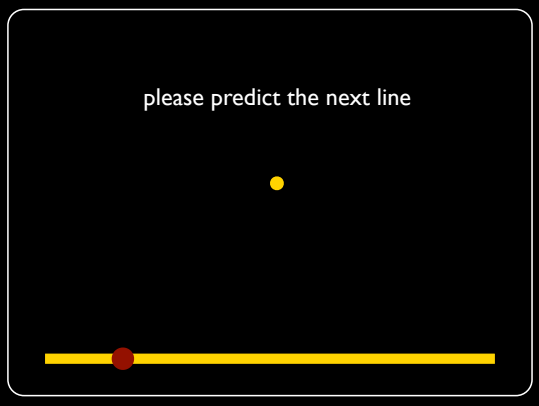

Angela Radulescu

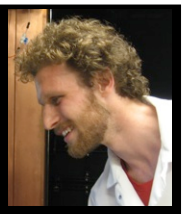
Sam Gershman

Gershman, Radulescu, Norman & Niv 2014

Inferred structure as memory traces

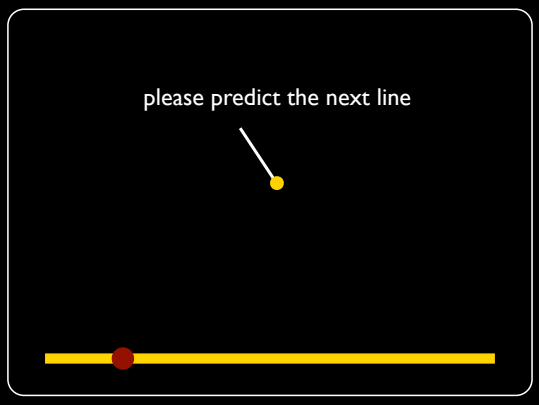

Gershman, Radulescu, Norman & Niv 2014

Inferred structure as memory traces

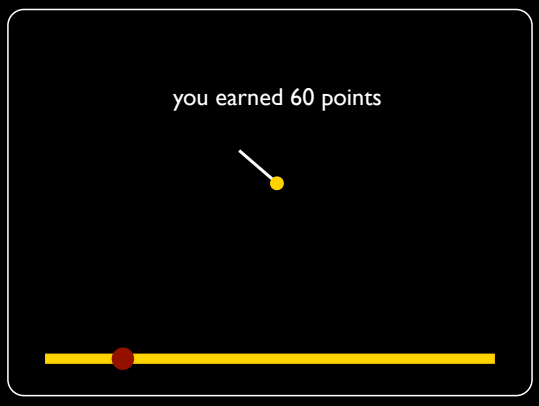

Gershman, Radulescu, Norman & Niv 2014

Inferred structure as memory traces

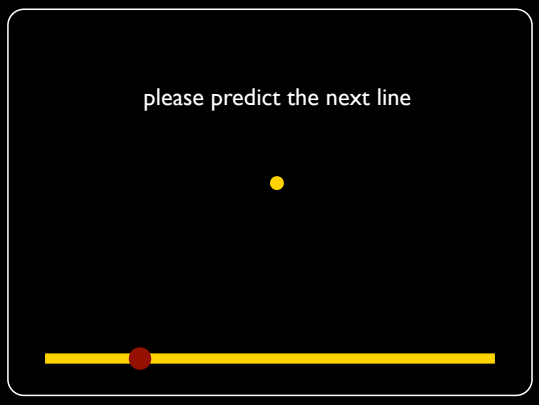

Gershman, Radulescu, Norman & Niv 2014

Inferred structure as memory traces

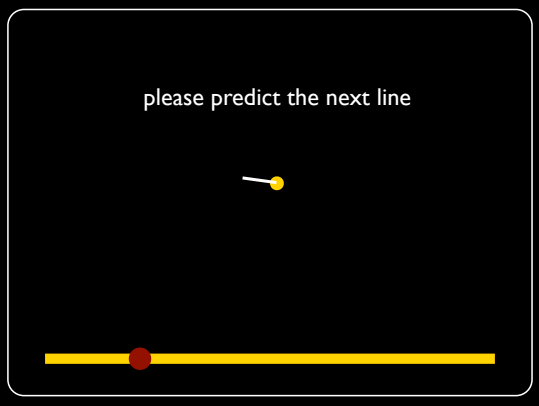

Gershman, Radulescu, Norman & Niv 2014

Inferred structure as memory traces

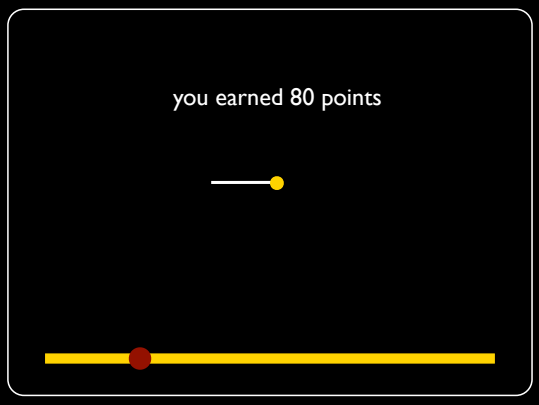

Gershman, Radulescu, Norman & Niv 2014

Inferred structure as memory traces

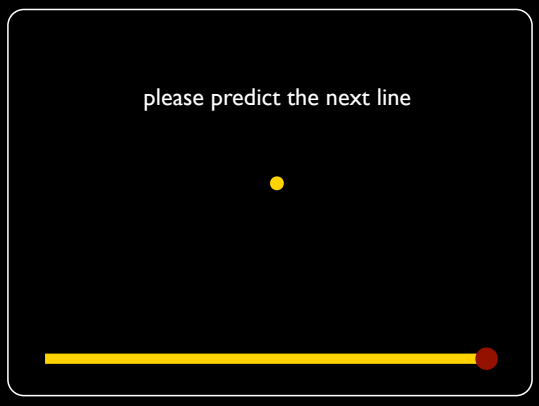

Gershman, Radulescu, Norman & Niv 2014

Inferred structure as memory traces

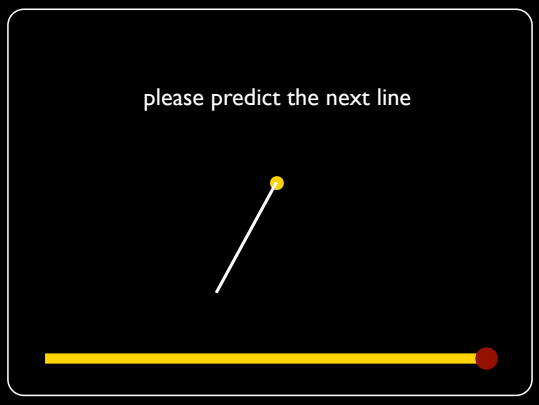

Gershman, Radulescu, Norman & Niv 2014

Inferred structure as memory traces

