The organization of the basal ganglia

Ilana Witten

Different parts of the brain contribute to different types of learning, on different timescales

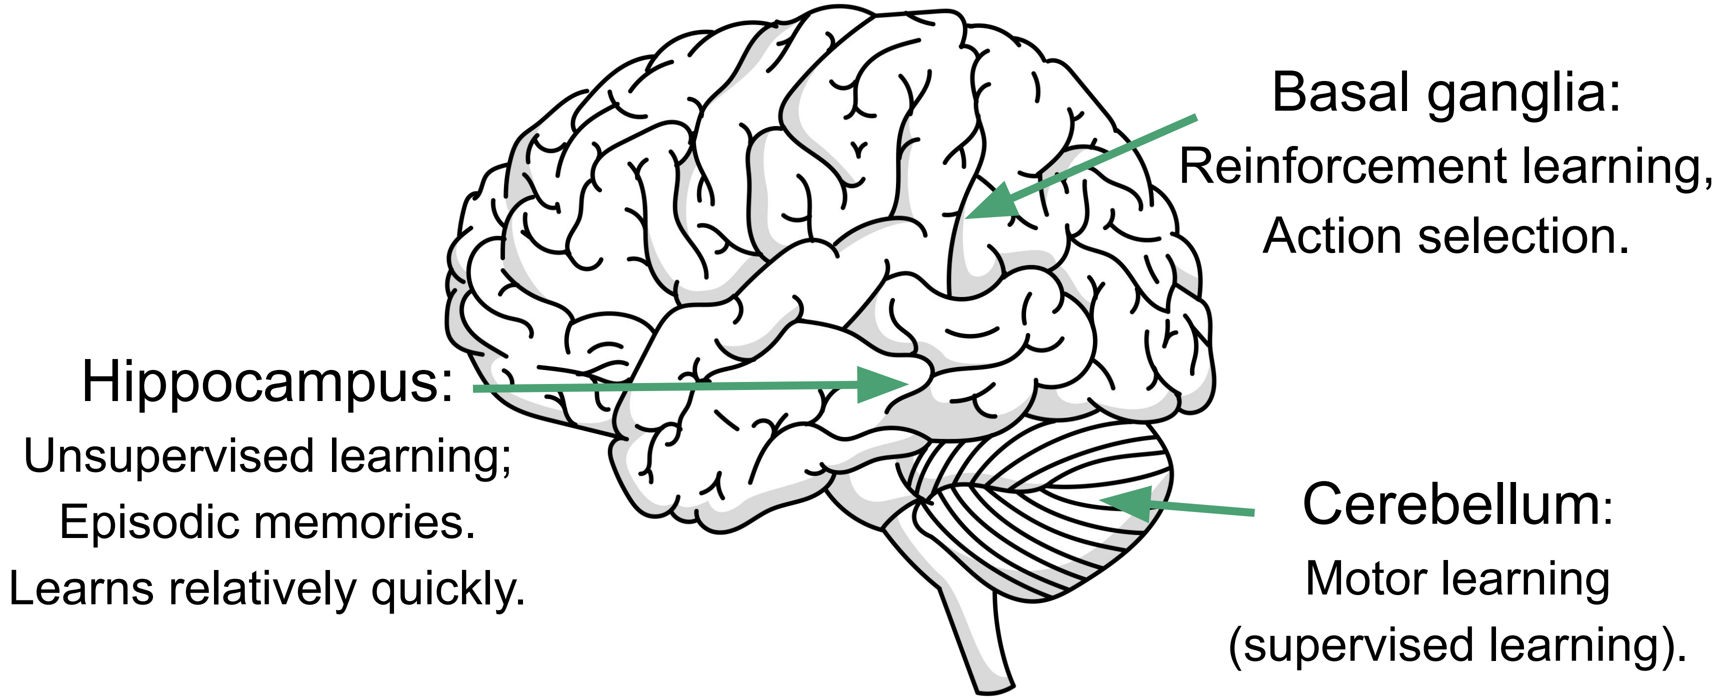

Different synaptic plasticity rules in different regions.

Basal ganglia: neural basis of reinforcement learning

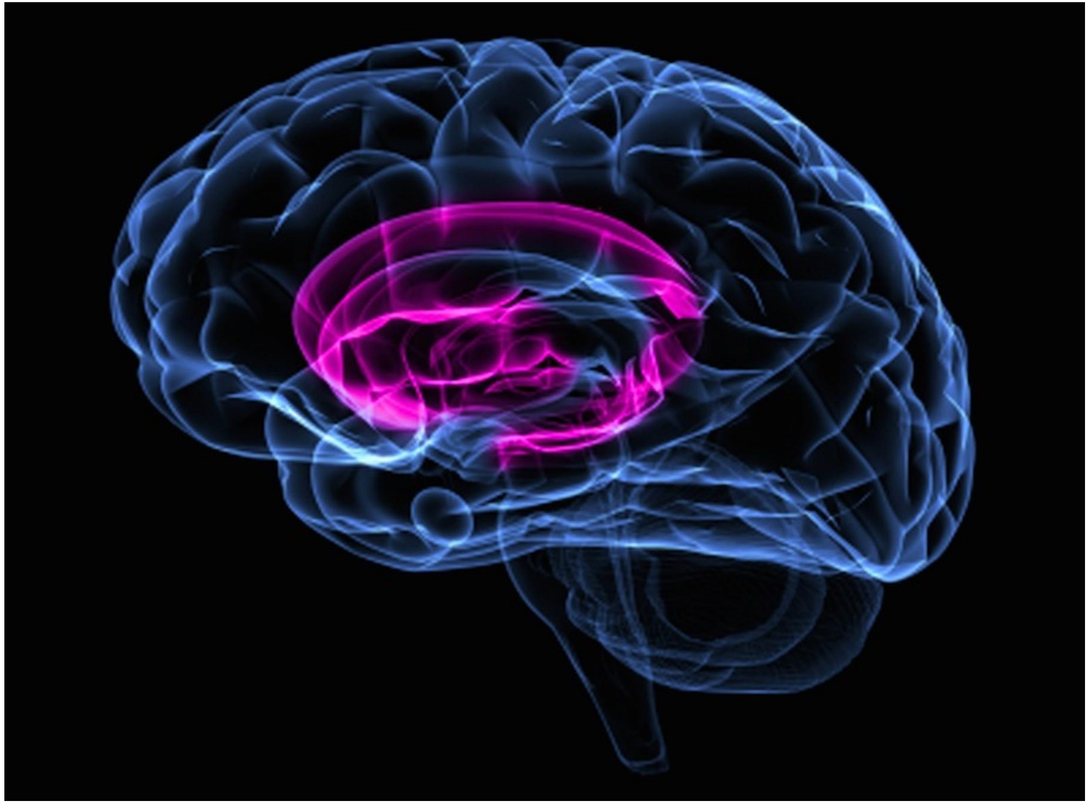

Outline for today:

1. Quick overview of **basal ganglia anatomy**
2. How does the **dopamine system** regulate the basal ganglia?
3. What is the function of the **direct** vs **indirect pathways** of the basal ganglia?

Outline for today:

1. **Quick overview of basal ganglia anatomy**
2. How does the **dopamine system** regulate the basal ganglia?
3. What is the function of the **direct** vs **indirect pathways** of the basal ganglia?

Basal ganglia function

- Different parts of the basal ganglia are connected to different parts of the cortex, and therefore have different functions
- The part of basal ganglia that is connected to motor cortex is involved in learning and selecting movements

Basal ganglia dysfunction

- Parkinson's and Huntington's are both motor diseases involving the basal ganglia
  - Parkinson's → difficult w/ movement initiation
  - Huntington's → too much movement

- Most psychiatric disorders have a major basal ganglia component
  - e.g. addiction, depression, OCD, schizophrenia

The basal ganglia forms a “cortico-striatal-thalamic loop”: cortical neurons -> striatal neurons --> thalamic neurons -> cortex

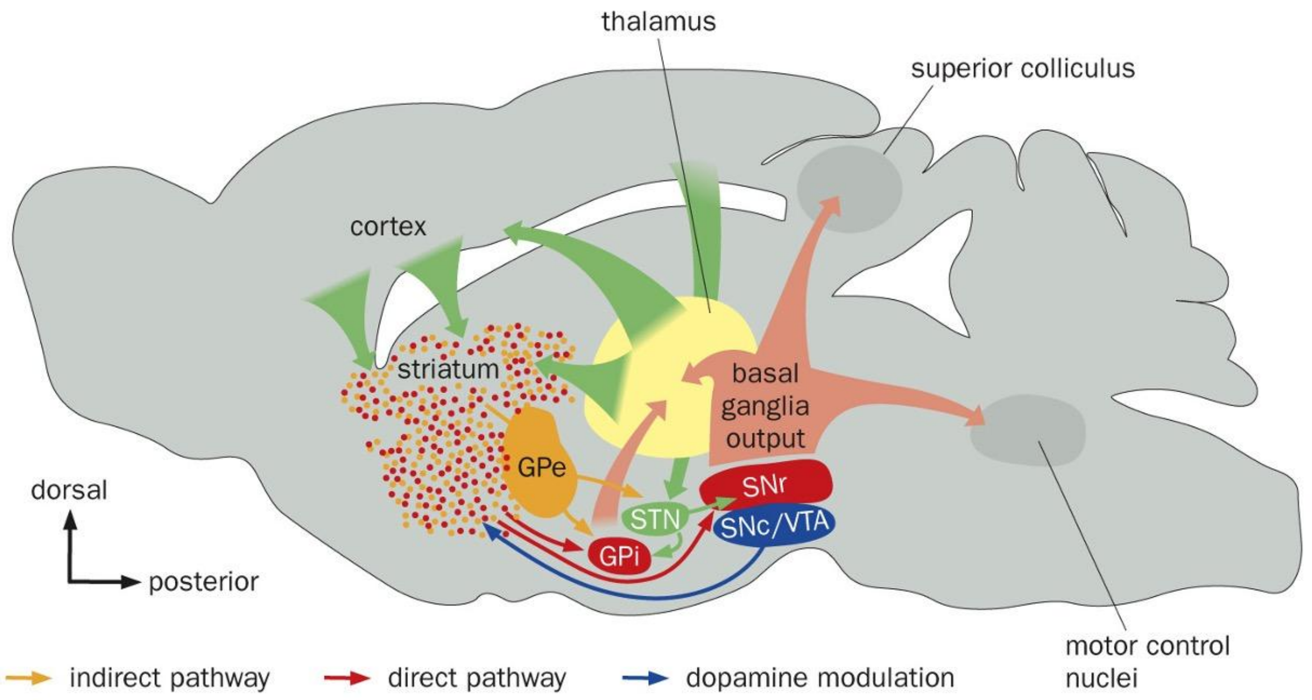

striatum == caudate/putamen

There are 2 outputs from the striatum: direct pathway (D1R SPNs) and indirect pathway (D2R SPNs) neurons

gaba

glutamate

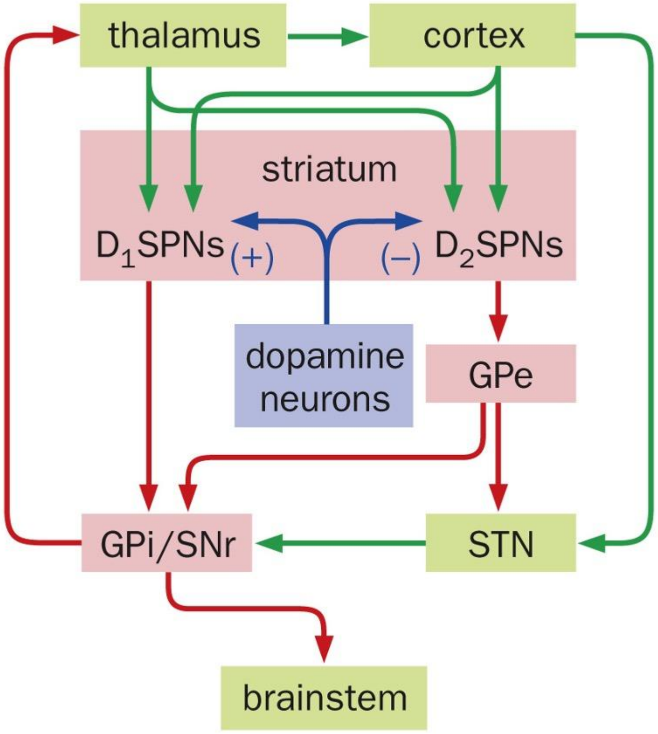

Figure 8-22c Principles of Neurobiology (© Garland Science 2016)

The basal ganglia forms MULTIPLE parallel “cortico-striatal-thalamic loops”: cortical neurons -> striatal neurons --> thalamus neurons -> cortex

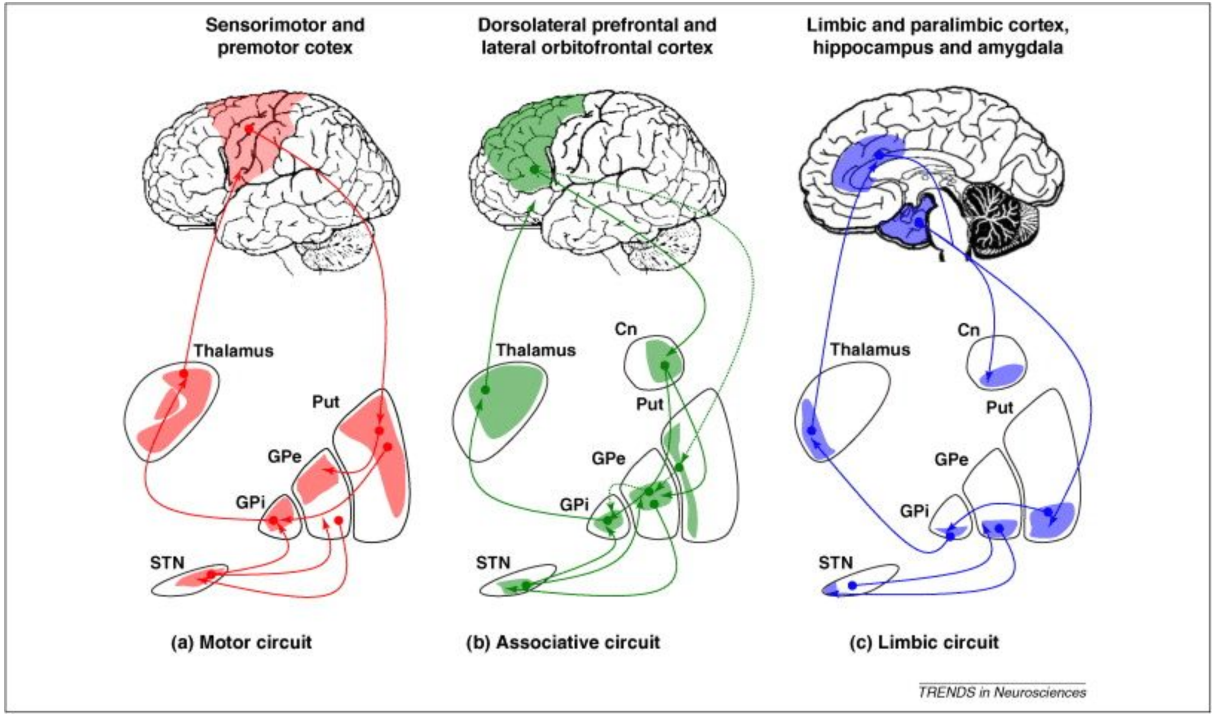

Different dopamine neurons project to different striatal subregions

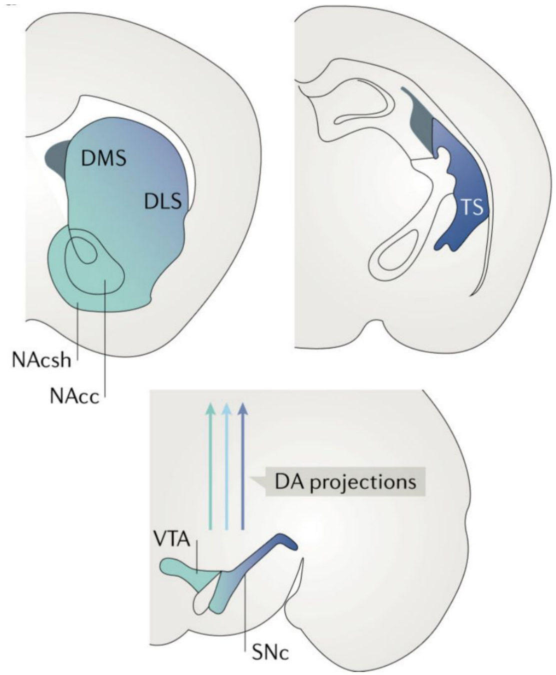

Different striatal subregions that have different specializations (and are part of different basal ganglia loops)

Dorsolateral striatum (DLS): motor sequences; habits

Dorsomedial striatum (DMS): goal-directed behavior; visual decision-making

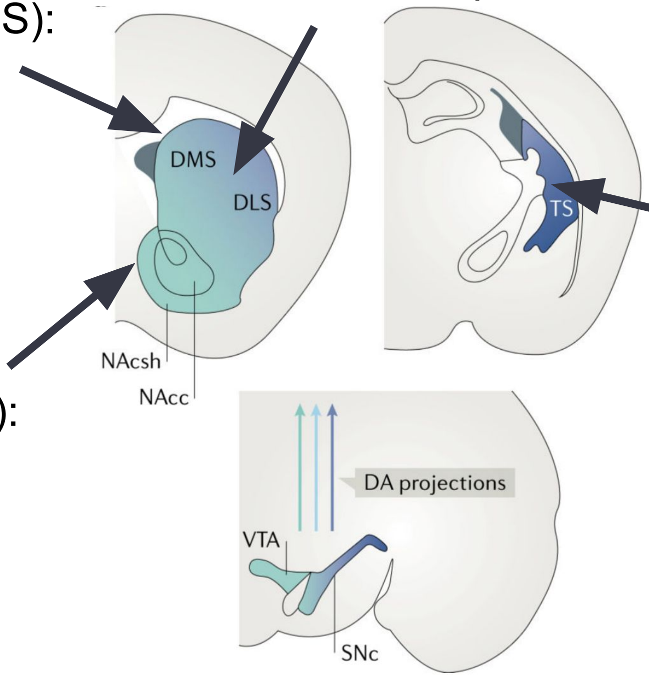

Tail of striatum (TS): threat learning; auditory decision-making

Nucleus accumbens (NAc): Pavlovian associations Emotional regulation Approach behavior

Outline for today

1. **Anatomy** overview
2. How does the **dopamine system** regulate the basal ganglia?
3. What is the function of the **direct** vs **indirect pathways** of the basal ganglia?

VTA dopamine neurons projecting to striatum: a reward prediction error (RPE)

- RPE is central to reinforcement learning algorithms

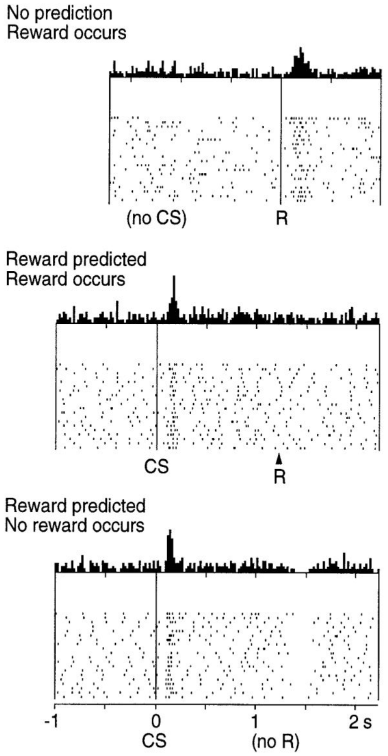

Schultz 1998 *J Neurophysiol*
In mice: Cohen & Uchida Nature *2012*

VTA dopamine neurons projecting to striatum: a reward prediction error (RPE)

- RPE is central to reinforcement learning algorithms

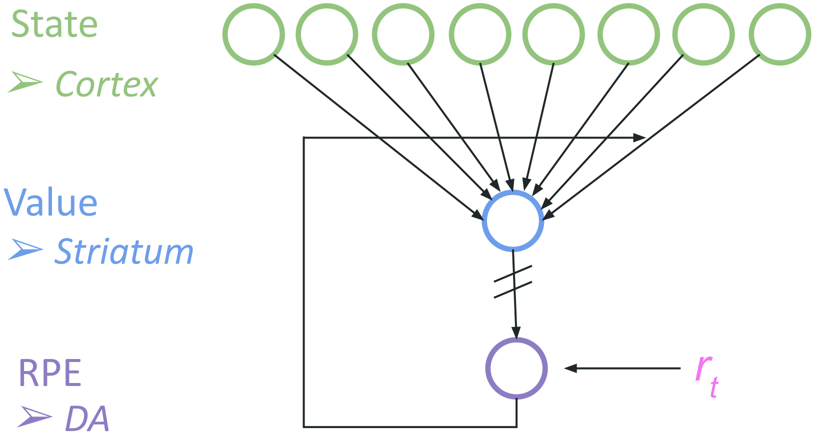

$$\delta_t = r_t + \gamma V(s_{t+1}) - V(s_t)$$

VTA dopamine neurons projecting to striatum: a reward prediction error (RPE)

- RPE is central to reinforcement learning algorithms
- Extensive evidence for this mapping
  - Value correlates in target regions (e.g. Ito & Doya 2009, Shin et al 2021)
  - Dopamine regulates synaptic plasticity (e.g. Reynolds and Wickens 2002; Calabresi et al. 2007; Fisher et al. 2017)
  - Optogenetic activation of dopamine neurons is reinforcing (e.g. Tsai et al. 2009; Witten et al. 2011; Steinberg et al. 2013; Stauffer et al. 2016, Parker et al 2016)

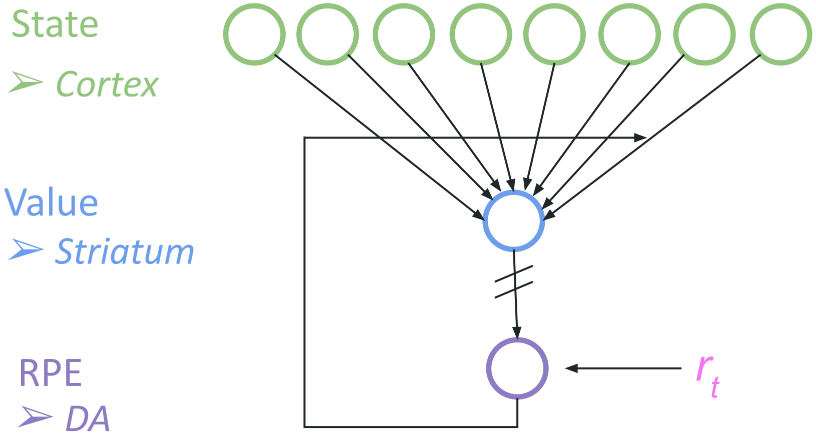

$$\delta_t = r_t + \gamma V(s_{t+1}) - V(s_t)$$

This circuit model assumes a single, scalar dopamine signal

- Experimental support in early recordings during low-dimensional tasks (i.e. Pavlovian conditioning)

- Conceptually, makes sense: a scalar prediction of the expected reward of the current state (value) can be used to guide a decision

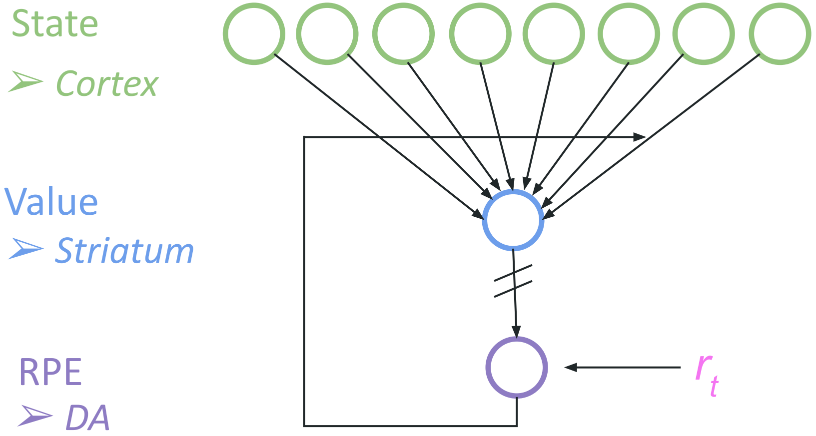

$$\delta_t = r_t + \gamma V(s_{t+1}) - V(s_t)$$

Heterogeneity in the dopamine system

Across projections:

