Constrained plasticity can compensate for ongoing drift in the parietal cortex

AR Loback, ME Rule, DV Raman, LN Driscoll, CD Harvey, TS O'Leary

NCCD, Capbreton, France, 2019

23 September 2019

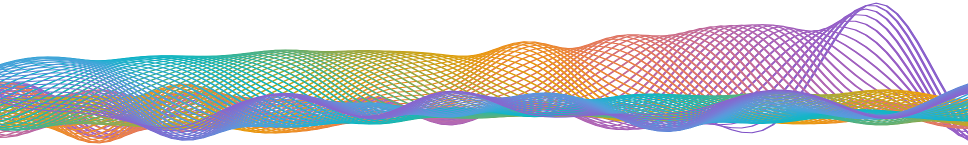

‘Drift’

Neurons in some sensorimotor areas reconfigure their tunings even after the task is learned (‘drift’).

‘Drift’

Neurons in some sensorimotor areas reconfigure their tunings even after the task is learned (‘drift’).

How can we interpret continuously changing neural codes?

‘Drift’

Neurons in some sensorimotor areas reconfigure their tunings even after the task is learned (‘drift’).

How can we interpret continuously changing neural codes?

What prevents this drift from disrupting task performance?

Driscoll et al. 2017

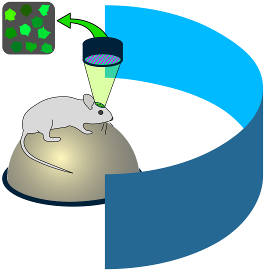

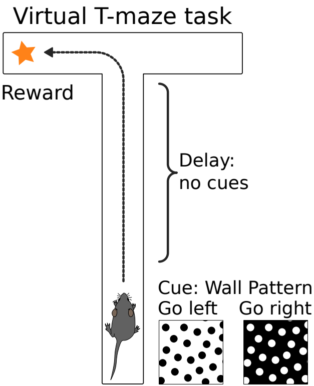

Driscoll et al. 2017

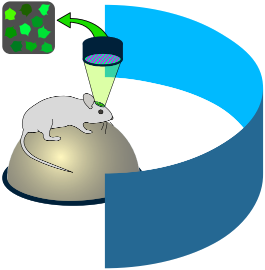

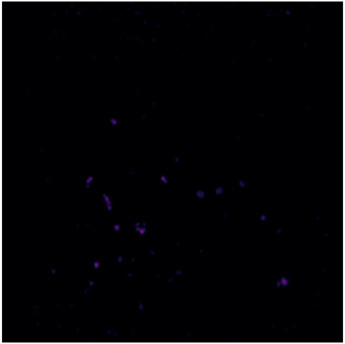

Normalized and filtered Ca2+ fluorescence

4× real-time

Neural tunings tile the task

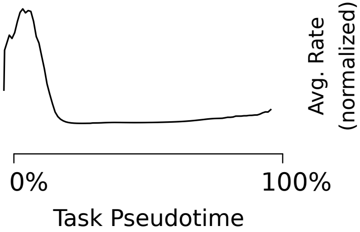

(cartoon example)

Neural tunings tile the task

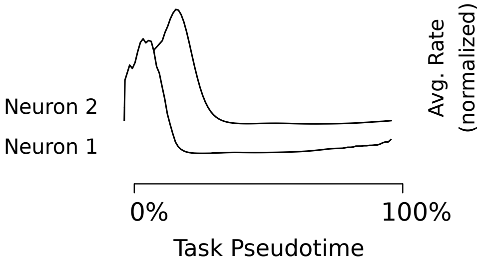

(cartoon example)

Neural tunings tile the task

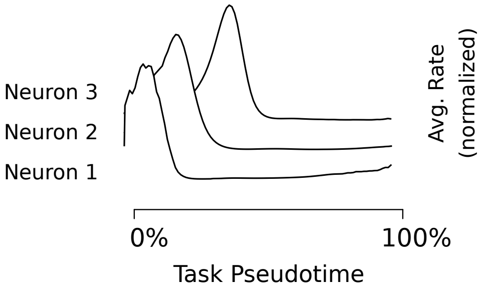

(cartoon example)

Neural tunings tile the task

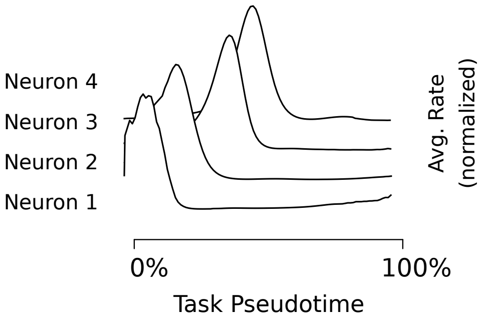

(cartoon example)

Neural tunings tile the task

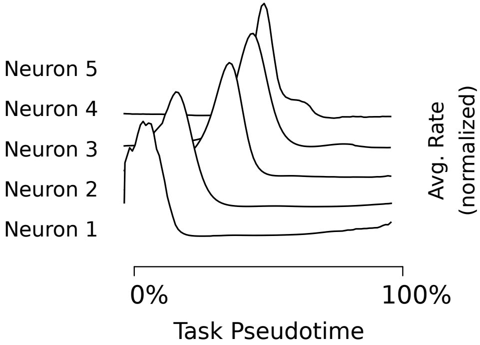

(cartoon example)

Neural tunings tile the task

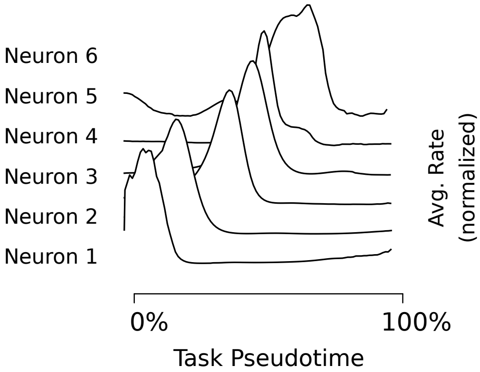

(cartoon example)

