Why measure from the brain during decision-making?

Two ideas running in the background:

Idea 1: multi-neuron recordings for better moment-by-moment estimates of internal signals

Idea 2: most neural activity is of unknown function, yet is coordinated across neurons and brain regions. What's going on with this "dark matter of the brain?"

electrophysiological recordings in **PPC**

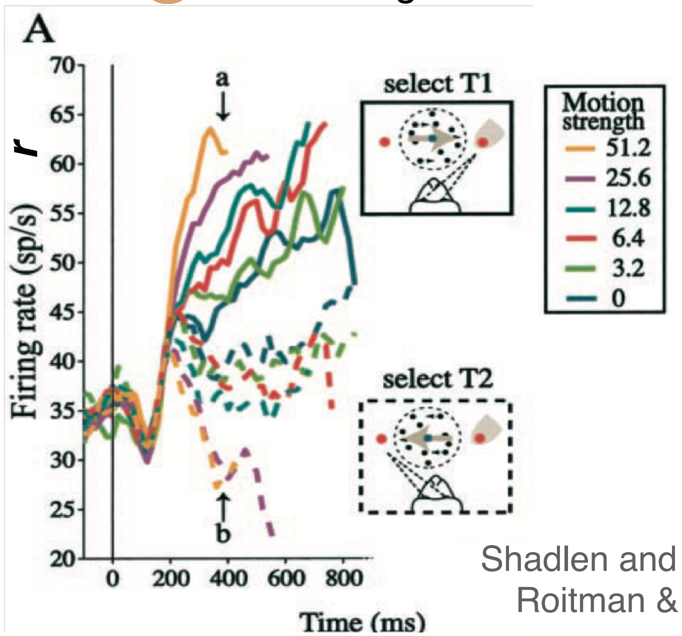

individual trials could be quite noisy

averages within a stimulus strength class will show key trends

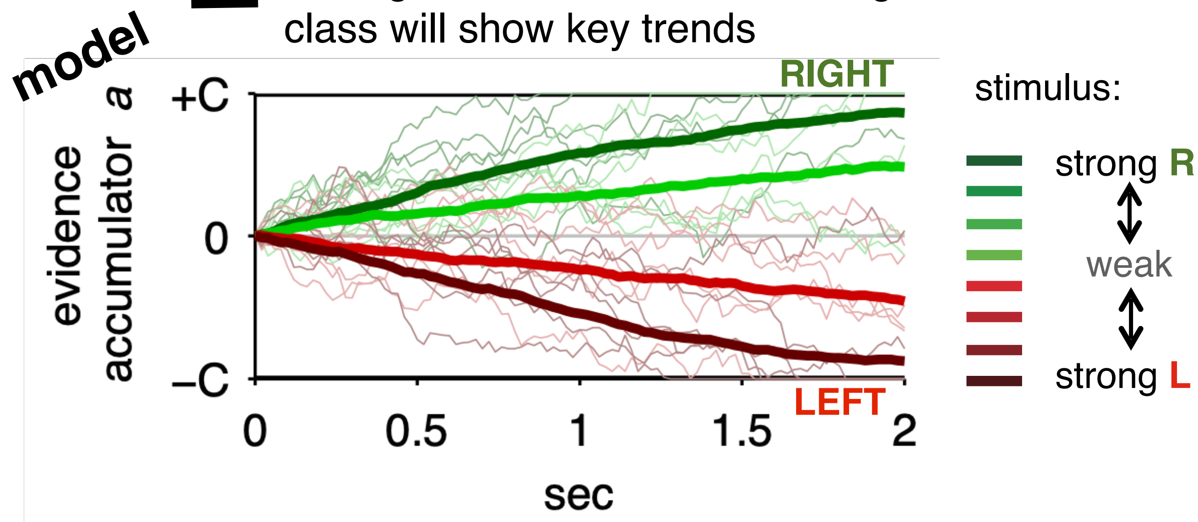

Shadlen and Newsome (1996);
Roitman & Shadlen (2002)

similarity suggests PPC firing rates *r* encode value of *a*

but this is based on averages over trials — what happens on single trials?

stepping model

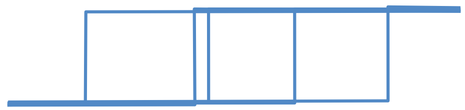

The stepping versus ramping controversy

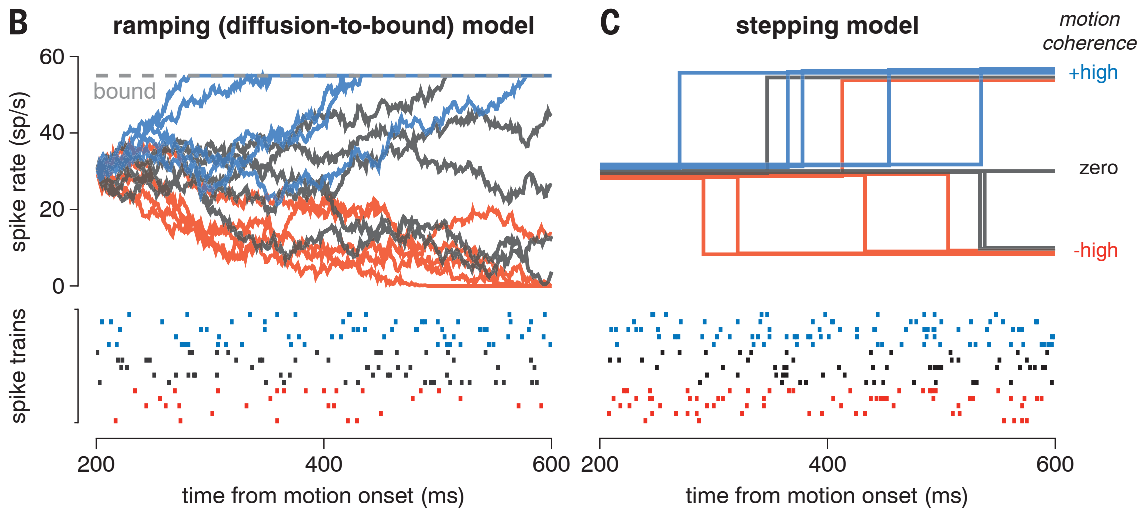

Latimer, … Pillow *Science* 2015

Comment on "Single-trial spike trains in parietal cortex reveal discrete steps during decision-making"

Michael N. Shadlen,[1*] Roozbeh Kiani,[2] William T. Newsome,[3] Joshua I. Gold,[4] Daniel M. Wolpert,[5] Ariel Zylberberg,[6] Jochen Ditterich,[7] Victor de Lafuente,[8] Tianming Yang,[9] Jamie Roitman[10]

Latimer *et al*. (Reports, 10 July 2015, p. 184) claim that during perceptual decision formation, parietal neurons undergo one-time, discrete steps in firing rate instead of gradual changes that represent the accumulation of evidence. However, that conclusion rests on unsubstantiated assumptions about the time window of evidence accumulation, and their stepping model cannot explain existing data as effectively as evidence-accumulation models.

Science 2016

It's both: some neurons look more like ramping, some neurons more like stepping
Zoltowski … Pillow, *Neuron* 2019

But this was still all single neurons !!!

Using multielectrode recordings to greatly improve prediction of behavior in single trials

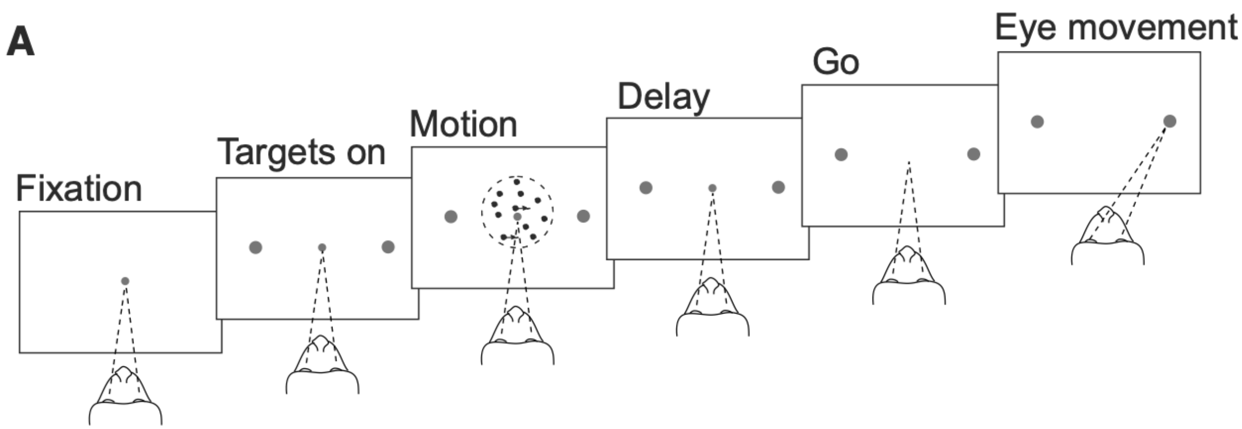

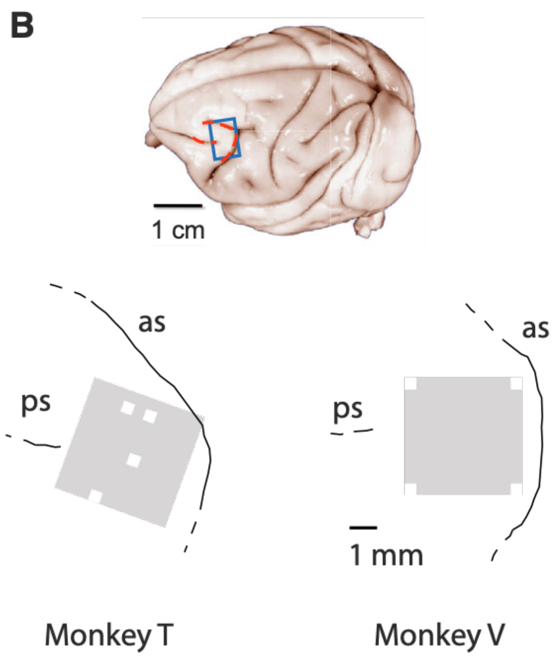

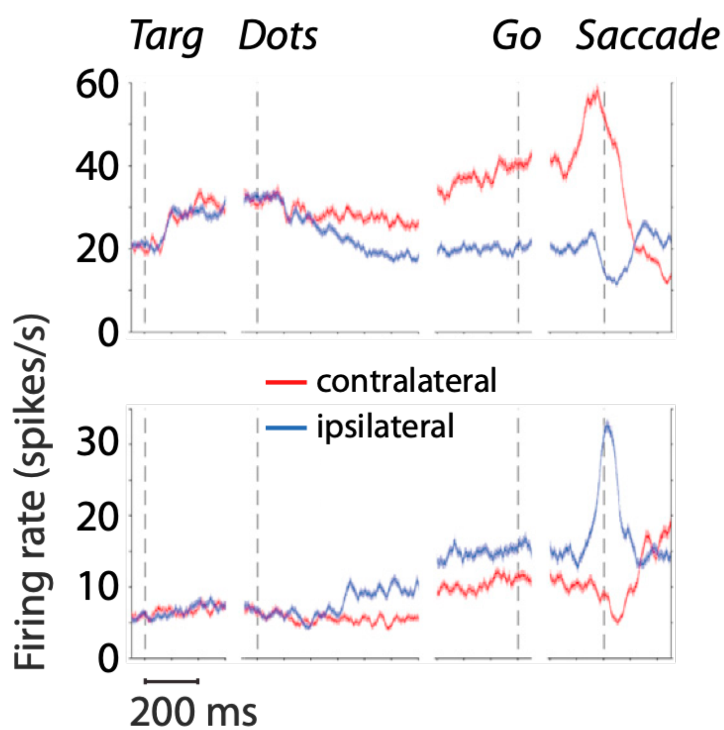

(Kiani, …, Newsome *Current Biology* 2015)

A simple linear model predicts behavior very well

Firing rates r(t) of N neurons

$$\mathrm{DV} = \log \frac{P(T_1|\vec{\mathbf{r}})}{P(T_2|\vec{\mathbf{r}})} = \beta_0(t) + \sum_{i=1}^{n} \beta_i(t) \times r_i(t)$$

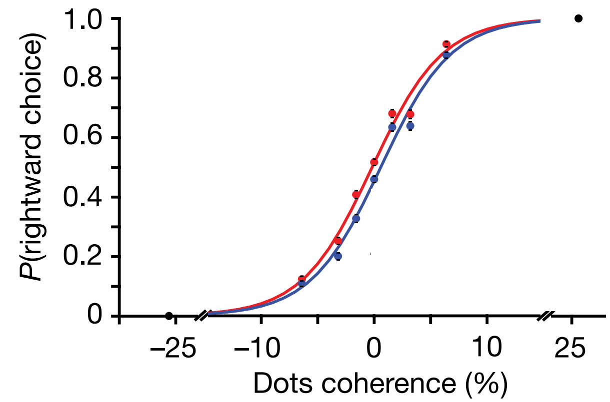

$$\mathrm{DV} = \log \frac{P(T_1|\mathbf{r})}{P(T_2|\mathbf{r})}$$
$$= \log \frac{P(T_1)}{1 - P(T_1)}$$
$$\Rightarrow$$
$$e^{\mathrm{DV}} = \frac{P(T_1)}{1 - P(T_1)}$$
$$\Rightarrow$$
$$P(T_1) = \frac{1}{1 + e^{-\mathrm{DV}}}$$

DV very positive: monkey chooses T1 almost always
DV = 0 : 50/50
DV very negative: monkey chooses T2 almost always

(Kiani, …, Newsome *Current Biology* 2015)

A simple linear model predicts behavior very well

Firing rates r(t) of N neurons

$$\mathrm{DV} = \log \frac{P(T_1|\vec{\boldsymbol{r}})}{P(T_2|\vec{\boldsymbol{r}})} = \beta_0(t) + \sum_{i=1}^{n} \beta_i(t) \times r_i(t)$$

optimize model params for 90% of the data,
test on remaining 10%

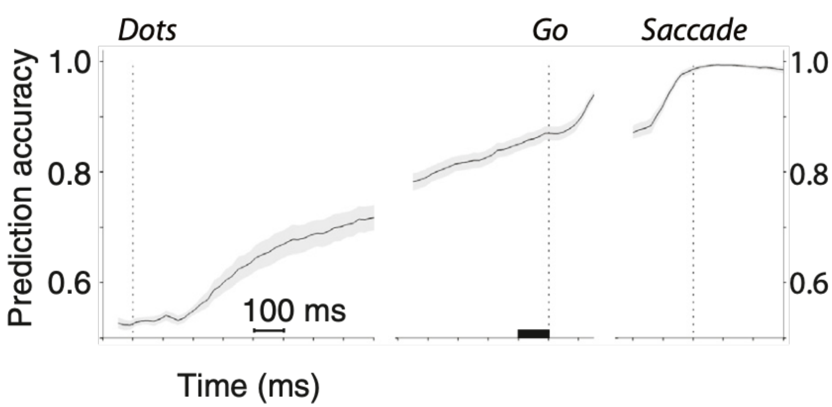

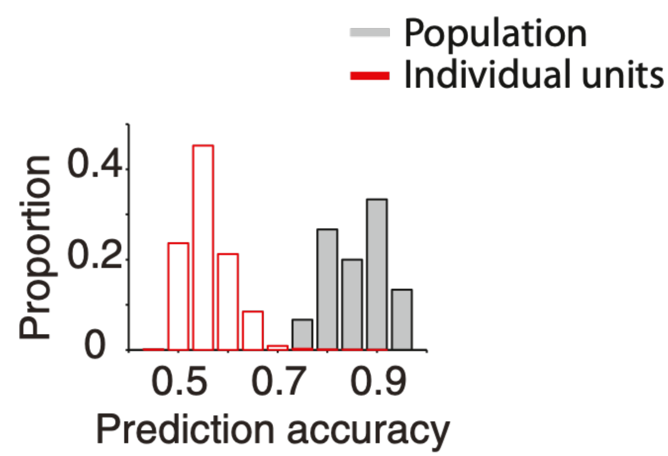

(Kiani, …, Newsome *Current Biology* 2015)

DV(t) in four example trials

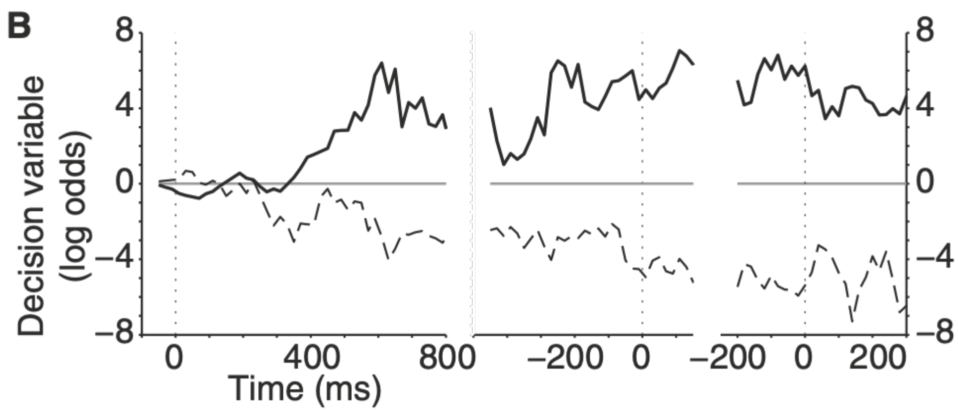

one trial

another trial

Arrows indicate “changes of mind” ?

(Kiani, …, Newsome *Current Biology* 2014)

Two ideas running in the background:

Idea 1: multi-neuron recordings for better moment-by-moment estimates of internal signals

Idea 2: most neural activity is of unknown function, yet is coordinated across neurons and brain regions. What's going on with this "dark matter of the brain?"

Idea 2: the “dark matter” of the brain …

Neural activity variance:

1.5 %

explained by known
task variables, i.e., what we study

(International Brain Lab 2023; our data)

~10 % correlated with
uninstructed movements
(but we don’t know why)

(Musall…Churchland 2019,
Wang… Svoboda Druckmann 2023;
our data)

~88 % ???

we… don’t know

- most neural activity looks like noise but is coordinated across neurons and regions.

(Arieli 1996, Fiser 2004, Stringer 2019, Manley 2024)

Two ideas running in the background:

Idea 1: multi-neuron recordings for better moment-by-moment estimates of internal signals

Idea 2: most neural activity is of unknown function, yet is coordinated across neurons and brain regions. What's going on with this "dark matter of the brain?"