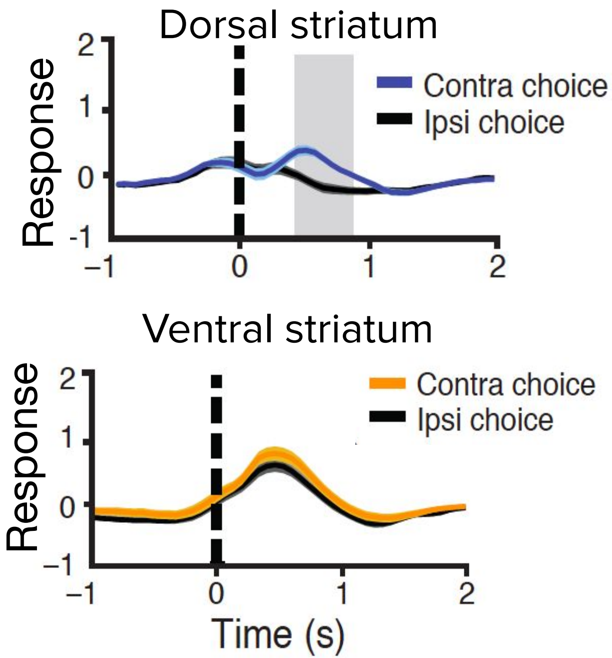

Parker *et al*, *Nature Neuro* 2016
Willmore *et al, Nature* 2022

Across neurons:

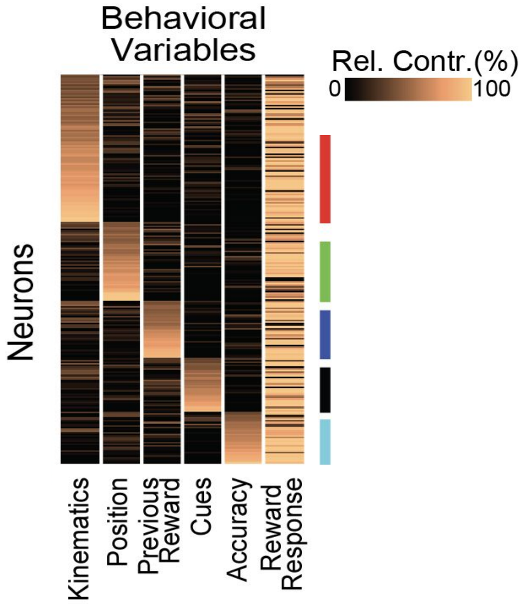

Engelhard *et al*, *Nature* 2019
Willmore *et al*, *Neuron* 2023

Causal effects:

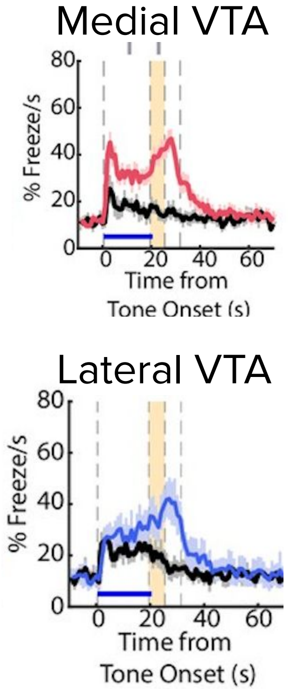

Cai et al, *eLife* 2020
Choi *et al, Cell Reports*, 2020

Mice choose the side with more cues in a virtual maze

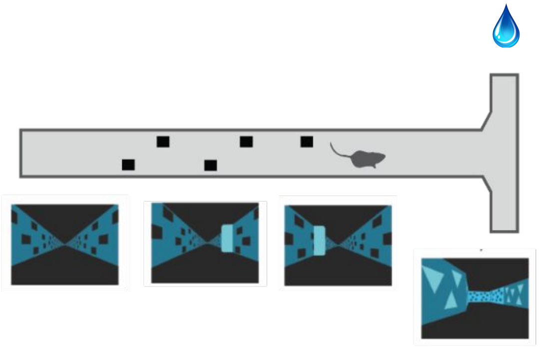

Engelhard et al. *Nature* 2019

Mice choose the side with more cues in a virtual maze

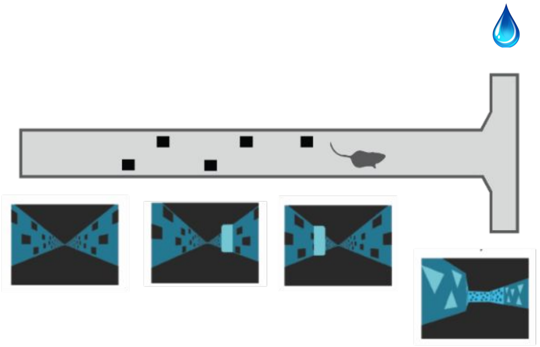

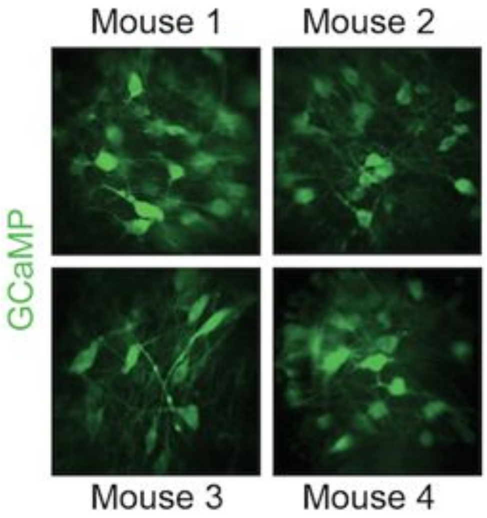

Engelhard et al. *Nature* 2019

Heterogeneity in the maze stem, homogeneity during outcome period

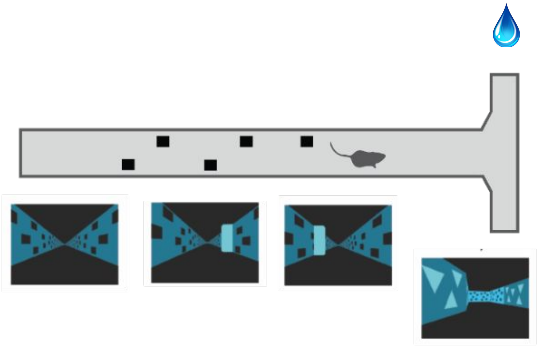

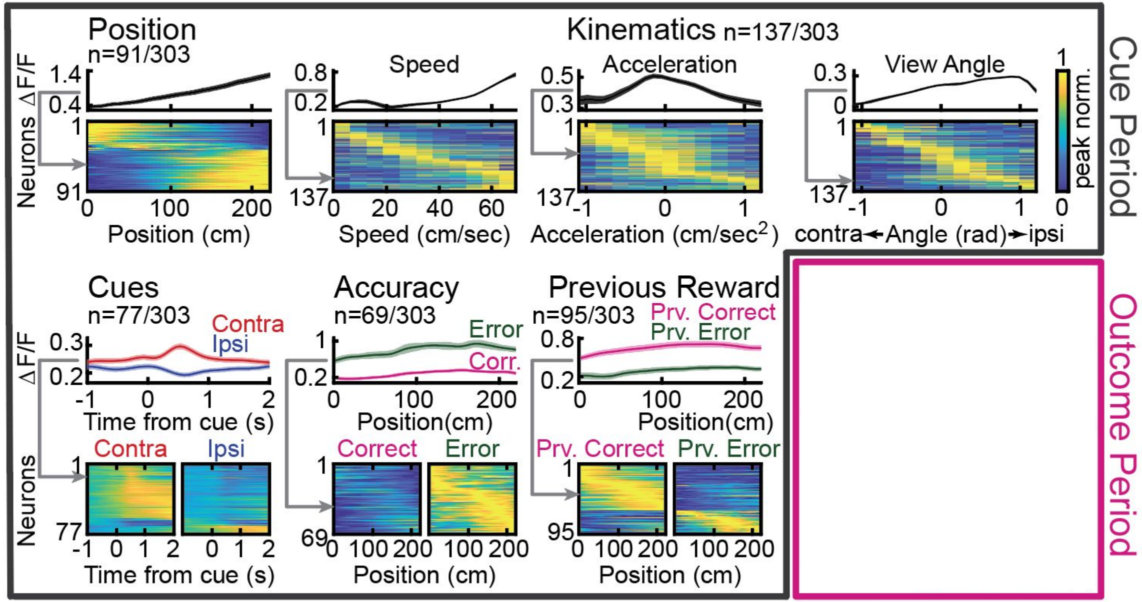

Engelhard et al. *Nature* 2019

Heterogeneity in the maze stem, homogeneity during outcome period

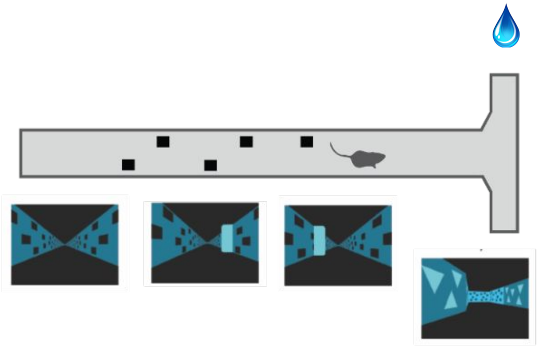

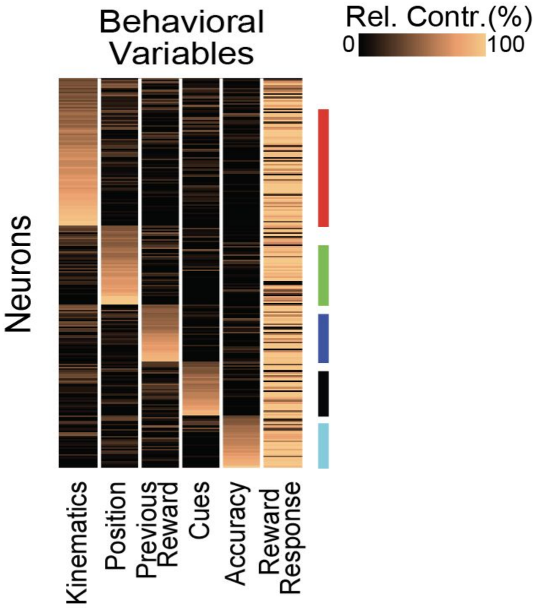

Engelhard *et al. Nature* 2019

Similar takeaway: Kremer *et al J Neuro* 2020

Scalar RPE Model

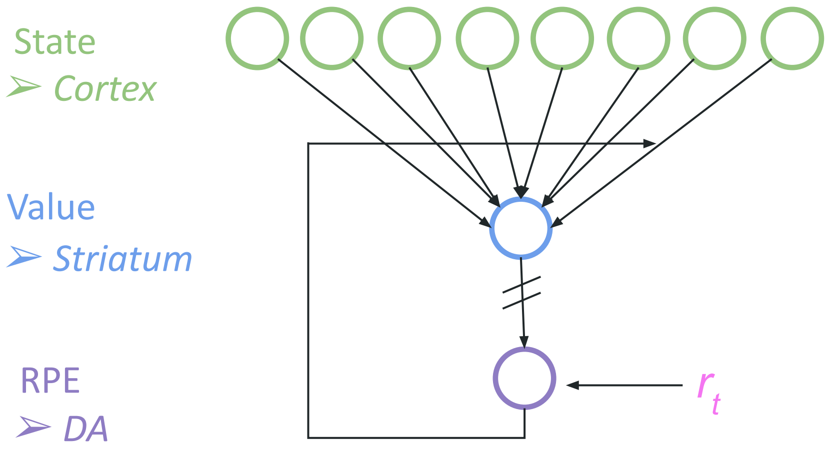

$$\delta_t = r_t + \gamma V(s_{t+1}) - V(s_t)$$

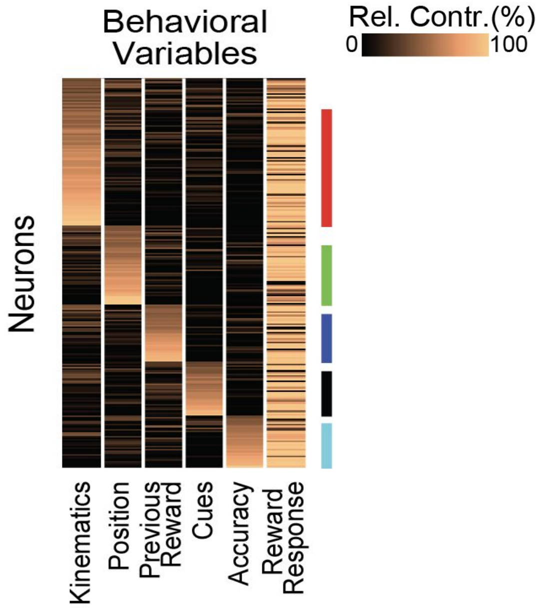

Engelhard *et al*. *Nature* 2019

Similar takeaway: Kremer *et al J Neuro* 2020

Previous accounts of DA heterogeneity: different projection-defined DA neurons for different “types” of prediction errors

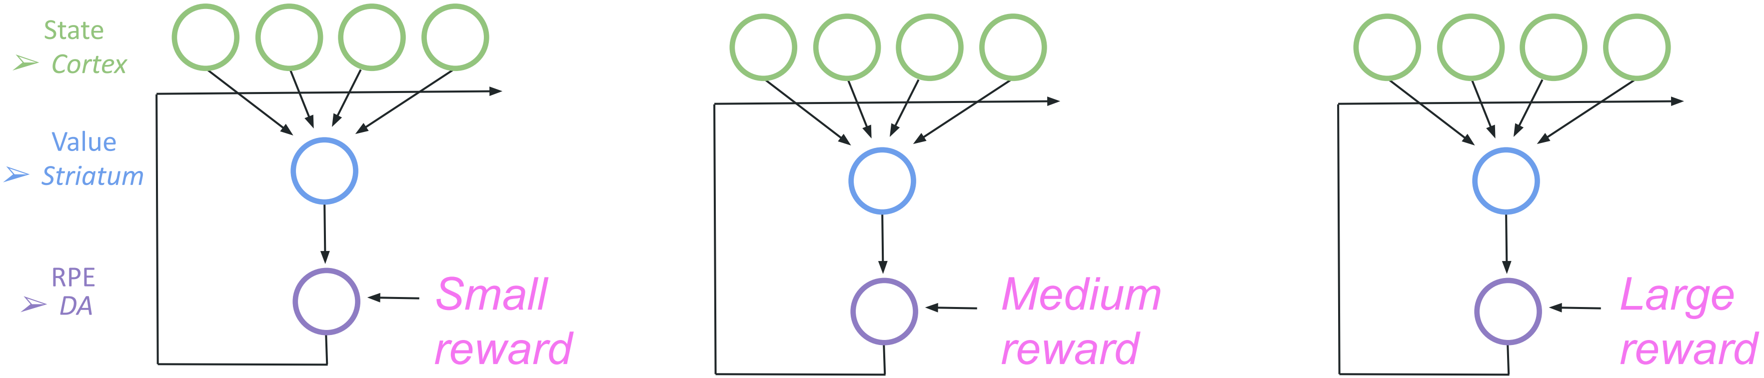

Dabney *et al Nature* 2020

Previous accounts of DA heterogeneity: different projection-defined DA neurons for different “types” of prediction errors

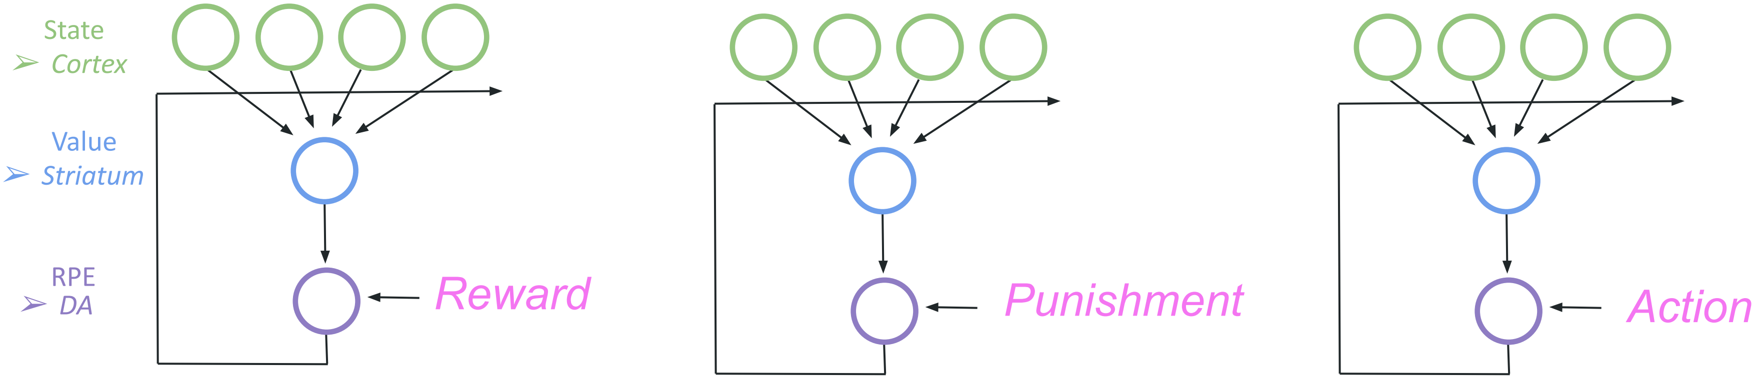

Daw *et al Neural Network*, 2002; Gershman *et al J Neuro* 2009; Lloyd & Dayan *Brain Behav Func* 2016; Gardner et al *Proc Biol Sci* 2018; Akiti *et al Neuron* 2022; Greenstreet *et al bioRxiv* 2022; Lindsey & Litwin-Kumar *arXiv* 2022

Previous accounts of DA heterogeneity: different projection-defined DA neurons for different "types" of prediction errors

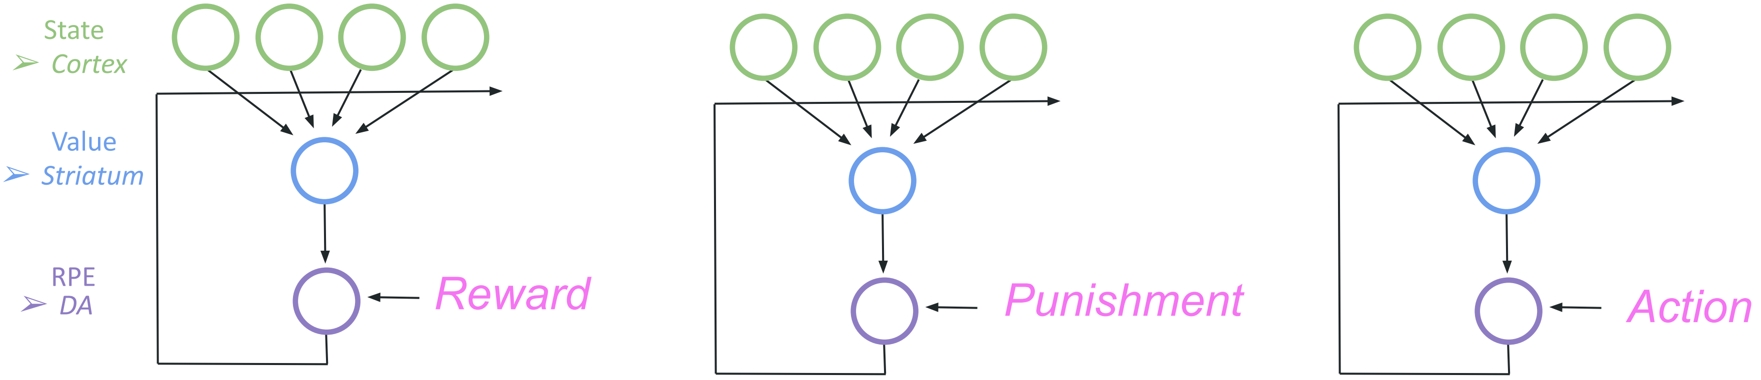

Does not account for:

1. uniform reward response
2. heterogeneous coding of task variables

Previous accounts of DA heterogeneity: different projection-defined DA neurons for different “types” of prediction errors

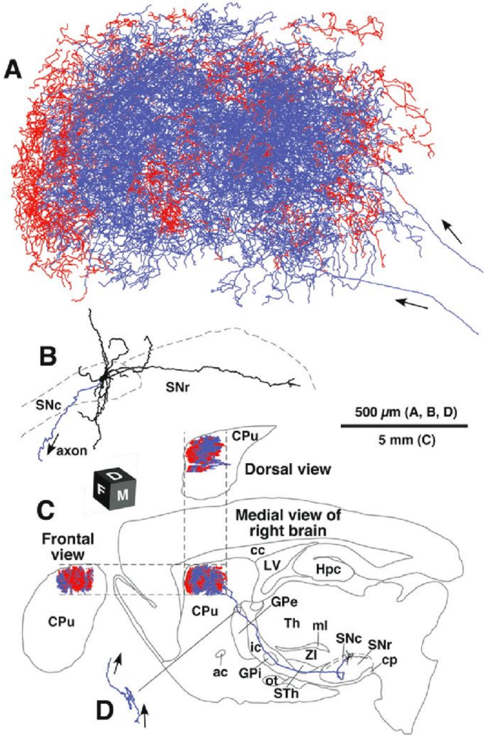

Does not account for anatomy:

3. Broad dopamine arborization
4. Volume transmission

A different mapping between learning algorithms and the brain: Feature-specific RPE model

Scalar RPE Model

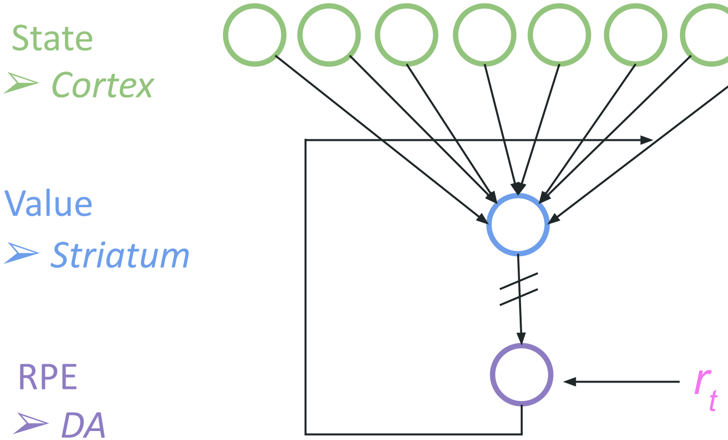

$$\delta_t = r_t + \gamma V(s_{t+1}) - V(s_t)$$

Feature-specific RPE Model

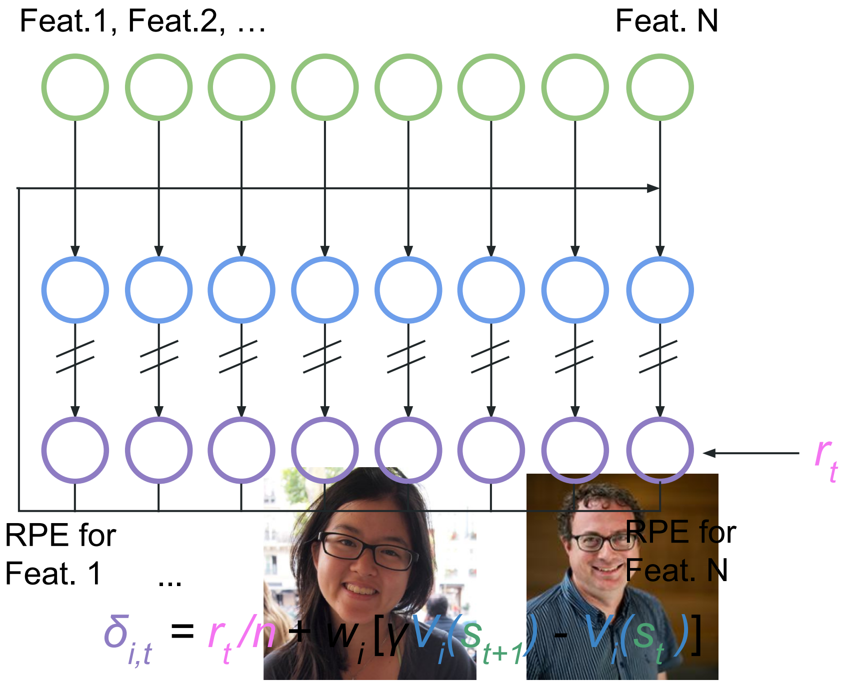

$$\delta_{i,t} = r_t / n + w_i [\gamma V_i(s_{t+1}) - V_i(s_t)]$$

Rachel Lee , Nathaniel Daw

Implies that dopaminergic heterogeneity tells us the features being used to predict reward ("state representation")