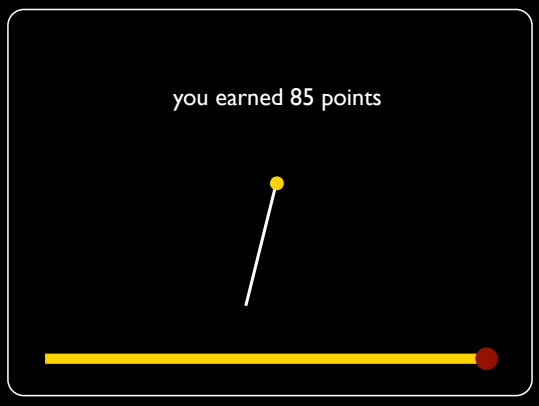

Gershman, Radulescu, Norman & Niv 2014

Inferred structure as memory traces

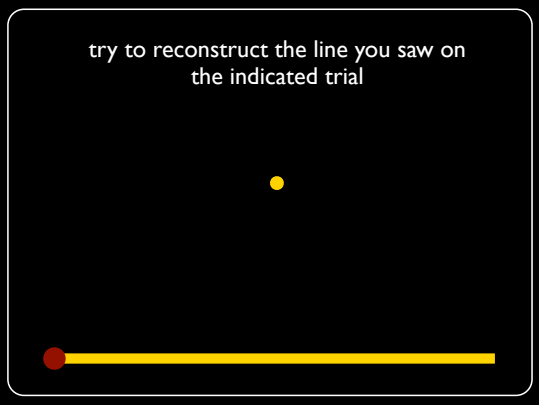

Gershman, Radulescu, Norman & Niv 2014

Inferred structure as memory traces

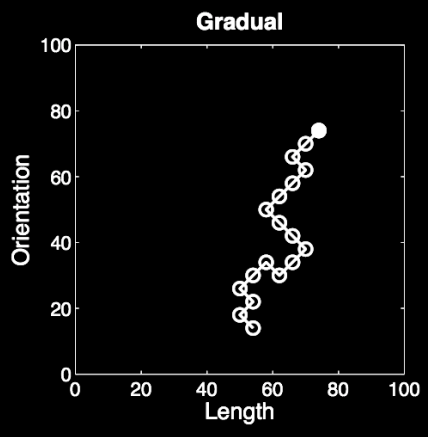

Gershman, Radulescu, Norman & Niv 2014

Inferred structure as memory traces

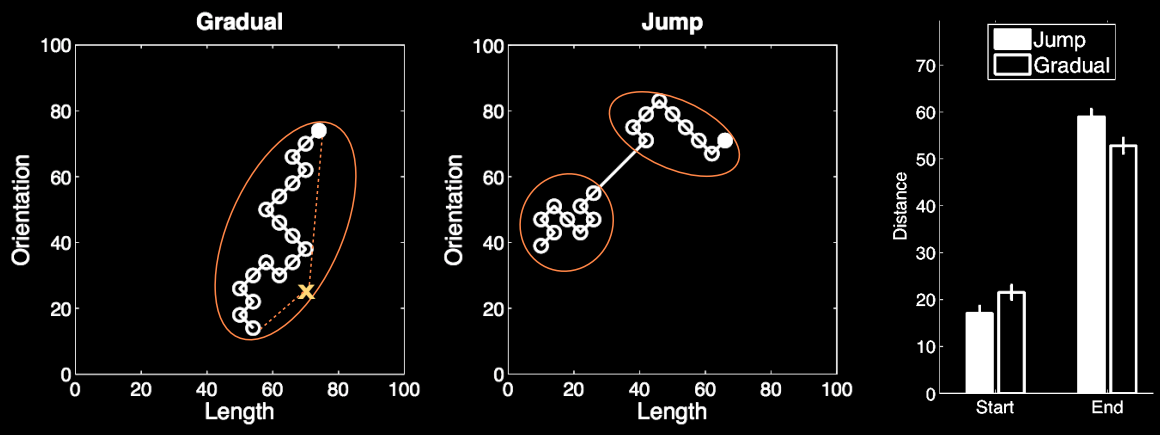

Gershman, Radulescu, Norman & Niv 2014

how wide can this simple theory reach? quite wide. one example:

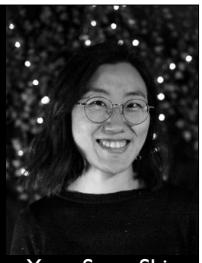

Yeon Soon Shin

"Happy families are all alike; every unhappy family is unhappy in its own way."
- Leo Tolstoy

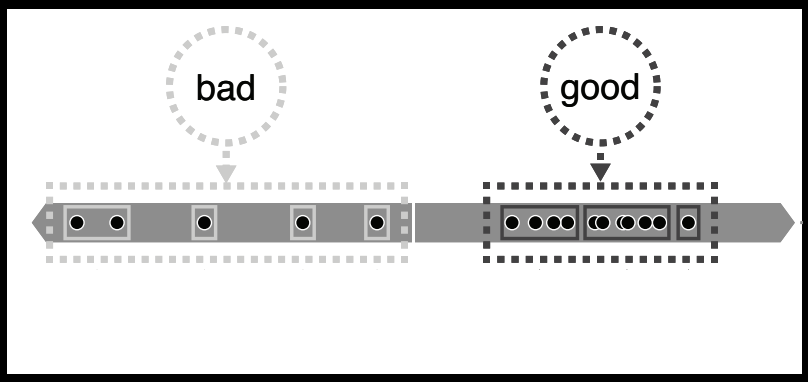

Shin & Niv 2021

how wide can this simple theory reach? quite wide. one example:

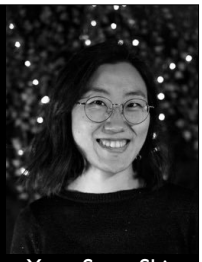

Yeon Soon Shin

"Happy families are all alike; every unhappy family is unhappy in its own way."
- Leo Tolstoy

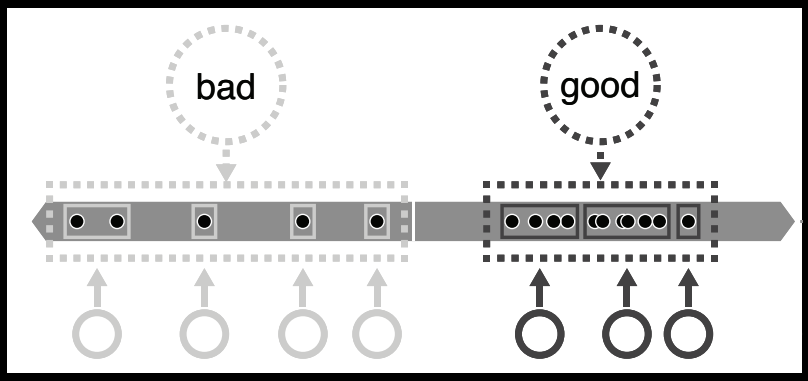

Shin & Niv 2021

how wide can this simple theory reach? quite wide. one example:

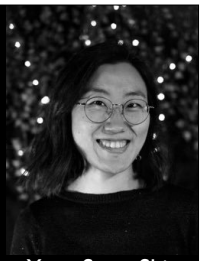

Yeon Soon Shin

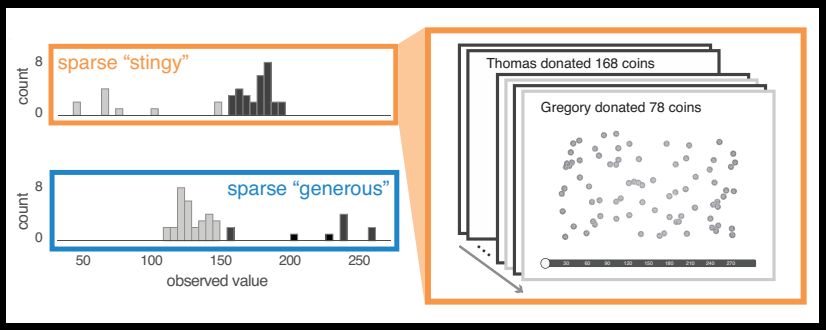

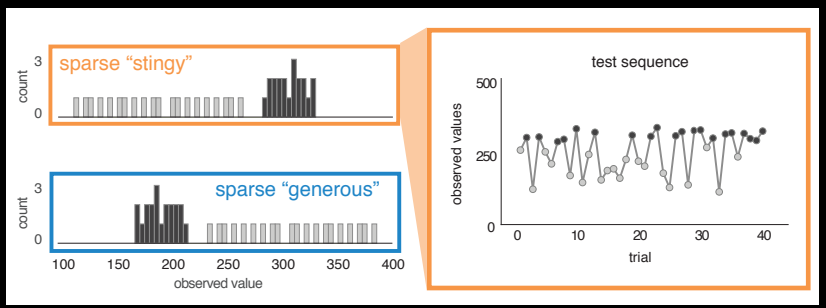

Shin & Niv 2021

latent-cause inference explains the negativity bias as a sparsity bias

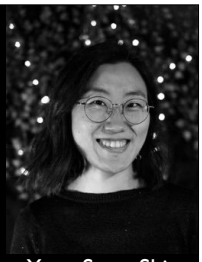

Yeon Soon Shin

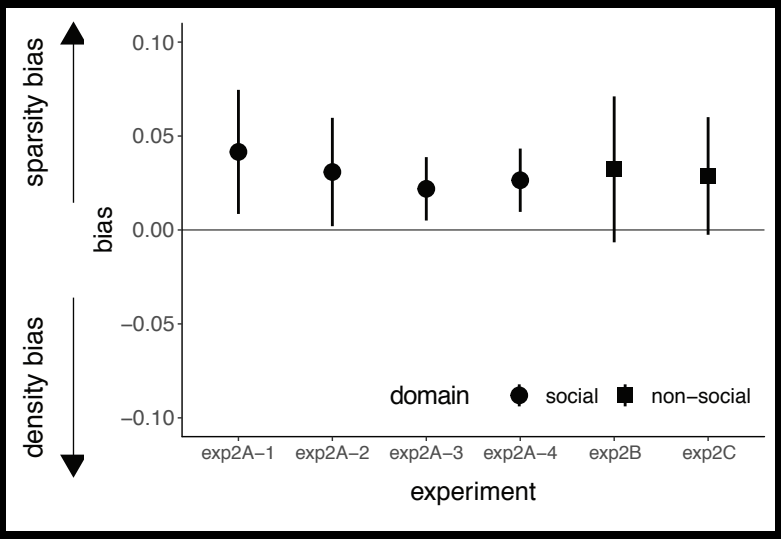

Shin & Niv 2021

"Battle" between learning and memory

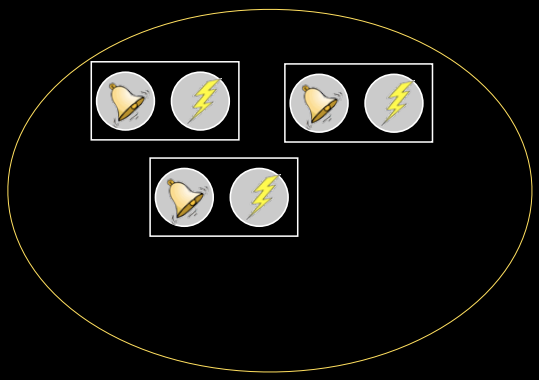

reinforcement learning
(modify existing cluster)

"Battle" between learning and memory

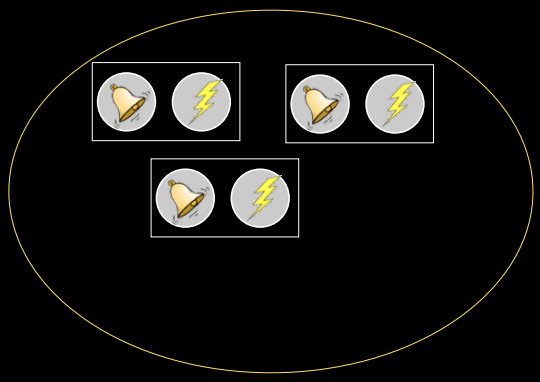

structural learning
(create new state)

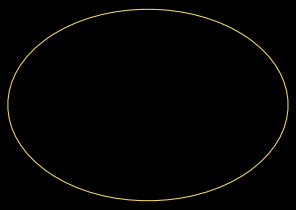

how can we erase a fear memory?

what if we make extinction more similar to acquisition?