Neural tunings tile the task

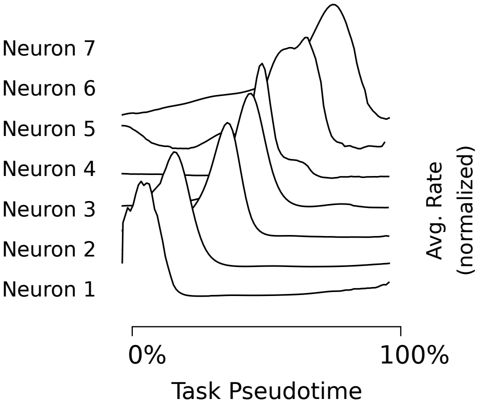

(cartoon example)

Neural tunings tile the task

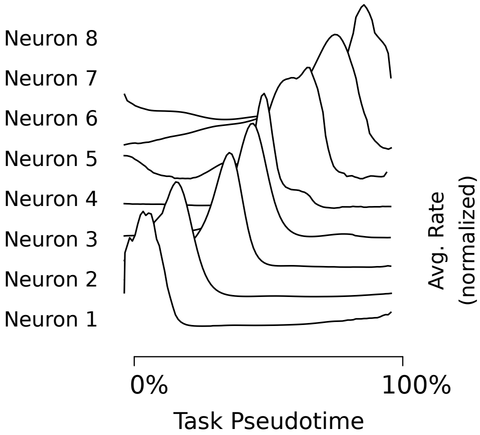

(cartoon example)

Neural tunings tile the task

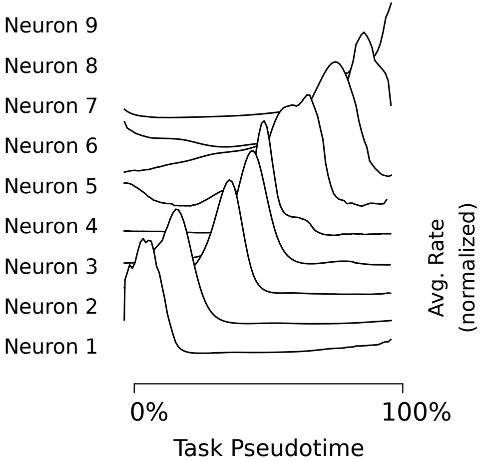

(cartoon example)