On to the main talk

Luo ’23

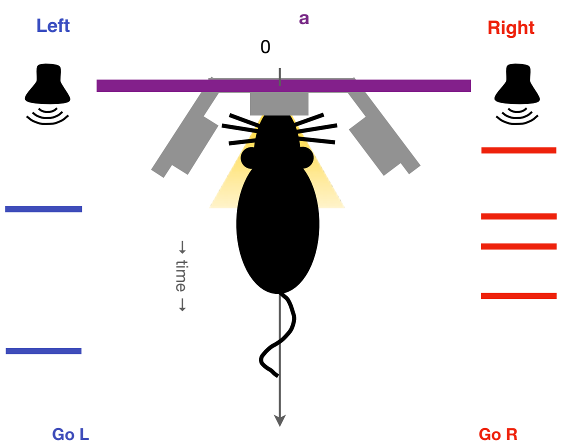

decision-making Luo ’23

Behavior level:
the “drift-diffusion model” (DDM)

accumulation
inputs drive dynamics

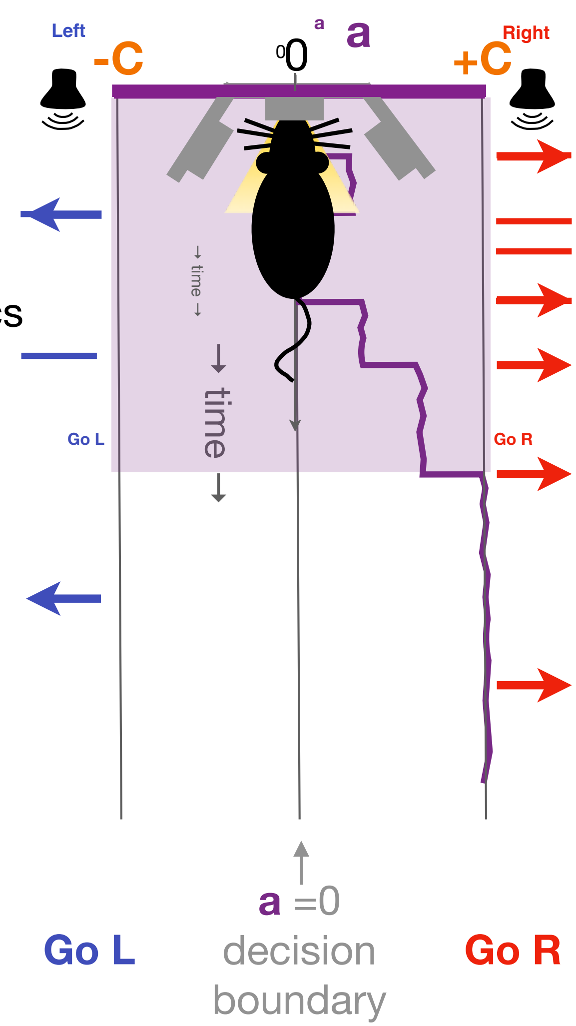

rightwards evidence

leftwards evidence

accounts for data in:

- social decisions (e.g., Krajbich 2012)
- sensory decisions (e.g., Newsome, 1989)
- economic decisions (e.g., Gluth 2012)
- gambling decisions (e.g., Busemeyer, 1993)
- memory decisions (e.g., Ratcliff, 1978)
- visual search decisions (e.g., Purcell, 2010)
- value decisions (e.g., Milosavljevic 2012)

Behavior level:
the "drift-diffusion model" (DDM)

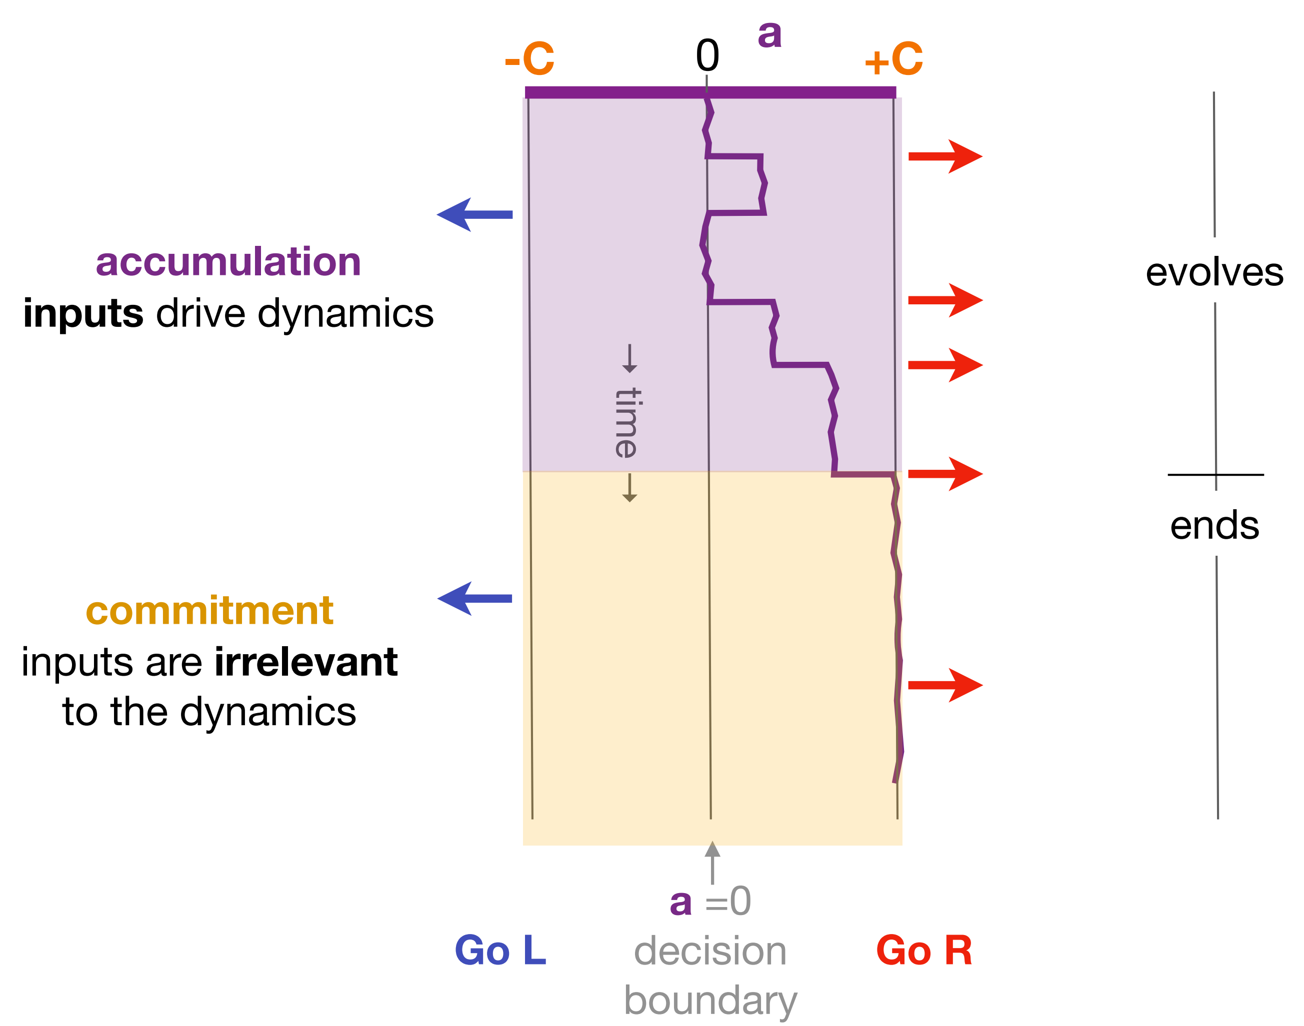

at the neural level:

how do decisions evolve?
and how do they end?

Behavior level:
the “drift-diffusion model” (DDM)

Neural level:
the “line attractor” hypothesis

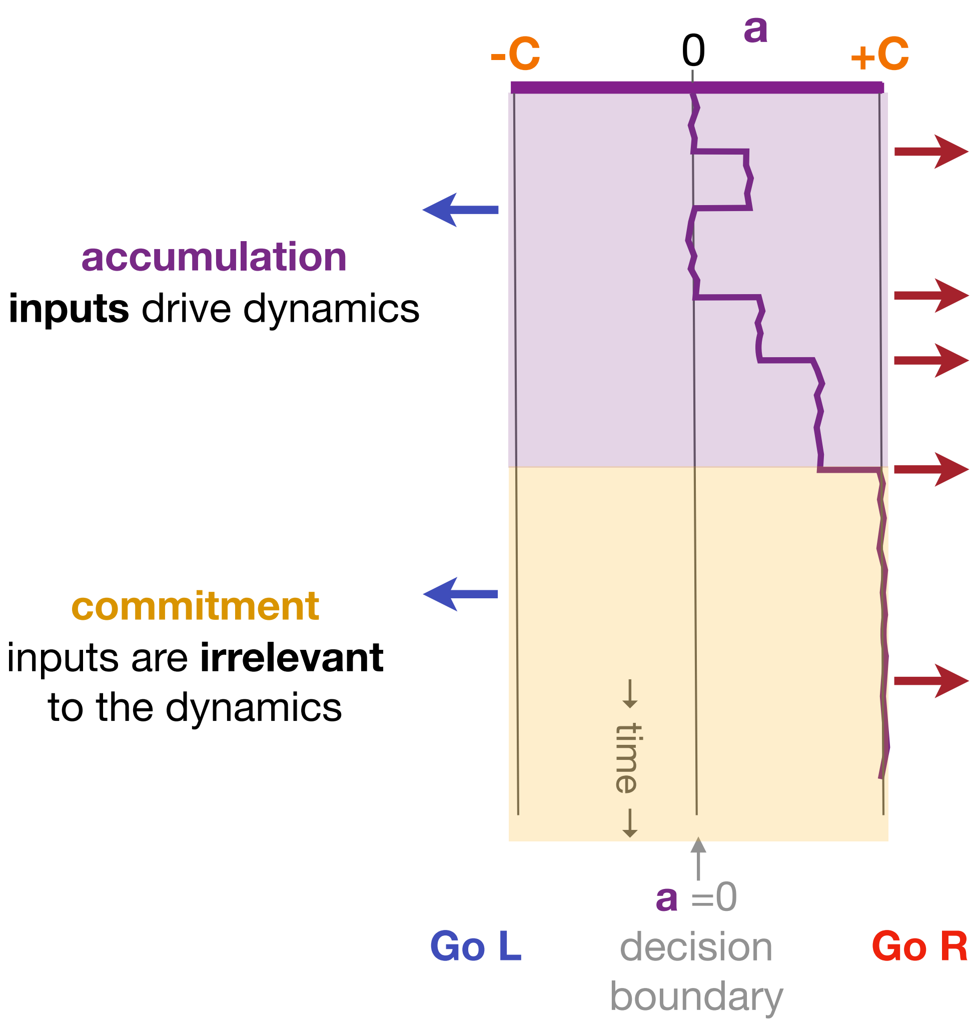

Behavior level:
the “drift-diffusion model” (DDM)

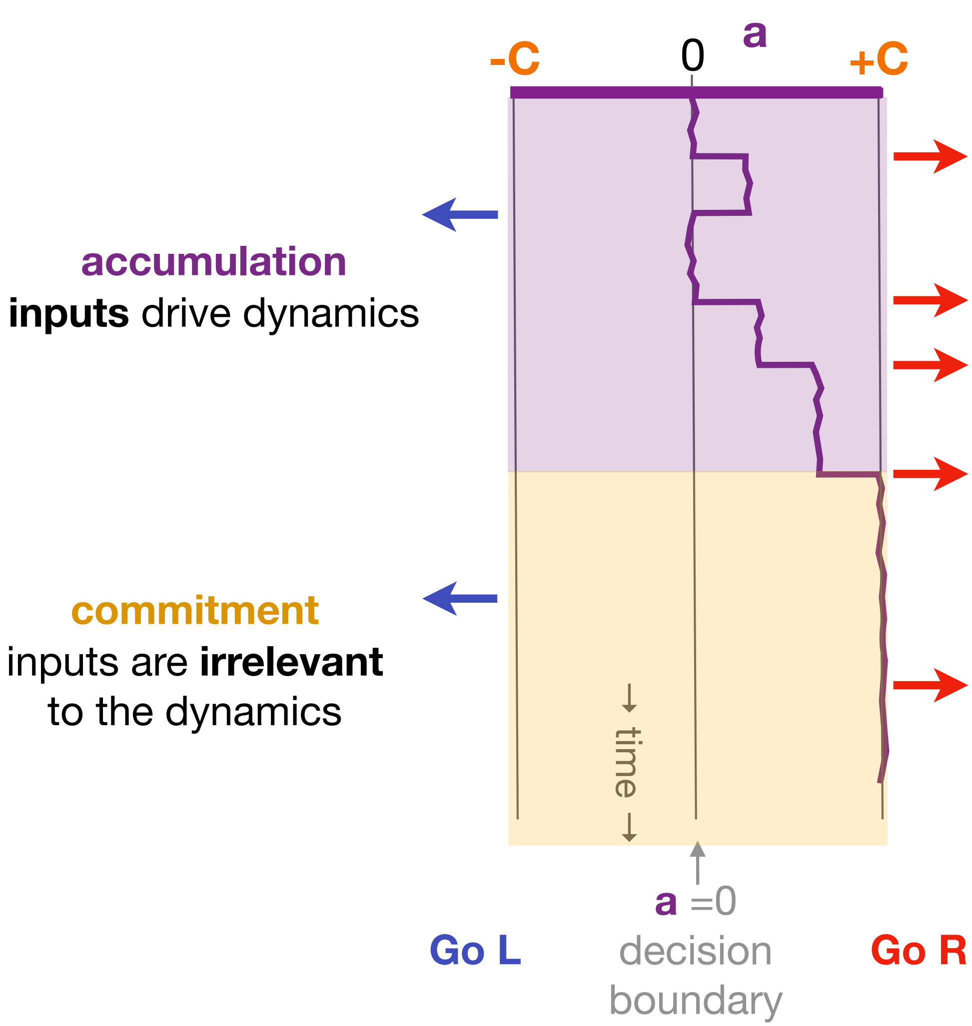

Neural level:
the “line attractor” hypothesis

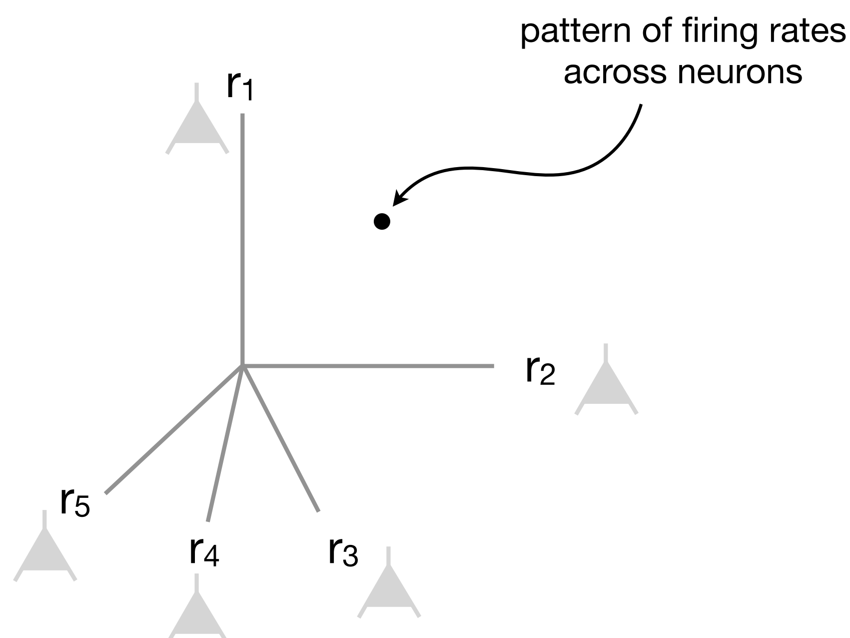

Behavior level:
the "drift-diffusion model" (DDM)

accumulation
inputs drive dynamics

commitment
inputs are **irrelevant**
to the dynamics

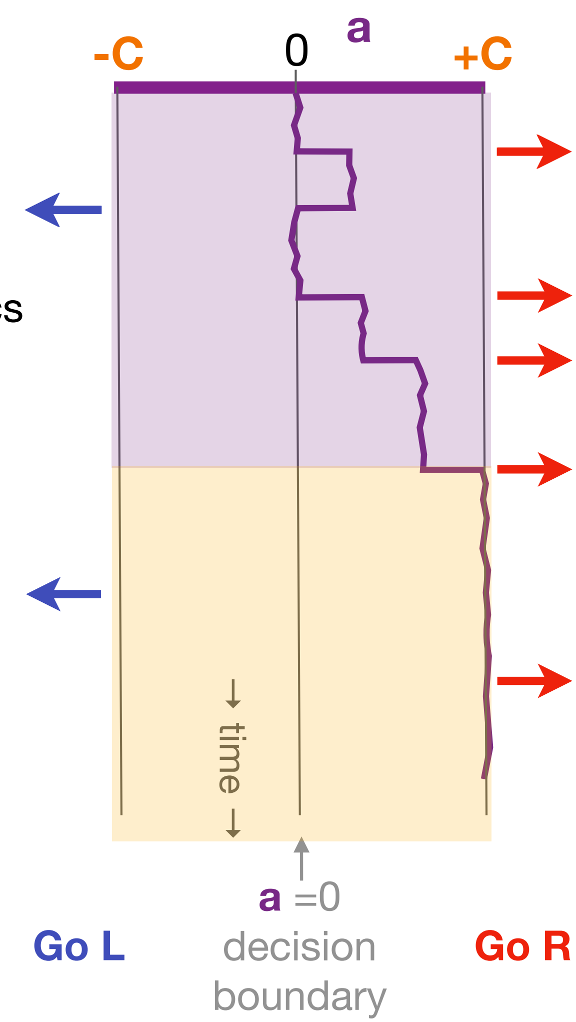

Neural level:
the "line attractor" hypothesis

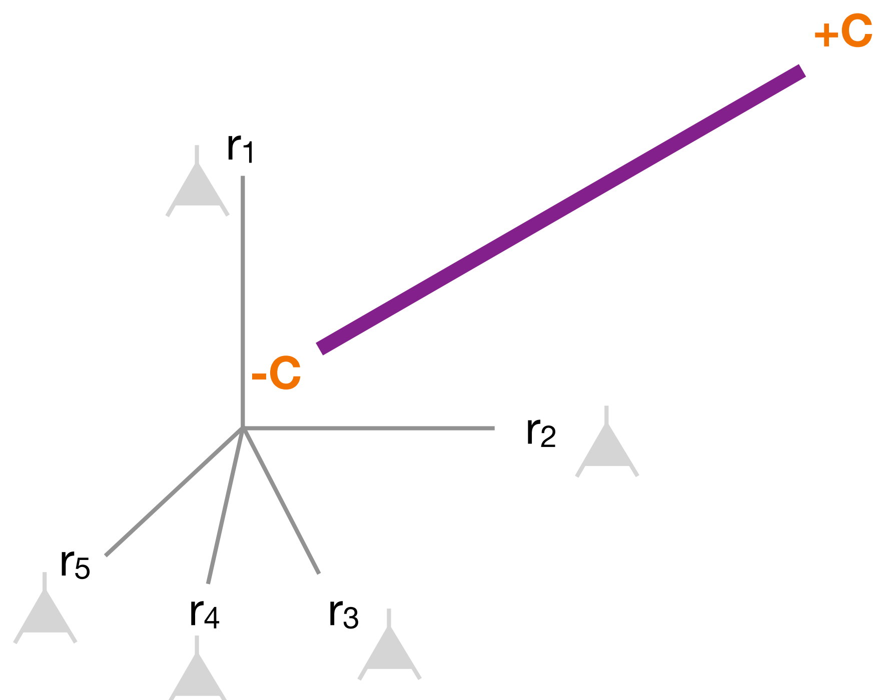

Behavior level:
the “drift-diffusion model” (DDM)