Erasing fear through gradual extinction

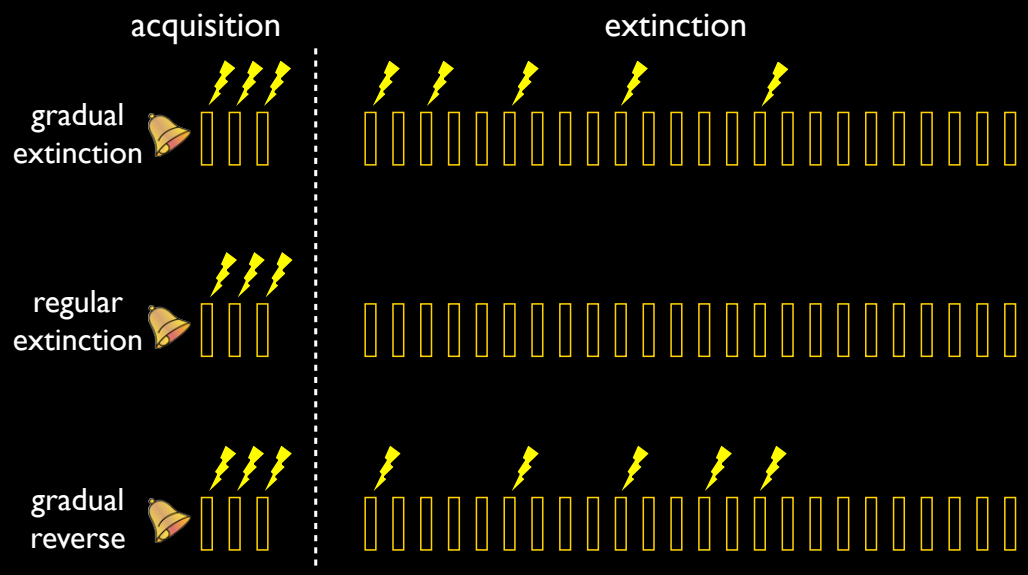

Gershman, Jones, Norman, Monfils & Niv 2013

Erasing fear through gradual extinction

Mingyu Song

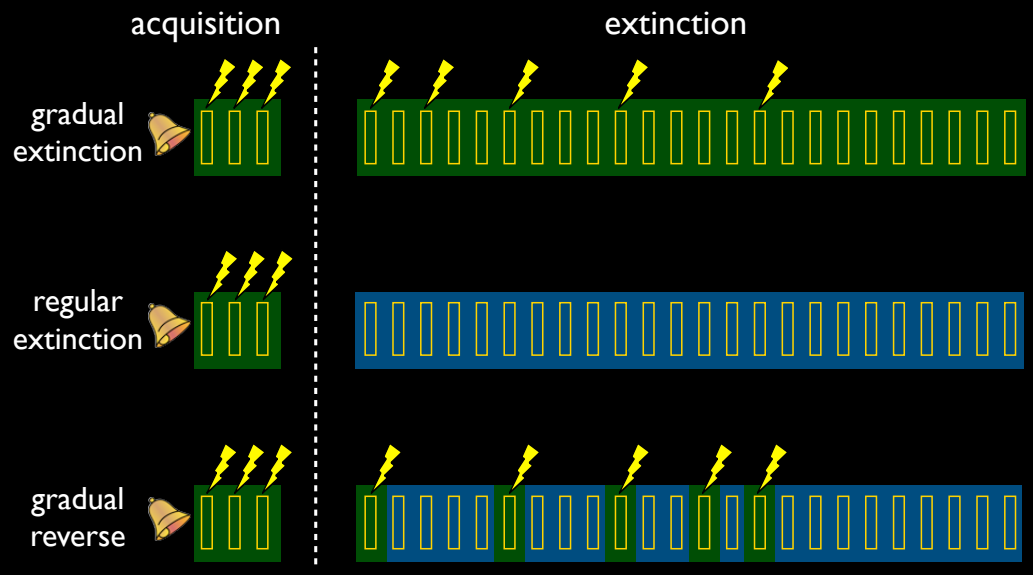

Model: Song, Jones, Monfils & Niv 2021

Erasing fear through gradual extinction

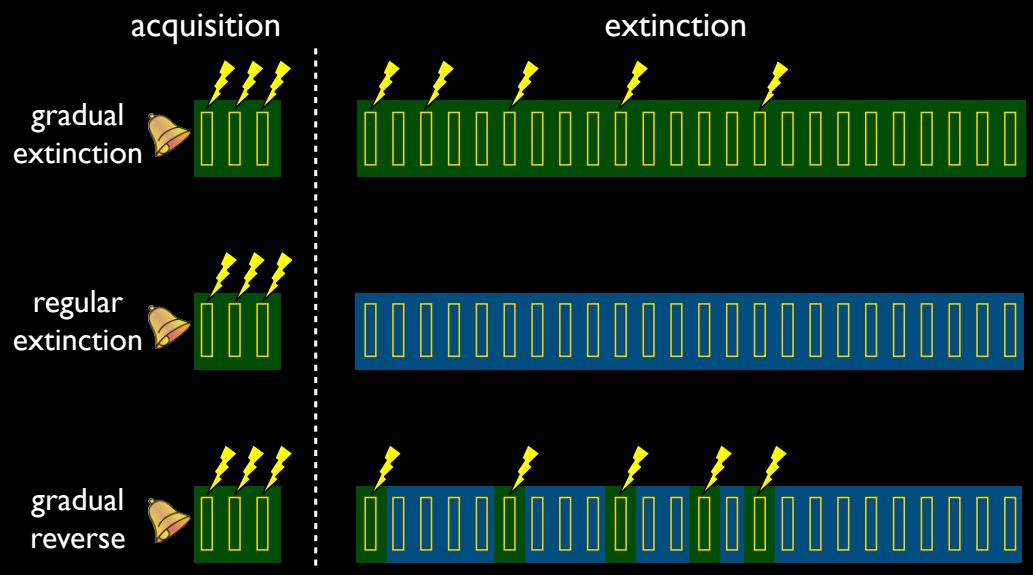

test one day (reinstatement) or 30 days later (spontaneous recovery)

No recovery of fear in gradual extinction group!

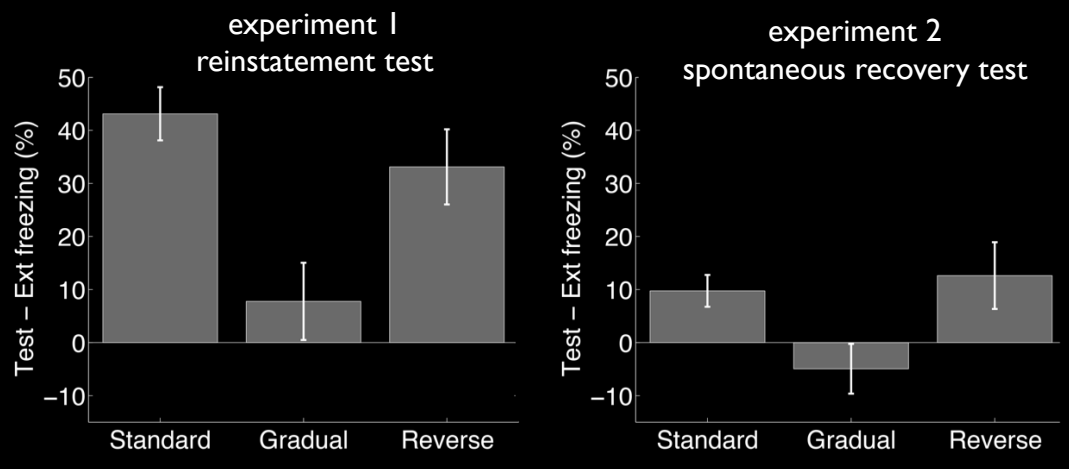

Gershman, Jones, Norman, Monfils & Niv 2013

Inferred structure controlled by titrating prediction errors

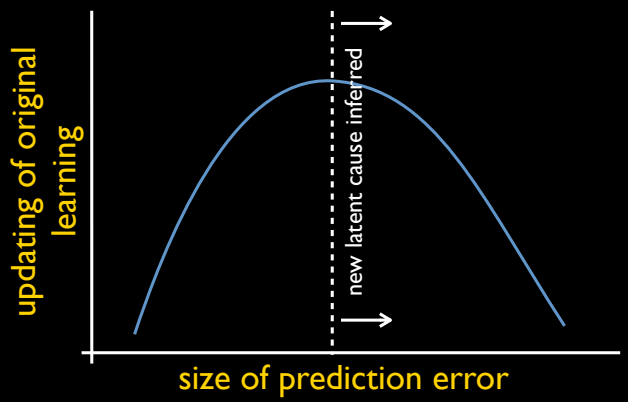

5 min break

Where in the brain are these states stored?