Task representation is not fixed

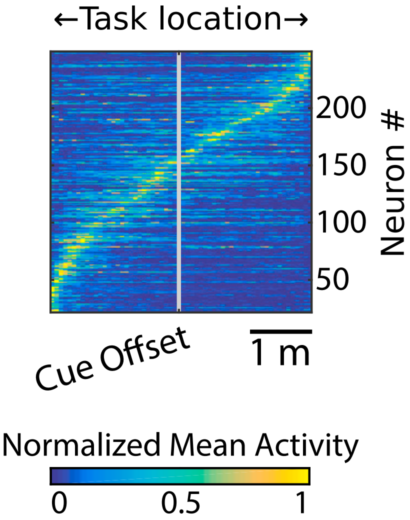

Day 1
Driscoll et al. 2017

Task representation is not fixed

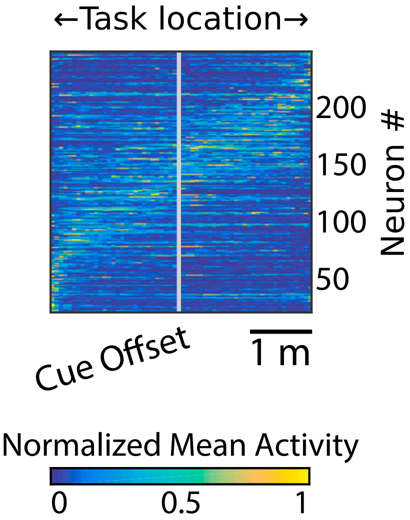

Day 10
Driscoll et al. 2017

Task representation is not fixed

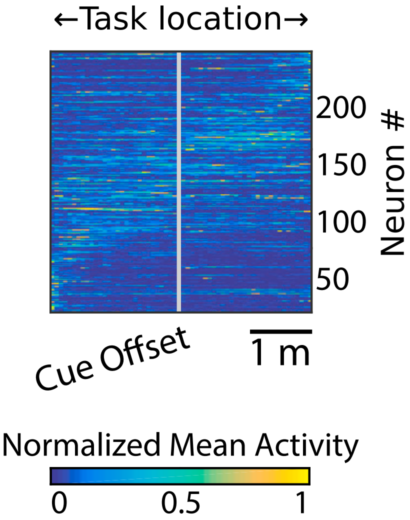

Day 20
Driscoll et al. 2017

Neural tunings change over days-weeks

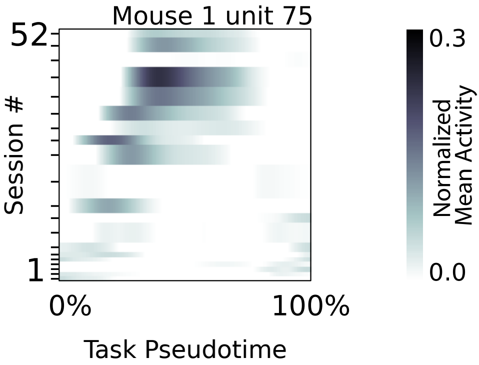

Neural tunings change over days-weeks

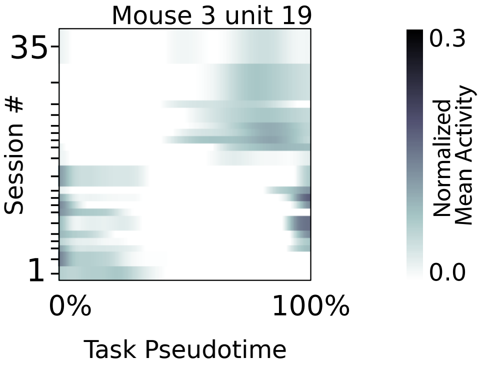

Neural tunings change over days-weeks

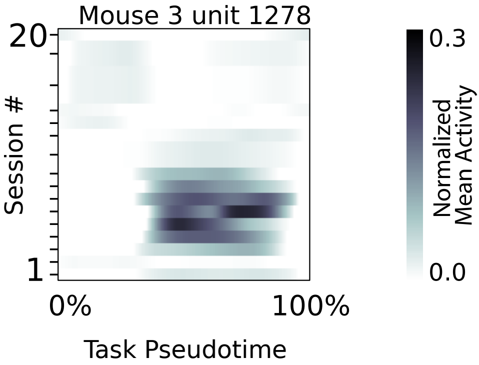

Neural tunings change over days-weeks

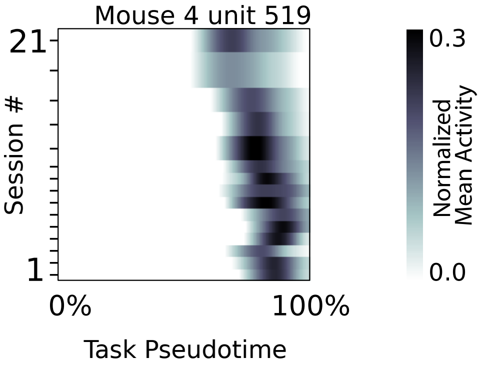

Neural tunings change over days-weeks

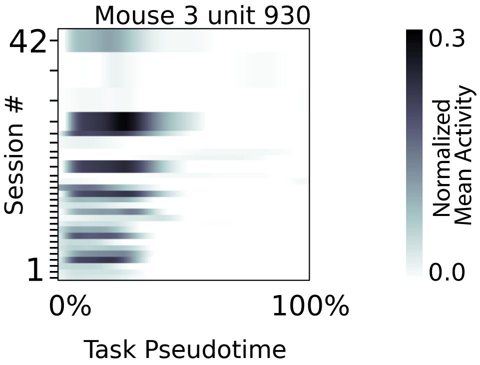

Task structure persists

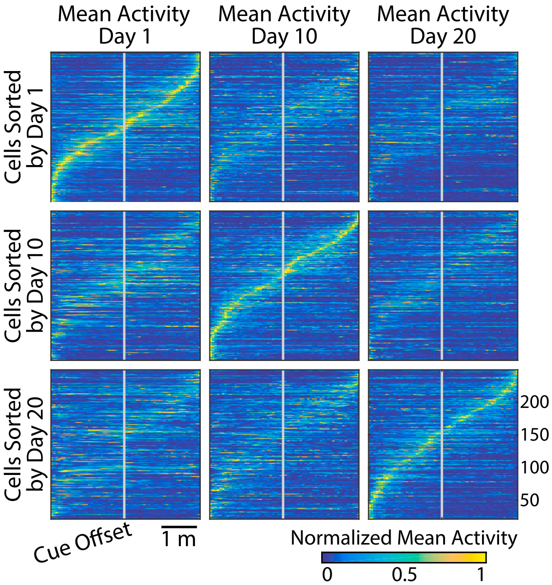

Task structure persists

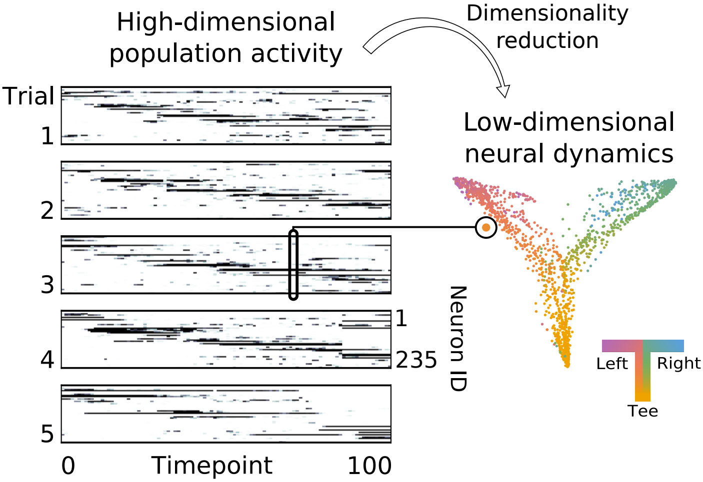

Task structure persists

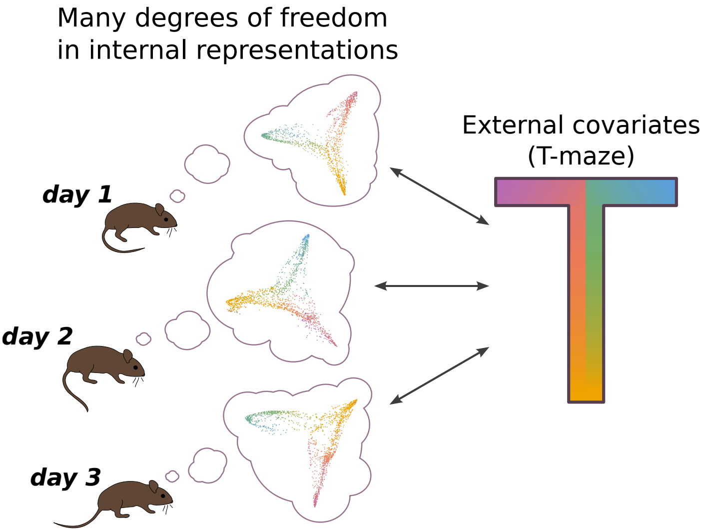

Reconcile drift with stable performance:

Task-relevant neural representation in parietal cortex **change**

Reconcile drift with stable performance:

Task-relevant neural representation in parietal cortex **change**

This change does not disrupt behavior

Reconcile drift with stable performance:

Task-relevant neural representation in parietal cortex **change**

This change does not disrupt behavior

What are the implications for neural coding?