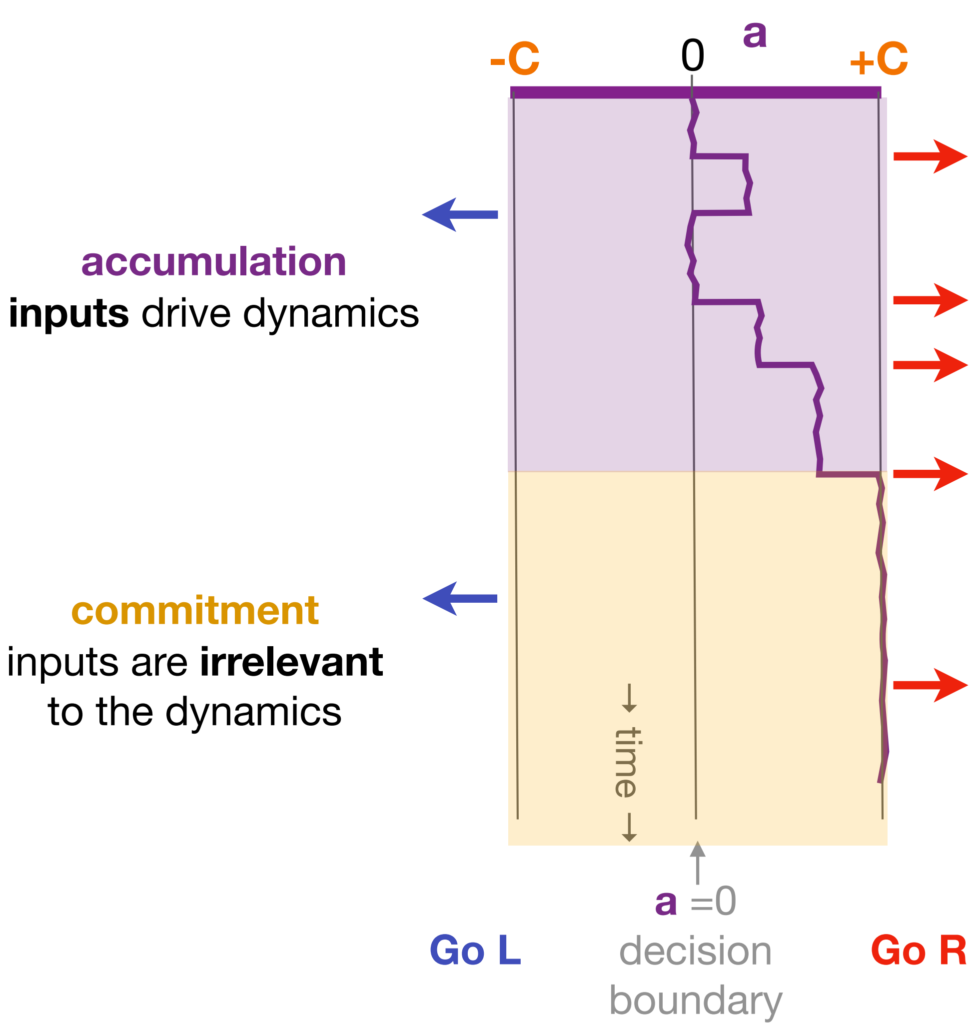

Neural level:
the “line attractor” hypothesis

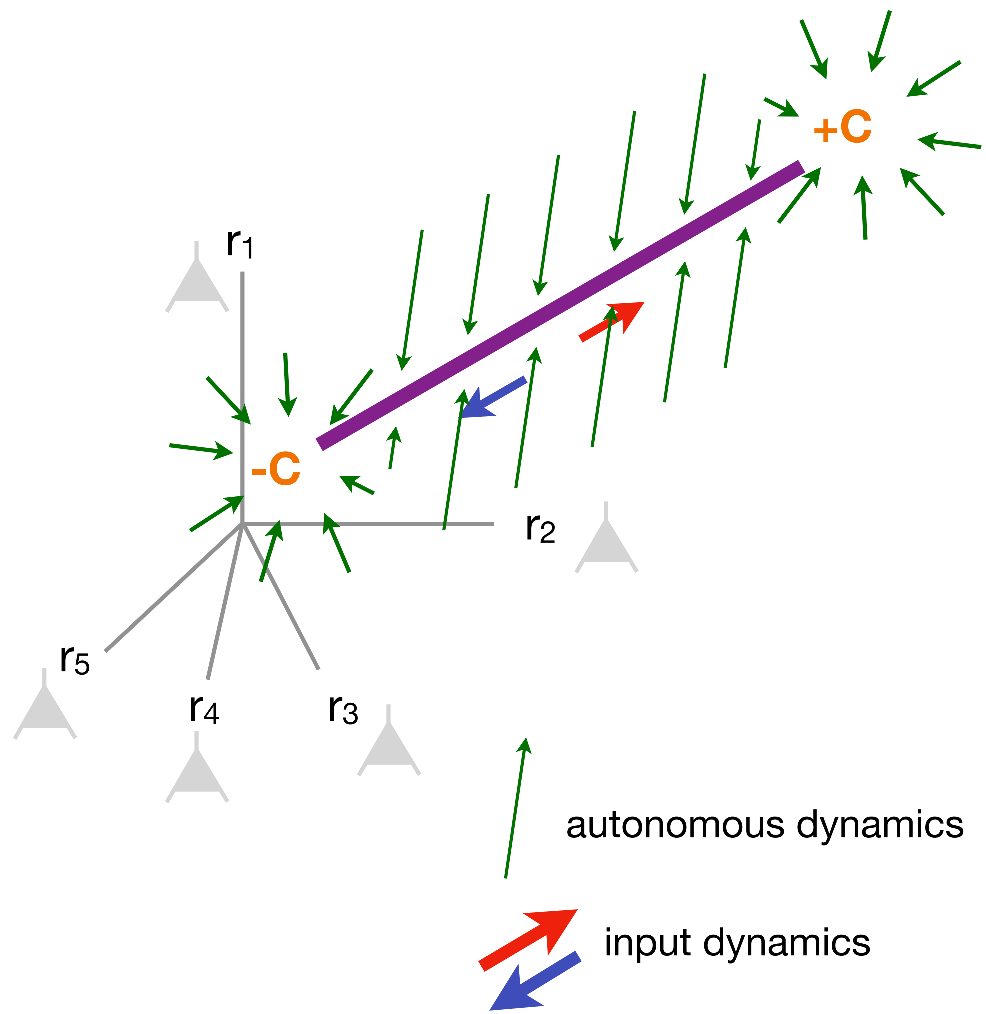

Why not just measure the flow lines and find what the data says?

theoretical model variants in the literature:

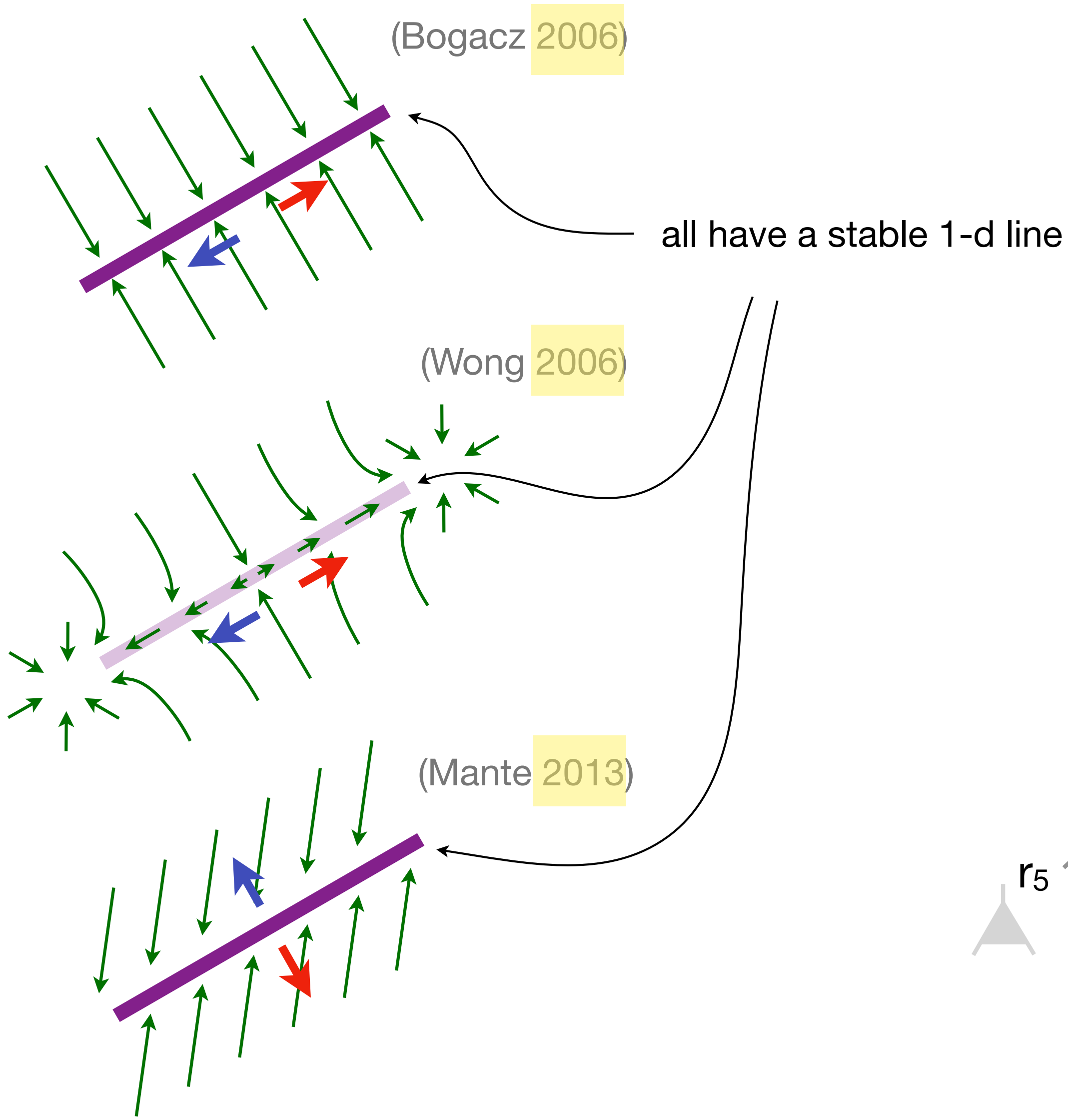

Neural level:
the "line attractor" hypothesis

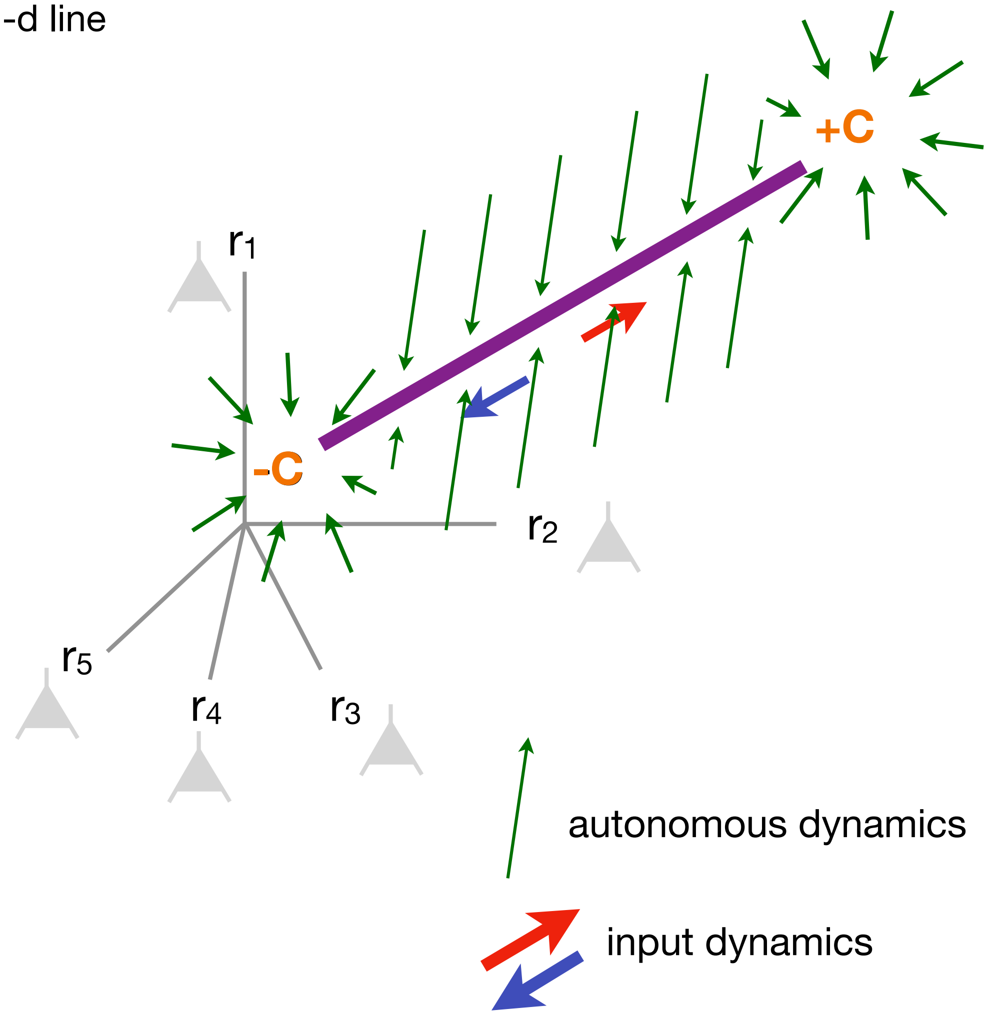

very noisy data:

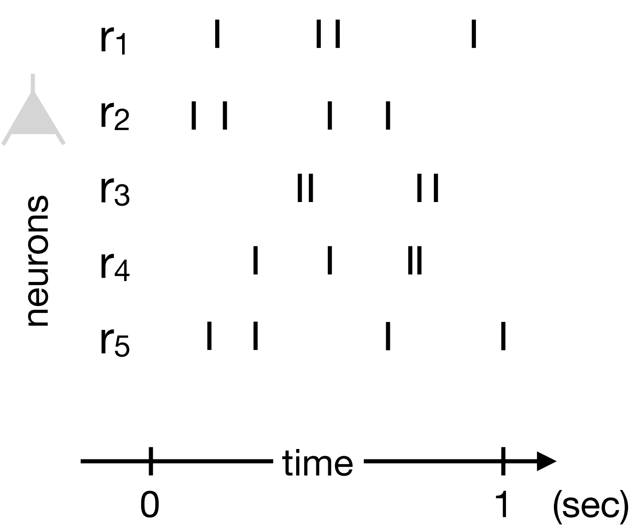

stimulus presentation:
auditory clicks ~ 1.0 sec

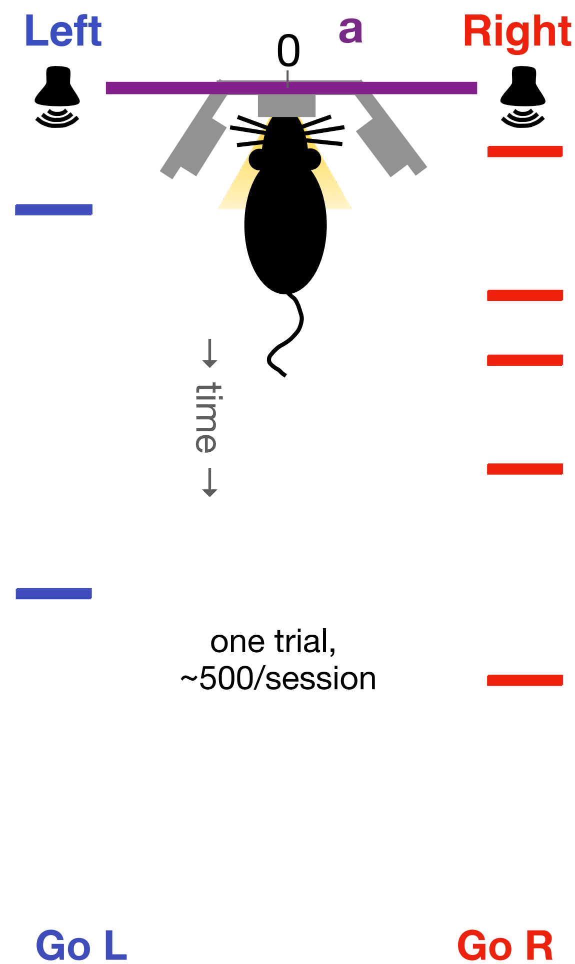

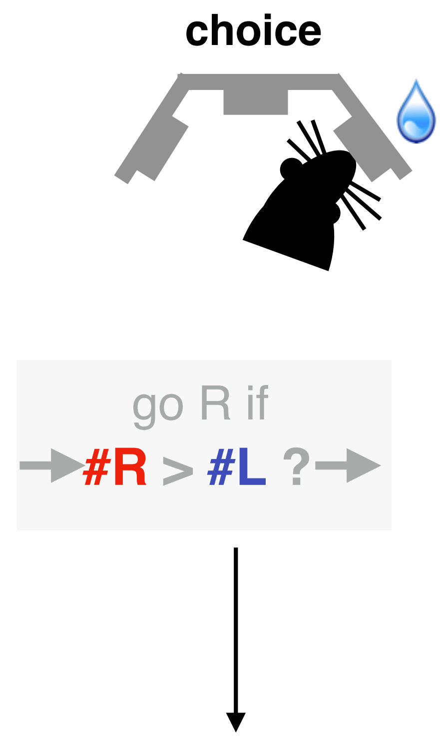

subjects must accumulate sensory evidence over time to form their decision

very noisy data:

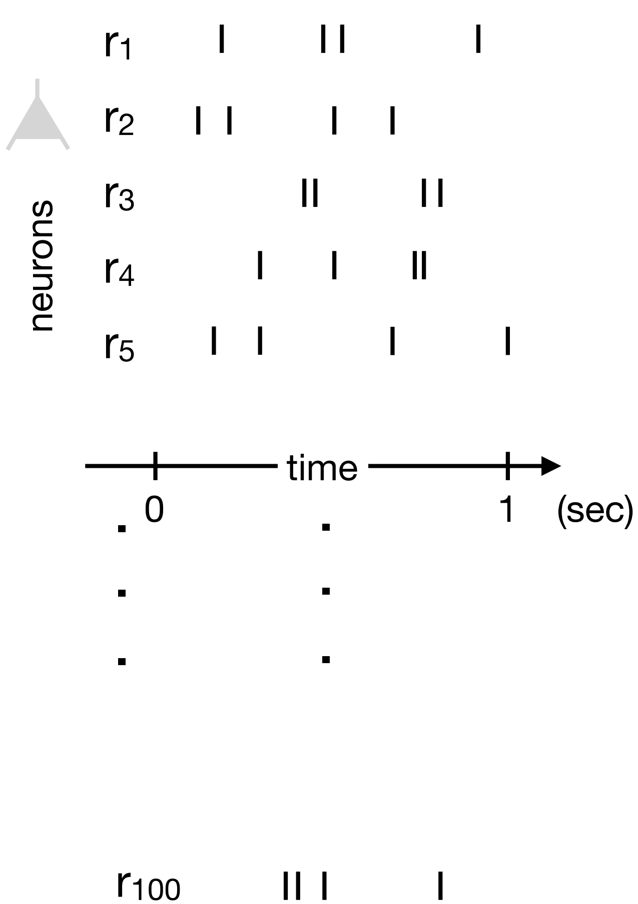

High-yield Neuropixels recordings (1 probe, 2-3 regions/rat)