The mysterious orbitofrontal cortex

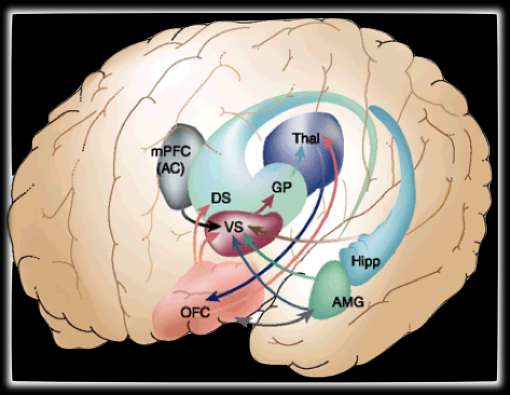

pervasive but subtle role in decision making in the lab...

involved in making good "life choices" in the real world

- after suffering OFC damage, "Elliot" divorced his wife, quit his job and married a prostitute (Damasio 1994, Eslinger & Damasio 1985)
- But performed well in lab tests of intelligence, reading and writing comprehension, verbal fluency, visuospatial abilities, and facial recognition

Hypothesis: OFC codes state space for tasks

Takahashi et al, 2012
Wilson et al. 2014

Robert Wilson

hypothesis

the orbitofrontal cortex codes the current location in a 'cognitive map' of task space

that is: it encodes *the current state of the task*

especially: inferred states that are not directly observable (partially observable states)

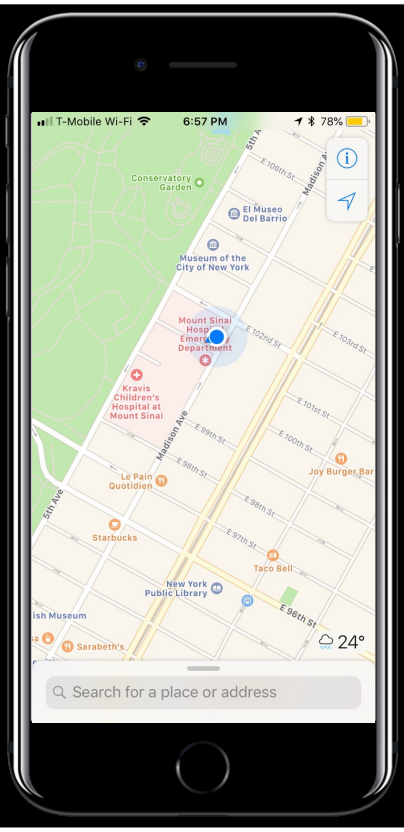

Sharpe et al., 2019

Case study I: reversal learning

Wilson et al. 2014

conditioning	**reversal 1**	**reversal 2**
R → reward	R → no reward	R → reward
L → no reward	L → reward	L → no reward

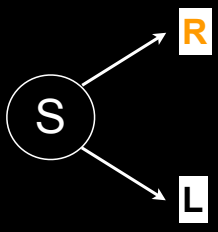

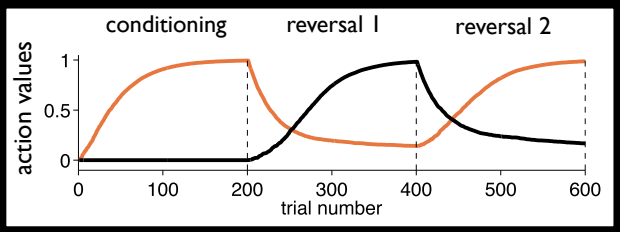

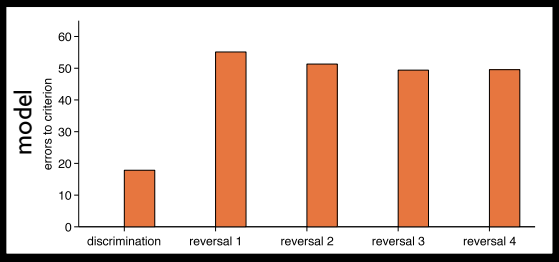

Case study I: reversal learning

Wilson et al. 2014

conditioning
R → reward
L → no reward

reversal 1
R → no reward
L → reward

reversal 2
R → reward
L → no reward

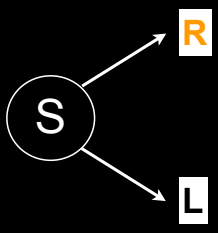

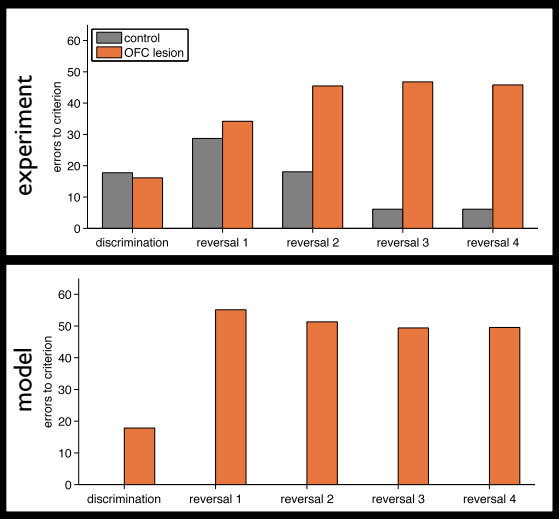

Case study I: reversal learning

Wilson et al. 2014

conditioning
R → reward
L → no reward

reversal 1
R → no reward
L → reward

reversal 2
R → reward
L → no reward

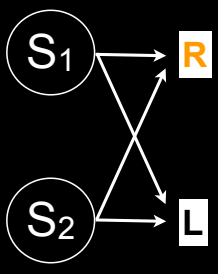

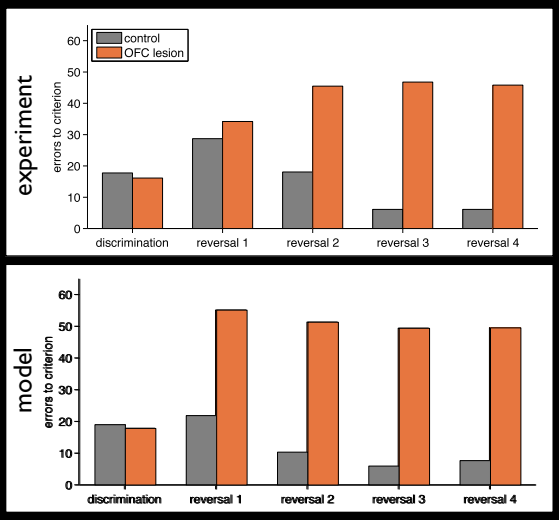

Case study I: reversal learning

Wilson et al. 2014

IDEA:

- OFC is critical for representing *inferred* states that are not stimulus-bound (internal states, memory-dependent states)
- without OFC the animal can only rely on stimulus-bound states

Case study II: extinction

Wilson et al. 2014

conditioning
P → reward
N → no reward

extinction
P → no reward
N → no reward

reacquisition
P → reward
N → no reward

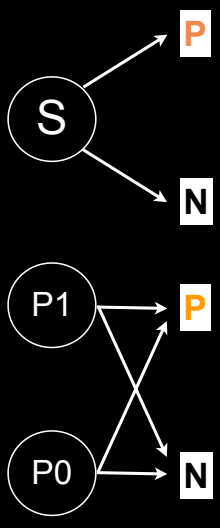

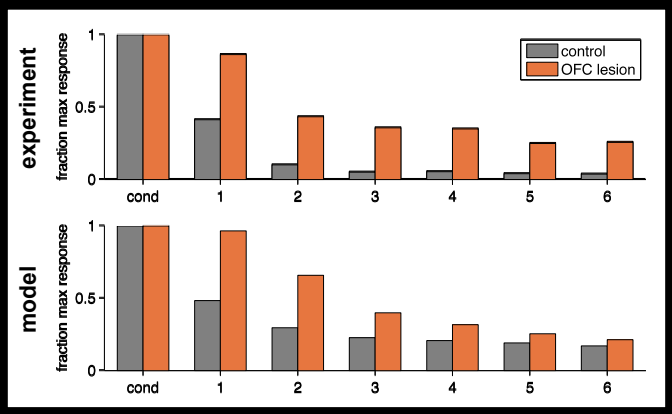

Butter et al. (1963)