Reconcile drift with stable performance:

Task-relevant neural representation in parietal cortex **change**

This change does not disrupt behavior

What are the implications for neural coding?

We will show that:

Reconcile drift with stable performance:

Task-relevant neural representation in parietal cortex **change**

This change does not disrupt behavior

What are the implications for neural coding?

We will show that:

Reconcile drift with stable performance:

Task-relevant neural representation in parietal cortex **change**

This change does not disrupt behavior

What are the implications for neural coding?

We will show that:

1. Shifts in mean activity are mostly irrelevant: they resemble noise

Reconcile drift with stable performance:

Task-relevant neural representation in parietal cortex **change**

This change does not disrupt behavior

What are the implications for neural coding?

We will show that:

1. Shifts in mean activity are mostly irrelevant: they resemble noise
2. Redundancy allows many codes, and there is a stable subspace

Reconcile drift with stable performance:

Task-relevant neural representation in parietal cortex **change**

This change does not disrupt behavior

What are the implications for neural coding?

We will show that:

1. Shifts in mean activity are mostly irrelevant: they resemble noise
2. Redundancy allows many codes, and there is a stable subspace
3. Modest plasticity would be required to track remaining drift

Reconcile drift with stable performance:

Task-relevant neural representation in parietal cortex **change**

This change does not disrupt behavior

What are the implications for neural coding?

We will show that:

1. **Shifts in mean activity are mostly irrelevant: they resemble noise**
2. Redundancy allows many codes, and there is a stable subspace
3. Modest plasticity would be required to track remaining drift

Shifts in mean activity are mostly irrelevant

Not all drift is disruptive

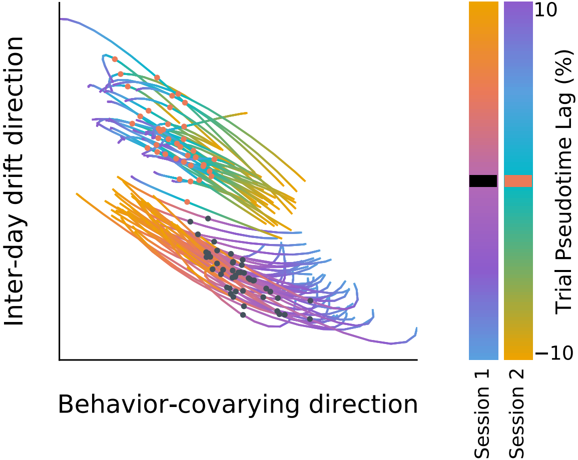

Mouse 3 session 2/3 pseudotime 40%±10

previous turn right, next turn right

Not all drift is disruptive

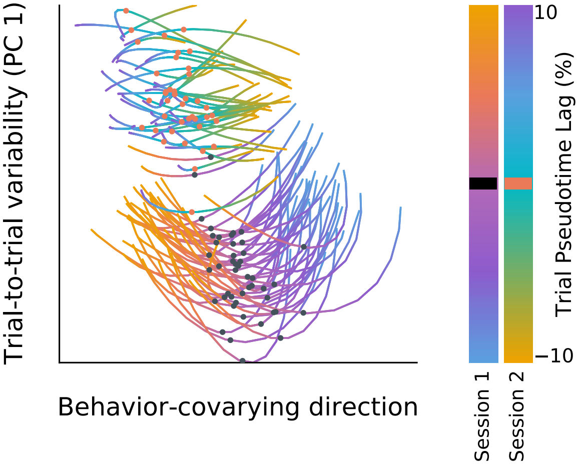

Mouse 3 session 2/3 pseudotime 40%±10

previous turn right, next turn right

Not all drift is disruptive

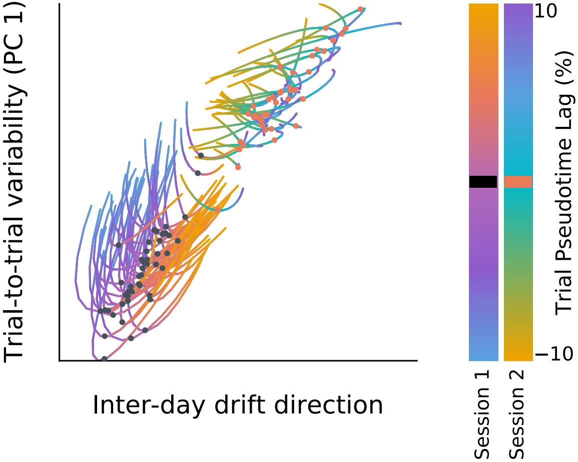

Mouse 3 session 2/3 pseudotime 40%±10

previous turn right, next turn right

Does drift resemble noise?

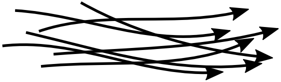

$z(x)$: neural population activity

Does drift resemble noise?

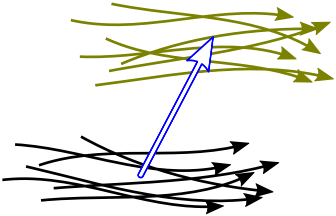

$z(x)$: neural population activity

Drift : $\Delta\mu(x) = \langle z(x) \rangle_{\text{Day 2}} - \langle z(x) \rangle_{\text{Day 1}}$

Does drift resemble noise?