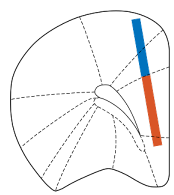

dmFC
mPFC
4.0 mm AP
4 rats

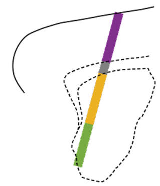

M1
dStr
vStr
2.4 mm ML
5 rats

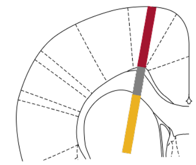

FOF
dStr
1.9 mm AP
3 rats

very noisy data:

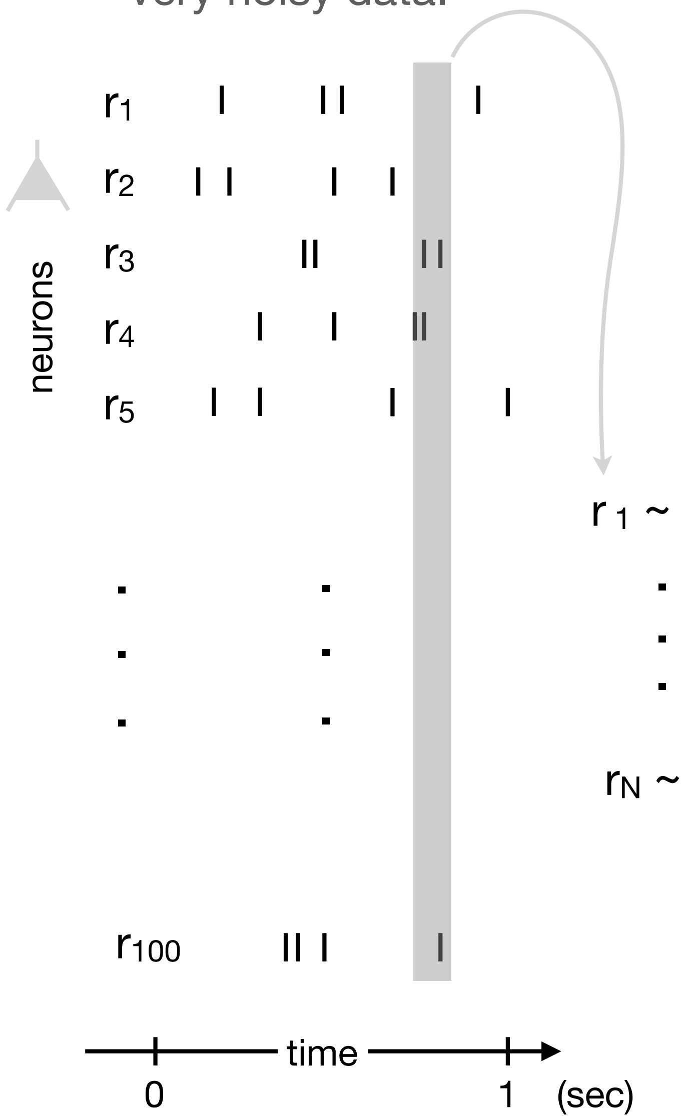

High-yield Neuropixels recordings (1 probe, 2-3 regions/rat)

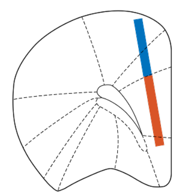

dmFC
mPFC
4.0 mm AP
4 rats

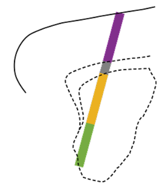

M1
dStr
vStr
2.4 mm ML
5 rats

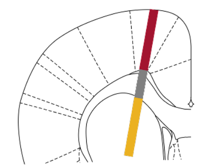

FOF
dStr
1.9 mm AP
3 rats

neural data well-described by few "latent" variables z

rectifying nonlinearity

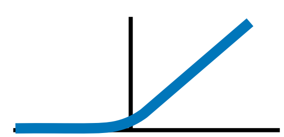

very noisy data:

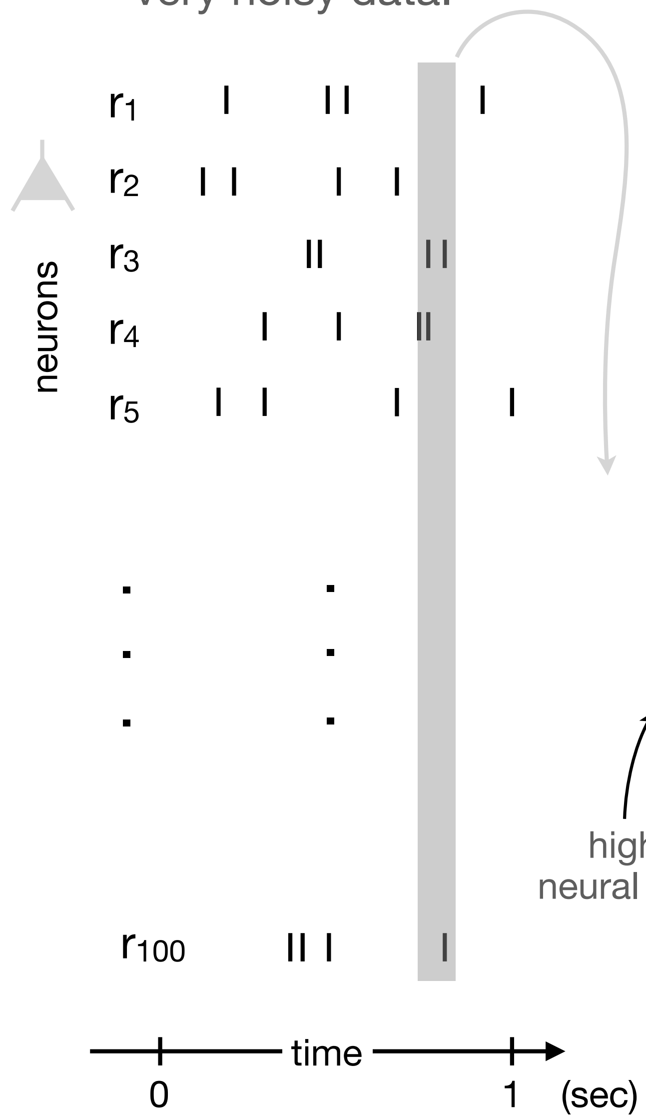

High-yield Neuropixels recordings (1 probe, 2-3 regions/rat)

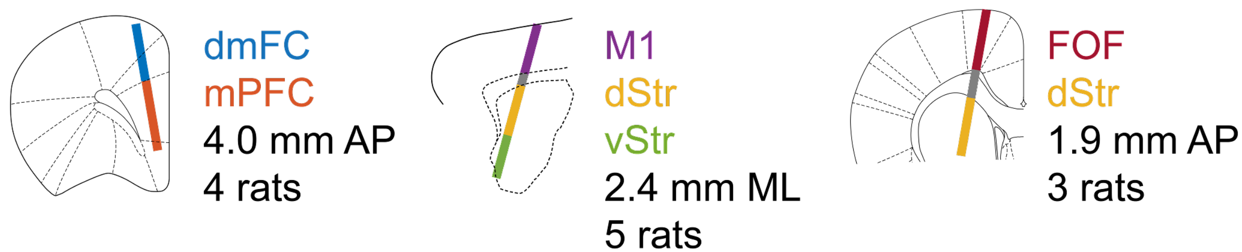

neural data well-described by few “latent” variables z

fit to data

$$\mathbf{r} \sim \text{Poisson}\left[h\left(\mathbf{C}\cdot\mathbf{z}+\mathbf{b}\right)\right]$$

high-D
neural space

pointwise nonlinearity

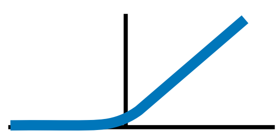

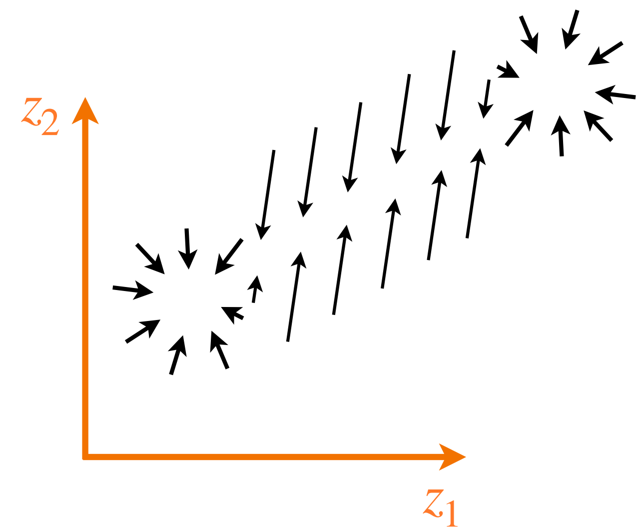

fit to data

$$\mathbf{r} \sim \text{Poisson}\left[h\left(\mathbf{C} \cdot \mathbf{z} + \mathbf{b}\right)\right]$$

high-D
neural space

inputs

$$\dot{\mathbf{z}} \equiv \frac{d\mathbf{z}}{dt} = F(\mathbf{z}, u = \uparrow \downarrow)$$

autonomous dynamics : $F(\mathbf{z}, u = 0)$

input dynamics : $F(\mathbf{z}, u = \uparrow \downarrow) - F(\mathbf{z}, u = 0)$

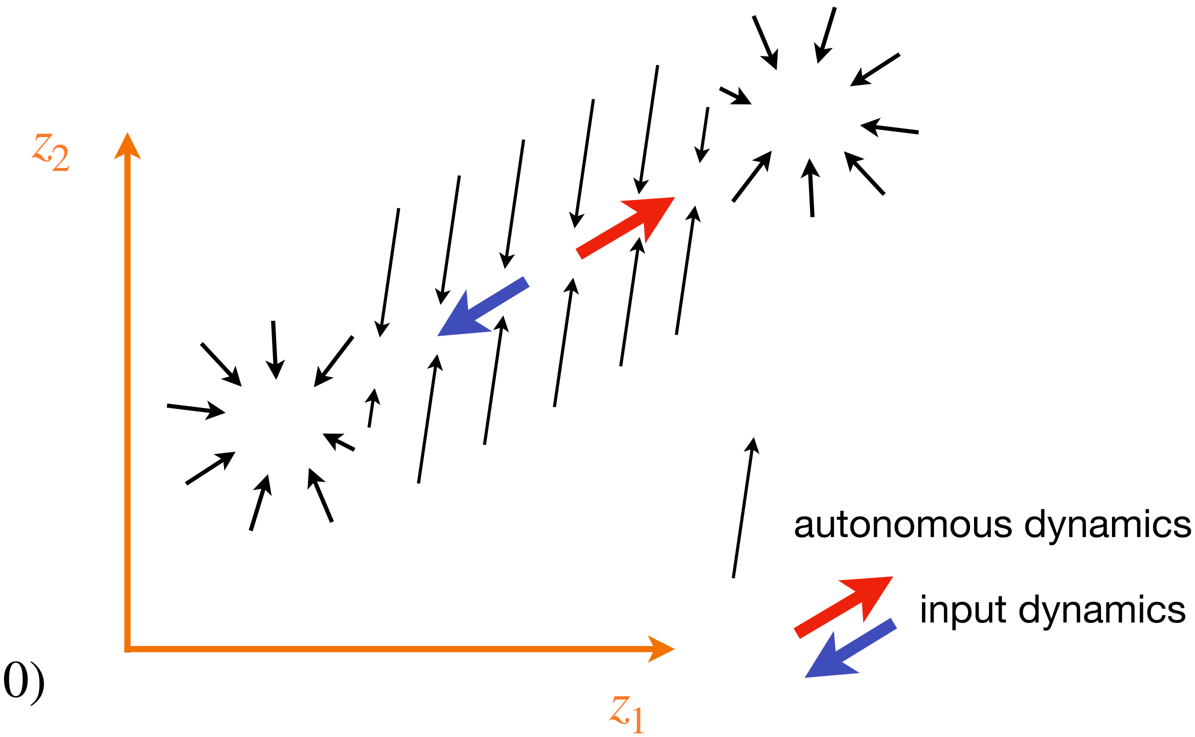

fit to data

$$\mathbf{r} \sim \text{Poisson}\left[h\left(\mathbf{C}\cdot\mathbf{z} + \mathbf{b}\right)\right]$$

high-D
neural space

fit to data   inputs

$$\dot{\mathbf{z}} = F\,(\,\mathbf{z}, u = \uparrow\downarrow\,) + \eta$$

autonomous dynamics : $F\,(\,\mathbf{z}, u = 0)$

input dynamics : $F\,(\,\mathbf{z}, u = \uparrow\downarrow\,) \;-\; F\,(\,\mathbf{z}, u = 0)$

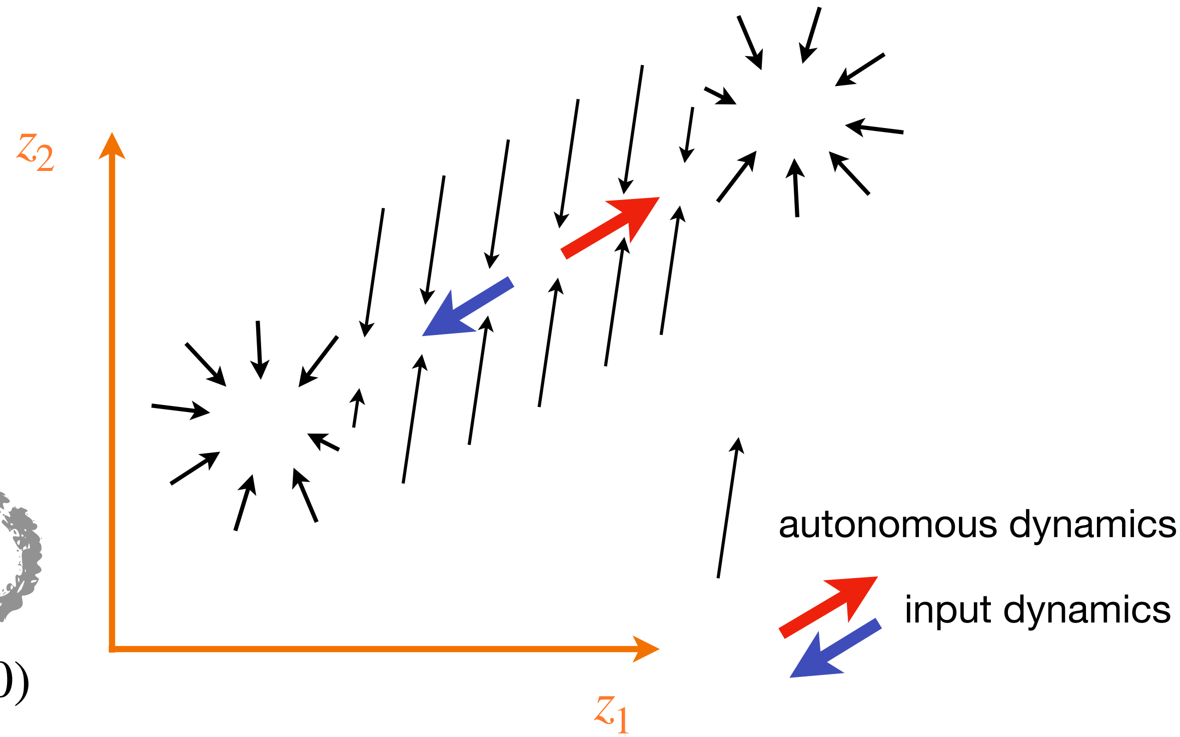

- LFADS (Pandarinath, … Sussillo, 2018) : *F()* is a 500-neuron RNN
- Genkin, … Engel (2023), Duncker, … Sahani (2019) : not yet equipped to take time-dependent inputs
- rSLDS (Scott Linderman's group, used in Nair …. Anderson 2023) : fit our data poorly
- Kim et al., "FINDR" (2023) : *F()* is parametrized by a deep FFNN — $\mathbf{z}$ stays low-d