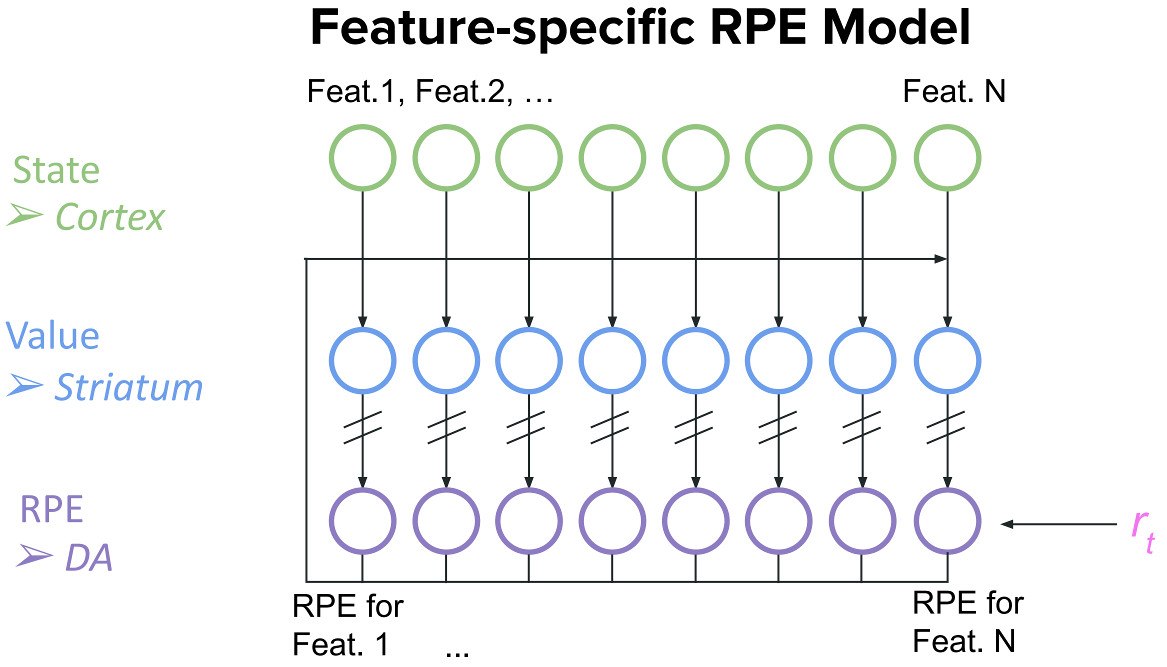

✅ Heterogenous coding of task variables
✅ Task variable coding as feature-specific RPE
✅ Homogeneous response to reward

Across striatal subregions, different corticostriatal "modules" with different "reward" inputs (that calculate different **types** of prediction errors)

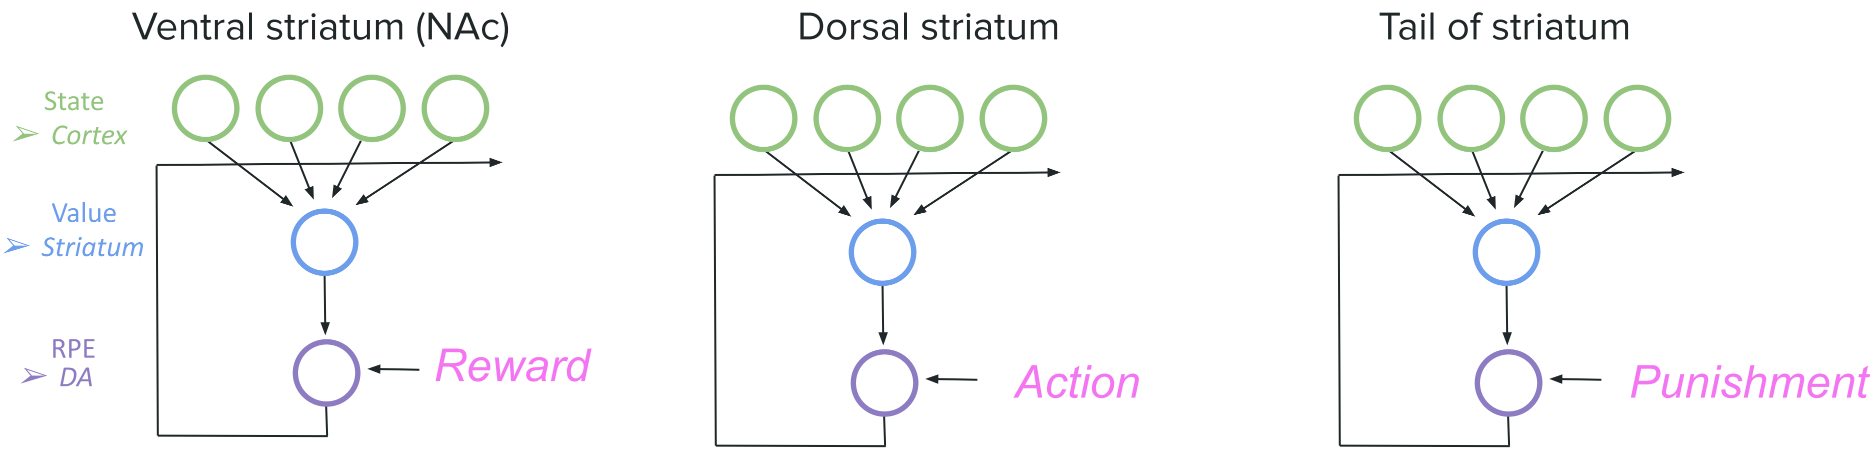

Daw *et al Neural Network*, 2002; Gershman *et al J Neuro* 2009; Lloyd & Dayan *Brain Behav Func* 2016; Gardner et al *Proc Biol Sci* 2018; Akiti *et al Neuron* 2022; Greenstreet *et al bioRxiv* 2022; Lindsey & Litwin-Kumar *arXiv* 2022

Within a striatal subregion, dopaminergic heterogeneity reflects the state representation

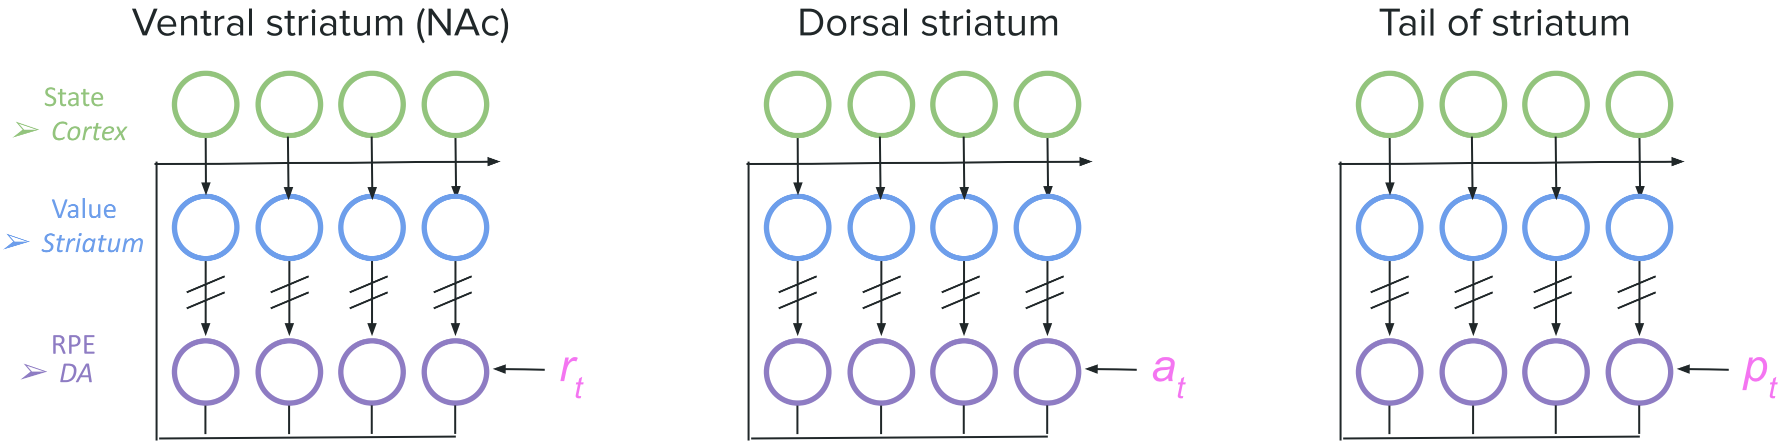

Two “axes” of dopamine heterogeneity:

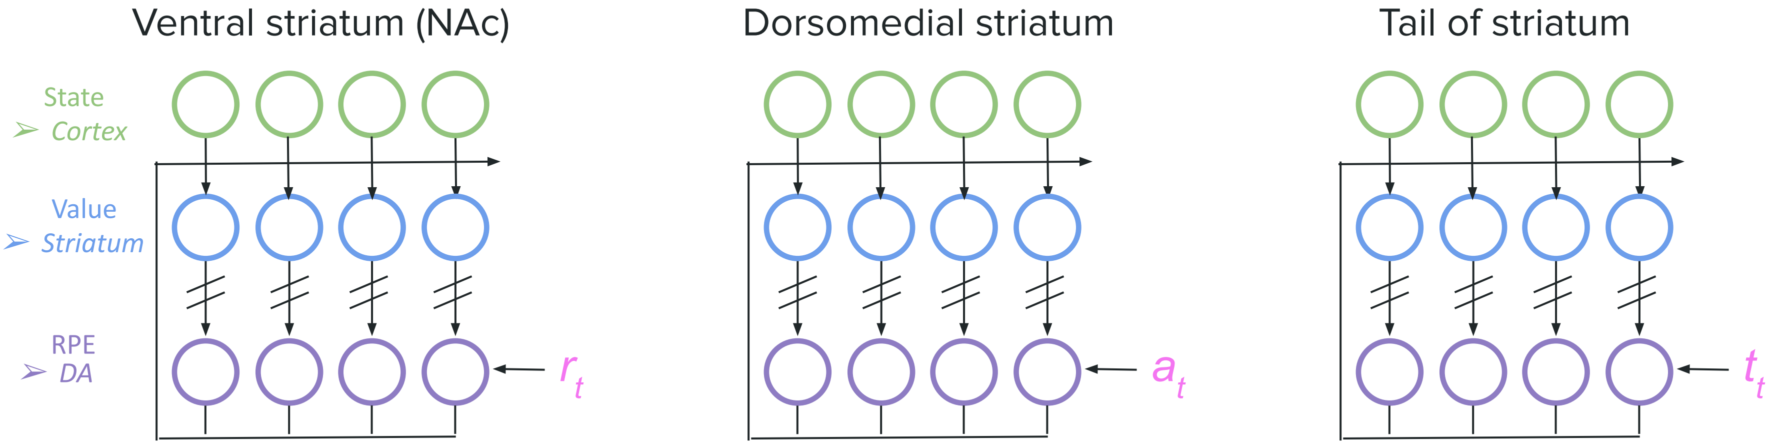

Lee *et al. Nature neuro* 2024

Two “axes” of dopamine heterogeneity:
- Within a module: reflect the state representation

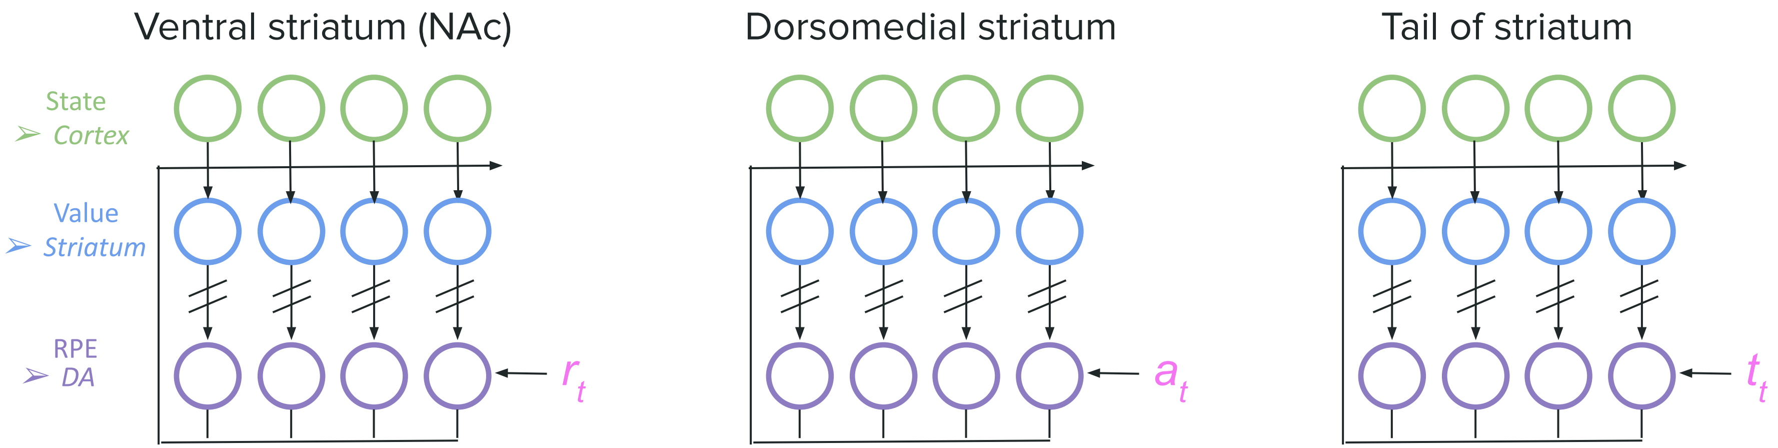

Lee *et al. Nature neuro* 2024

Two "axes" of dopamine heterogeneity:

- Within a module: reflect the state representation
- Across modules: different "types" of prediction errors

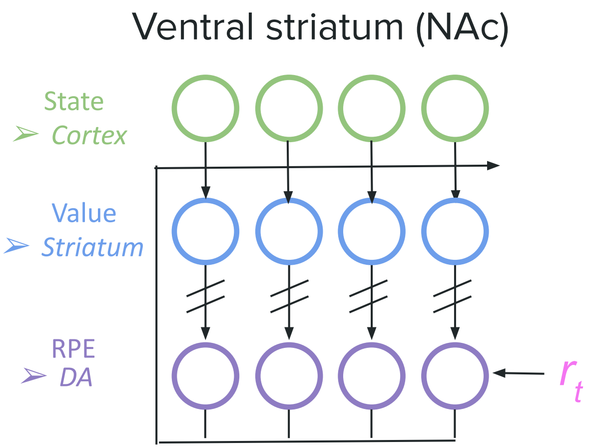

Lee *et al. Nature neuro* 2024

Outline for today

1. **Anatomy** overview
2. How does the **dopamine system** regulate the basal ganglia?
3. What is the function of the **direct** vs **indirect pathways** of the basal ganglia?

Classic view: striatal pathways oppositely modulate behavior

direct pathway

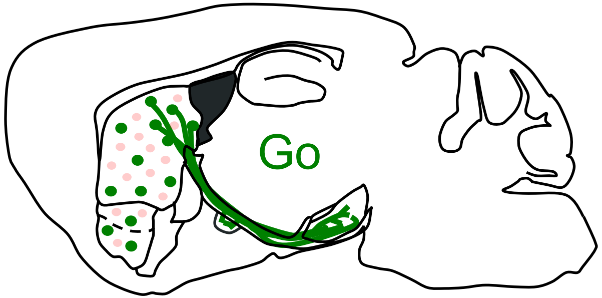

Direct pathway promotes behavioral output ("Go")

indirect pathway

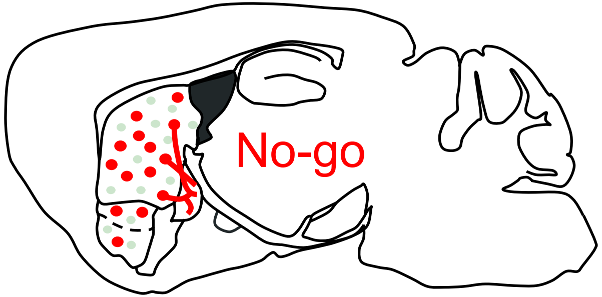

Indirect pathway suppresses behavioral output ("No-go")

Many influential optogenetic studies support the go/ no-go model

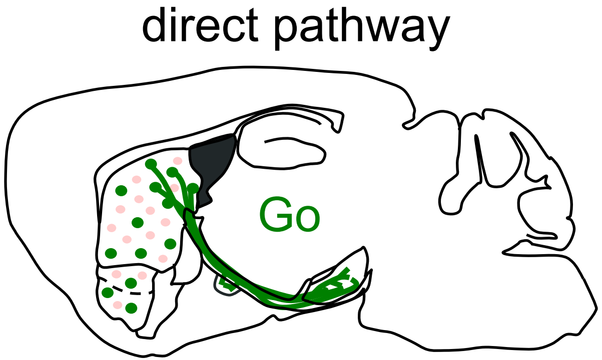

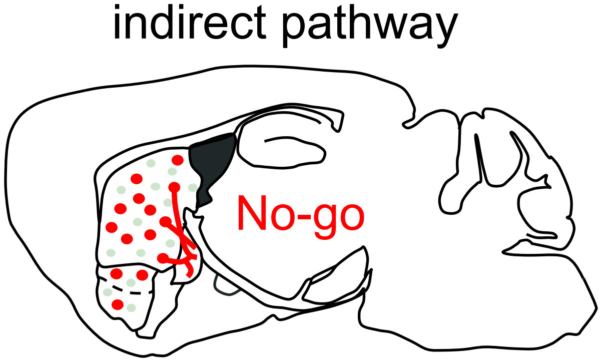

Movement (go/no-go)
- Kravitz *et al.* (2010)
- Roseberry *et al.* (2016)
- Bartholomew *et al.* (2016)
- Bakhurin *et al.* (2019)

Reinforcement (repeat/cease)
- Lobo *et al.* (2010)
- Kravitz *et al.* (2012)
- Yttri *et al.* (2016)

Rotations (left/right)
- Tai *et al.* (2012)
- Nonomura *et al.* (2018)
- Lee *et al.* (2020)

Despite these opposing effects on behavior, both pathways are <u>usually</u> <u>coactive</u> during movement

direct pathway

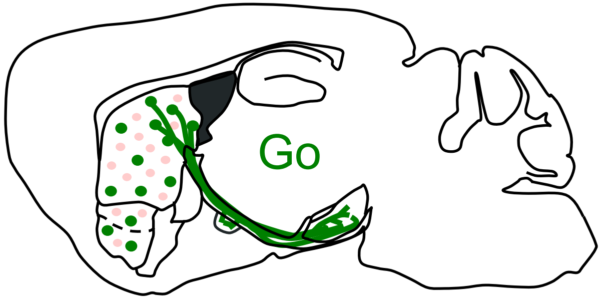

indirect pathway

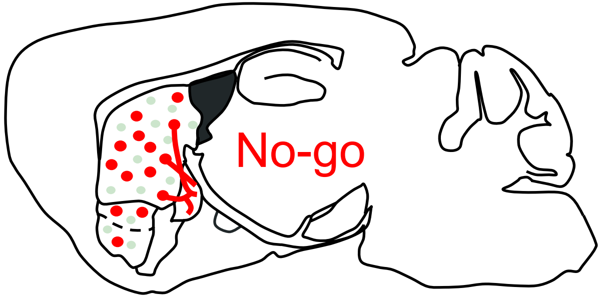

Cui, G. *et. al.* (2013)
Parker, J. G. *et. al*. (2018)
Markowitz, J.E. *et al*. (2018)

However, the pathways may transiently have divergent activity

direct pathway

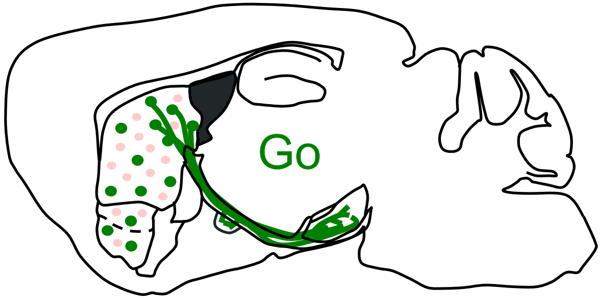

indirect pathway

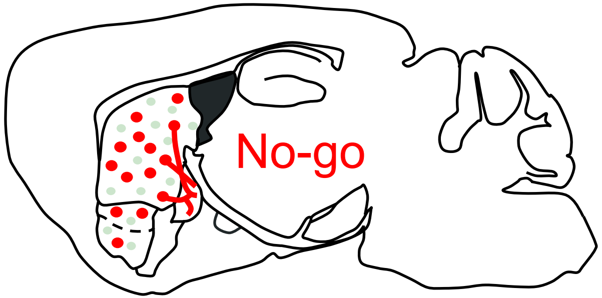

Nonomura *et al.* (2018)
Cruz *et al.* (2022)
Fearey *et al.* (2024)
Bruce *et al.* (2024)
Faust *et al.* (2025)

Why similar activity and transient divergence?