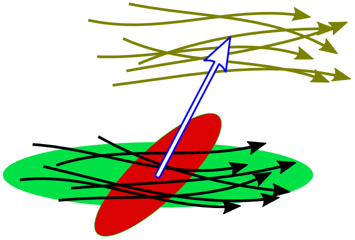

$z(x)$: neural population activity

Drift : $\Delta\mu(x) = \langle z(x) \rangle_{\text{Day 2}} - \langle z(x) \rangle_{\text{Day 1}}$

$\nabla_{z(x)}$: task-co-varying activity

$\Sigma_{z(x)}$: trial-to-trial variability

Does drift resemble noise?

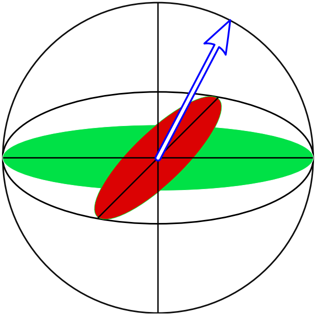

$z(x)$: neural population activity

Drift : $\Delta\mu(x) = \langle z(x) \rangle_{\text{Day 2}} - \langle z(x) \rangle_{\text{Day 1}}$

$\nabla_{z(x)}$: task-co-varying activity

$\Sigma_{z(x)}$: trial-to-trial variability

Drift & task-covarying directions: $\left\langle \| \Delta\mu(x)^\top \nabla_{z(x)} \|^2 \right\rangle$

Drift & noise directions: $\left\langle \Delta\mu(x)^\top \Sigma_{z(x)} \Delta\mu(x) \right\rangle$

(normalize for expected alignment)

… (some) drift resembles noise

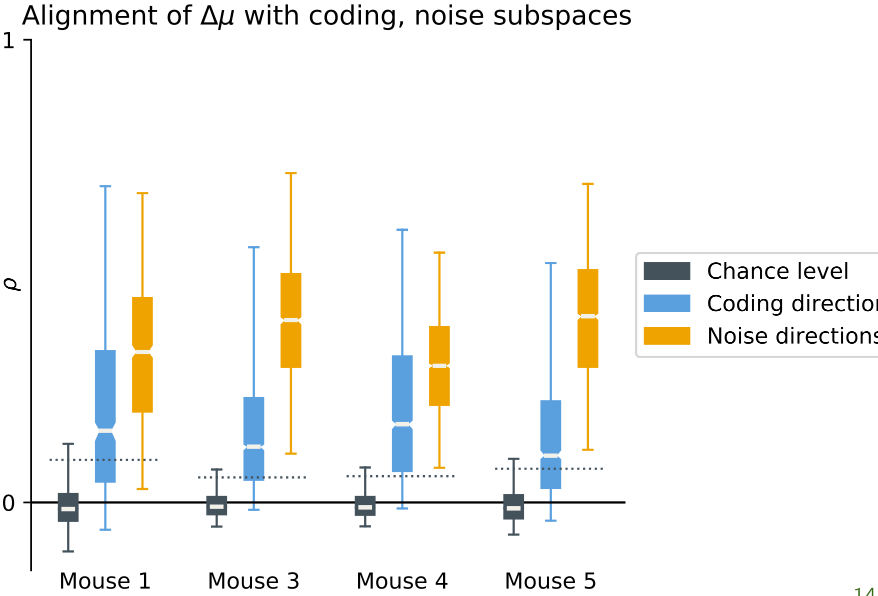

... (some) drift resembles noise

Neural populations exhibit structured trial-to-trial variability

... (some) drift resembles noise

Neural populations exhibit structured trial-to-trial variability

Drift between days is concentrated in these directions

... (some) drift resembles noise

Neural populations exhibit structured trial-to-trial variability

Drift between days is concentrated in these directions

Decoders that ignore trial-to-trial fluctuations can ignore this drift

… (some) drift resembles noise

Neural populations exhibit structured trial-to-trial variability

Drift between days is concentrated in these directions

Decoders that ignore trial-to-trial fluctuations can ignore this drift

Reconcile drift with stable performance:

Task-relevant neural representation in parietal cortex **change**

This change does not disrupt behavior

What are the implications for neural coding?

We will show that:

1. Shifts in mean activity are mostly irrelevant: they resemble noise
2. **Redundancy allows many codes, and there is a stable subspace**
3. Modest plasticity would be required to track remaining drift

Can we find a long-term stable encoding subspace?

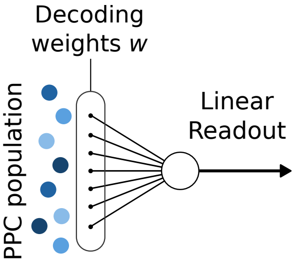
Decoding
weights w
PPC population
Linear
Readout

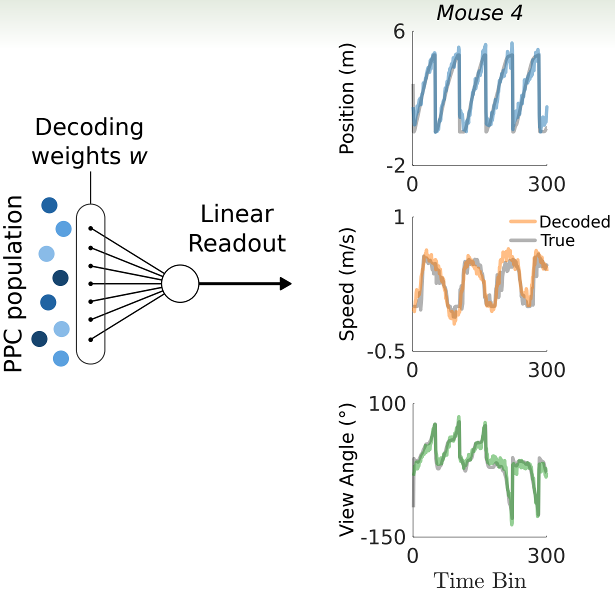
Decoding
weights w
PPC population
Linear
Readout
Mouse 4
Position (m)
6
-2
0
300
Speed (m/s)
1
-0.5
0
300
Decoded
True
View Angle (°)
100
-150
0
300
Time Bin