"Neural ODEs" : Weinan 2017; Chen 2018

+ Stochastic diff. equ. (Li, … Duveneaud 2020) : stochastic dynamics for $\mathbf{z}$
+ Poisson observations
+ Non-differentiable pulsatile inputs (clicks)
+ *F()* is a gated NN (Kim et al. 2023)

$$\dot{\mathbf{z}} = F\,(\,\mathbf{z}, u = \uparrow\downarrow\,)$$

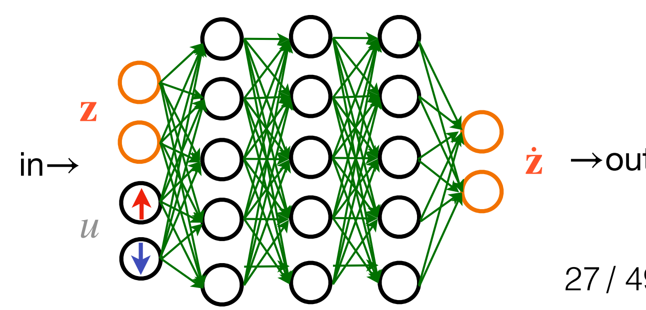

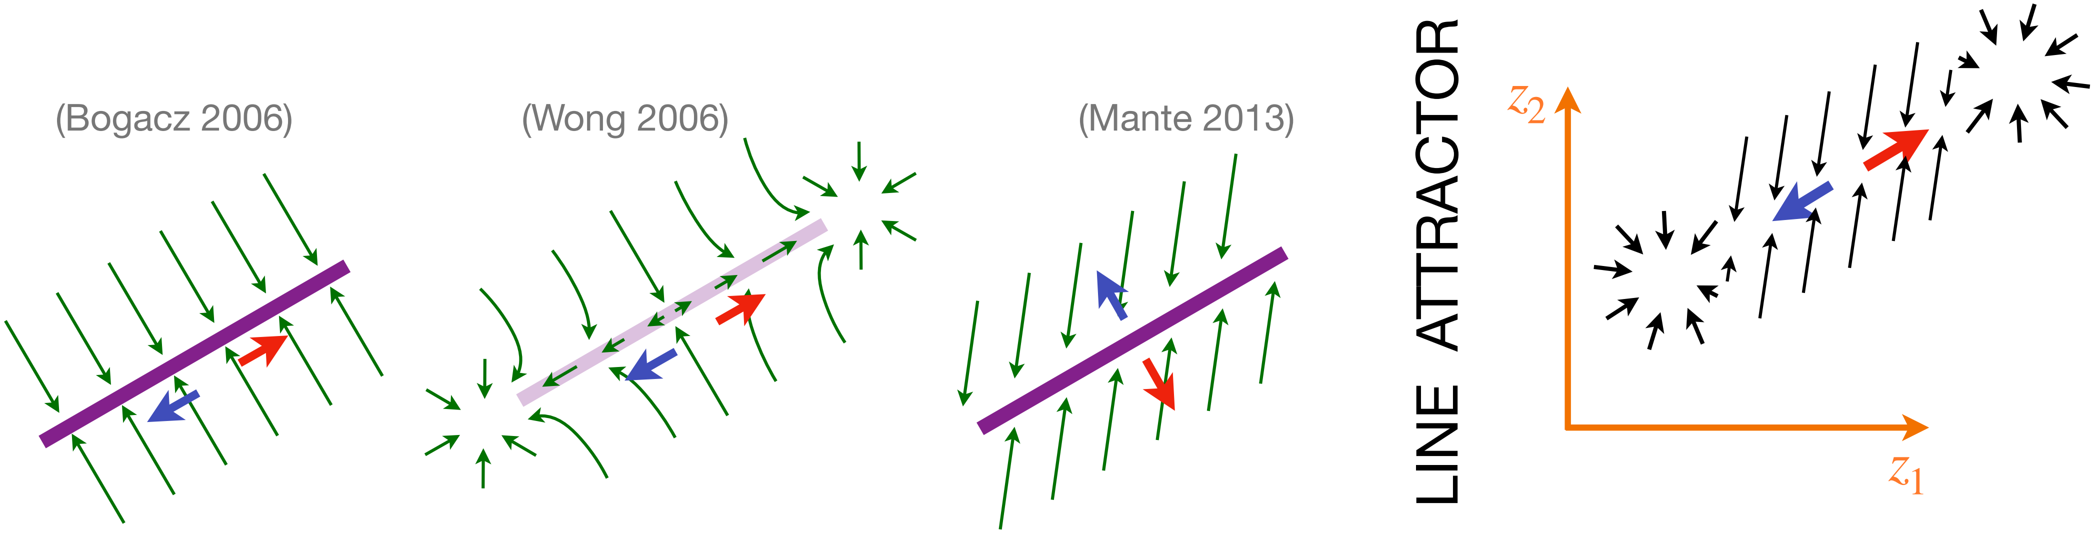

OUR DATA:

autonomous weak, inputs strong

inputs weak, autonomous strong

autonomous dynamics : $F(\mathbf{z}, u = 0)$

input dynamics : $F(\mathbf{z}, u = \uparrow\downarrow) - F(\mathbf{z}, u = 0)$

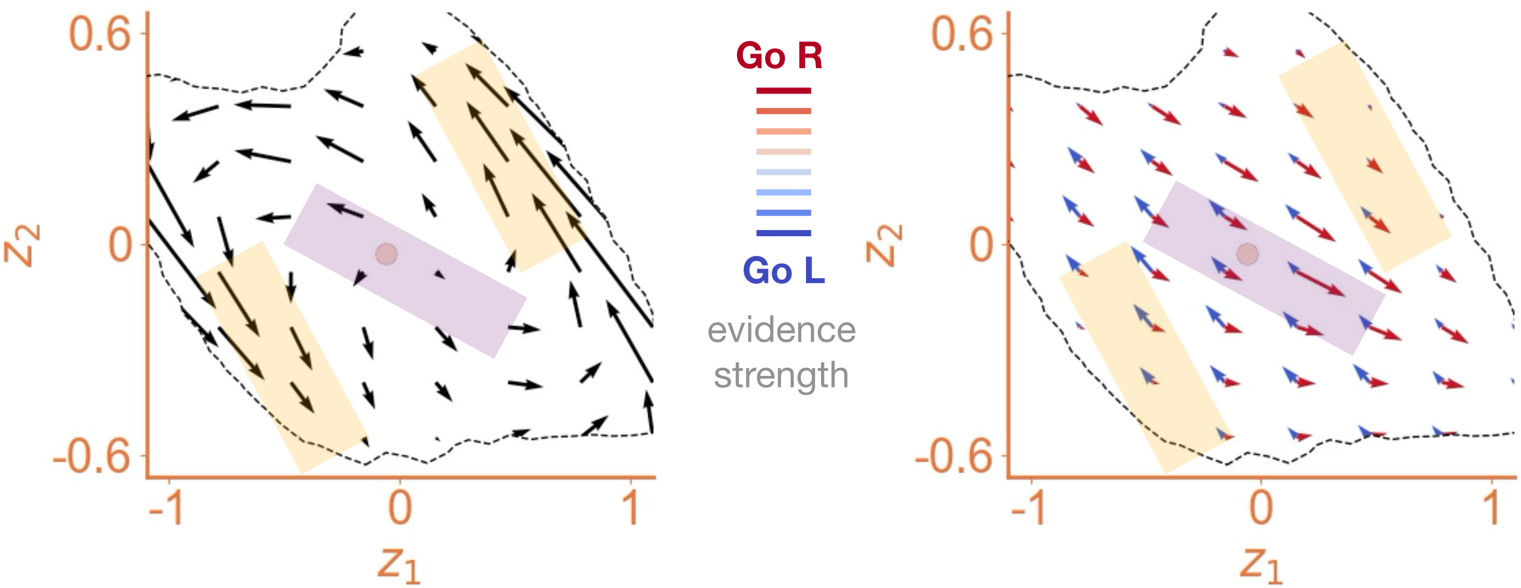

behavioral DDM:

accumulation
inputs drive dynamics

commitment
inputs are **irrelevant** to the dynamics

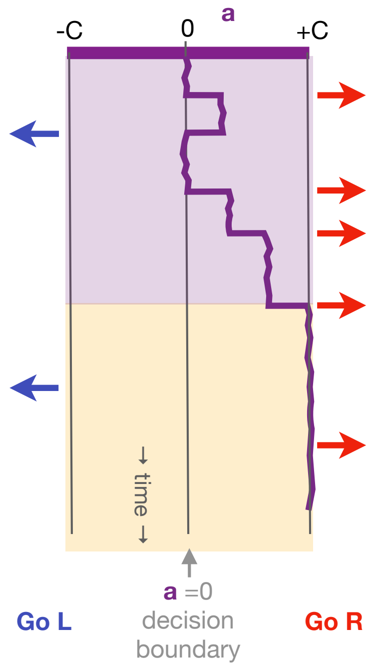

subject has made up their mind and committed to a decision?

OUR DATA:

autonomous weak, inputs strong

inputs weak, autonomous strong

autonomous dynamics : $F(\mathbf{z}, u = 0)$

input dynamics : $F(\mathbf{z}, u = \uparrow\downarrow) - F(\mathbf{z}, u = 0)$

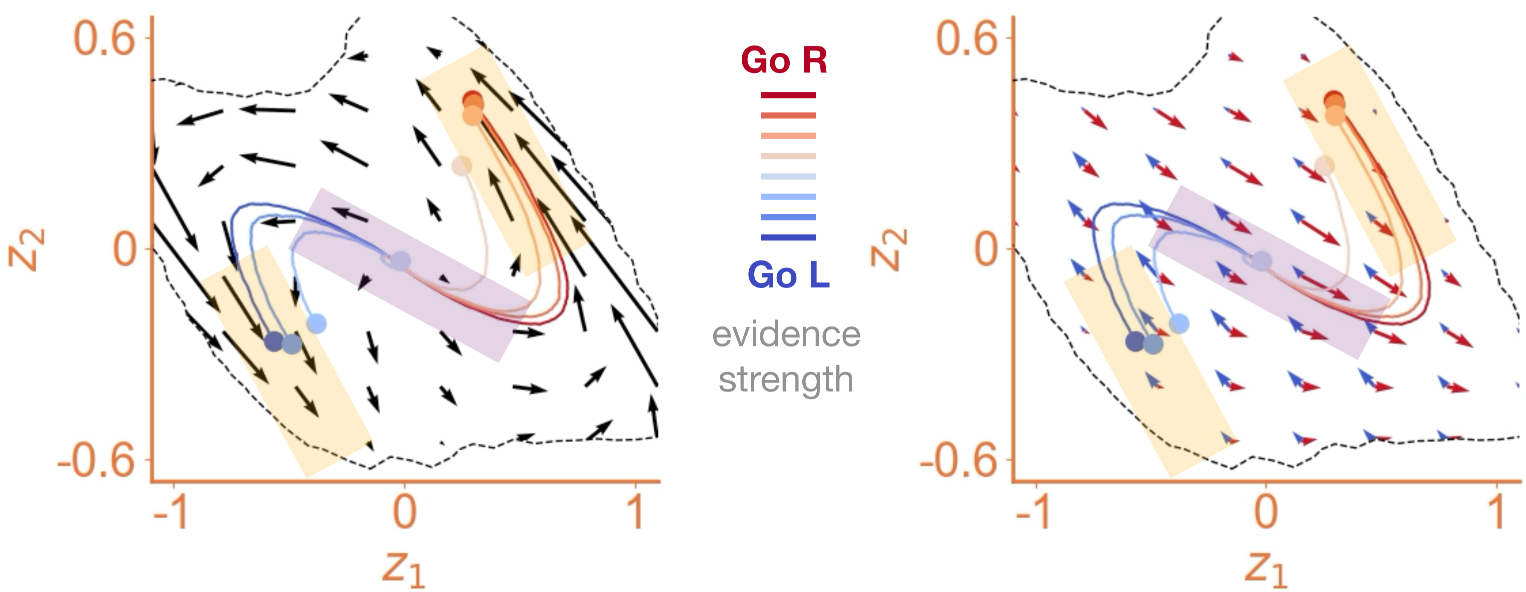

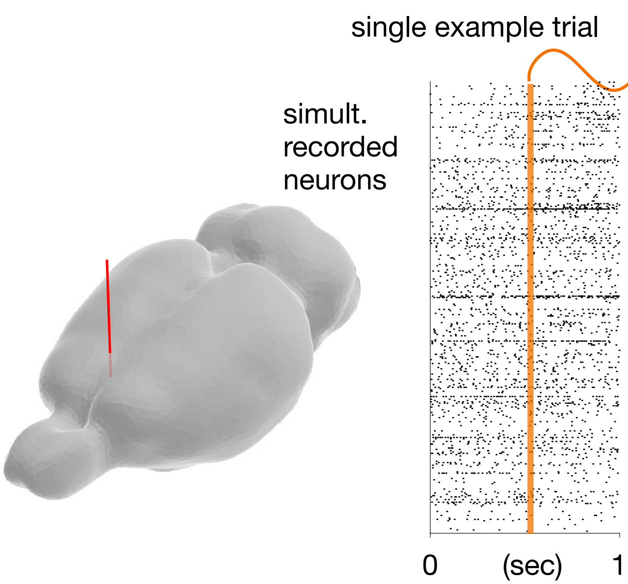

“neurally-inferred Time of commitment” (nTc)

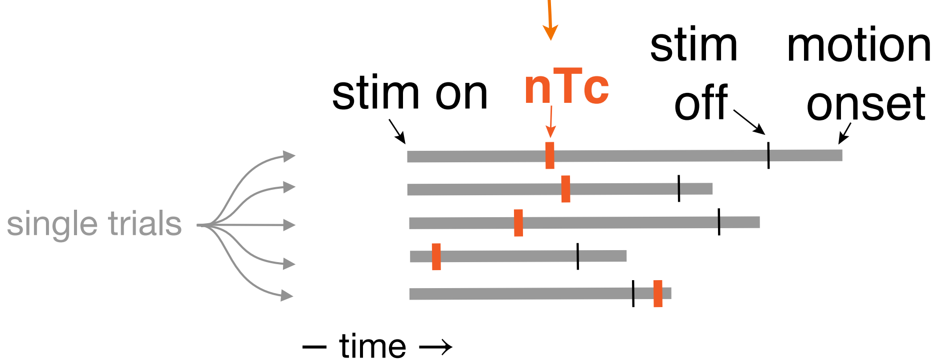

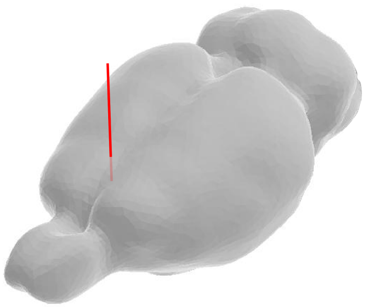

“neurally-inferred Time of commitment” (nTc)

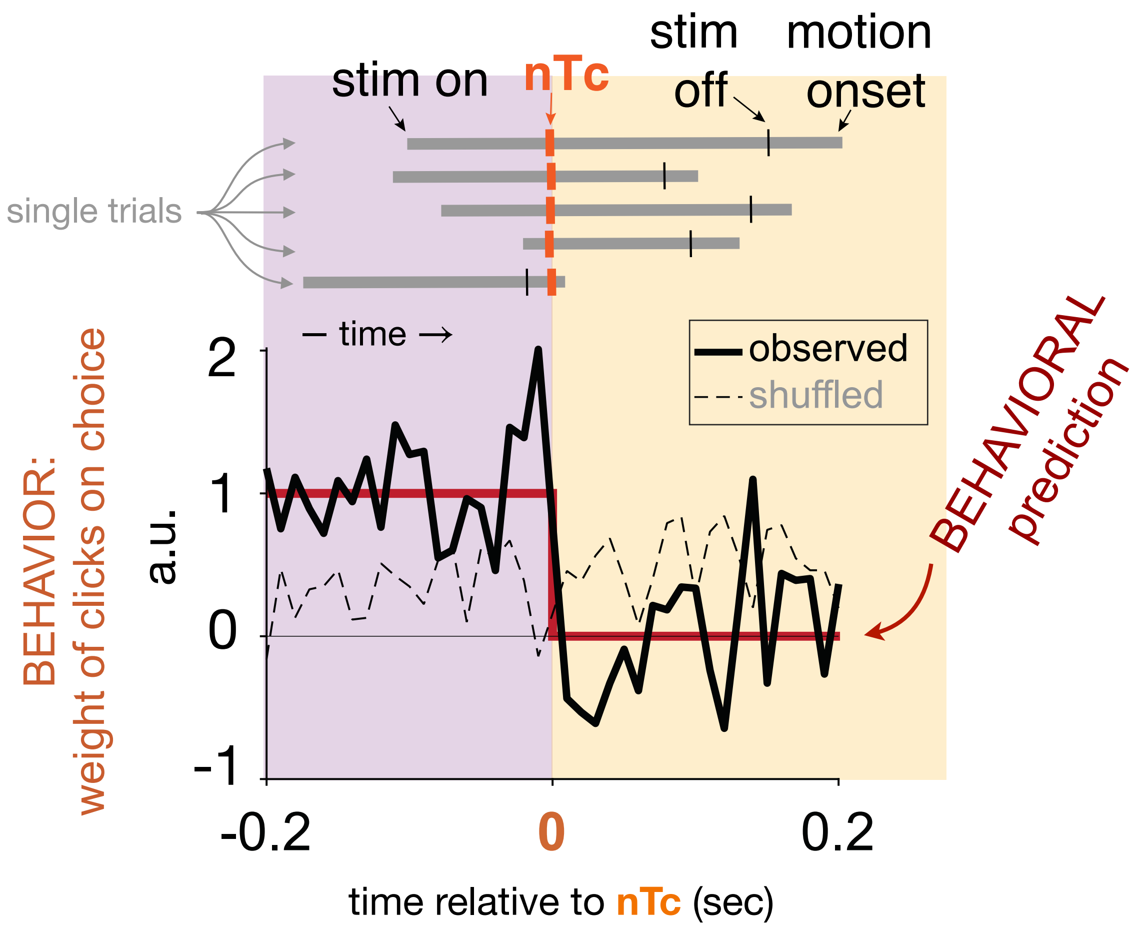

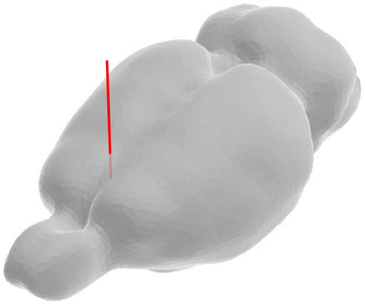

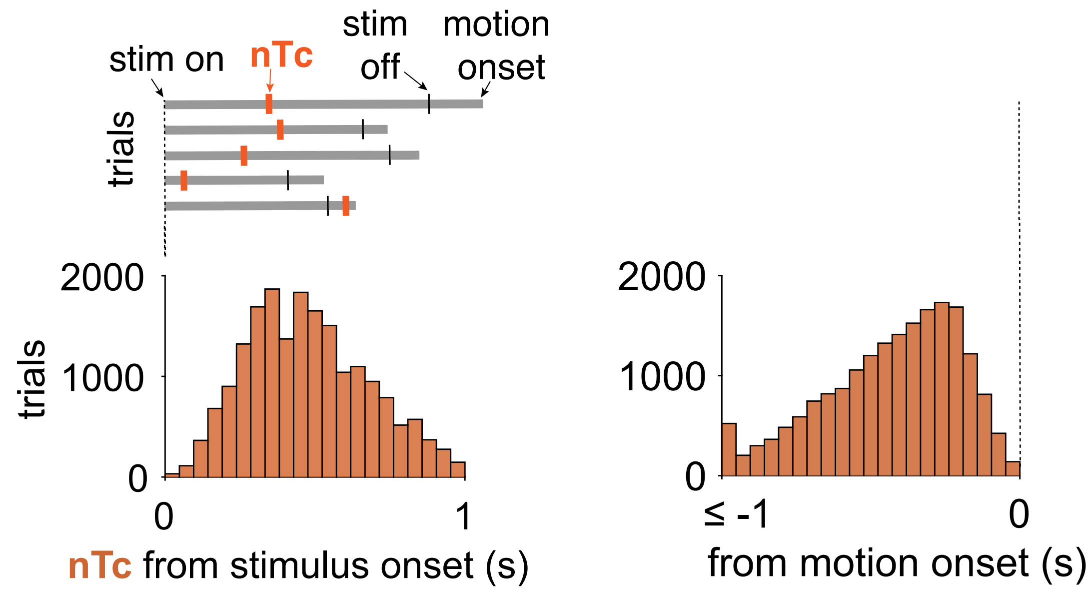
stim on
nTc
stim off
motion onset
trials
2000
1000
0
0
1
nTc from stimulus onset (s)
2000
1000
0
≤ -1
0
from motion onset (s)