One idea: Direct pathway promotes the chosen action, indirect pathway suppresses un-chosen actions

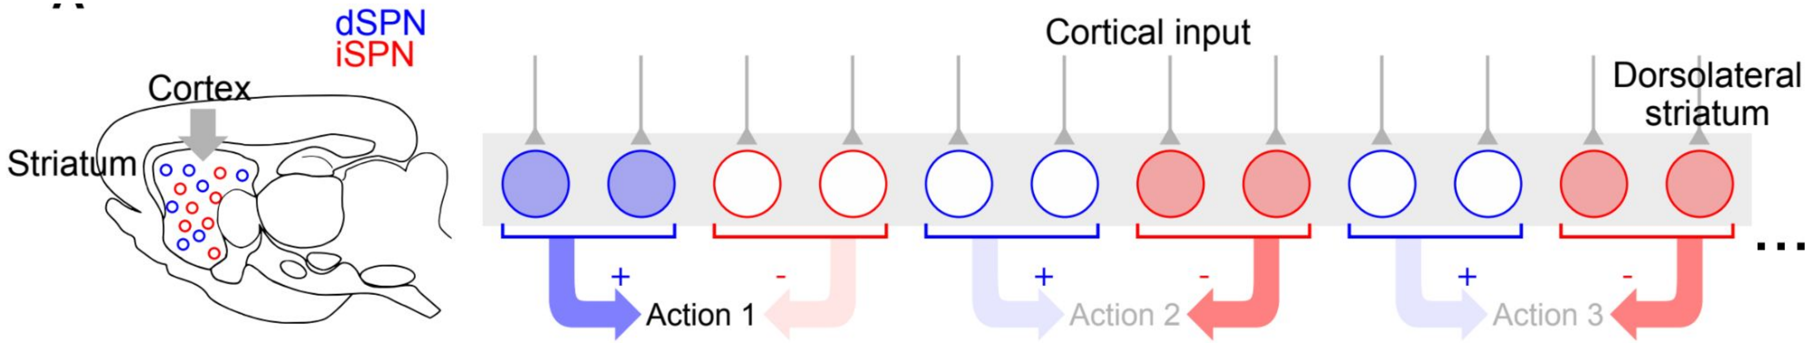

Mink 1996

Background info: opposite dopamine-mediated plasticity in direct & indirect pathway

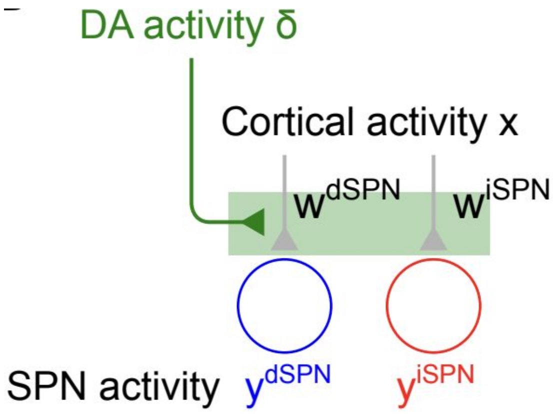

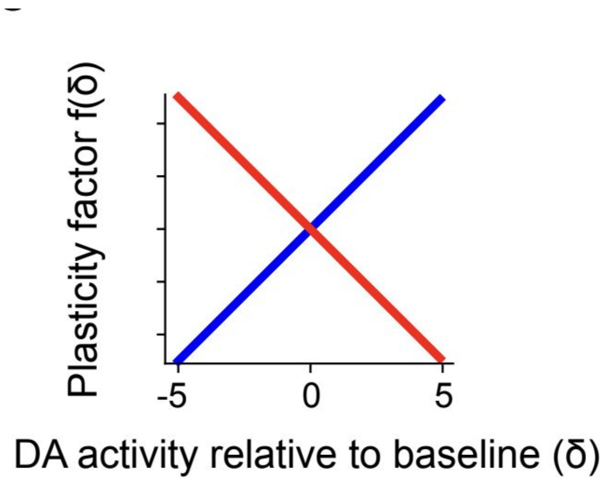

Background info: D1R are needed for LTP in direct pathway neurons (through an AMPA/NMDA mechanism)

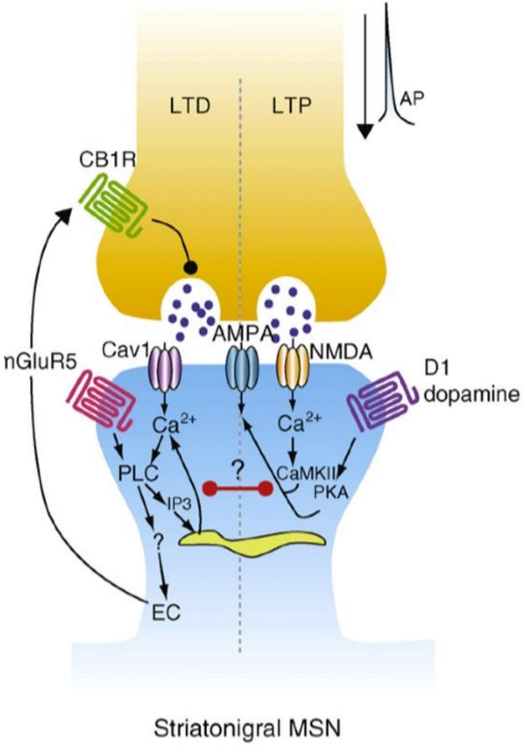

Background info: D2R are needed for LTD in indirect pathway neurons (through a presynaptic endocannabinoid mechanism)

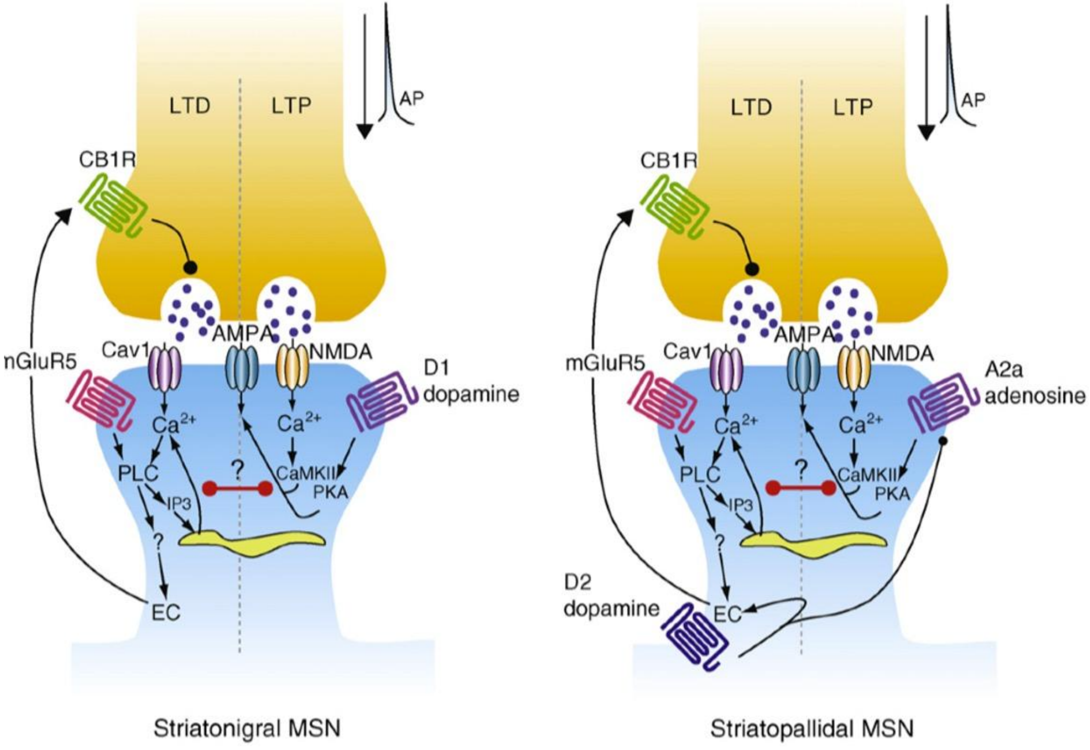

Dopamine increases (after reward) will reinforce the selected action

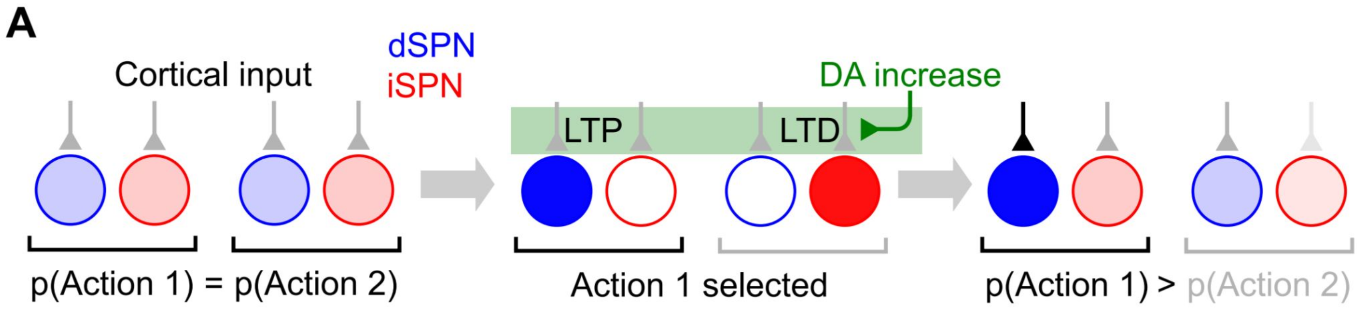

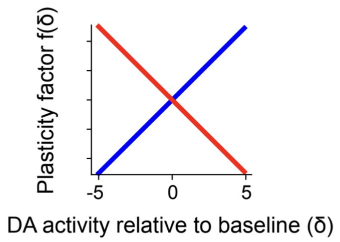

Lindsey. *et. al.* (2024)

Problem: dopamine <u>de</u>creases (after failure) will also reinforce the selected action

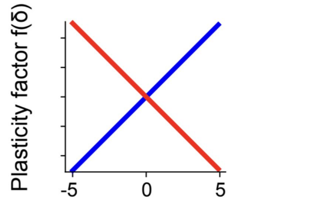

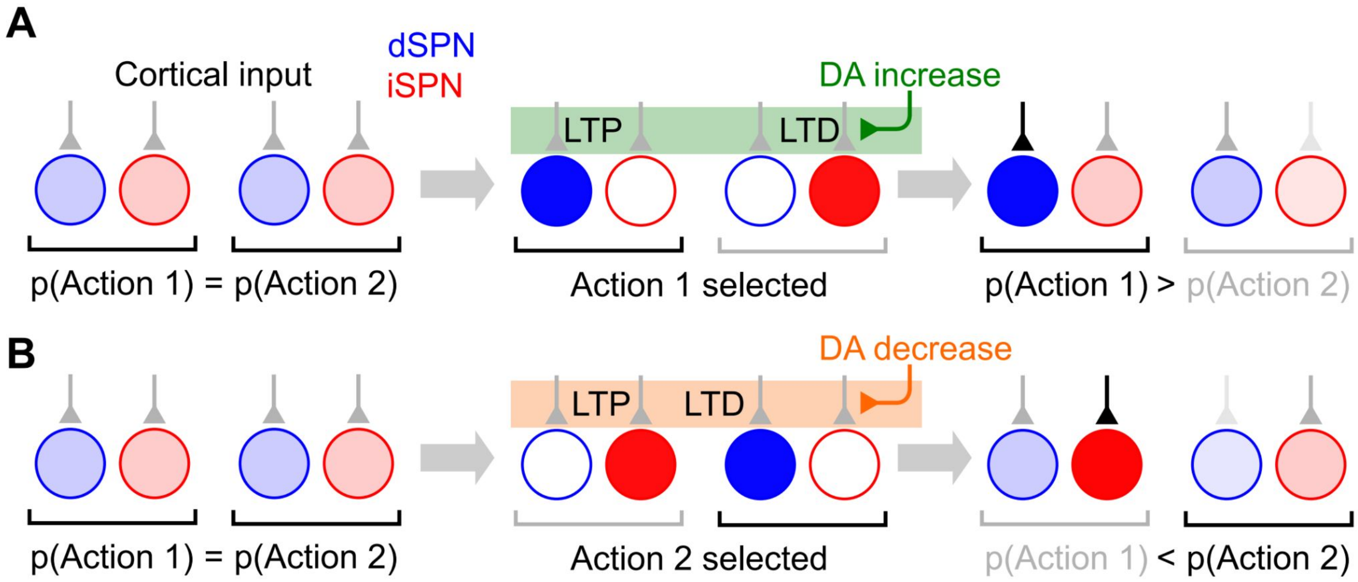

Lindsey. *et. al.* (2024)

Solution: efference copy of the <u>chosen</u> action is encoded in each pathway AFTER action is performed

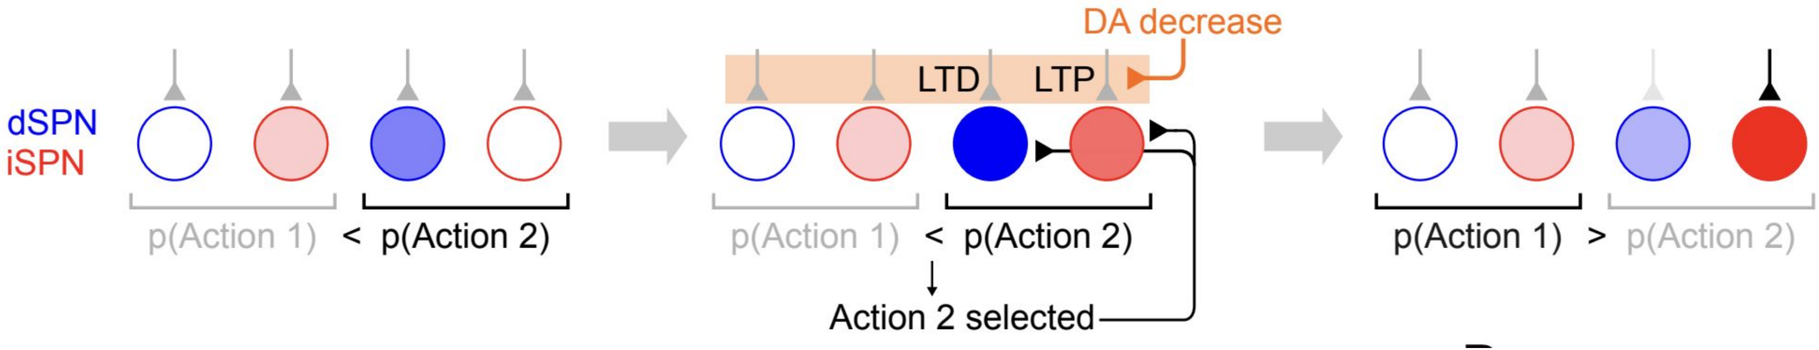

- Co-active after the action is chosen (efference copy to support action value update)
- Direct pathway will show stronger activity than indirect pathway for the chosen action before it is chosen (to allow it to happen)
    - *opposite coding of action value*

Does optogenetic stimulation support the go/ no go model?

Many influential optogenetic studies support the go/ no-go model

direct pathway

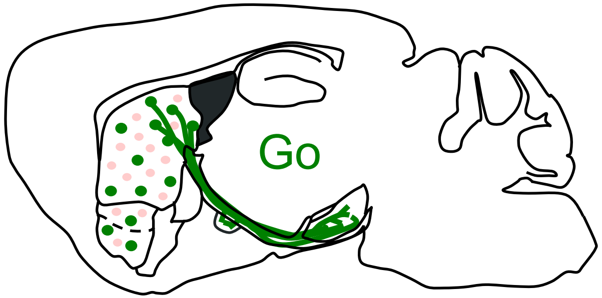

indirect pathway

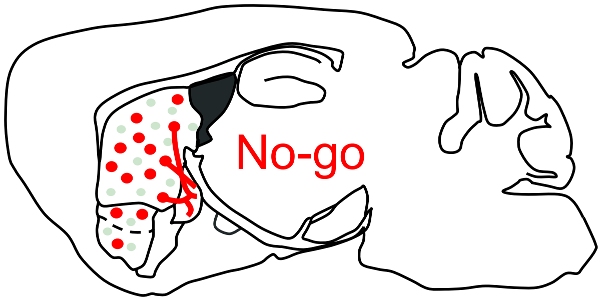

Movement (go/no-go)
- Kravitz *et al.* (2010)
- Roseberry *et al.* (2016)
- Bartholomew *et al.* (2016)
- Bakhurin *et al.* (2019)

Reinforcement (repeat/cease)
- Lobo *et al.* (2010)
- Kravitz *et al.* (2012)
- Yttri *et al.* (2016)

Rotations (left/right)
- Tai *et al.* (2012)
- Nonomura *et al.* (2018)
- Lee *et al.* (2020)

But no (or few) reports of opposing effects of inhibition

direct pathway

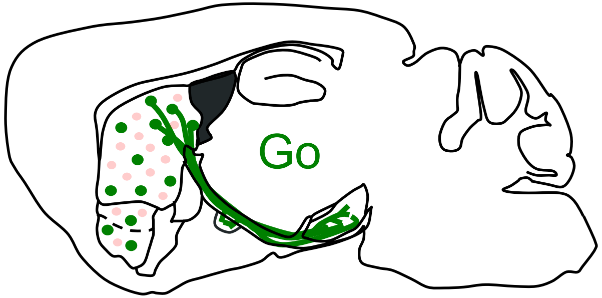

indirect pathway

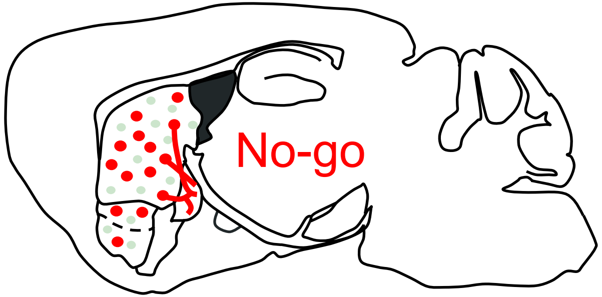

Movement (go/no-go)
- Kravitz *et al.* (2010)
- Roseberry *et al.* (2016)
- Bartholomew *et al.* (2016)
- Bakhurin *et al.* (2019)

Reinforcement (repeat/cease)
- Lobo *et al.* (2010)
- Kravitz *et al.* (2012)
- Yttri *et al.* (2016)

Rotations (left/right)
- Tai *et al.* (2012)
- Nonomura *et al.* (2018)
- Lee *et al.* (2020)

Why?

Pathway specific optogenetic inhibition in striatum (dorsomedial striatum, DMS)

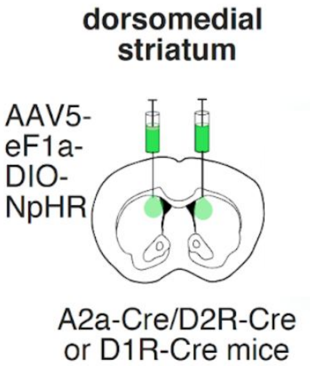

Bolkan*, Stone* *et al.*
Nature Neuro 2022

No behavior effect while running in a 2-D corridor

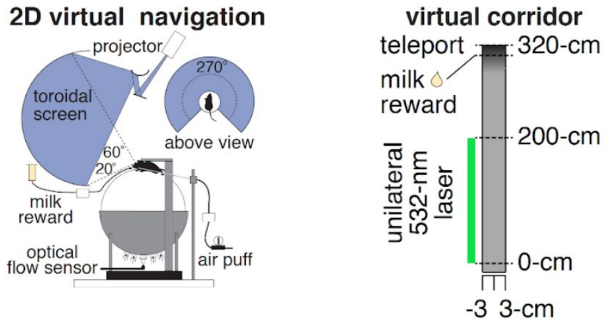

Bolkan*, Stone* *et al.*
Nature Neuro 2022

Hypothesis: pathways may only have opposing control of behavior in the context of the "function" of each subregion

Dorsolateral striatum (DLS):
sensory-motor learning

Dorsomedial striatum (DMS):
goal-directed behavior;
visual decision-making

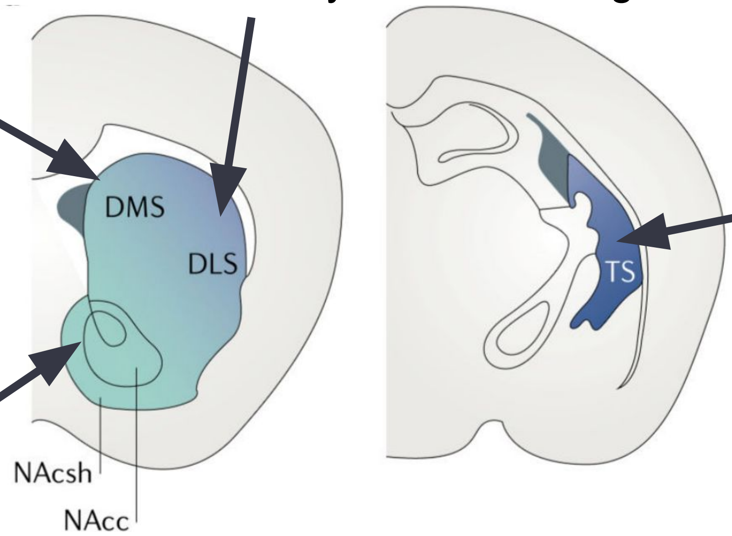

Tail of striatum (TS):
threat learning

Nucleus accumbens (NAc):
Pavlovian associations

Cox & Witten
NRN 2019

Could DMS pathways oppositely be controlling visual decision-making, not movements?

Could DMS pathways oppositely be controlling visual decision-making, not movements?

accumulation-of-evidence (AoE) task

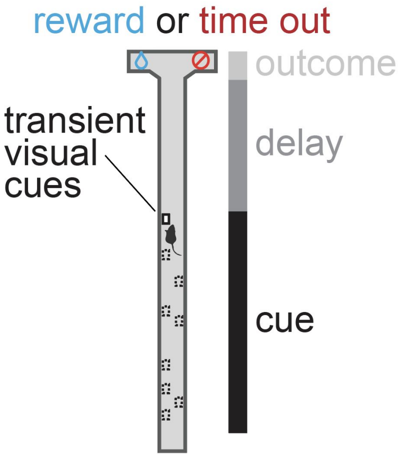

Pinto*, Koay*,...., Witten, Tank, Brody 2019

Large and opposing effects of DMS pathway-specific inhibition during accumulation of evidence

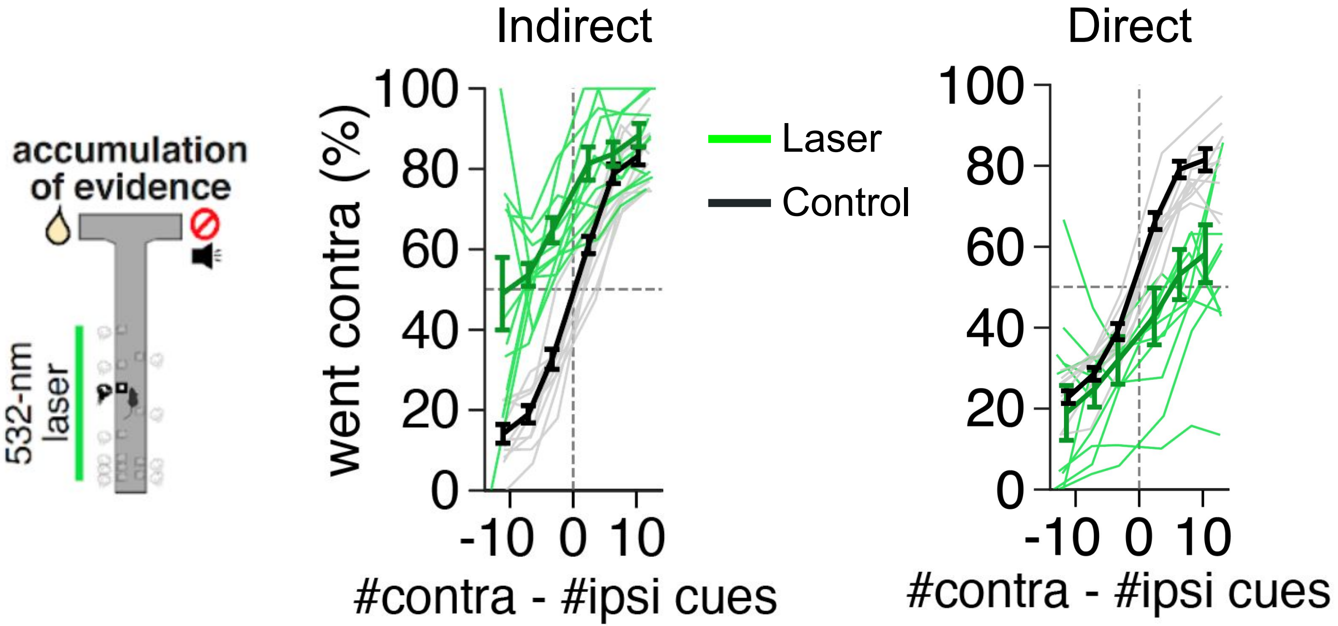

Little effect of inhibition of either pathways in NAc

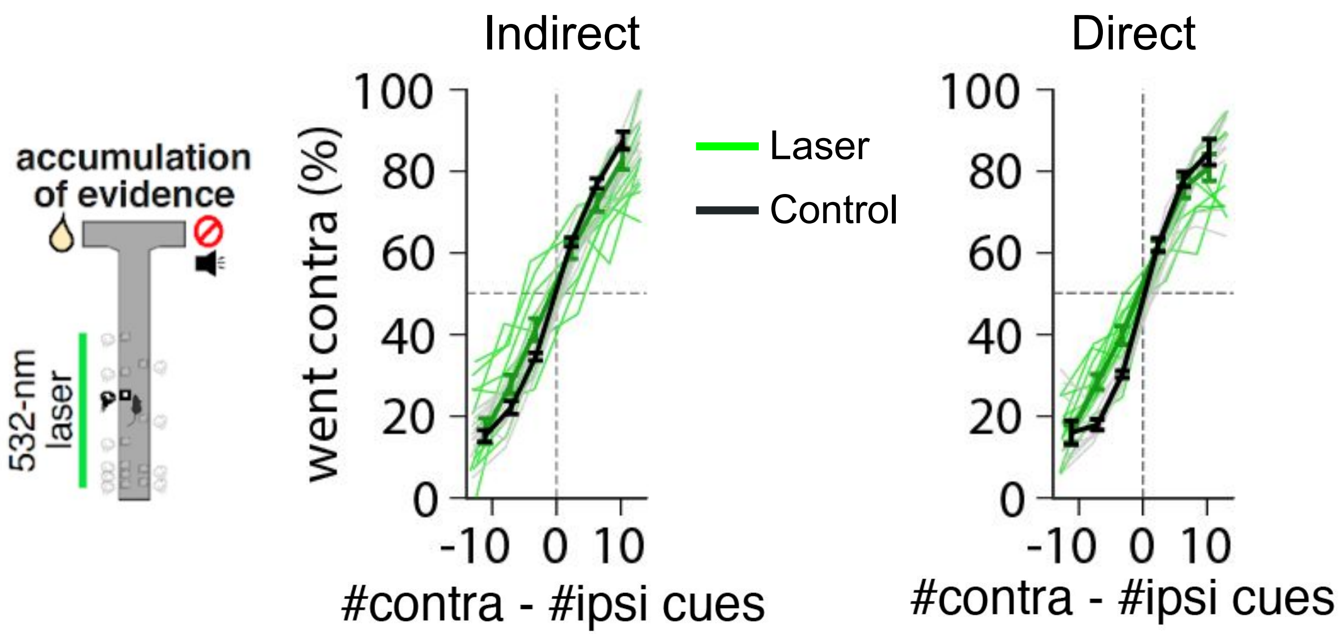

Little effect of DMS inhibition in sensory-guided tasks that do not involve gradual evidence accumulation