Case study II: extinction

Wilson et al. 2014

- can also explain why rapid reacquisition is not seen after OFC lesions
- predicts no reinstatement or spontaneous recovery without OFC

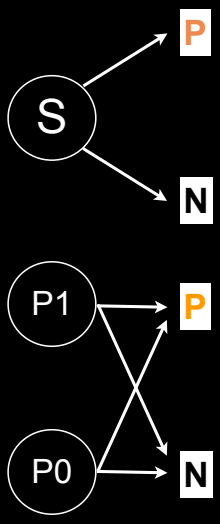

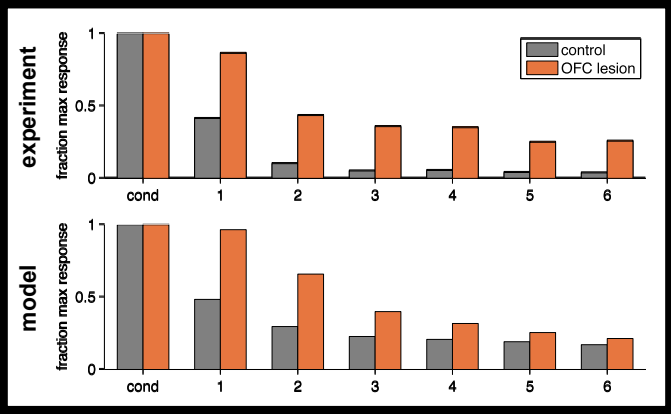

Butter et al. (1963)

How do we test this hypothesis?

Two conditions for 'task state' area:

1) ***Representation*** - all variables relevant to the task (stimuli, actions, context) *must* be encoded
2) ***Specificity*** - variables that are irrelevant for the current task should *not* be encoded

Wilson et al. 2014

Navigating a hidden-state task

Nico Schuck Ming Bo Cai

Task: is it Young or Old?
Perform task on either Face or House
Switch category when age switches

Schuck et al. 2016

Testing the OFC hypothesis:

a new block begins now.
You will start with
FACES
good luck!

young

young

old

young

old

+

Markov states for the Young/Old task

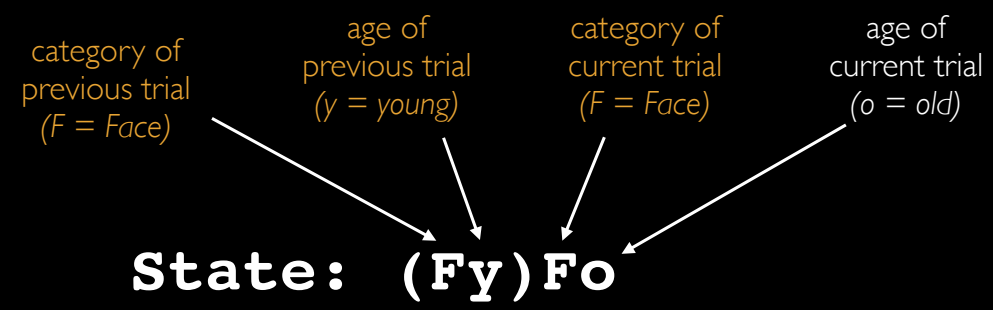

Schuck et al. 2016

Can state information be decoded from the OFC?

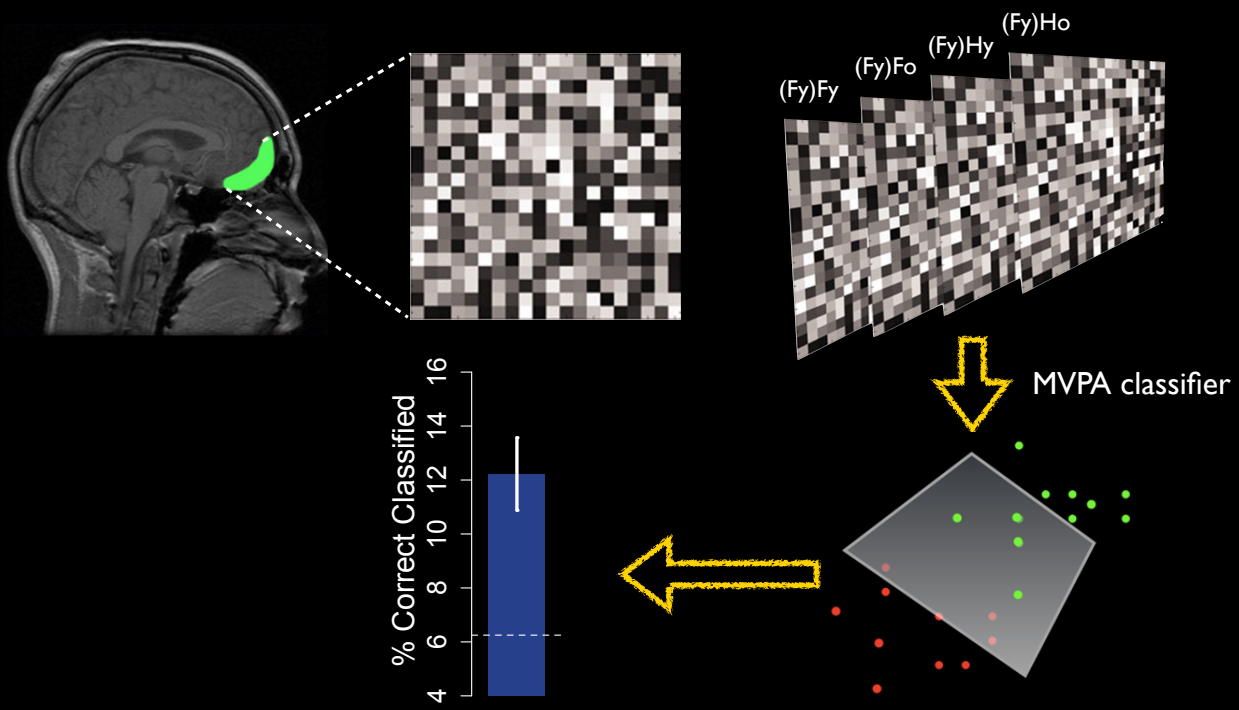

Schuck et al. 2016

Can state information be decoded from the OFC?