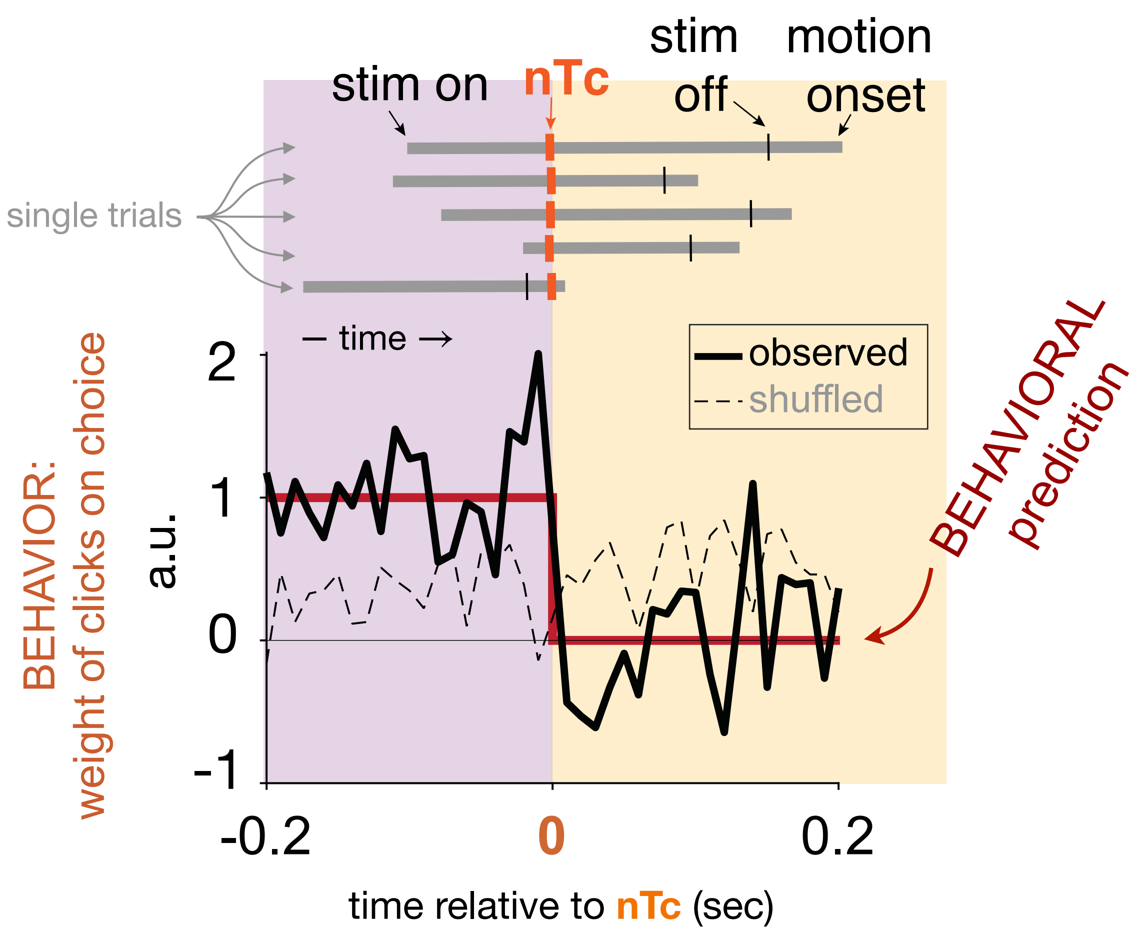
stim on
nTc
stim off
motion onset
single trials
— time →
observed
shuffled
BEHAVIOR:
weight of clicks on choice
a.u.
2
1
0
-1
-0.2
0
0.2
time relative to nTc (sec)
BEHAVIORAL prediction

“**neurally-inferred Time of commitment” (nTc)** : a neural biomarker for covertly making up one’s mind

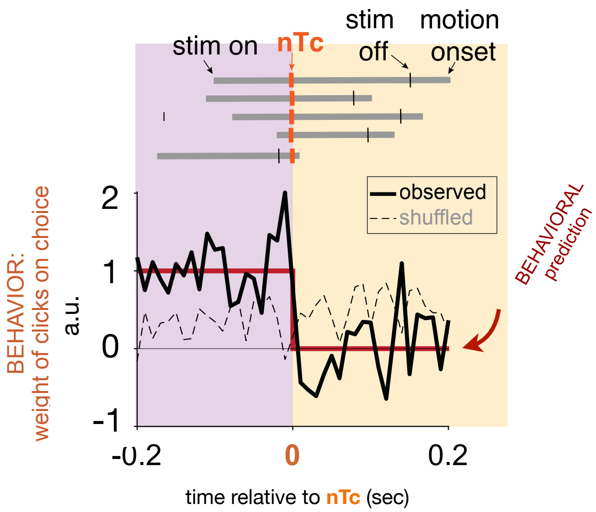

Luo*, Kim* ’23

- new flow line methods led us to **nTc**
- timing of **nTc** appears to be internally determined
- Can now read neural activity and tell if and when a subject *secretly* makes up their mind

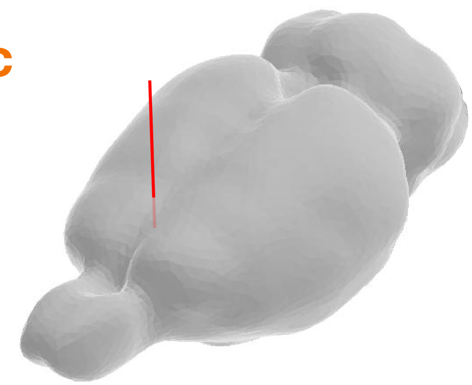

(Idea 1: multi-neuron recordings for better moment-by-moment estimates of internal signals)

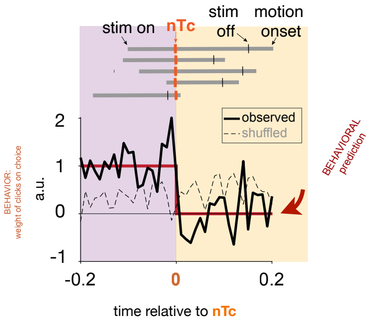

Luo*, Kim* '23

- new flow line methods led us to **nTc**
- Can now read neural activity and tell *if* and *when* a subject secretly makes up their mind

8 *simultaneous*
Neuropixels probes

~2,700 neurons/session

for animals engaged in a task, world record # simultaneous ephys recorded neurons

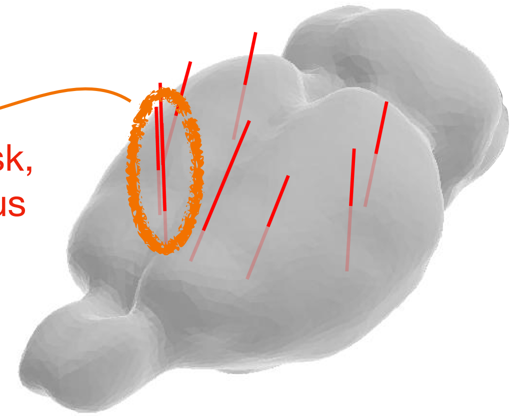

("**n**eurally-inferred **T**ime of **c**ommitment")

Bondy*, Charlton, Luo* '24

frontally-estimated

How does nTc affect the rest of the brain?

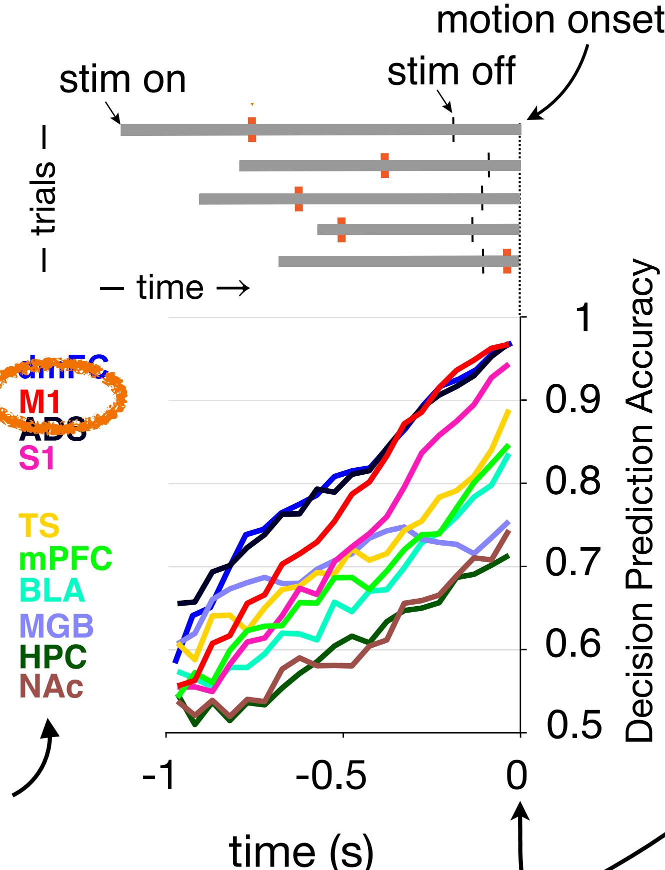

brain region

onset of motion to report decision

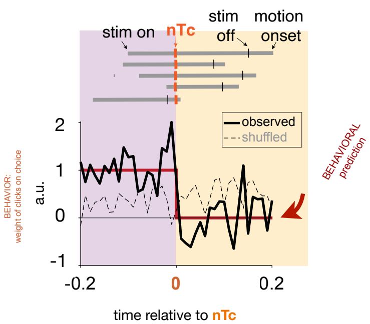

Luo*, Kim* '23

- new flow line methods led us to **nTc**
- Can now read neural activity and tell *if* and *when* a subject secretly makes up their mind

8 *simultaneous*
Neuropixels probes

~2,700 neurons/session

for animals engaged in a task,
world record # simultaneous
ephys recorded neurons

Bondy*, Charlton, Luo* '24

**("neurally-inferred
Time of commitment")**

How does nTc affect the rest of the brain?

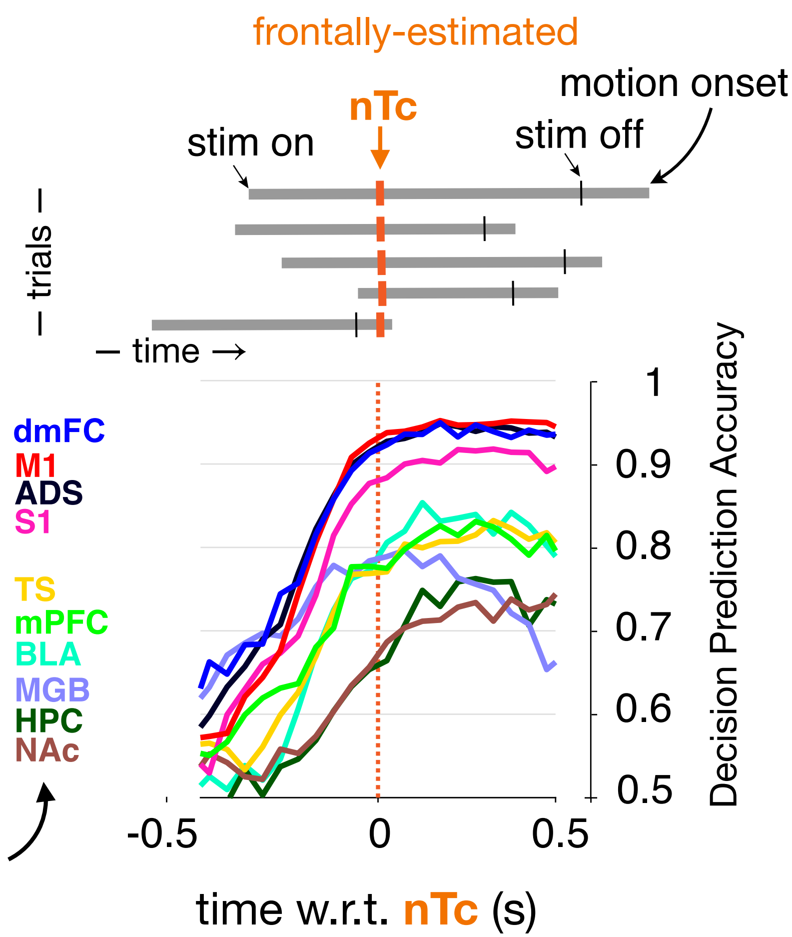

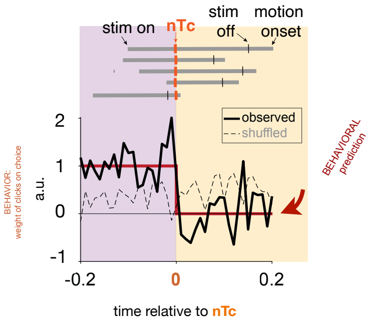

Luo*, Kim* ’23

- new flow line methods led us to **nTc**
- Can now read neural activity and tell *if* and *when* a subject secretly makes up their mind

8 *simultaneous* Neuropixels probes

~2,700 neurons/session

for animals engaged in a task, world record # simultaneous ephys recorded neurons

Bondy*, Charlton, Luo* ’24

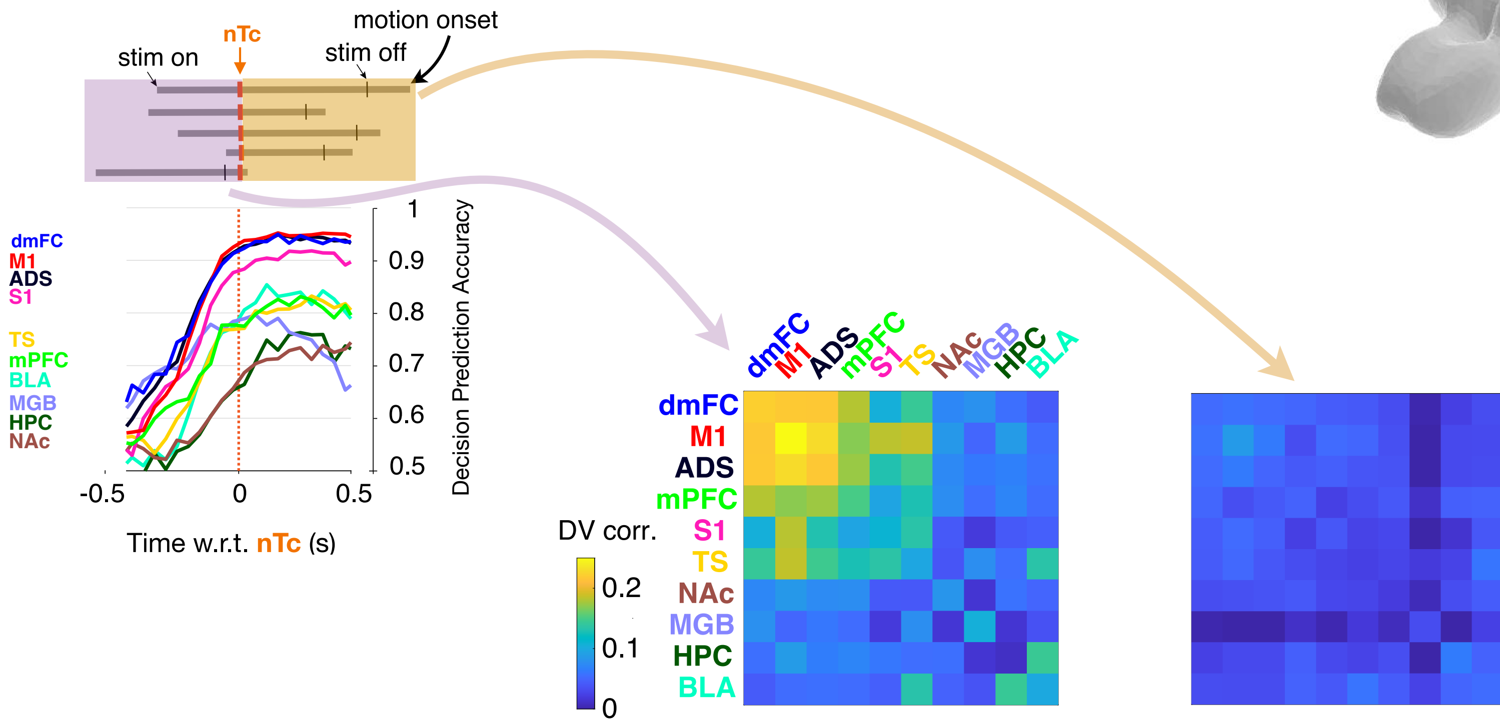

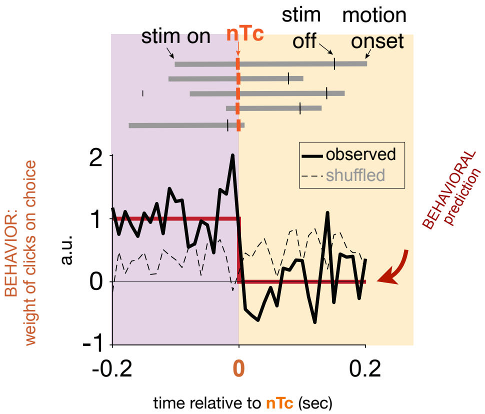

Luo*, Kim* ’23

- discovery of **nTc**, a neural biomarker for covert decision commitment

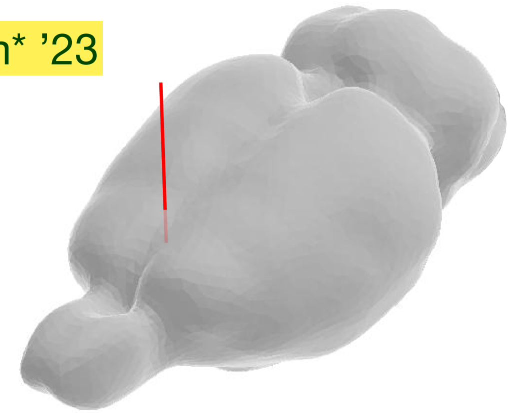

• nTc marks a sweeping state change *across* the brain

previous analyses, which were not sensitive to **nTc**, were blurring entirely disparate data together

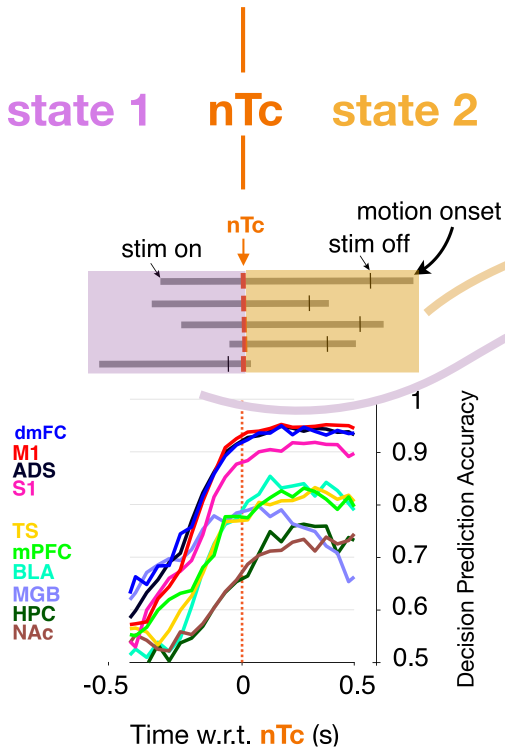

(note: uninstructed movements don’t affect these conclusions at all)

Bondy*, Charlton, Luo* ’24

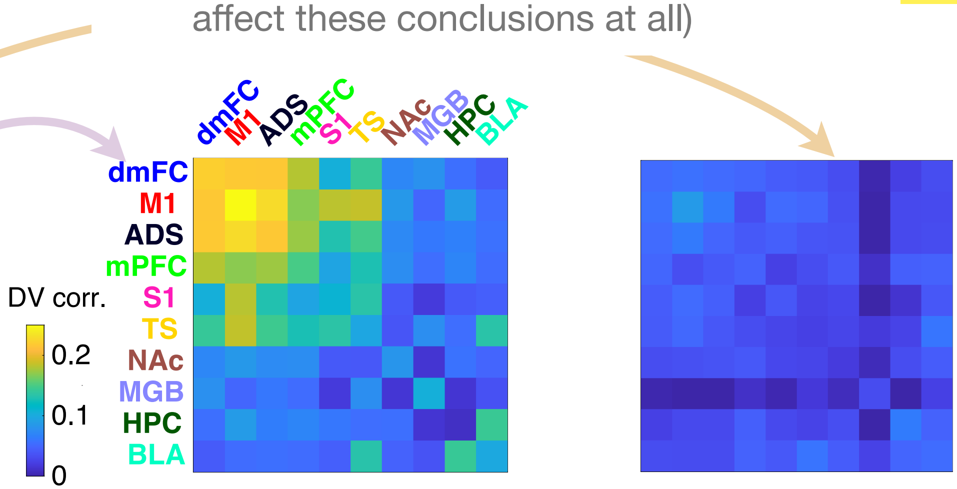

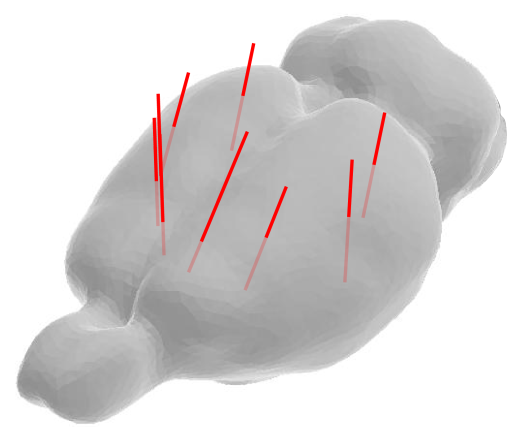

dmFC-ADS

DV corr. ρ=0.25 .