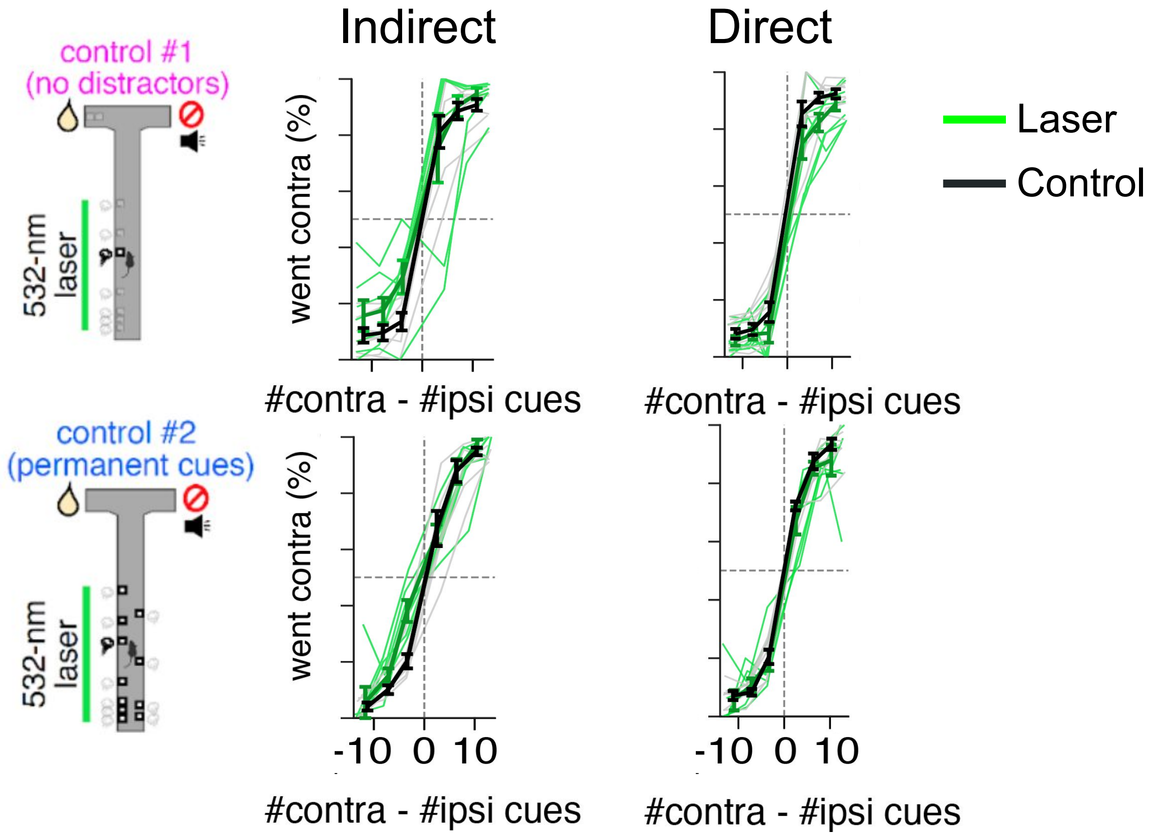

Striatal pathways in DMS oppositely control decisions based on gradually accumulating evidence

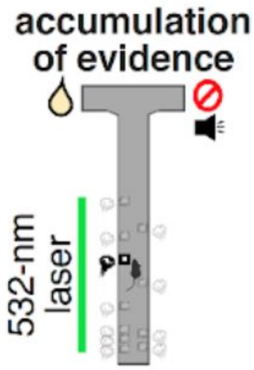

- Not directly on movement
- Not on "easy" decisions that don't involve evidence accumulation
- Not in a "task disengaged" state when animals aren't actually integrating evidence
- No effect in neighboring striatal subregion (NAc)

Interpretation: pathways have opposing control of behavior, in the context of the "function" of each subregion

Dorsolateral striatum (DLS):
sensory-motor learning

Cruz *et al* 2022

Dorsomedial striatum (DMS):
goal-directed behavior;
visual decision-making

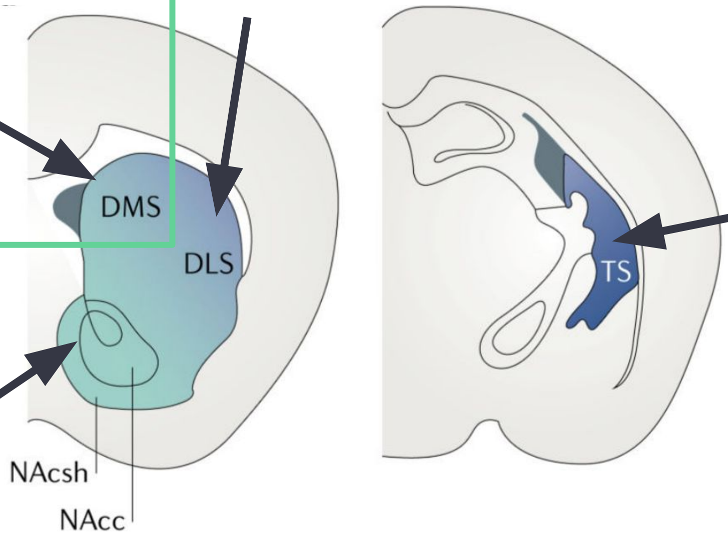

Tail of striatum (TS):
threat learning

Tsutsui-Kimura *et al* 2025

Nucleus accumbens (NAc):
Pavlovian associations

Lowet *et al* 2025

How does pathway-specific inhibition drive opposite effects on decisions? Do the pathways oppositely control cortical activity, as predicted by classic model?

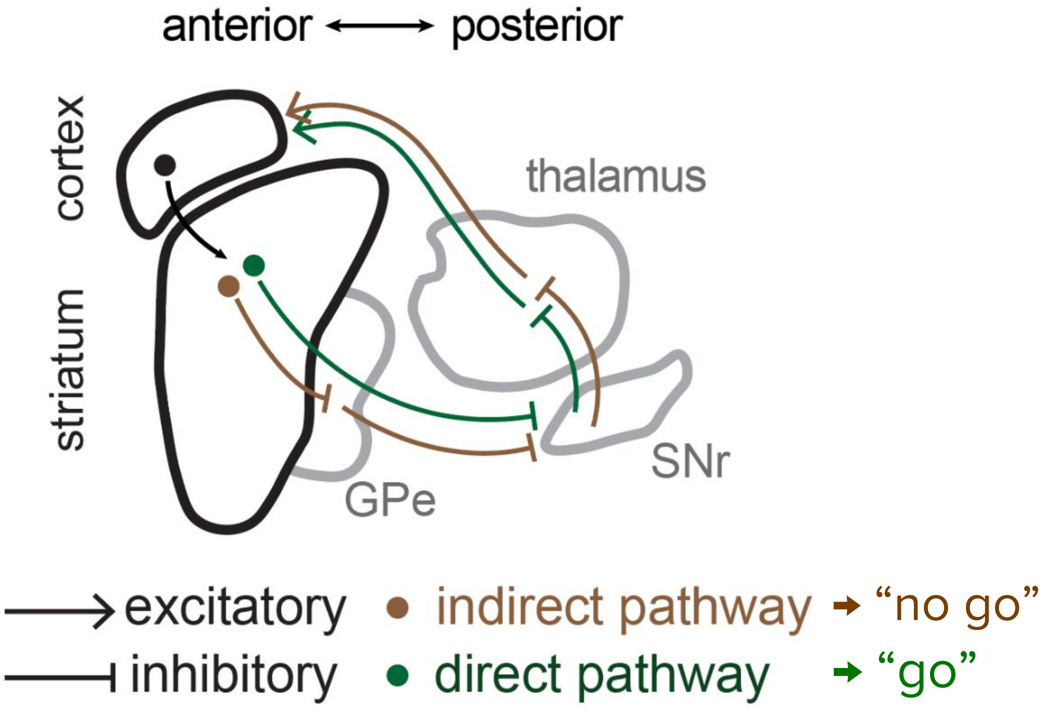

See also Oldenburg & Sabatini 2015

Record from cortex (ACC) while inhibiting direct or direct pathway neurons in DMS during evidence accumulation

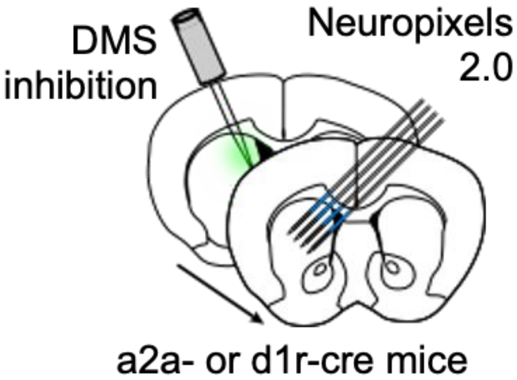

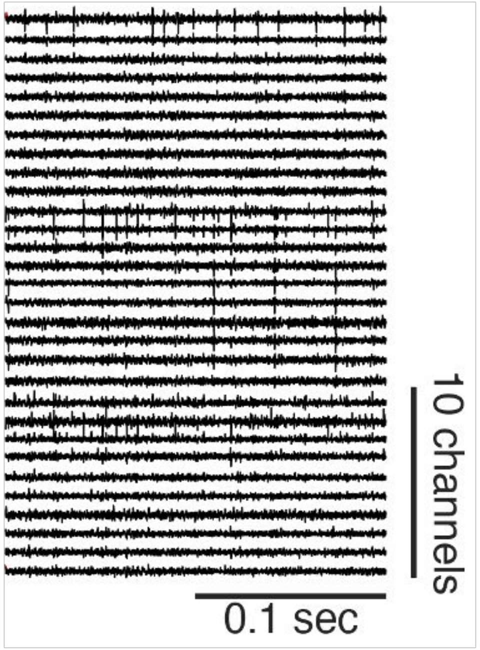

Cho & Bolkan & Brown, *bioRxiv 2025*

Heterogeneous effects of direct or indirect pathway inhibition during evidence accumulation on example cortical neurons

indirect pathway inhibition

example neuron #1 example neuron #2

Time (normalized)

direct pathway inhibition

example neuron #1 example neuron #2

Time (normalized)

Heterogeneous effects of direct or indirect pathway inhibition during evidence accumulation across cortical neurons

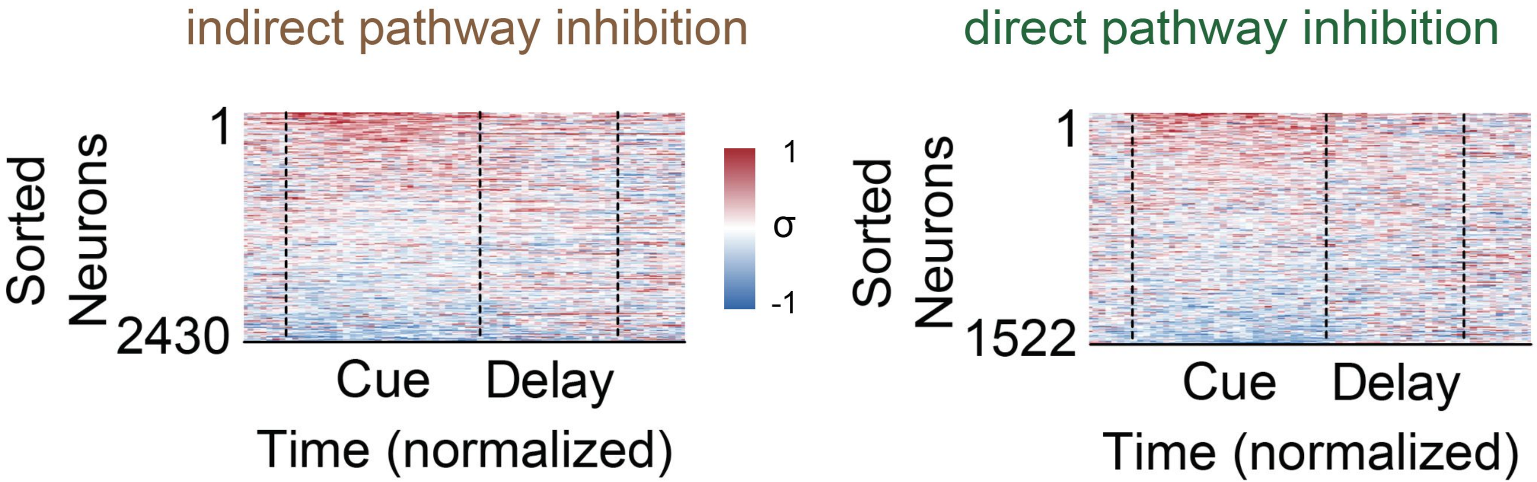

Contradicts the simple idea of opponent control of cortex by the 2 pathways

Hypothesis: opposing effects on decision-making depends on opposing effects specifically on decision-coding neurons

Identify decision-encoding ACC neurons with an encoding model

firing ~ b0 + b1*evidence + b2*choice +b3*laser

Identify decision-encoding ACC neurons with an encoding model

firing ~ b0 + b1*evidence + b2*choice +b3*laser

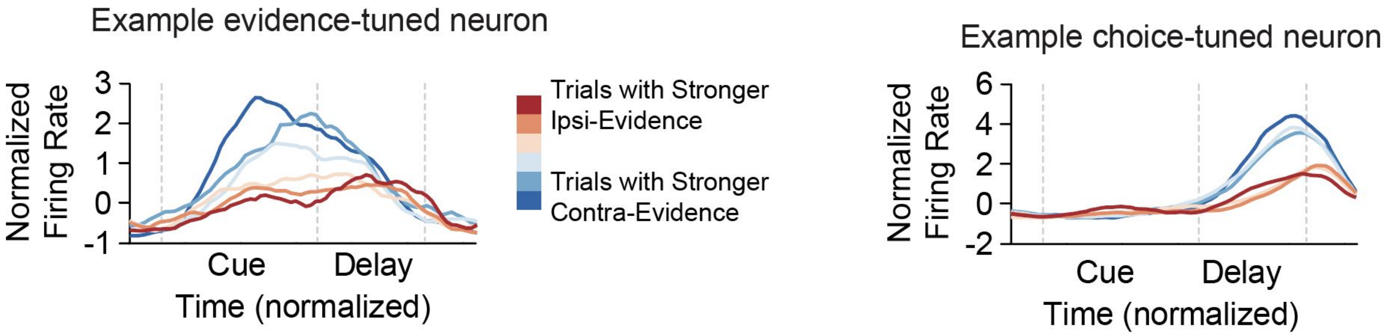

Identify decision-related ACC neurons with an encoding model

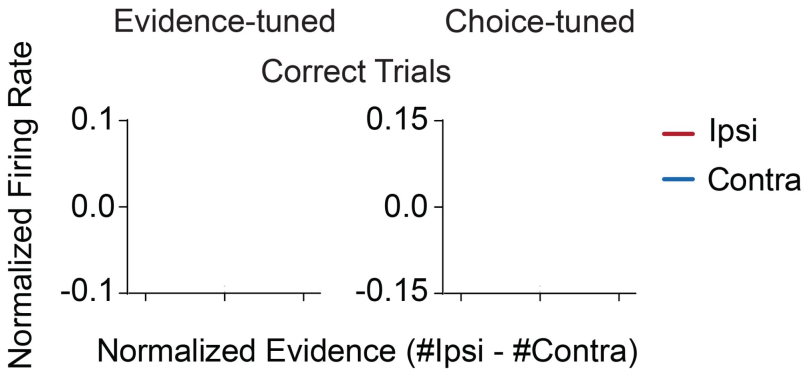

Identify decision-related ACC neurons with an encoding model

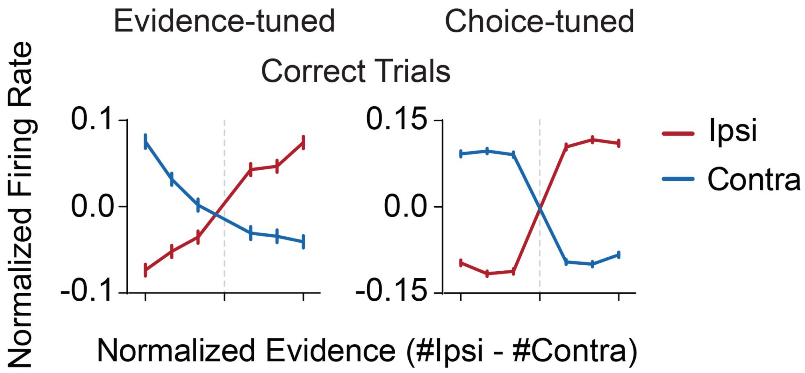

Indirect pathway inhibition activates contra-preferring and inhibits ipsi-preferring evidence coding neurons

Indirect pathway inhibition activates contra-preferring and inhibits ipsi-preferring evidence coding neurons

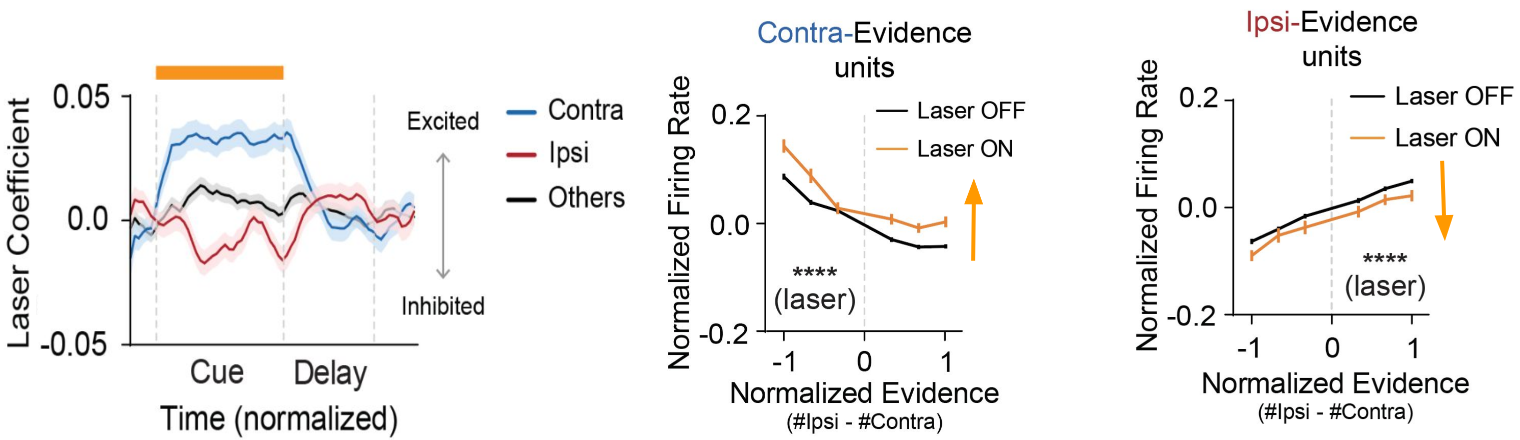

Project population activity onto the evidence decoder axis

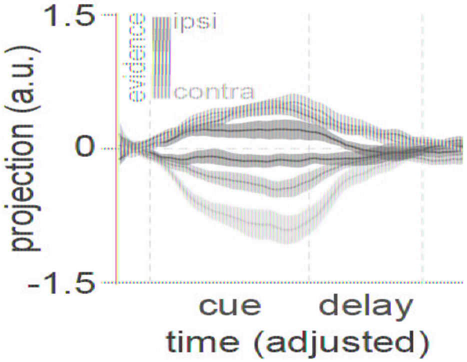

Population-level effect of indirect pathway inhibition is to shift evidence coding in the contralateral direction

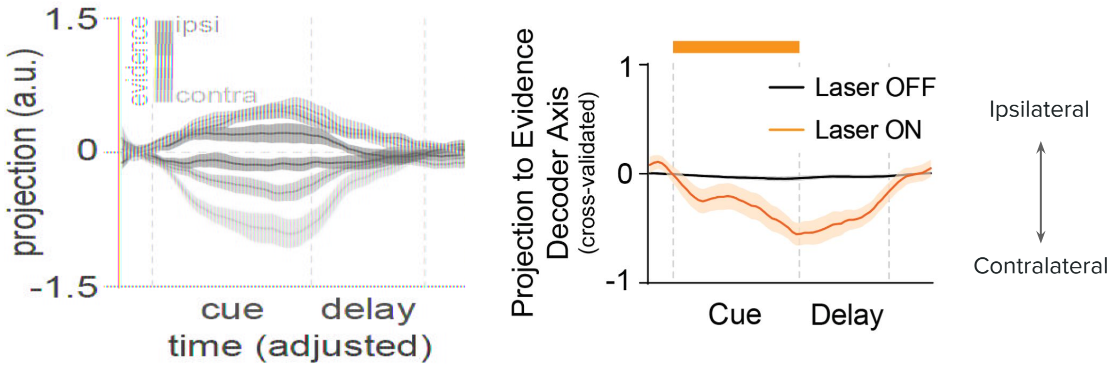

Population-level effect of indirect pathway inhibition is to shift evidence coding in the contralateral direction

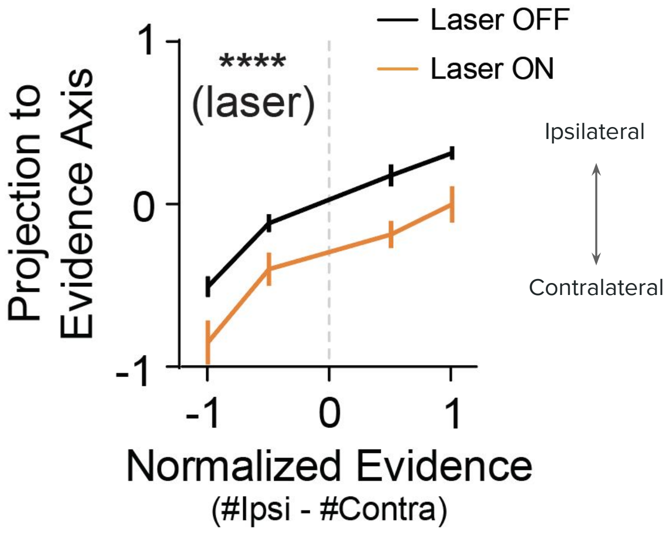

In contrast, no effect of laser on activity along 1st PC axis

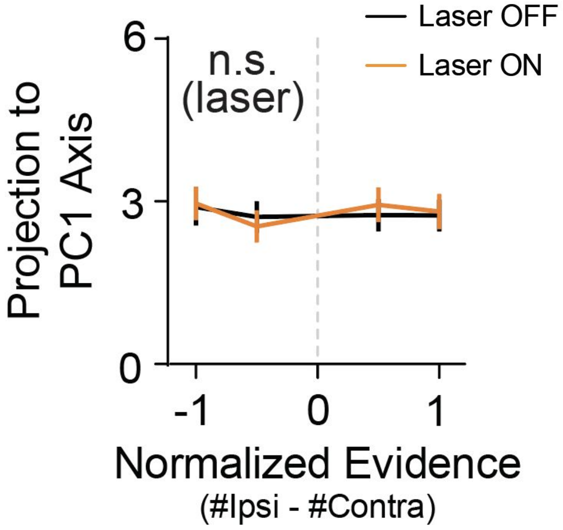

What about direct pathway inhibition?