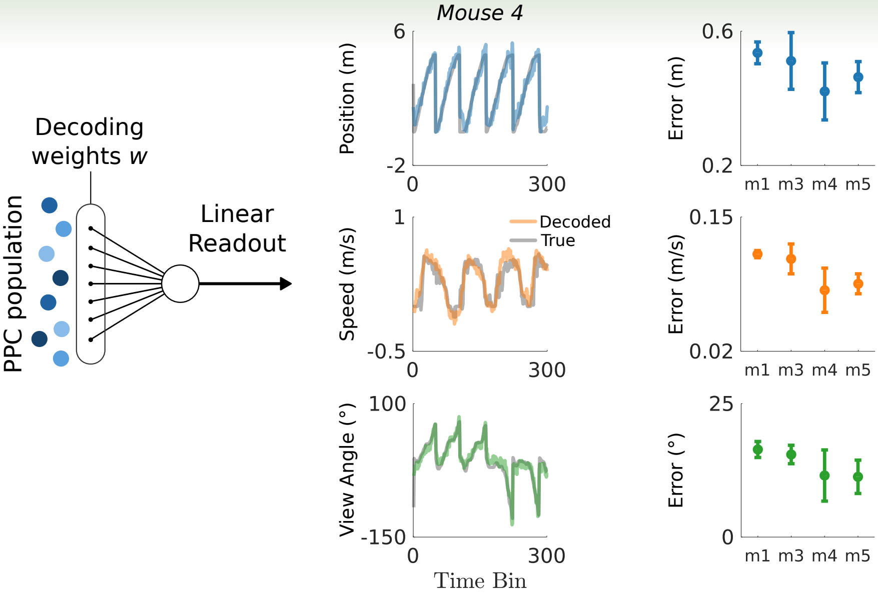

Decoding weights w
PPC population
Linear Readout
Mouse 4
Position (m)
6
-2
0
300
Speed (m/s)
1
-0.5
Decoded
True
View Angle (°)
100
-150
Time Bin
Error (m)
0.6
0.2
Error (m/s)
0.15
0.02
Error (°)
25
0
m1 m3 m4 m5

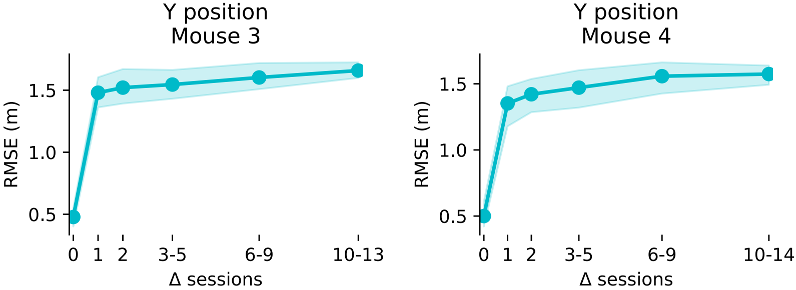

Y position
Mouse 3
RMSE (m)
1.5
1.0
0.5
0
1
2
3-5
6-9
10-13
Δ sessions
Y position
Mouse 4
RMSE (m)
1.5
1.0
0.5
0
1
2
3-5
6-9
10-14
Δ sessions

Concatenated decoder

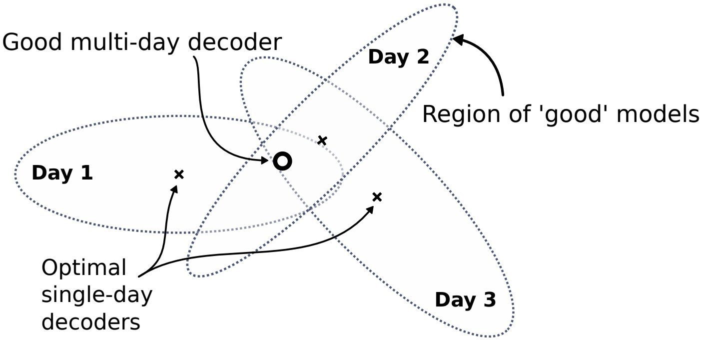

$$\sum_{d=1}^{n} \|x_d - M z_d\|^2$$

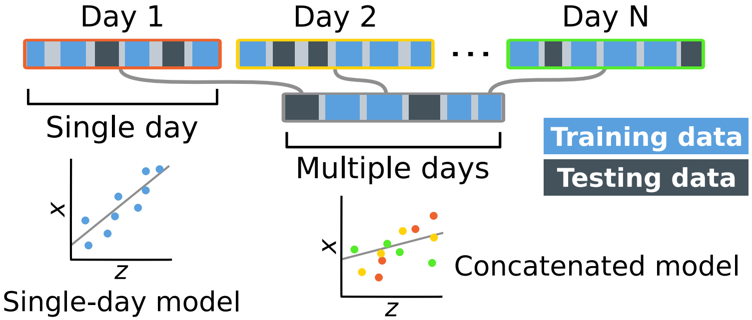
Day 1
Day 2
...
Day N
Single day
Multiple days
Training data
Testing data
x
z
Single-day model
x
z
Concatenated model