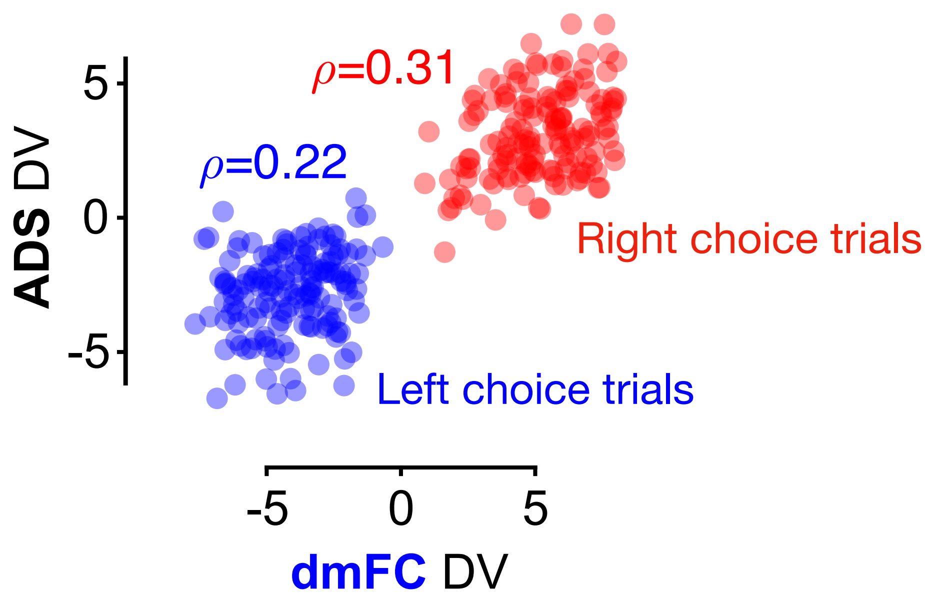

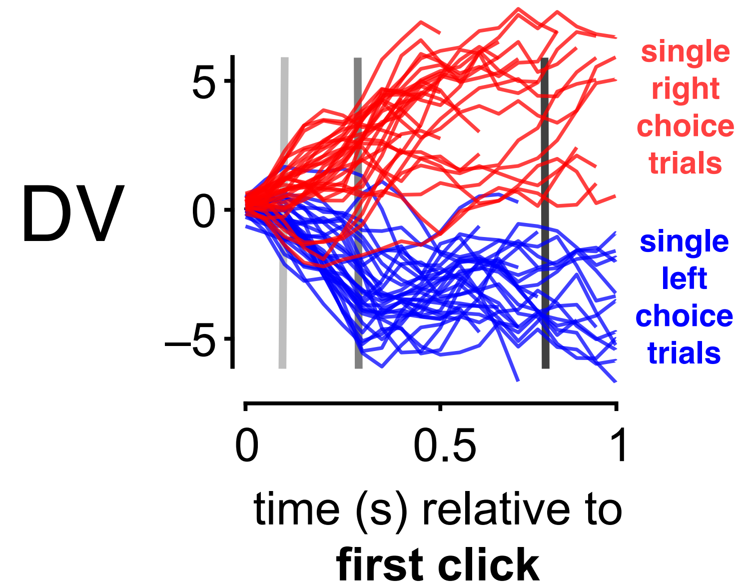

"Decision Variable" for each region

firing rates of neurons in one region

$$\log \frac{P(\text{go R})}{P(\text{go L})}(t) \equiv \text{DV}(t) = w_1^t\, r_1^t + w_2^t\, r_2^t + \ldots + w_N^t\, r_N^t$$

weights that best predict upcoming **R** v **L** choice

Kiani et al., 2014

“Frozen noise” task design

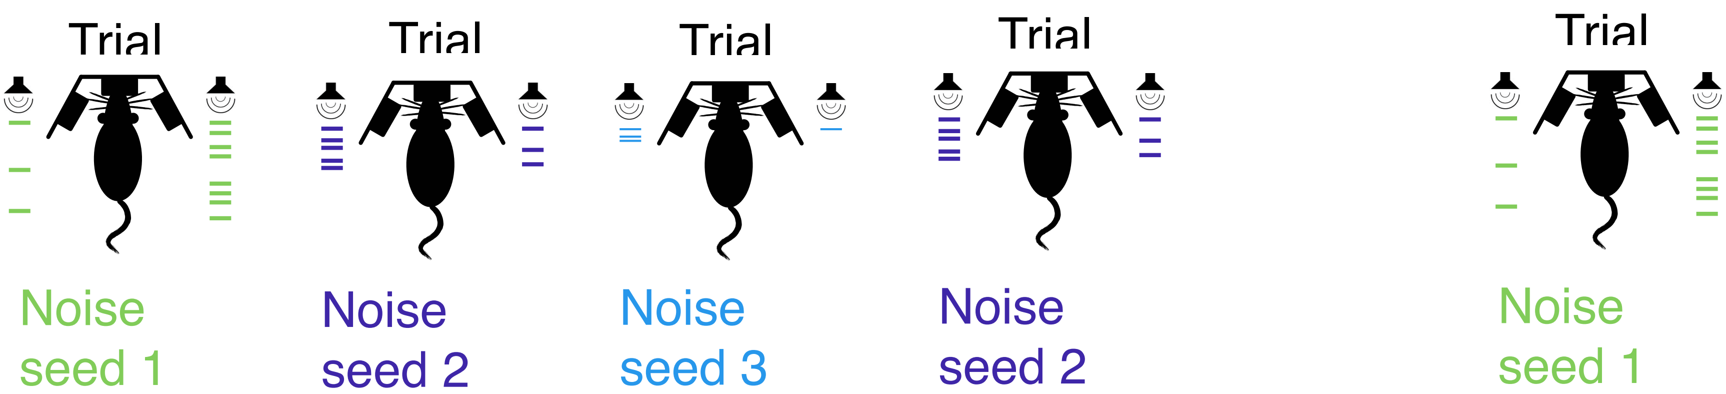

54 unique seeds,
~10 repeats each per session

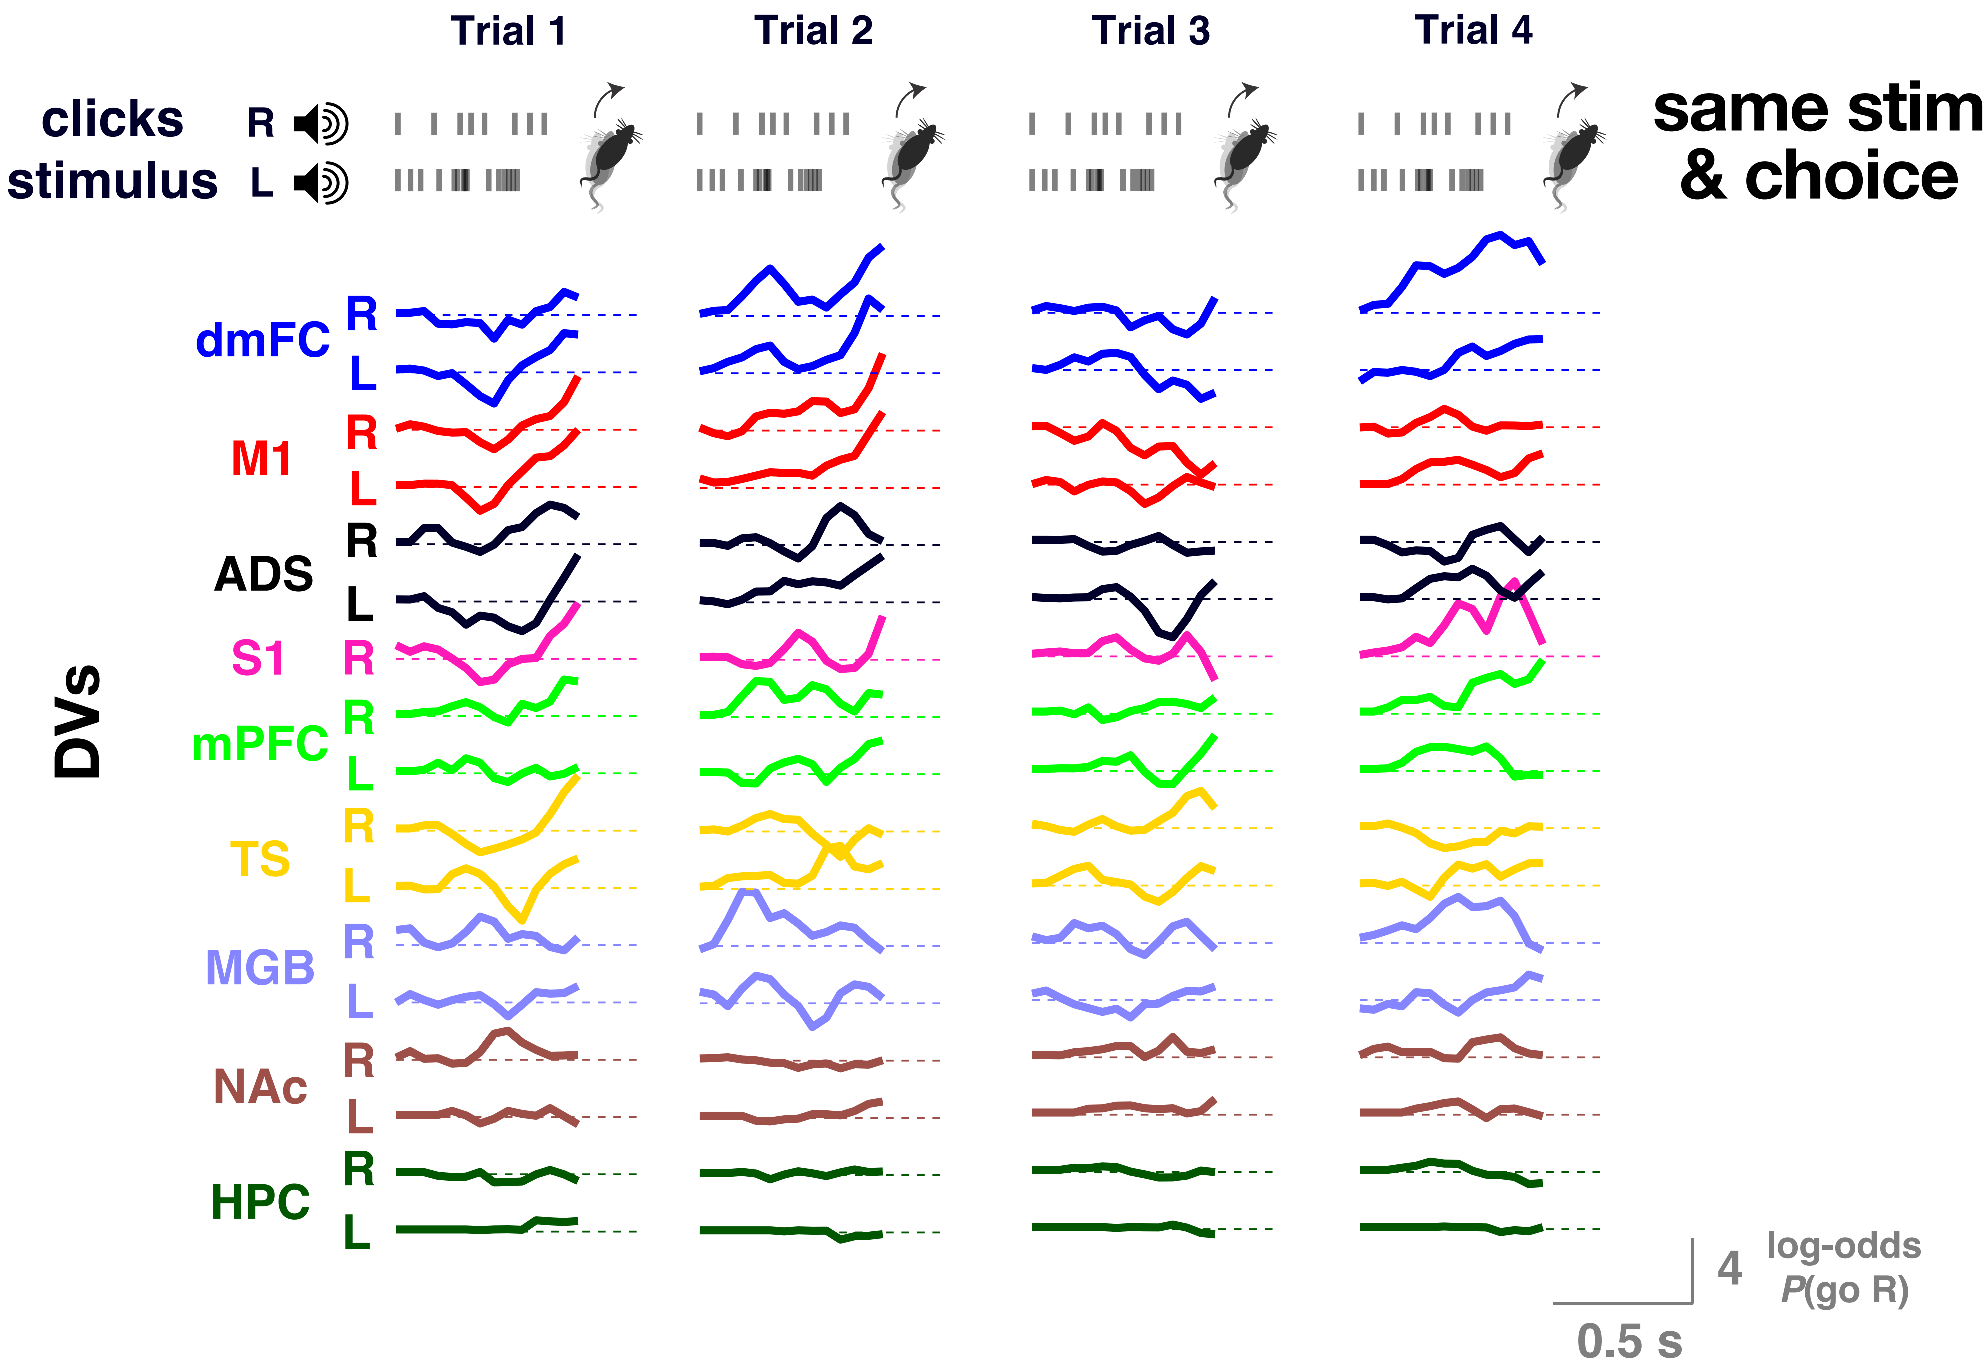

Trial 1
Trial 2
Trial 3
Trial 4
clicks
stimulus
R
L
same stim
& choice
DVs
dmFC
M1
ADS
S1
mPFC
TS
MGB
NAc
HPC
R
L
4
log-odds
P(go R)
0.5 s

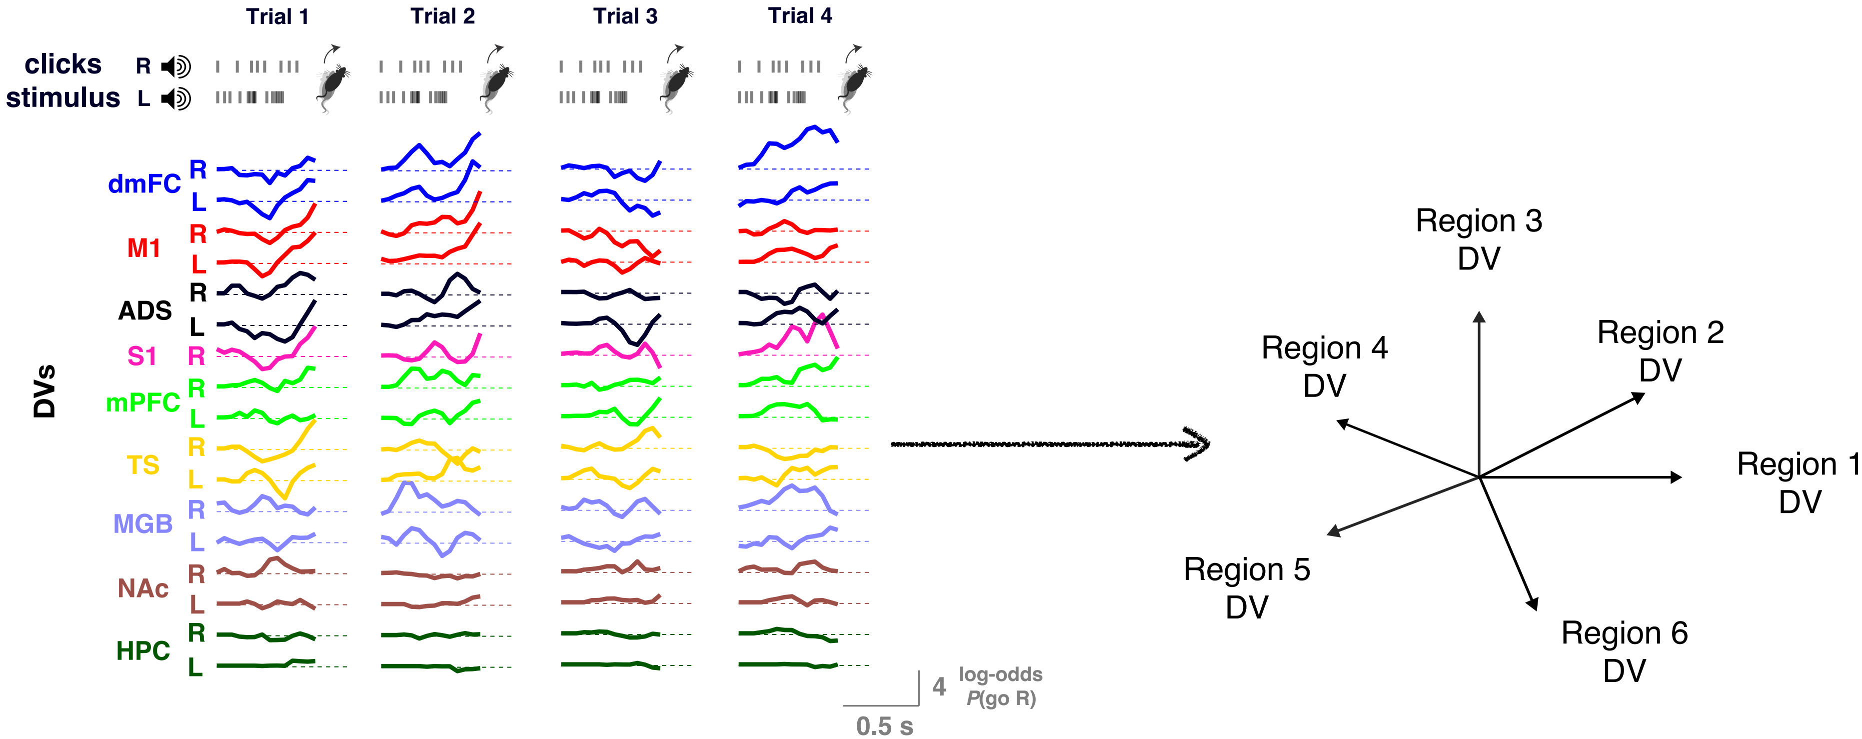
Trial 1
Trial 2
Trial 3
Trial 4
clicks stimulus
R
L
DVs
dmFC
M1
ADS
S1
mPFC
TS
MGB
NAc
HPC
4
log-odds P(go R)
0.5 s
Region 1 DV
Region 2 DV
Region 3 DV
Region 4 DV
Region 5 DV
Region 6 DV