Direct pathway inhibition has opposite effects on contra vs ipsi preferring evidence neurons

indirect pathway inhibition

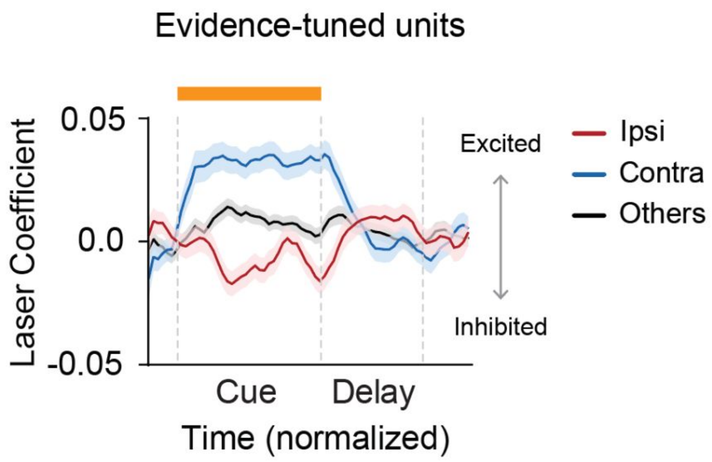

direct pathway inhibition

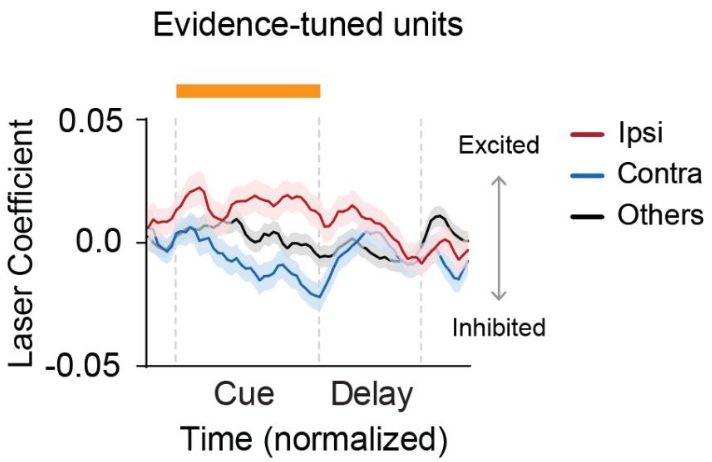

.... And on population coding along the evidence axis

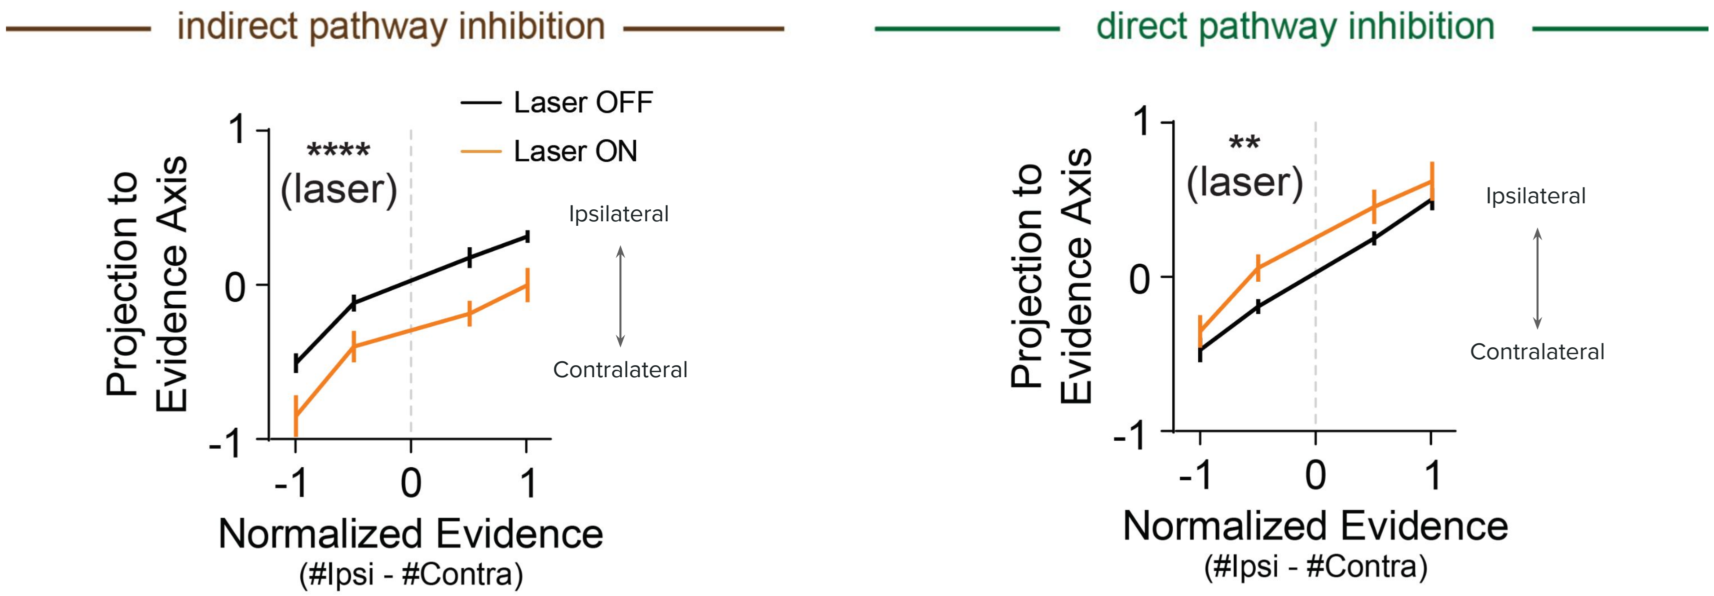

Direct and indirect pathways do not have opposite effects on overall cortical activity, but instead on the coding of accumulated evidence

Are opposite effects on cortical coding also present for choice-tuned neurons?

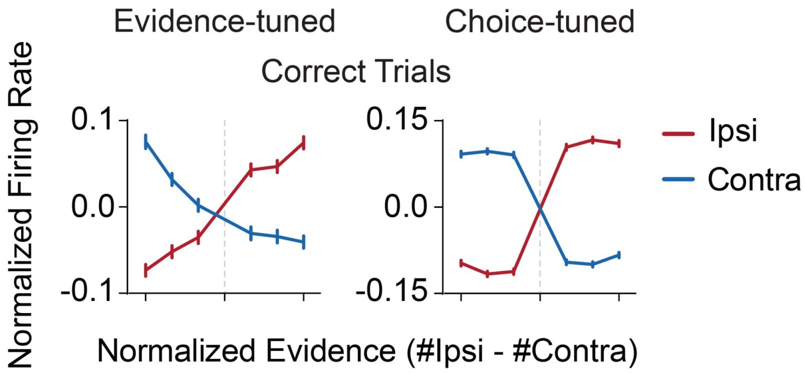

In contrast to evidence-coding neurons, little effect of inhibition on choice-selective neurons

In contrast to evidence-coding neurons, little effect of inhibition on choice-selective neurons

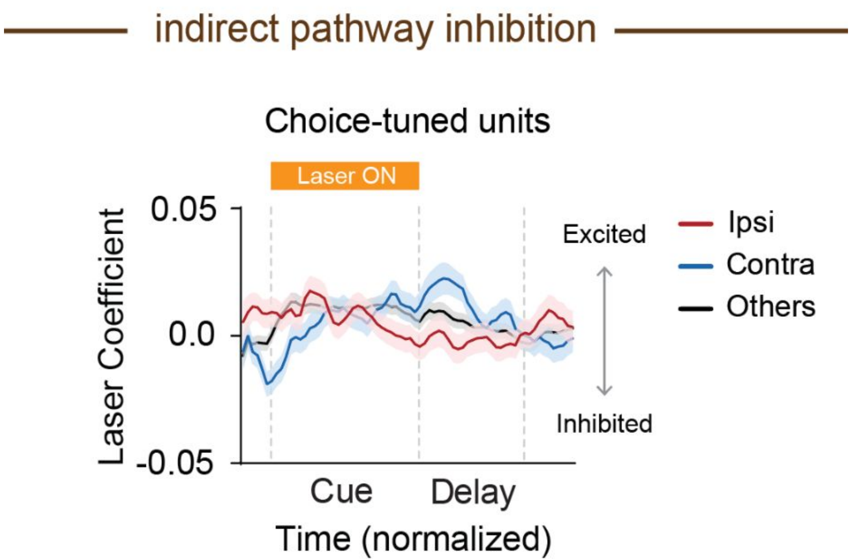

In contrast to evidence-coding neurons, little effect of inhibition on choice-selective neurons

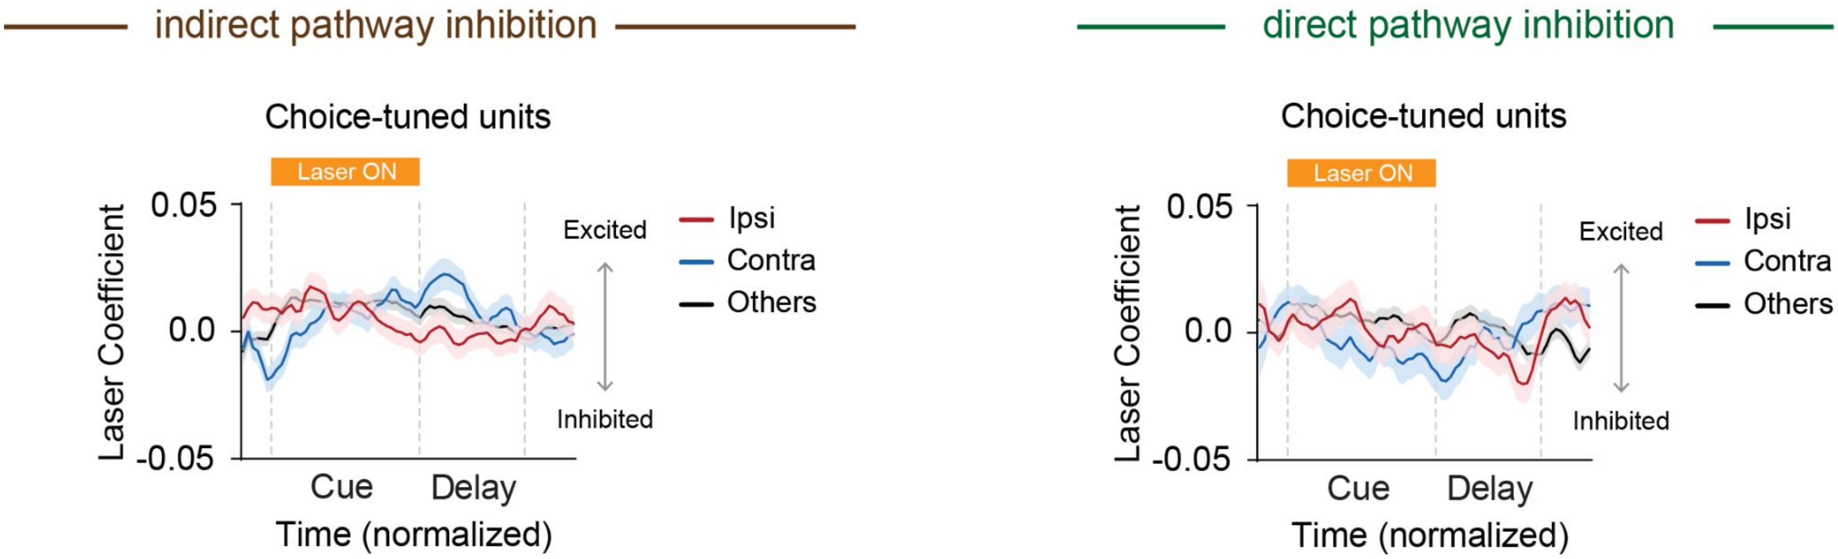

Note: this is after accounting for the selected choice

Classic model: the pathways provide opposite control of behavior, & of cortical activity

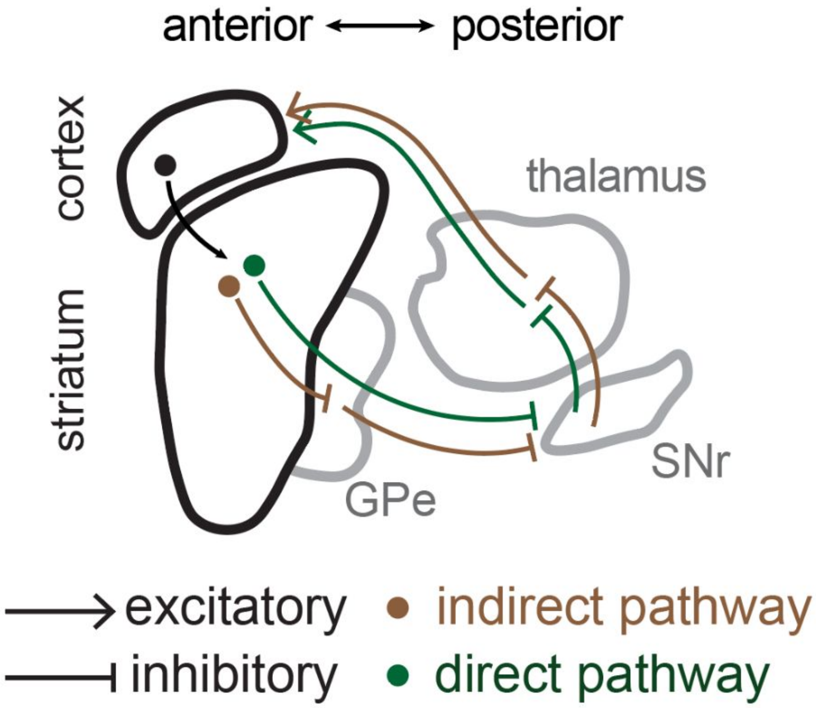

→ "no go"
→ "go"

1) Behavior: opposite effects of DMS pathways, <u>but only during evidence accumulation</u>

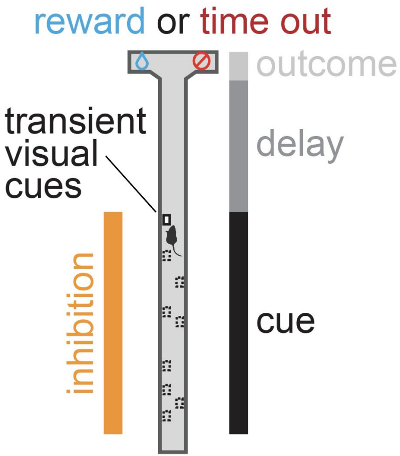

No effect:

- Directly on movement
- On control tasks w/o evidence accumulation
- Of NAc pathway inhibition

2) Neural: opposite effects of DMS pathways specifically <u>on the cortical coding of accumulated visual evidence</u>

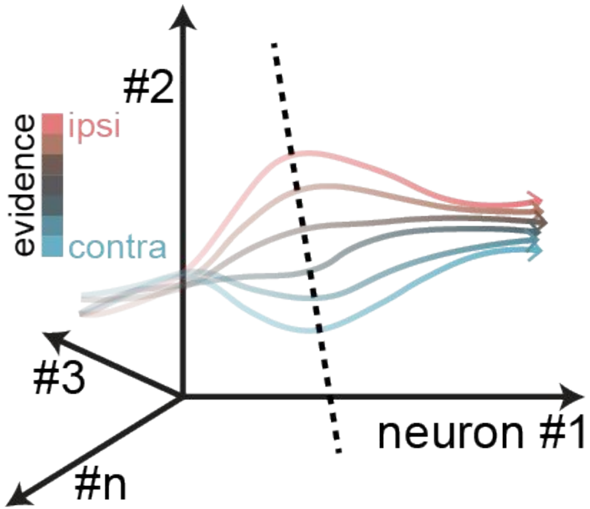

2) Neural: opposite effects of DMS pathways specifically <u>on the cortical coding of accumulated visual evidence</u>

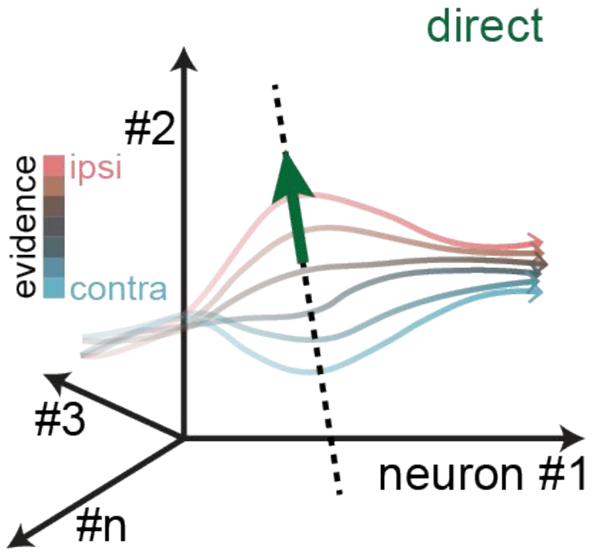

2) Neural: opposite effects of DMS pathways specifically on the cortical coding of accumulated visual evidence

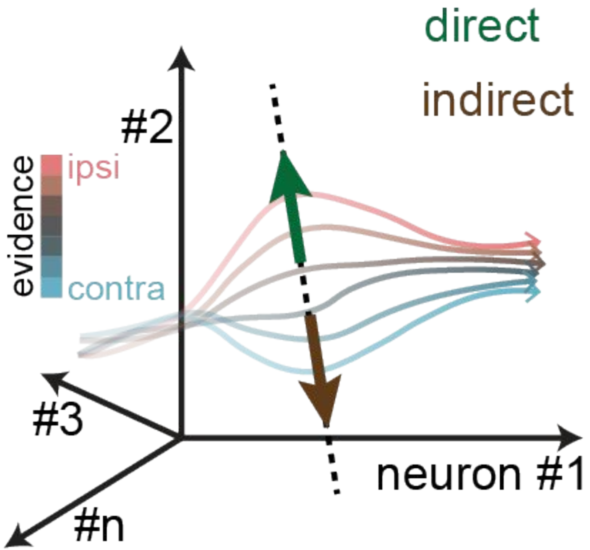

How (and why) are the pathways oppositely shifting
behavior & cortical coding
specifically along the decision-axis?

Evidence integration signals are present beyond ACC (& cortex): in the striatum as well

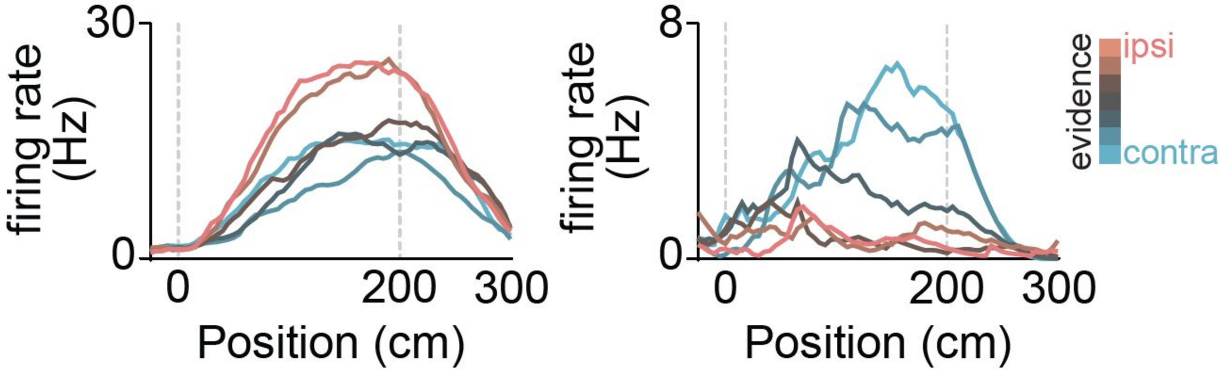

Ding & Gold, 2010
Yartsev et al, 2018
Gupta et al, 2026

Hypothesis: integration may be implemented by the DMS basal ganglia loop

Hypothesis: integration may be implemented by the DMS basal ganglia loop

canonical integration circuit

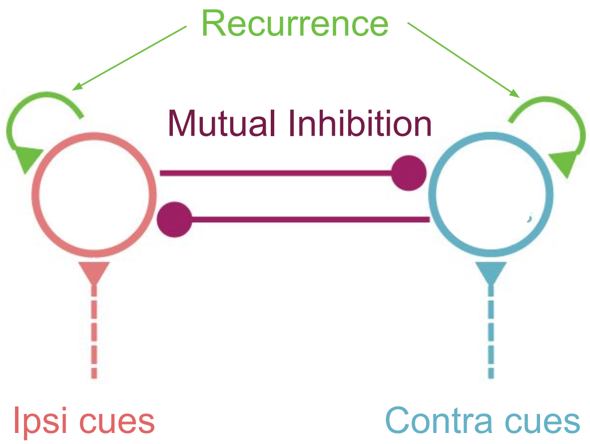

Wong & Wang 2006 J. Neurosci. ;
Usher & McClelland 2001 Psychol. Rev.

Ipsilateral & Contralateral cues

Positive Feedback

Mutual Inhibition

Remapping the integration circuit onto the DMS-thalamo-cortical loop

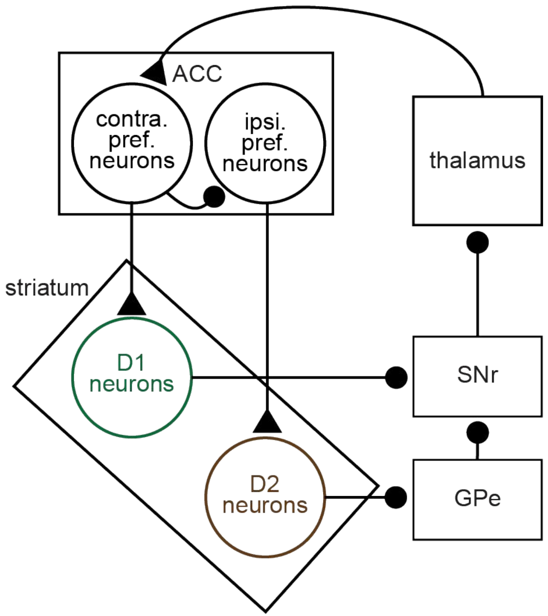

Ipsilateral & Contralateral cues

Positive Feedback

Mutual Inhibition

Remapping the integration circuit onto the DMS-thalamo-cortical loop

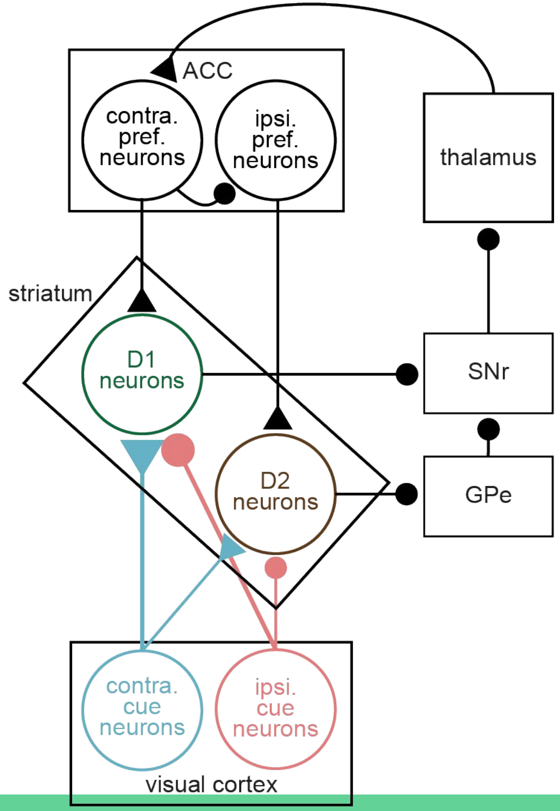

Ipsilateral & Contralateral cues

Positive Feedback

Mutual Inhibition

Remapping the integration circuit onto the DMS-thalamo-cortical loop

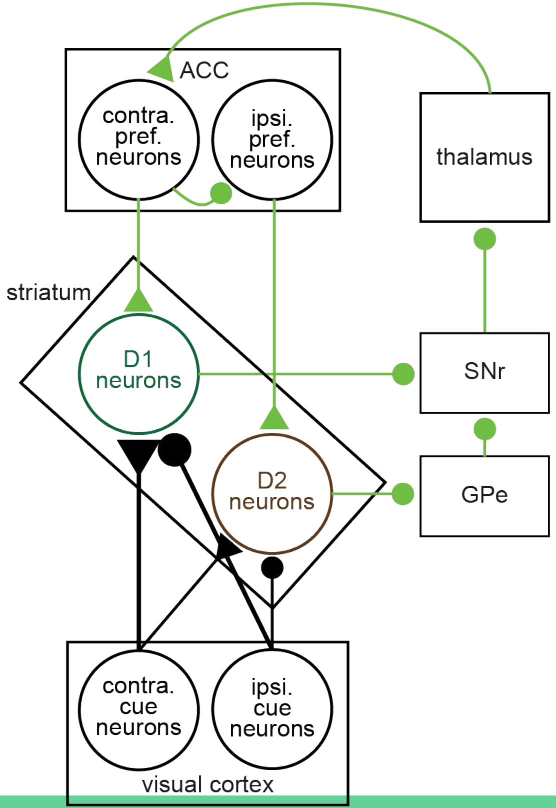

Ipsilateral & Contralateral Cues

Positive Feedback

Mutual Inhibition

Remapping the integration circuit onto the DMS-thalamo-cortical loop

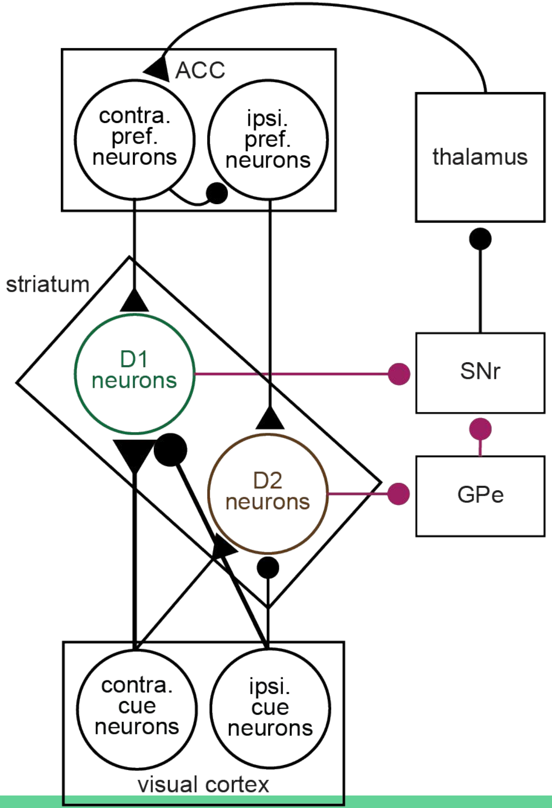

Ipsilateral & Contralateral Cues

Positive Feedback

Mutual Inhibition

This circuit model integrates evidence, producing neurons with contra or ipsi preference throughout the loop