… there is a (mostly) stable subspace

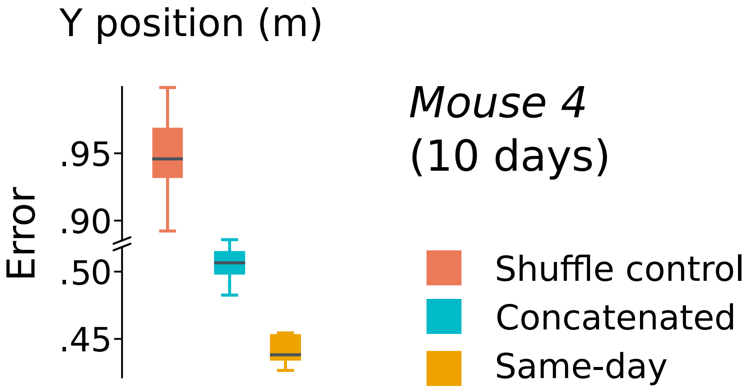

… there is a (mostly) stable subspace

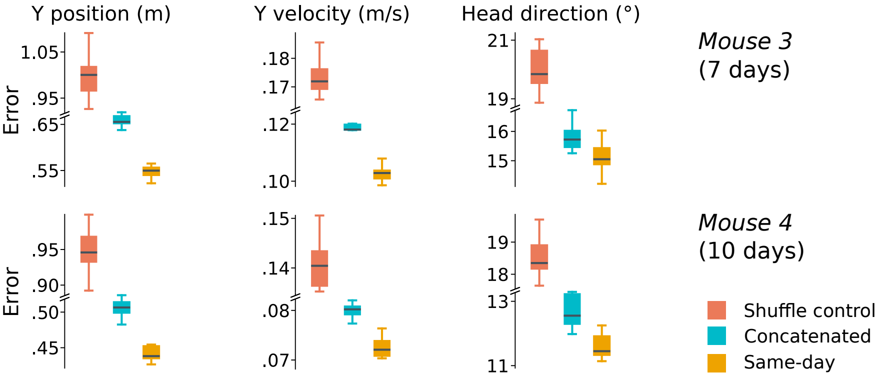

… there is a (mostly) stable subspace

Observation over time can identify a relatively stable linear subspace

… there is a (mostly) stable subspace

Observation over time can identify a relatively stable linear subspace

For 100-200 neurons, over 7-10 days, we can decode from this subspace with 10-20% error increase

… there is a (mostly) stable subspace

Observation over time can identify a relatively stable linear subspace

For 100-200 neurons, over 7-10 days, we can decode from this subspace with 10-20% error increase

Addressing this error increase:

... there is a (mostly) stable subspace

Observation over time can identify a relatively stable linear subspace

For 100-200 neurons, over 7-10 days, we can decode from this subspace with 10-20% error increase

Addressing this error increase:

… there is a (mostly) stable subspace

Observation over time can identify a relatively stable linear subspace

For 100-200 neurons, over 7-10 days, we can decode from this subspace with 10-20% error increase

Addressing this error increase:

- Accuracy required to perform task unclear?

… there is a (mostly) stable subspace

Observation over time can identify a relatively stable linear subspace

For 100-200 neurons, over 7-10 days, we can decode from this subspace with 10-20% error increase

Addressing this error increase:

- Accuracy required to perform task unclear?
- Larger population observations may be needed?

… there is a (mostly) stable subspace

Observation over time can identify a relatively stable linear subspace

For 100-200 neurons, over 7-10 days, we can decode from this subspace with 10-20% error increase

Addressing this error increase:

- Accuracy required to perform task unclear?
- Larger population observations may be needed?
- **Plasticity could track these changes?**

Reconcile drift with stable performance:

Task-relevant neural representation in parietal cortex **change**

This change does not disrupt behavior

What are the implications for neural coding?

We will show that:

1. Shifts in mean activity are mostly irrelevant: they resemble noise
2. Redundancy allows many codes, and there is a stable subspace
3. **Modest plasticity would be required to track remaining drift**

How much plasticity is needed to track an evolving code?

Constrained decoder

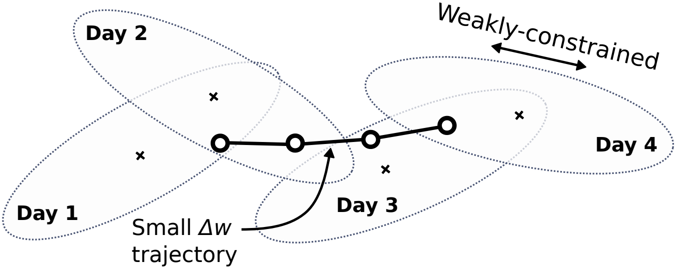

$$(1-\lambda)\sum_{d=1}^{n}\|x_d - M_d z_d\|^2 + \lambda\sum_{d=1}^{n-1}\|M_{d+1} - M_d\|^2$$