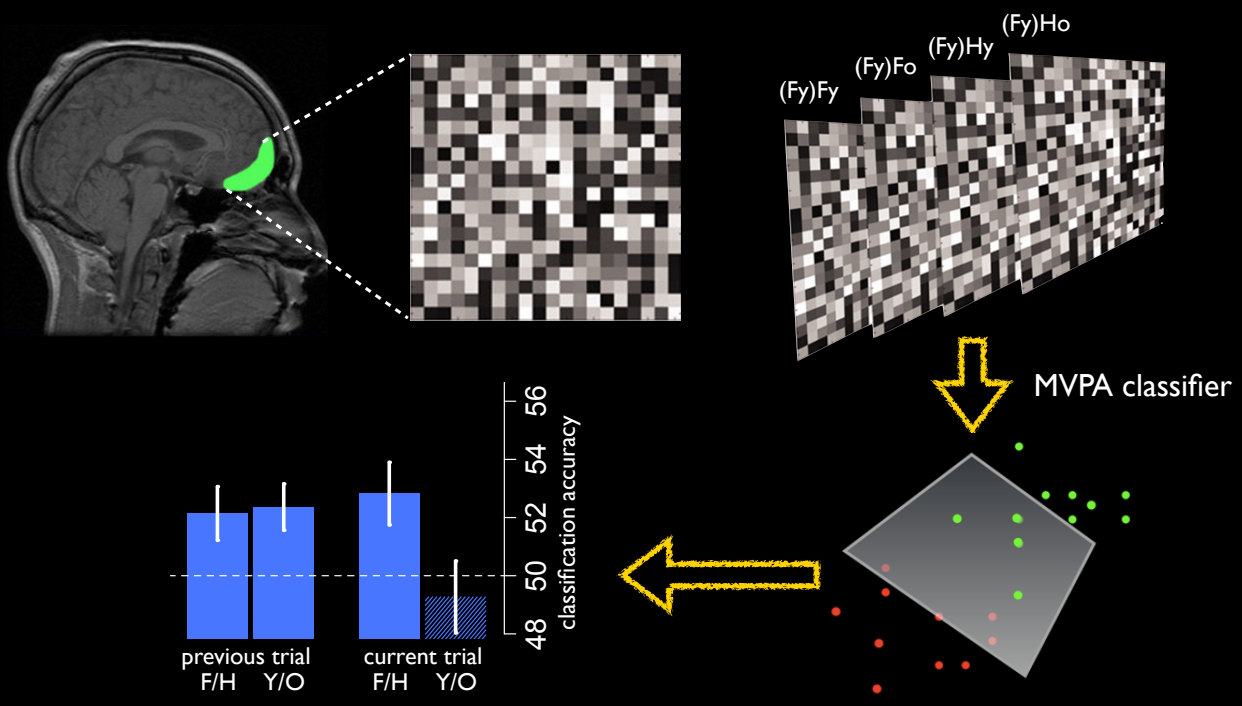

Schuck et al. 2016

...in fact, it is the only area that has all the information

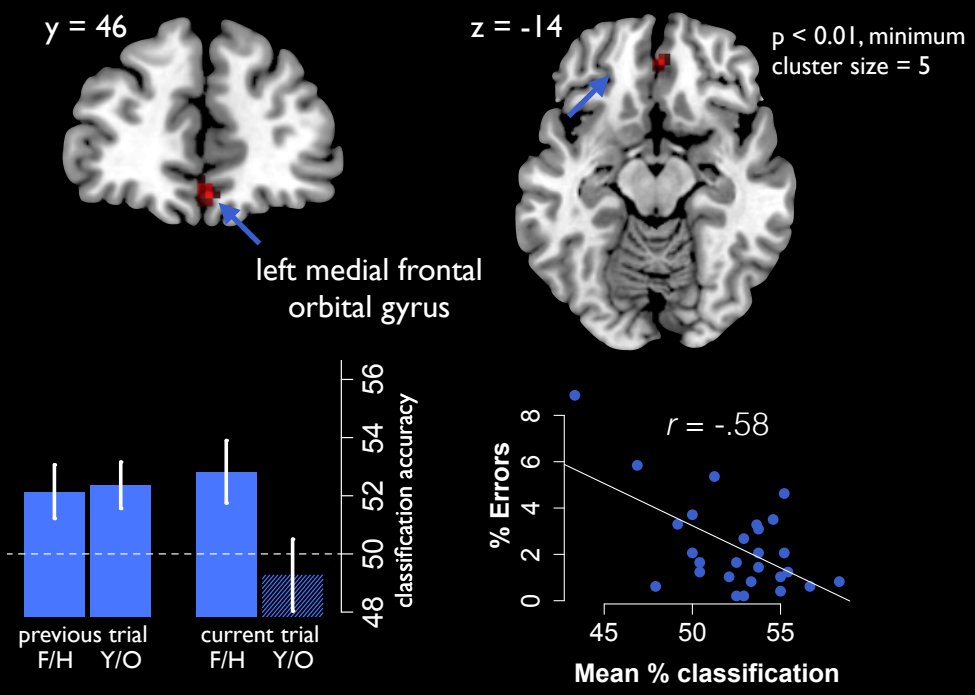

Schuck et al. 2016

The cognitive map in the OFC affects performance

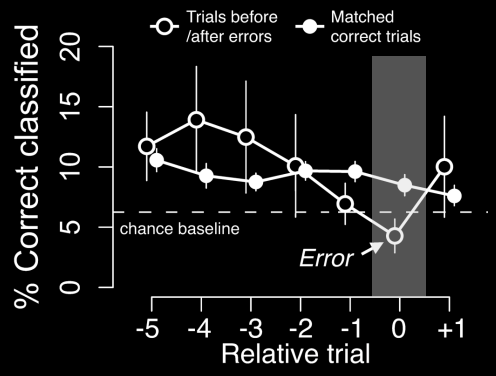

Schuck et al. 2016

How do we test this hypothesis?

Two conditions for 'task state' area:

✓ ***Representation*** - all variables relevant to the task (stimuli, actions, context) *must* be encoded

2) ***Specificity*** - variables that are irrelevant for the current task should *not* be encoded

OFC encodes only task-relevant information...

...likely building on information in other areas

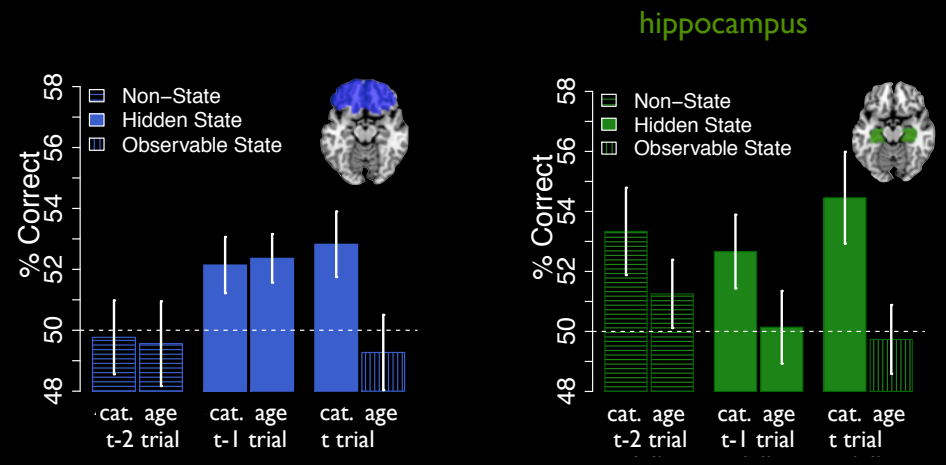

Schuck et al. 2016

How do we test this hypothesis?