This circuit model integrates evidence, producing neurons with contra or ipsi preference throughout the loop

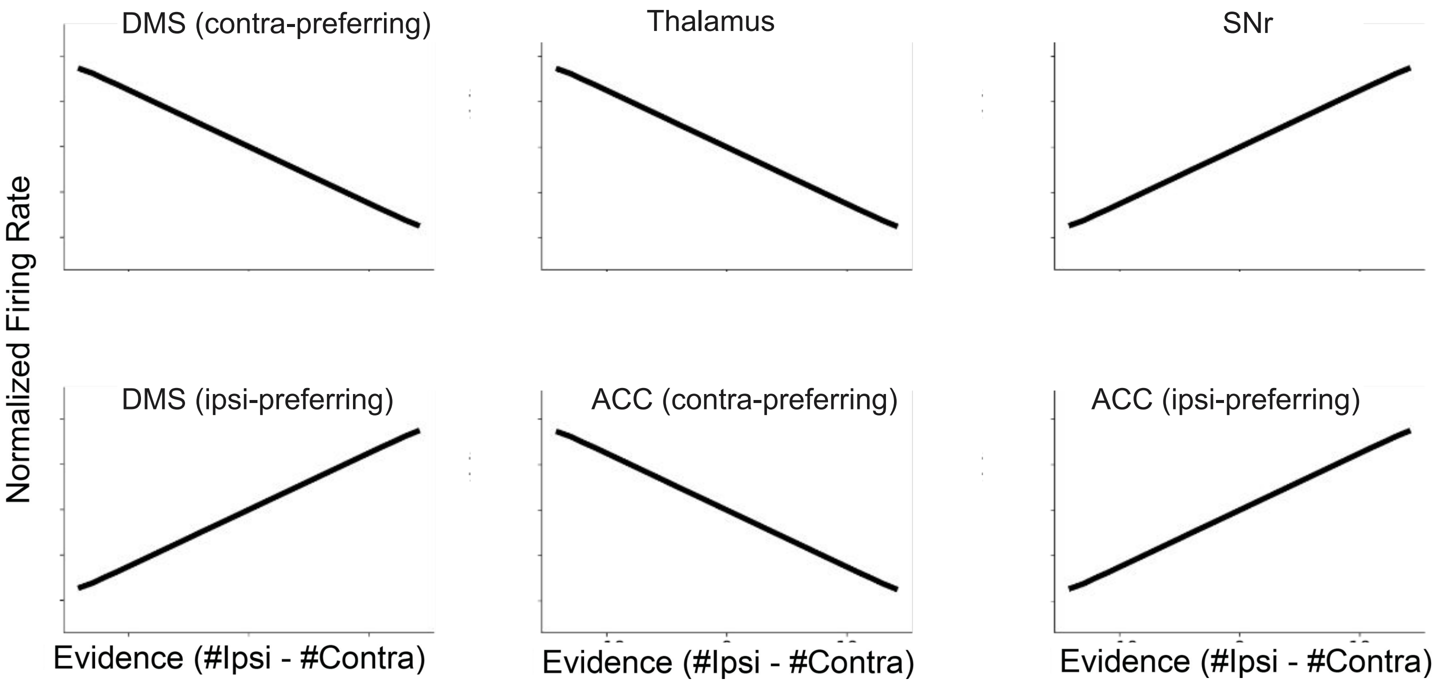

This model replicates our optogenetic data: shifts towards contralateral evidence representation for indirect pathway inhibition

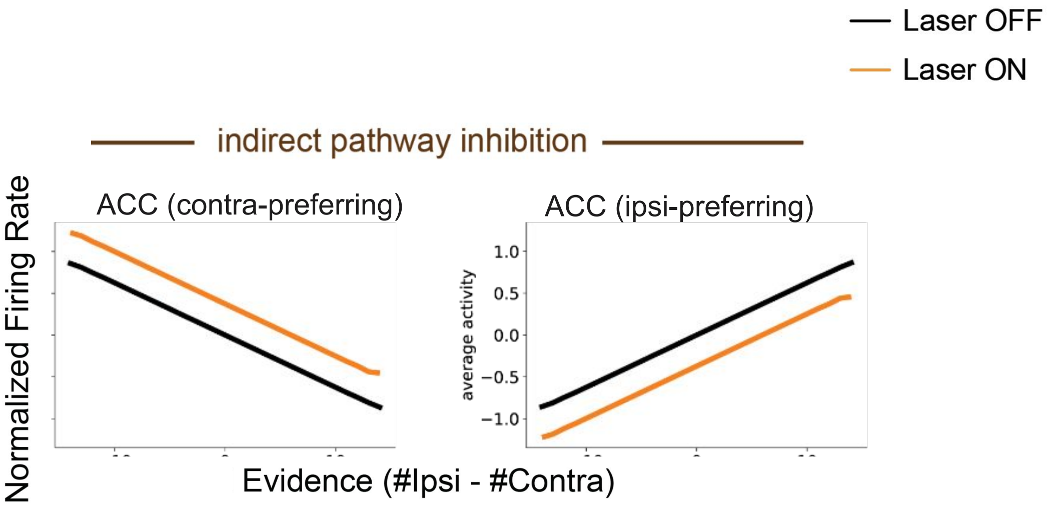

This model replicates our optogenetic data: opposite shifts for direct vs indirect pathway inhibition in ACC

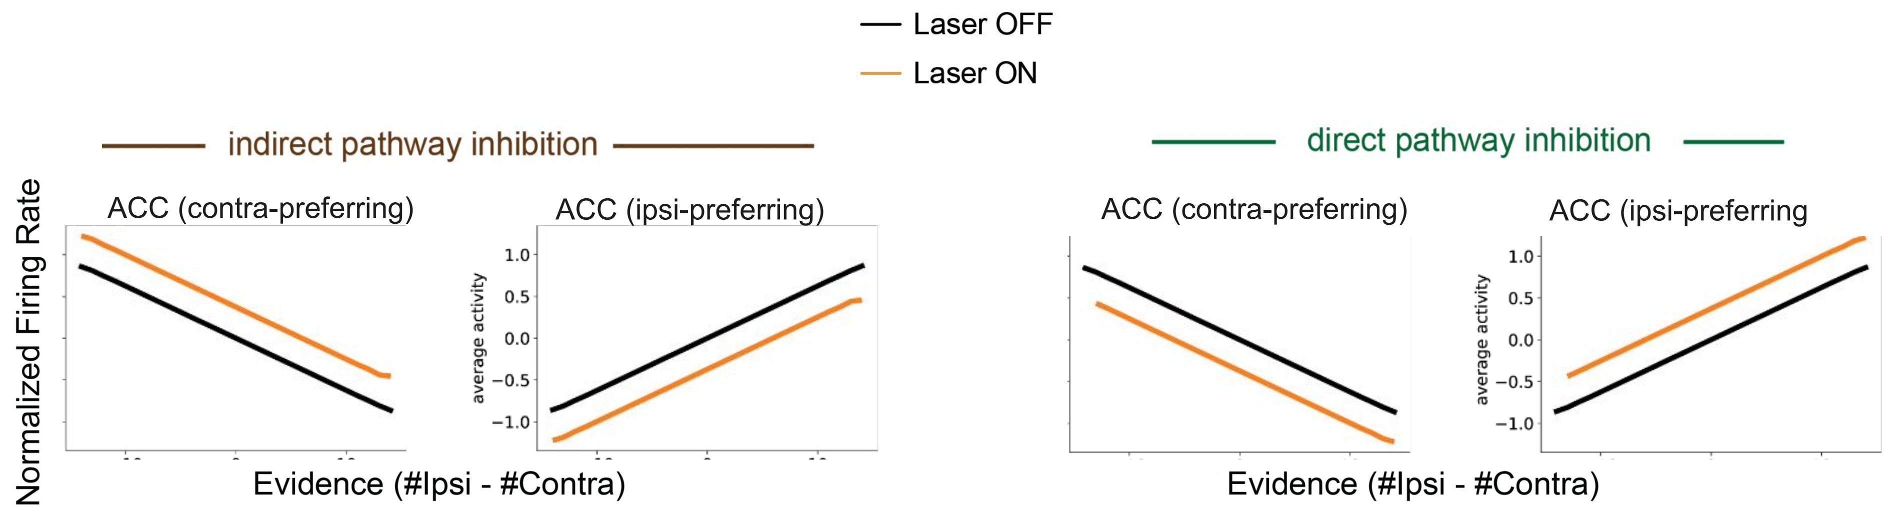

The key: the two DMS populations have opposite evidence preference

The key: the two DMS populations have opposite evidence preference

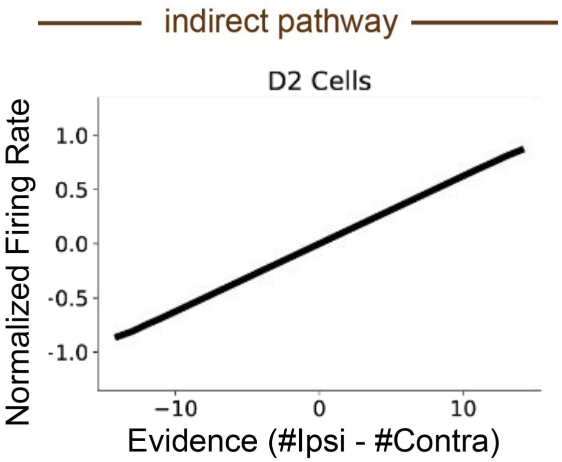

The key: the two DMS populations have opposite evidence preference

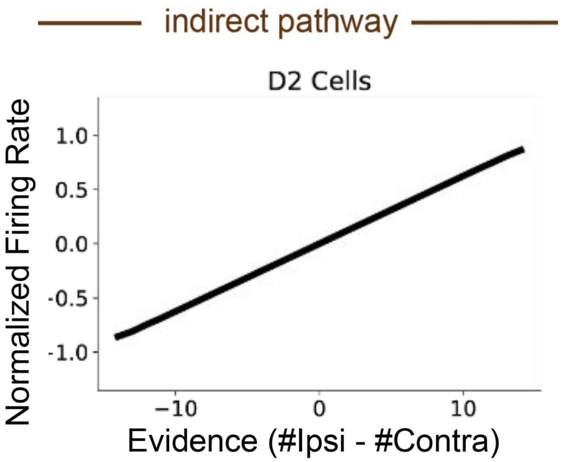

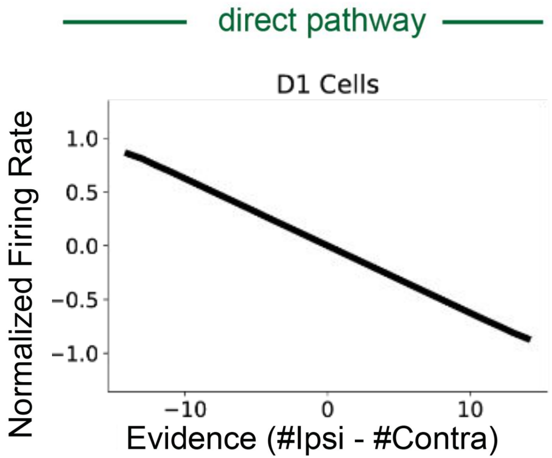

The model achieves opposite evidence preference in the two pathways based on stronger weights on the cues on the direct pathway

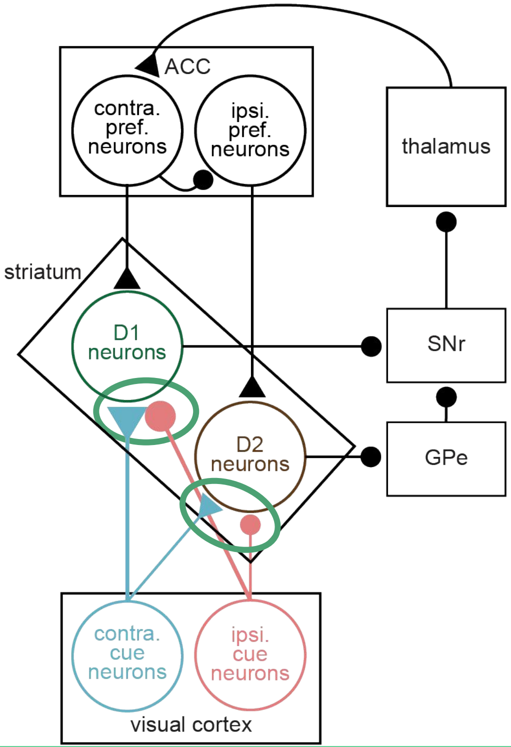

→ which could be produced by DA-mediated plasticity on the visual inputs to DMS

→ i.e. LTP in direct (D1R) vs LTD in indirect (D2R)

To recap: When, how, and why do the two pathways exert opponent effects on behavior, & neural activity?

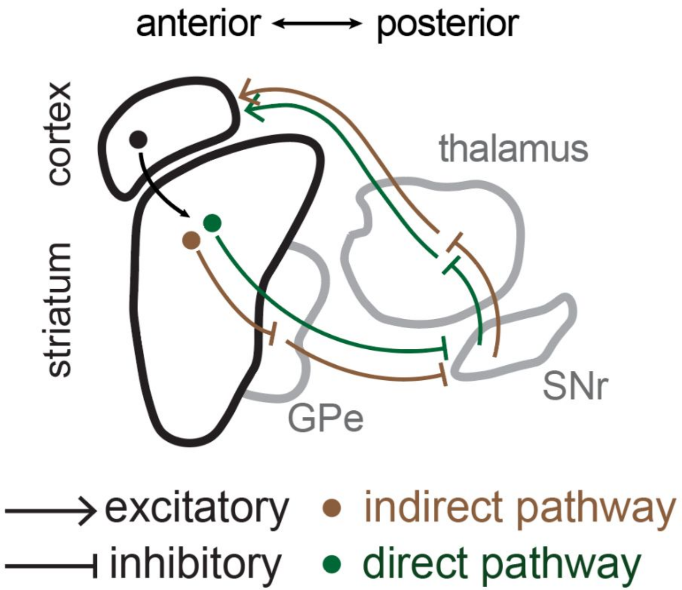

To recap: When, how, and why do the two pathways exert opponent effects on behavior, & neural activity?

- When? Opposing effects on behavioral and neural activity that are very specific <u>for the relevant decision-variable</u>

To recap: When, how, and why do the two pathways exert opponent effects on behavior, & neural activity?

- When? Opposing effects on behavioral and neural activity that are very specific for the relevant decision-variable

- How? These effects could result from a distributed cortical-BG loop that computes the decision-variable

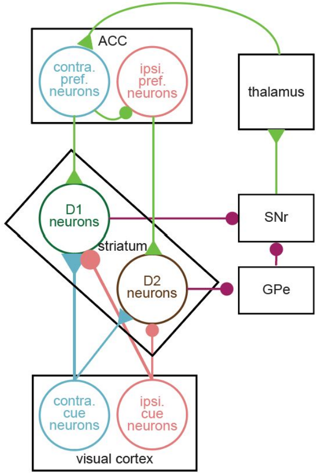

To recap: When, how, and why do the two pathways exert opponent effects on behavior, & neural activity?

- When? Opposing effects on behavior and neural activity that are very specific <u>for the relevant decision-variable</u>

- How? These effects could result from <u>a distributed cortical-BG loop that computes the decision-variable</u>

- Why?

To recap: When, how, and why do the two pathways exert opponent effects on behavior, & neural activity?

- When? Opposing effects on behavior and neural activity that are very specific for the relevant decision-variable

- How? These effects could result from a distributed cortical-BG loop that computes the decision-variable

- Why? The computational function of opposing pathways (at least in the model) is to produce competition between the choices (left vs right)

canonical integration circuit

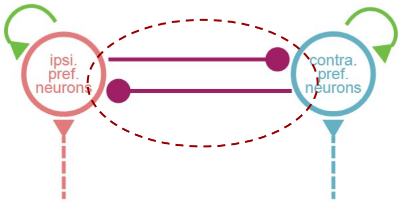

Summary for today

1. Anatomy overview

→ **multiple parallel cortico-striatal-thalamic loops**

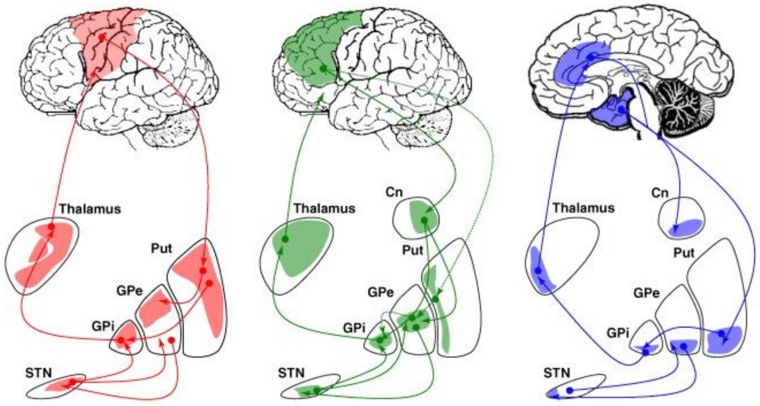

(a) Motor circuit (b) Associative circuit (c) Limbic circuit

Summary for today

1. Anatomy overview

→ **within each loop, direct & indirect pathways**

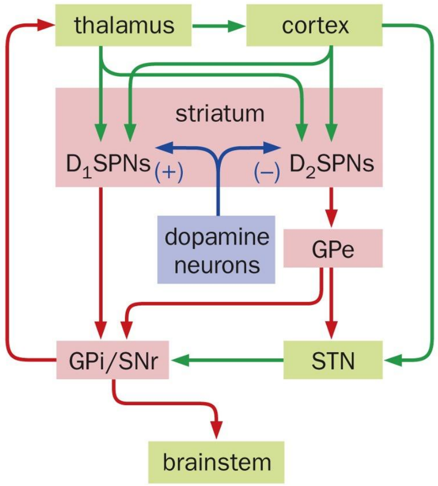

Figure 8-22c Principles of Neurobiology (© Garland Science 2016)

Summary for today

2. How does the dopamine system regulate the basal ganglia?

➔ **dopamine encodes a reward prediction error**

➔ **working model:**

- **across projections, different types of error signals**
- **within projection, feature-specific error signals**

Summary for today

3. What is the function of the direct and indirect pathway?

    ➜ **opposing effects on behavior, depending on the function of the striatal regions (e.g. decisions vs motor control vs threat learning)**

    ➜ **opposing effects on cortical activity, which may arise from a distributed cortico-thalamic-BG circuit calculating the decision variable**

Summary for today

4. Why are signals similar in the direct and indirect pathway?

→ **similar signals in the 2 pathways after an action could allow plasticity updating (efference copy)**

→ **opponent signals/ competition between the two pathways *before* an action could cause it to be selected (opponent coding of action values)**

Acknowledgments

Current lab:

*Ryan Cho
*Lindsey Brown
Alex Pan Vazquez
Bichan Wu
Brenna McMannon
Anna Zhukovskaya
Michael Siniscalchi
Rob Fetcho
Yotam Sagiv
Addie Minerva
Yousuf Al-Jayyousi
Yunchang Zhang
Vanessa Roser
Miranta Louka
Jordan Mearas-Garcia

Collaborators:

Yoel Sanchez Araujo
*Nathaniel Daw
Jonathan Pillow
Sam Wang
Annegret Falkner
David Tank
Carlos Brody
Cate Pena
Mark Goldman

Recent alumni:

Chris Zimmerman (Utah)
*Scott Bolkan (Iowa)
Lindsay Willmore (DeepMind)
Ben Engelhard (Technion)
Malu Murugan (Emory)
Julia Cox (Chicago)
Clare Choi (Northwestern)
Weston Fleming (UW)
Iris Stone (Allen)
*Rachel Lee (Allen)

PRINCETON Neuroscience INSTITUTE

SIMONS FOUNDATION
NIMH National Institute of Mental Health
BRAIN RESEARCH THROUGH ADVANCING INNOVATIVE NEUROTECHNOLOGIES BRAIN INITIATIVE
NIH DIRECTOR'S PIONEER AWARD

D1R are needed for LTP in D1R neurons (through an AMPA/NMDA mechanism)

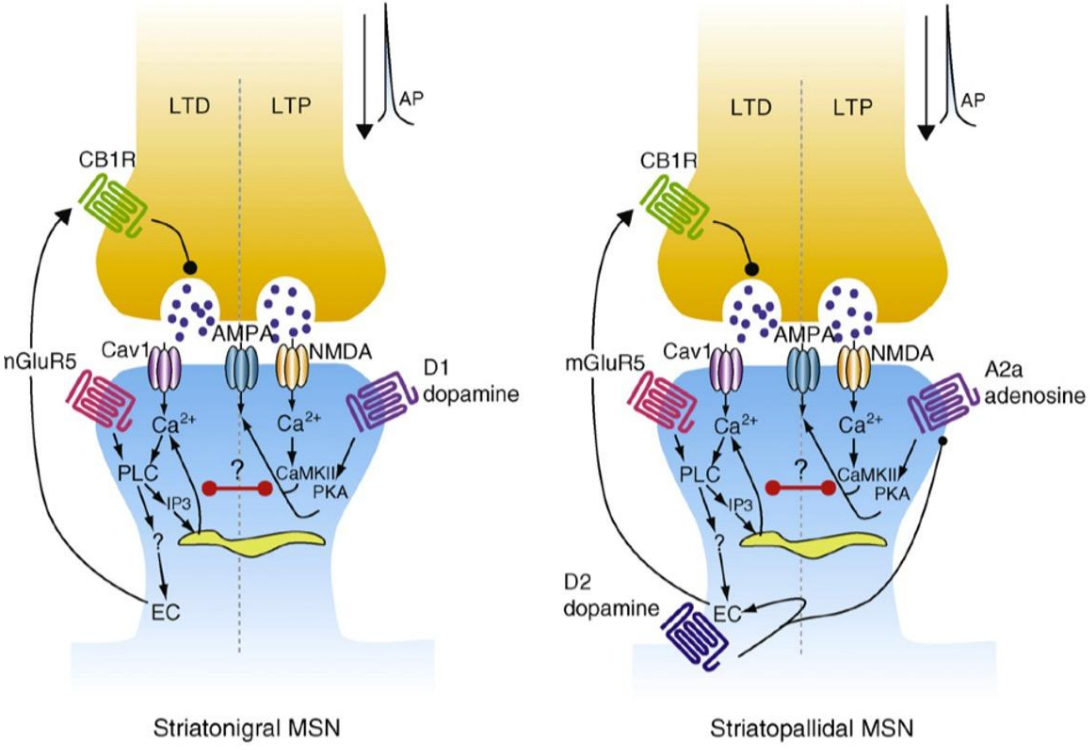

D2R are needed for LTD in D2R neurons (through a presynaptic endocannabinoid mechanism)