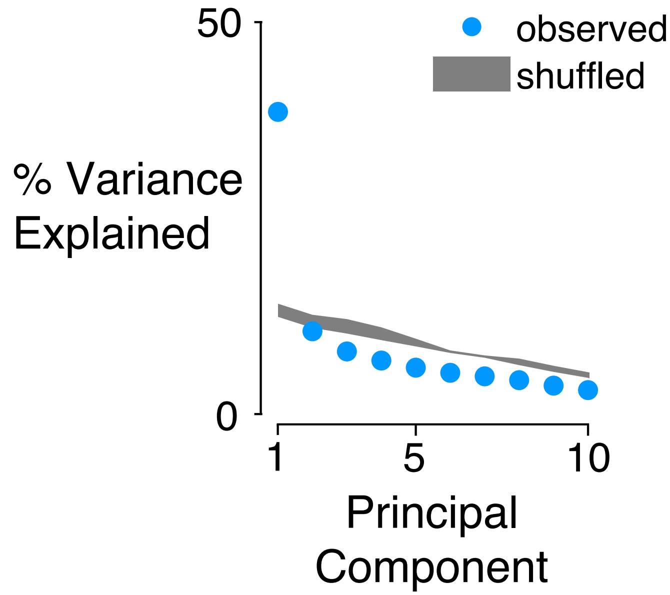
observed
shuffled
% Variance Explained
50
0
1
5
10
Principal Component

What did we discover with the multiprobe recordings?

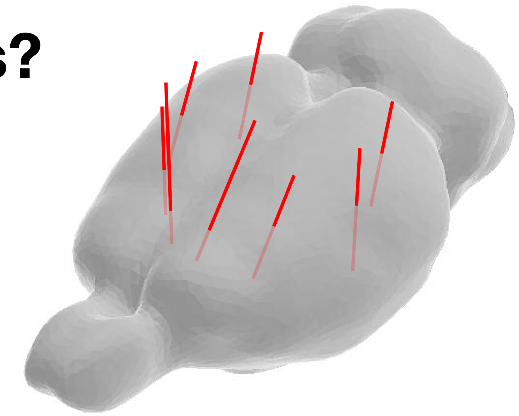

Can't see either one without the multiprobe recordings

1. Before **nTc**, decision-aligned activity is highly correlated across the brain, with correlations dominated by a single dimension.

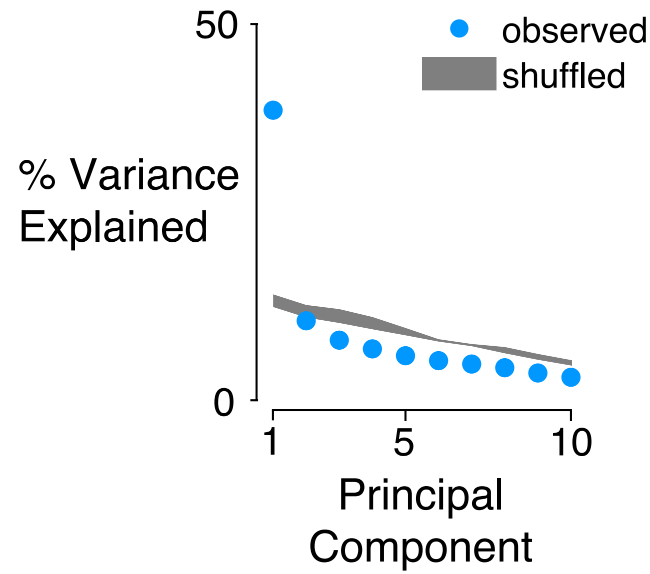

2. **nTc** marks a previously unknown, sweeping state change across the brain

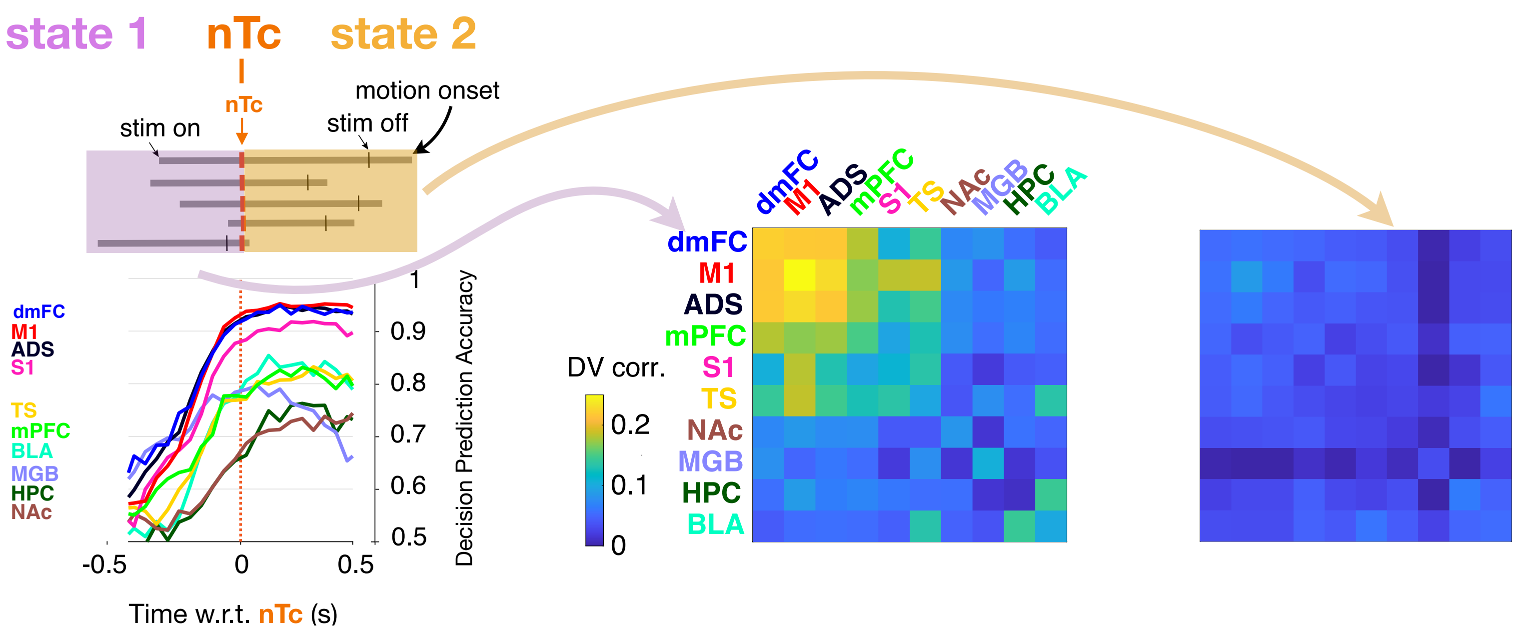

Two ideas running in the background:

Idea 1: multi-neuron recordings for better moment-by-moment estimates of internal signals

Idea 2: most neural activity is of unknown function, yet is coordinated across neurons and brain regions. What's going on with this "dark matter of the brain?"

ongoing and future work: **moving beyond decision-making …**

nTc is an example of something much broader

Neural activity variance:

1.5 %

explained by known
task variables, i.e., what we study

(International Brain Lab 2023; our data)

~10 % correlated with
uninstructed movements
(but we don't know why)

(Musall…Churchland 2019,
Wang… Svoboda Druckmann 2023;
our data)

~88 % ???

we… don't know

- most neural activity looks like noise but is coordinated across neurons and regions.

(Arieli 1996, Fiser 2004, Stringer 2019, Manley 2024)

what **nTc**, an *internal* covert signal,
looked like before we discovered it

hypothesis: could much neural activity consist of undiscovered internal signals?

- most neural activity looks like noise but is coordinated across neurons and regions.

hypothesis: could much neural activity consist of undiscovered internal signals?

~88 %

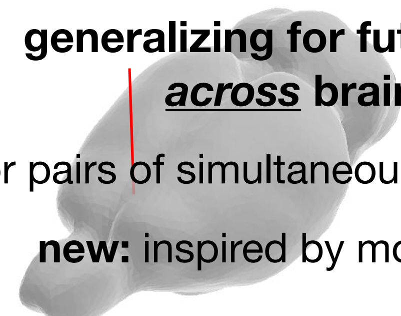

generalizing for future discoveries *across* brain regions

for pairs of simultaneously recorded regions:

new: inspired by modern *nonlinear* AI

how nTc was discovered:

- identify structure in simultaneous recordings
- characterize that structure

step 1

- characteristics → hypotheses of functional significance
- test hypotheses, find meaning

step 2

- state-of-the-art data-driven AI method to identify and separate private vs. shared latents
- infer intrinsic geometry of latents

step 1

- large-scale simultaneous recordings + new AI-based analysis methods give us access to these internal signals

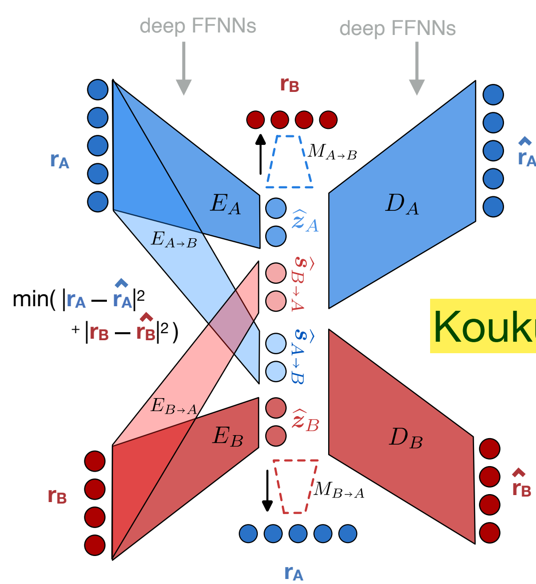

our "SPLICE" autoencoder architecture

Koukuntla, 2024

- most neural activity looks like noise but is coordinated across neurons and regions.

hypothesis: could much neural activity consist of undiscovered internal signals?

~88 %

how nTc was discovered:

- identify structure in simultaneous recordings
- characterize that structure

step 1

- characteristics → hypotheses of functional significance
- test hypotheses, find meaning

step 2

generalizing for future discoveries *<u>across</u>* brain regions

for pairs of simultaneously recorded regions:

new: inspired by modern *nonlinear* AI

- state-of-the-art data-driven AI method to identify and separate private vs. shared latents
- infer intrinsic geometry of latents

step 1

- large-scale simultaneous recordings + new AI-based analysis methods give us access to these internal signals

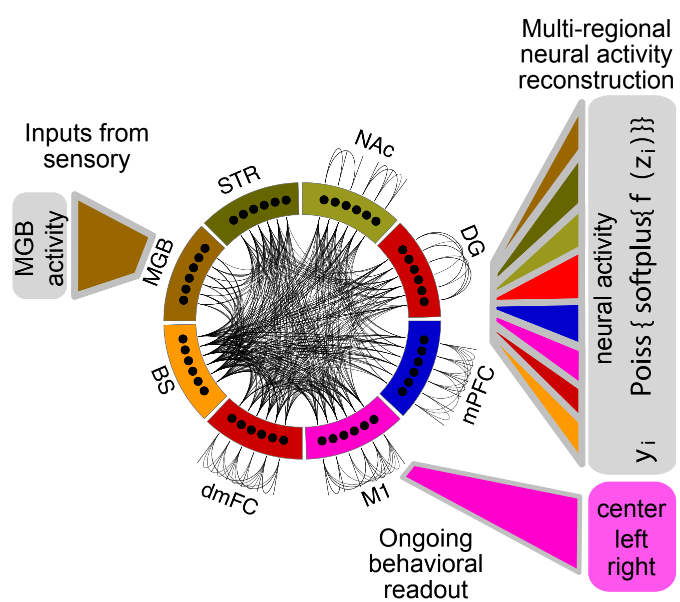

- dynamical models of *all* simultaneously recorded neurons across the brain, together, to understand cross-brain *single-trial* dynamics.

Summary from today

- most neural activity looks like noise but is coordinated across neurons and regions.
- discovered internal signal **nTc**, a neural biomarker for covert decision commitment
- **nTc** marks a sweeping state change across the brain

hypothesis: could much neural activity consist of undiscovered internal signals?

Quantitative cognitive rodent behavior

Pagan '22

Nieh*, Schottdorf* '21

- large-scale simultaneous recordings + new AI-based analysis methods give us access to these internal signals

Summary from today

- discovered internal signal **nTc**, a neural biomarker for covert decision commitment
- **nTc** marks a sweeping state change across the brain
- most neural activity looks like noise but is coordinated across neurons and regions.

  hypothesis: could much neural activity consist of undiscovered internal signals?
- large-scale simultaneous recordings + new AI-based analysis methods give us access to these internal signals

Quantitative cognitive rodent behavior

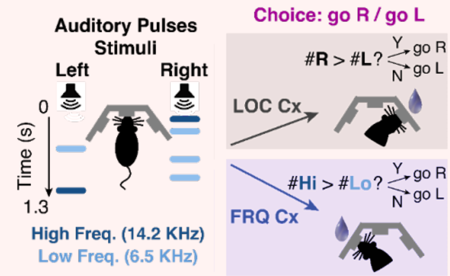

Pagan '22

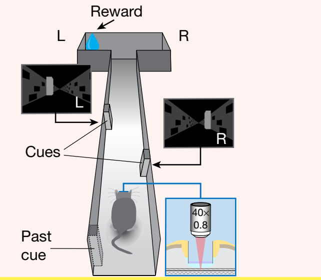

Nieh*, Schottdorf* '21

~88 % of activity

inner conversation of the mind?

Large-scale simultaneous recordings

Bondy*, Charlton, Luo* '24

AI-based large-scale data analysis

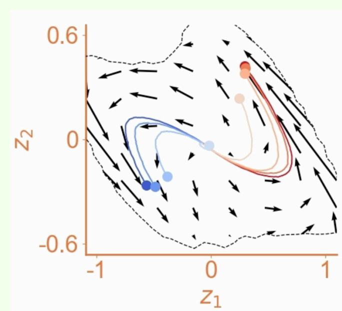

Luo*, Kim* '23

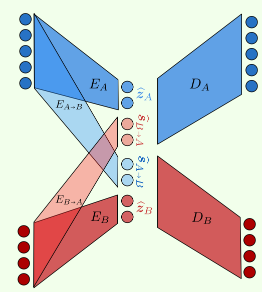

Koukuntla '24

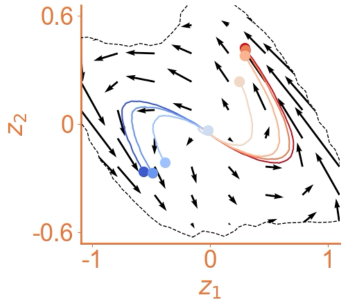

Luo*, Kim* '23

flow line estimation, discovery of **nTc**

Thomas Luo

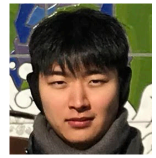

Tim Kim

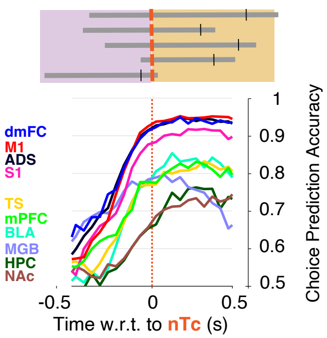

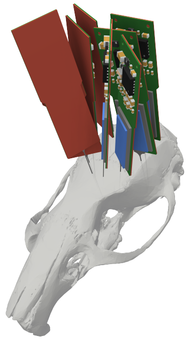

Bondy*, Charlton, Luo* '24

Cross-brain state change at the time of **nTc**

Adrian Bondy

Julie Charlton

Thomas Luo

Sarah Jo Venditto

Chuck Kopec

Wynne Stagnaro

Tim Harris (PI) @ Johns Hopkins

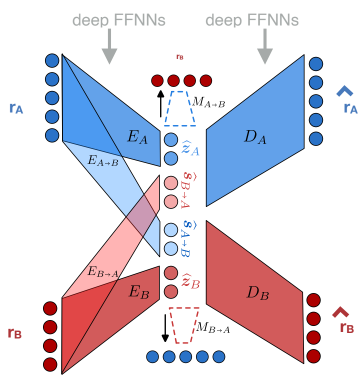

Koukuntla '24

AI-based self-supervised identification of latent structure in large-scale recordings

Sai Koukuntla

Adam Charles (PI)

@ Johns Hopkins

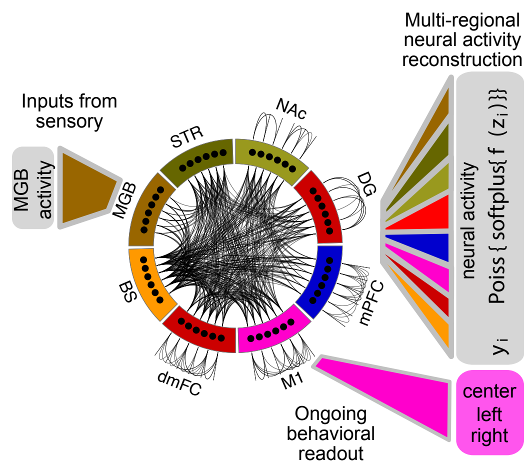

Lankow *in prep.*

large-scale dynamics models

Ben Lankow

Mark Goldman (PI)

@ UC Davis

Srdjan Ostojic (PI)

@ ENS

funding:

Howard Hughes Medical Institute

Simons Collaboration on the Global Brain

National Institutes Of Health