Two conditions for 'task state' area:

- ✓ ***Representation*** - all variables relevant to the task (stimuli, actions, context) *must* be encoded
- ✓ ***Specificity*** - variables that are irrelevant for the current task should *not* be encoded

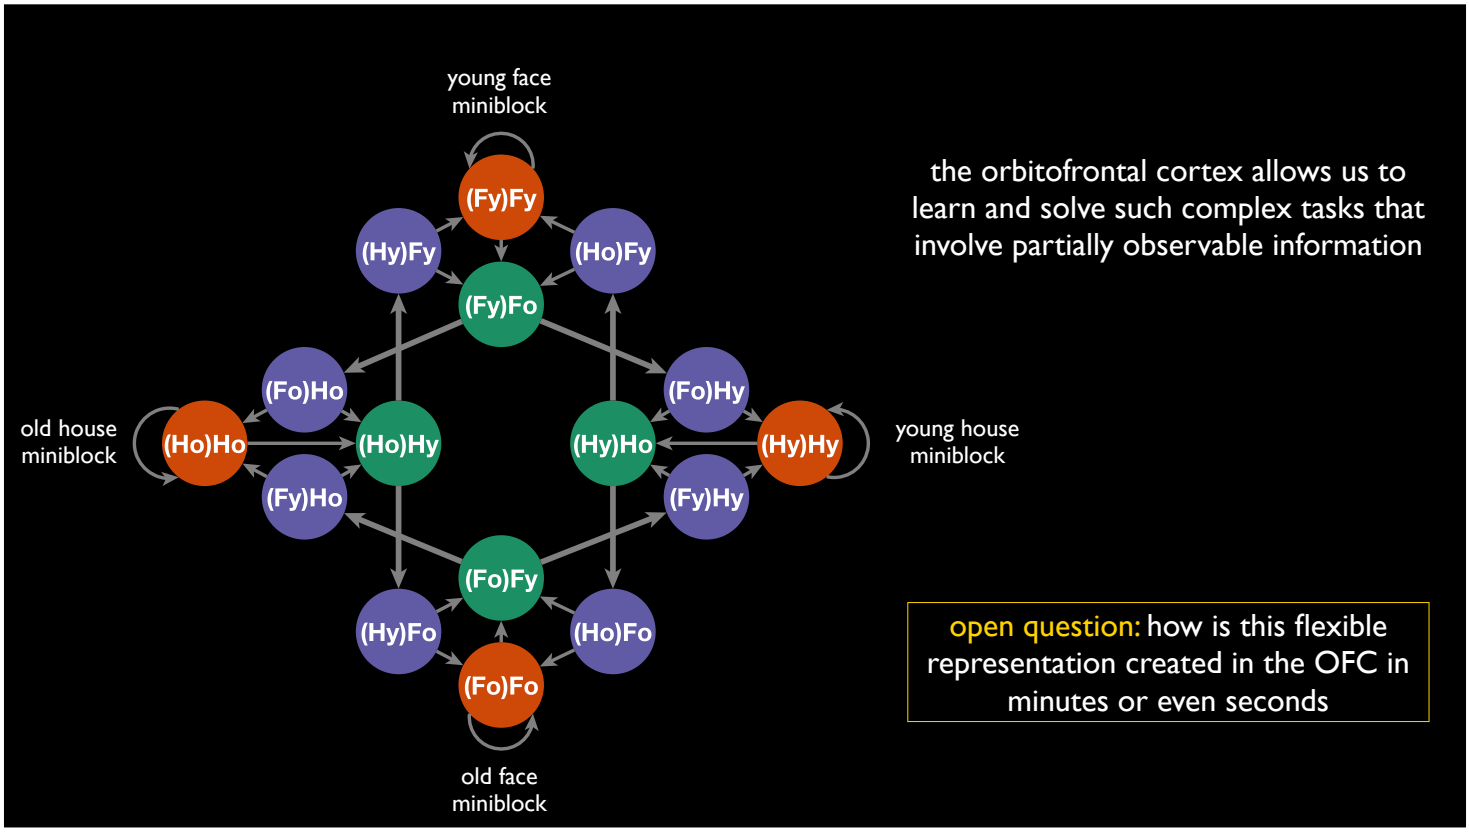
young face miniblock
(Fy)Fy
(Hy)Fy
(Ho)Fy
(Fy)Fo
(Fo)Ho
(Fo)Hy
old house miniblock
(Ho)Ho
(Ho)Hy
(Hy)Ho
(Hy)Hy
young house miniblock
(Fy)Ho
(Fy)Hy
(Fo)Fy
(Hy)Fo
(Ho)Fo
(Fo)Fo
old face miniblock
the orbitofrontal cortex allows us to learn and solve such complex tasks that involve partially observable information
open question: how is this flexible representation created in the OFC in minutes or even seconds

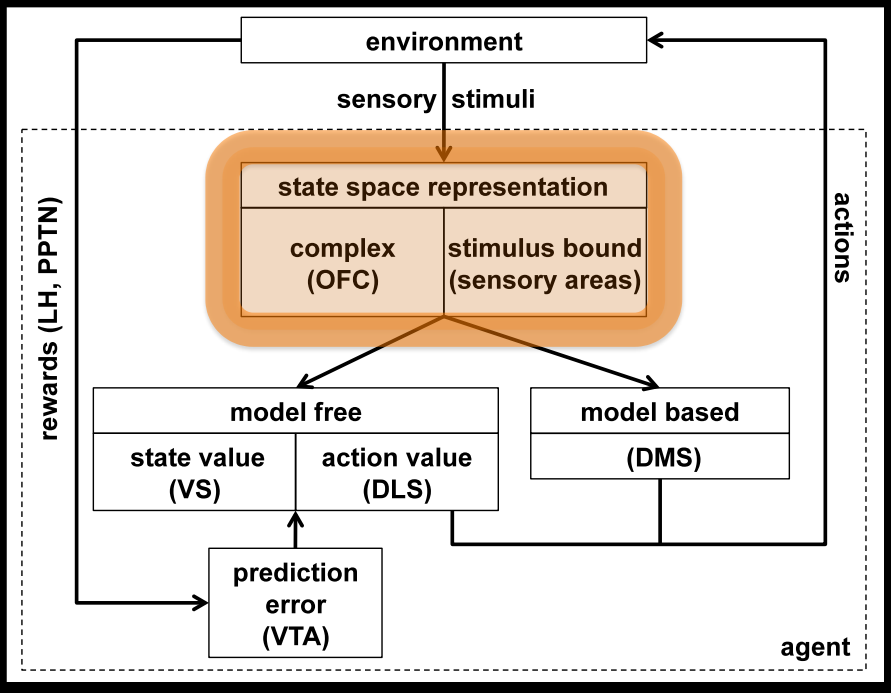
environment
sensory stimuli
state space representation
complex (OFC)
stimulus bound (sensory areas)
rewards (LH, PPTN)
actions
model free
state value (VS)
action value (DLS)
model based
(DMS)
prediction error (VTA)
agent

Shallow learning with deep representations

learning = generalization

reinforcement learning

e.g. basal ganglia dopamine

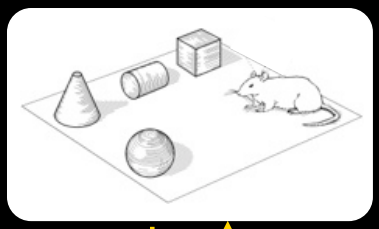

representation learning

e.g. hippocampus, frontoparietal orbitofrontal cortex

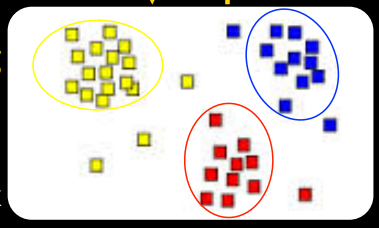

Implications for daily life

1. Slow learning can be good. If you want to change previous knowledge, gradual learning is better than abrupt large changes.
2. Memory/learning are tied to context. Consider changing contexts when you want to differentiate memories. Or go back to old context to recall memory (is this why time seems to move much slower when you are traveling? Lots of new contexts?)