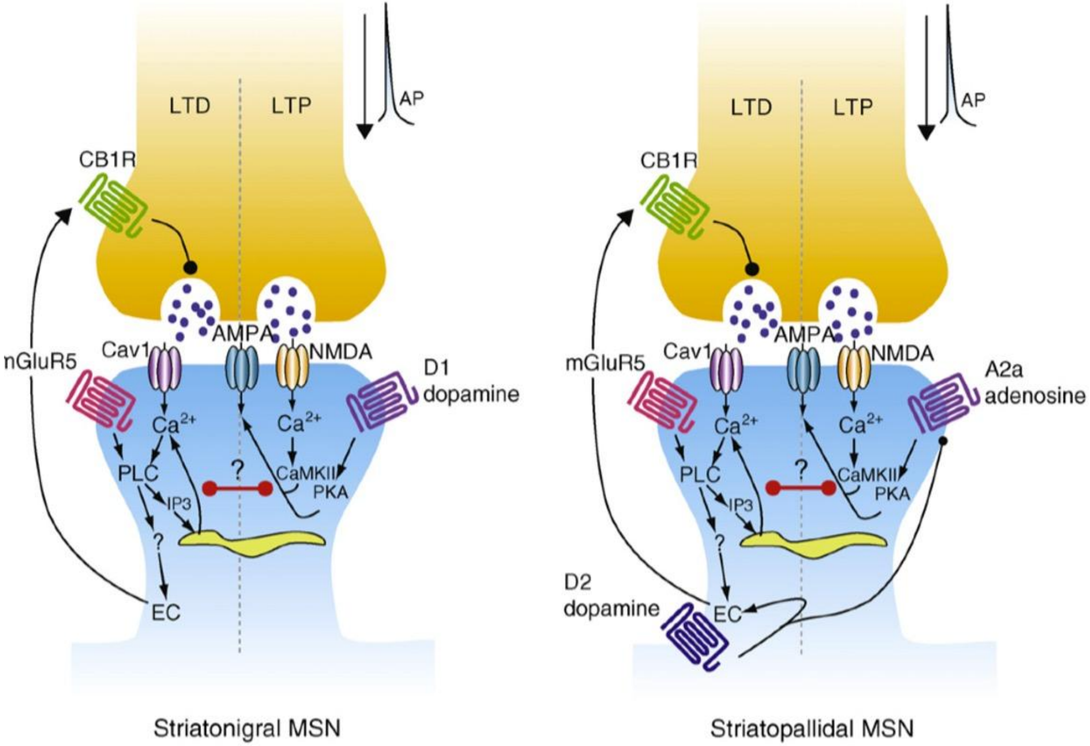

Train deep reinforcement learning network on task to simulate Vector RPE model

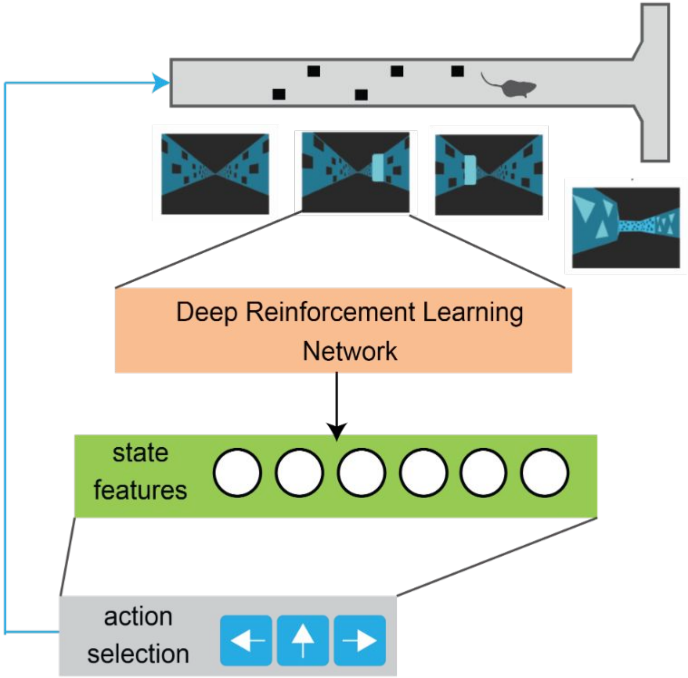

Lee *et al.* *Nature neuro* 2024

Train deep reinforcement learning network on task to simulate Vector RPE model

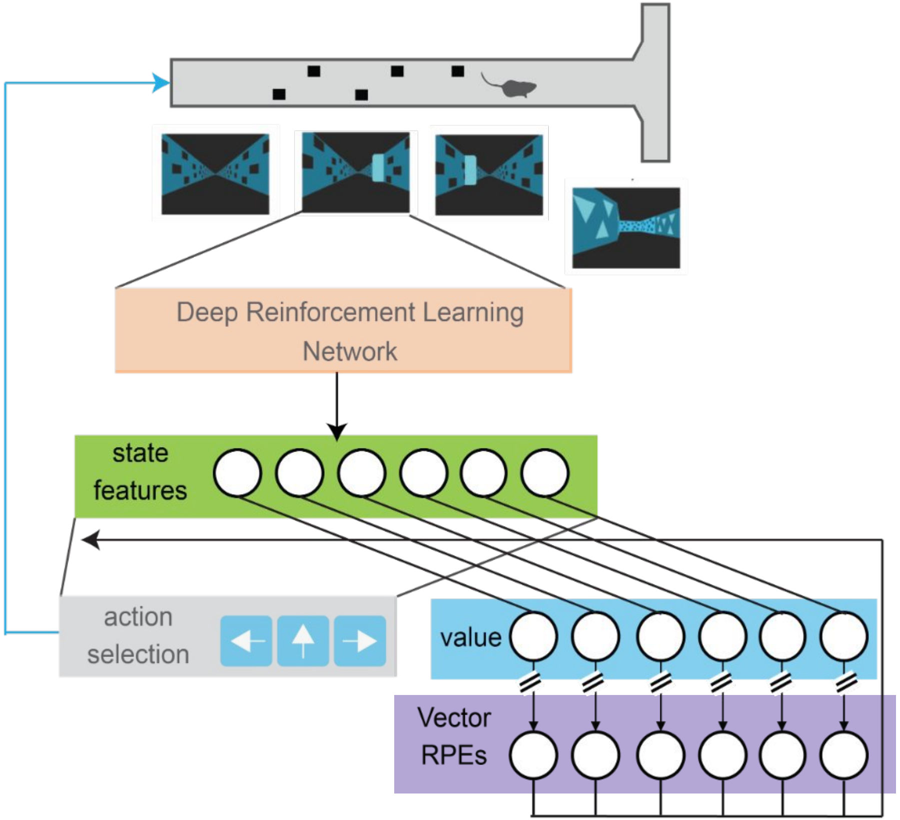

Lee *et al.* *Nature neuro* 2024

Train deep reinforcement learning network on task to simulate Vector RPE model

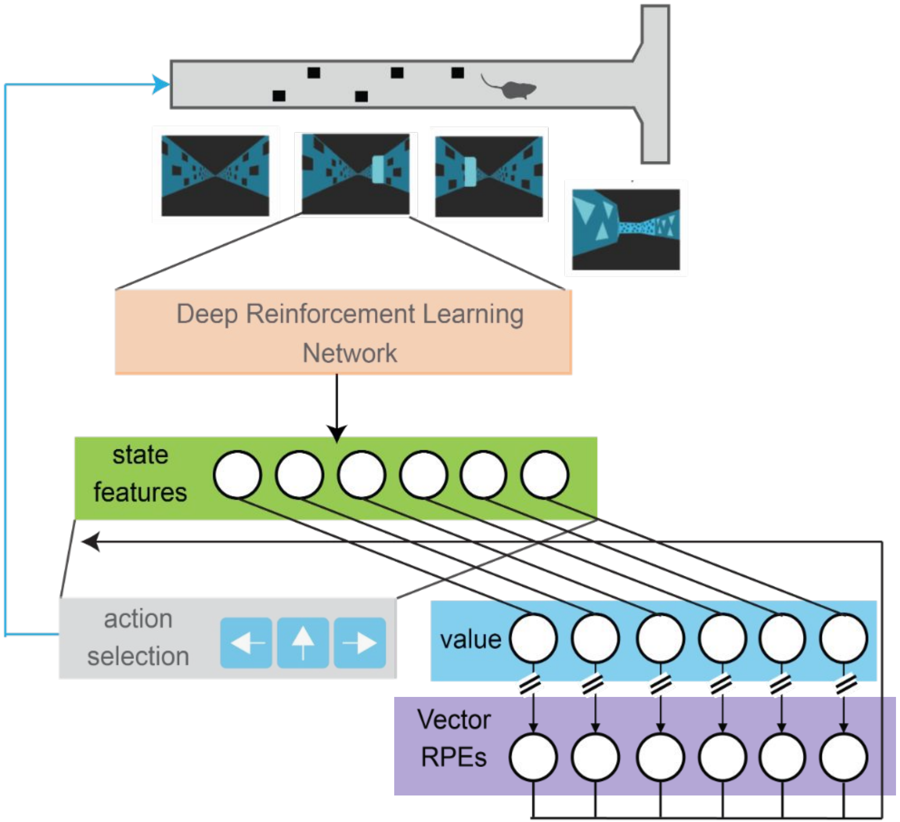

Lee *et al. Nature neuro* 2024

Heterogeneity in tuning during cue period across feature-specific RPEs

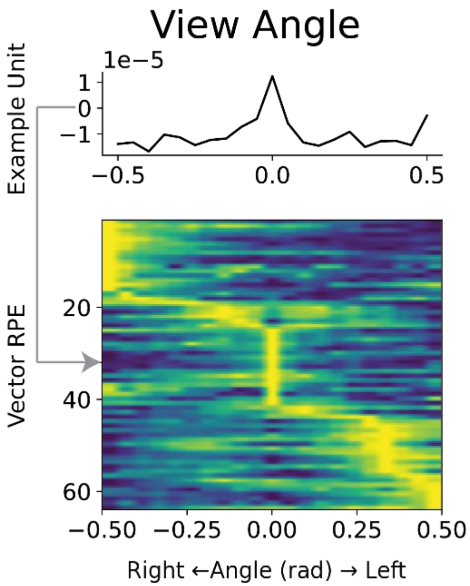

Heterogeneity in tuning during cue period across feature-specific RPEs

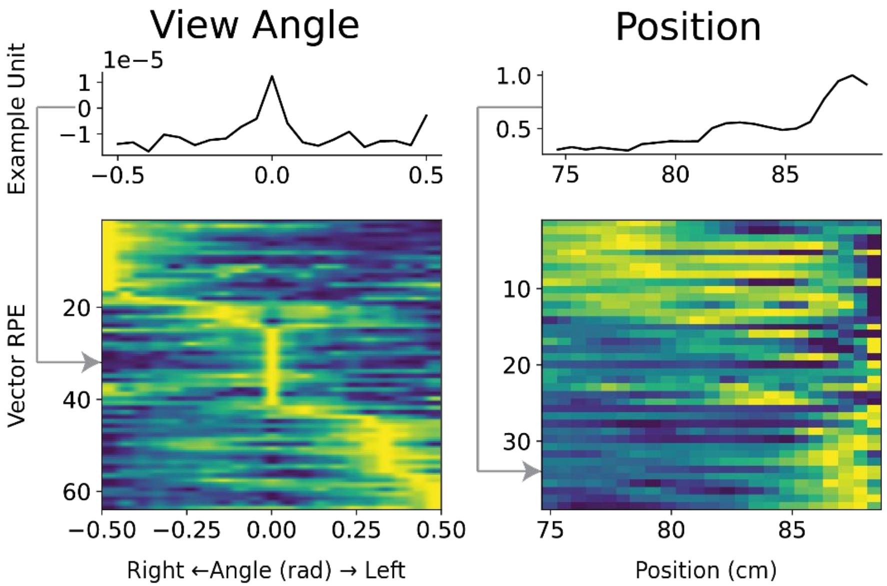

Heterogeneity in tuning during cue period across feature-specific RPEs

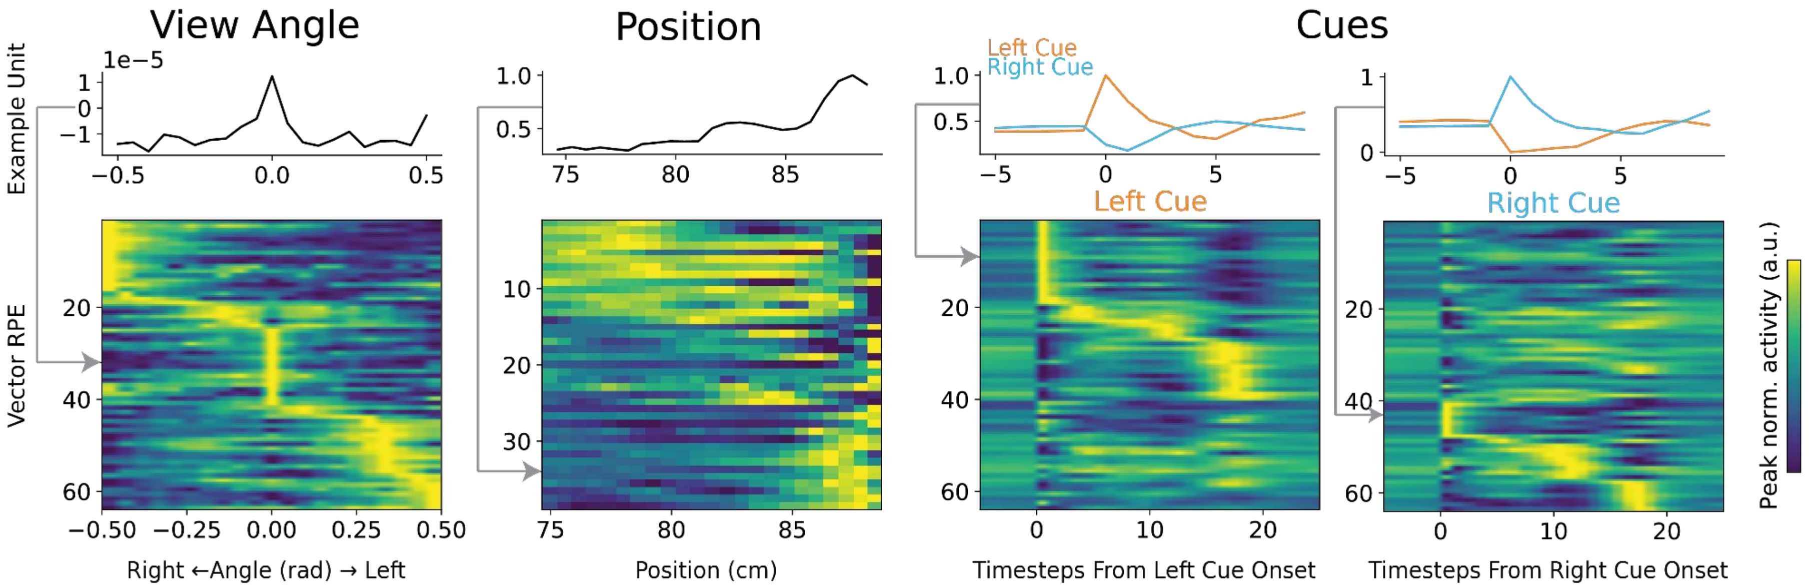

Cue responses as feature-specific RPEs

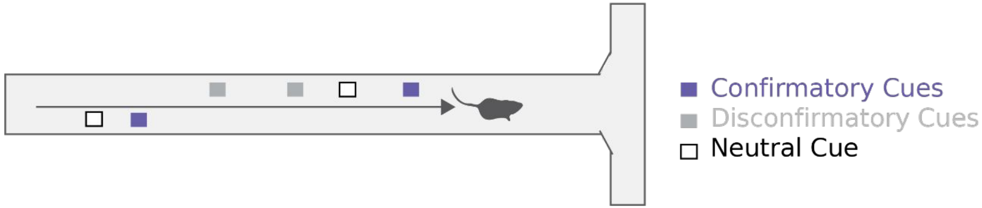

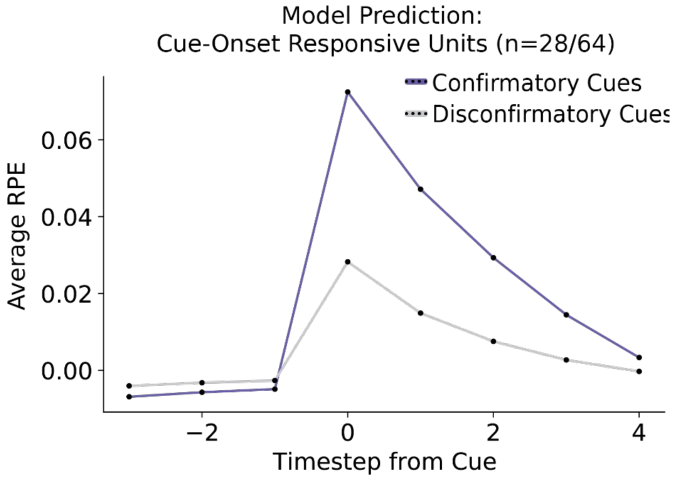

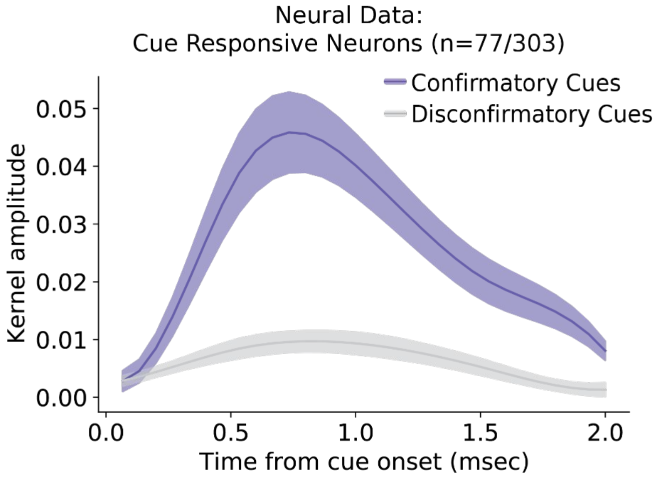

Homogeneity in the reward response

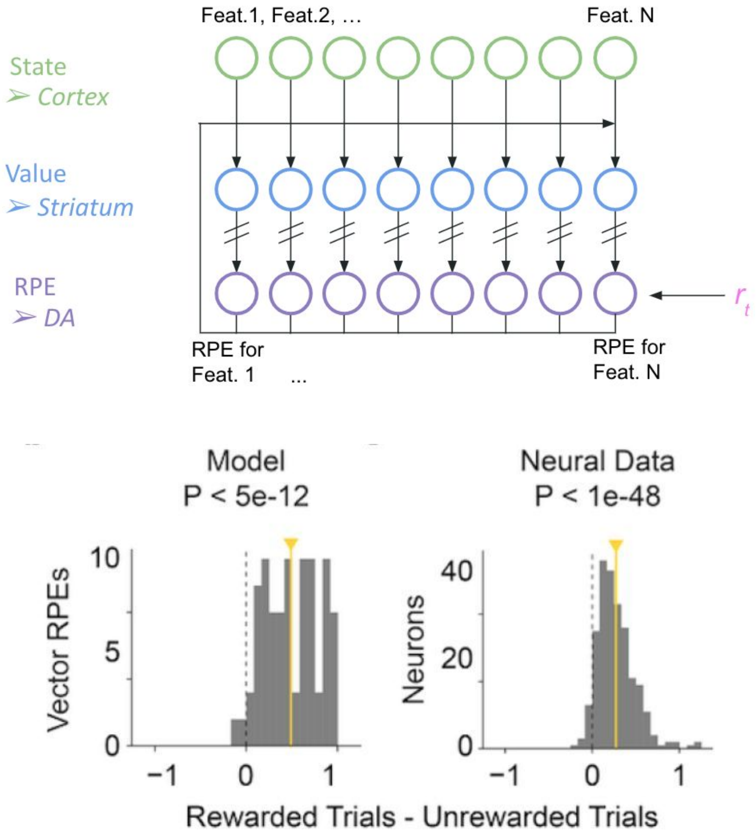

Implies that dopaminergic heterogeneity tells us the features being used to predict reward ("state representation")

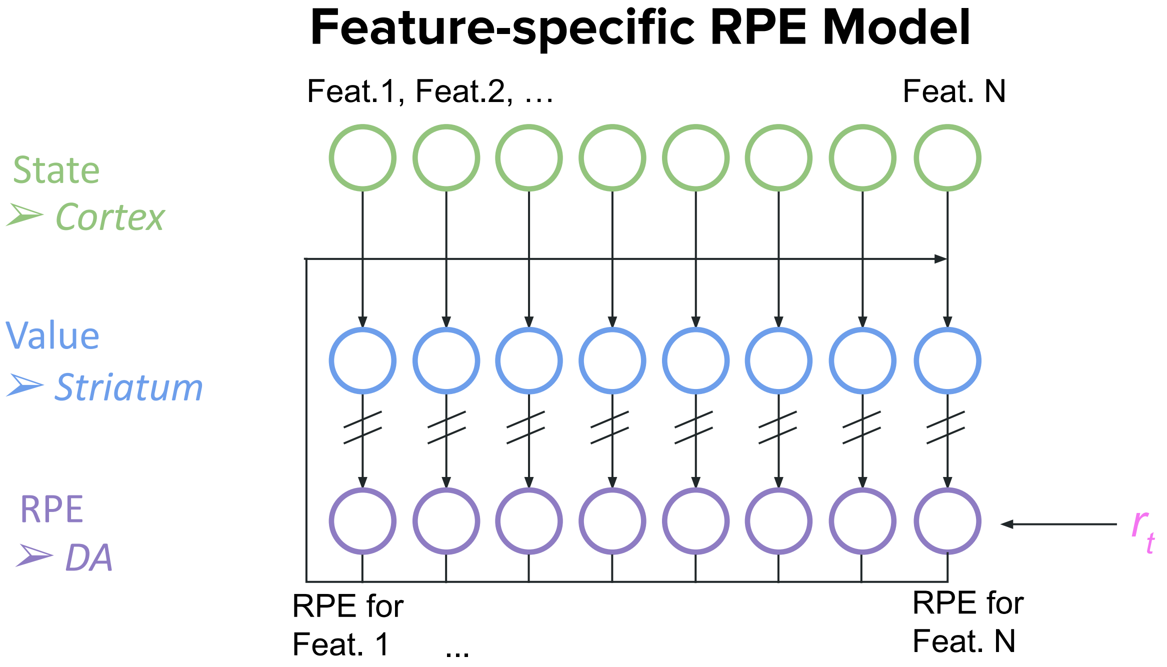

✅ Heterogenous coding of task variables
✅ Task variable coding as feature-specific RPE
✅ Homogeneous response to reward

Different striatal subregions that have different specializations

Dorsolateral striatum (DLS):
motor learning; habits

Dorsomedial striatum (DMS):
goal-directed behavior;
visual decision-making

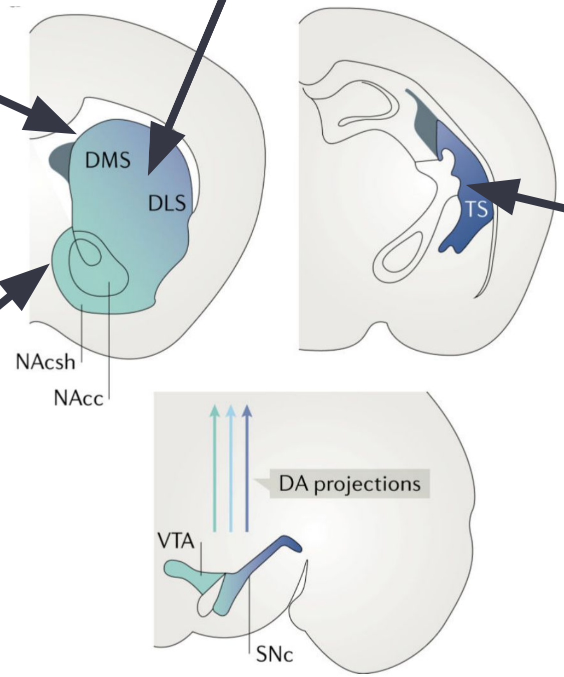

Tail of striatum (TS): threat
learning; auditory
decision-making

Nucleus accumbens (NAc):
Pavlovian associations
Emotional regulation
Approach behavior