Small changes are enough

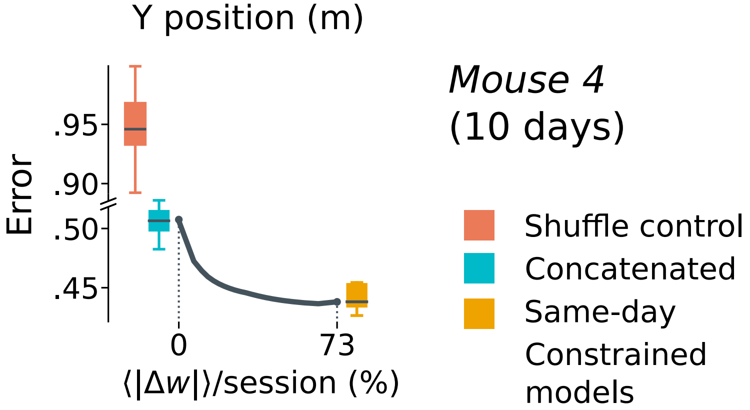

Small changes are enough

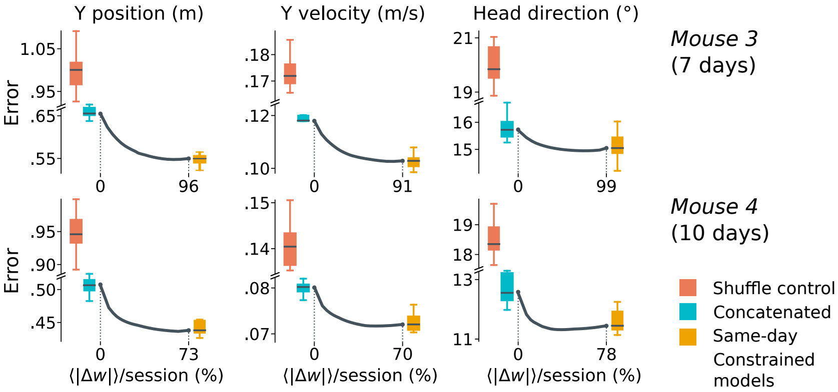

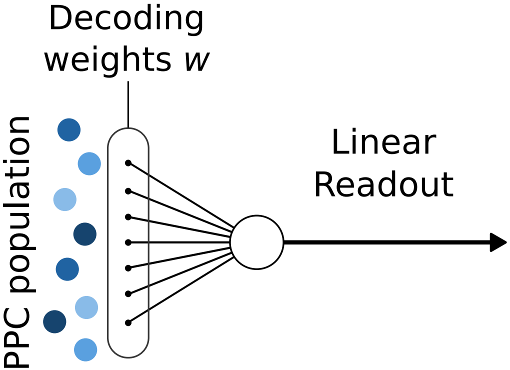
Decoding
weights *w*
PPC population
Linear
Readout

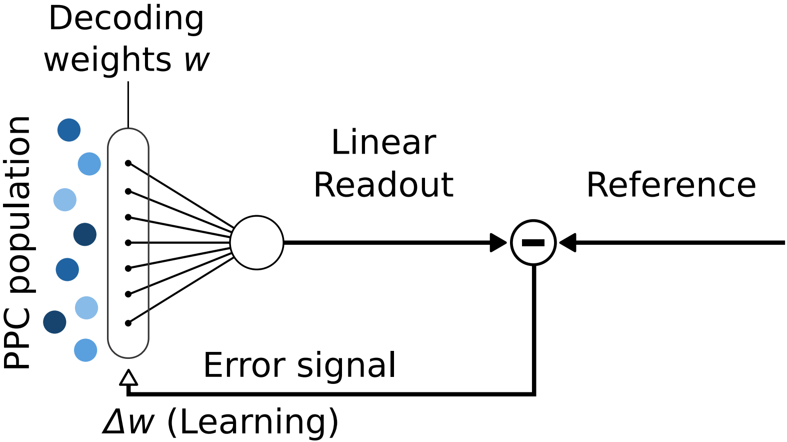
Decoding
weights w
PPC population
Linear
Readout
Reference
Error signal
Δw (Learning)

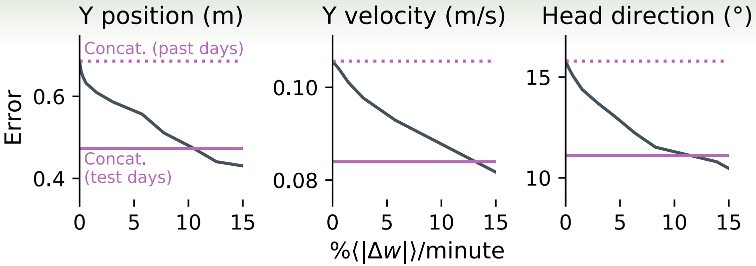
Y position (m)
Y velocity (m/s)
Head direction (°)
Error
Concat. (past days)
Concat.
(test days)
0.6
0.4
0.10
0.08
15
10
0
5
10
15
%⟨|Δw|⟩/minute

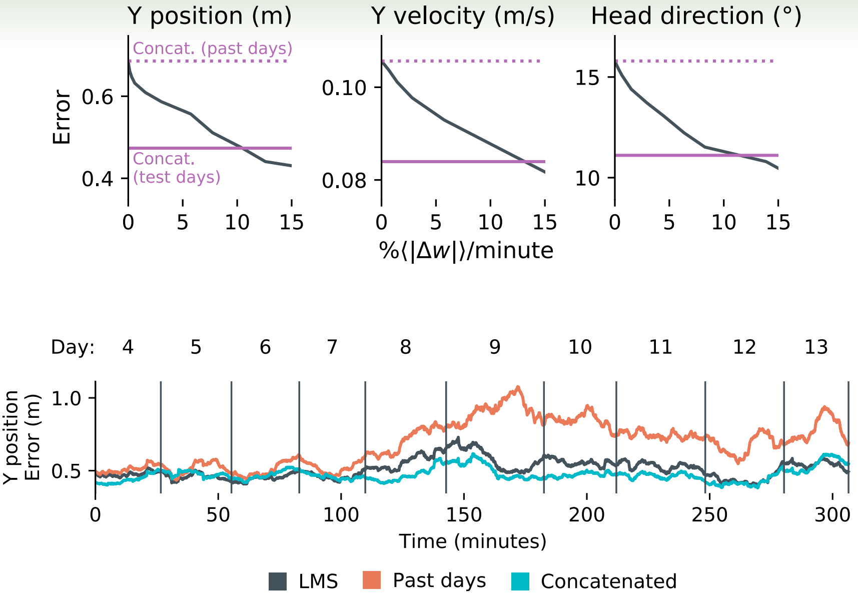
Y position (m)
Y velocity (m/s)
Head direction (°)
Error
Concat. (past days)
Concat.
(test days)
0.6
0.4
0.10
0.08
15
10
0
5
10
15
%⟨|Δw|⟩/minute
Day: 4 5 6 7 8 9 10 11 12 13
Y position
Error (m)
1.0
0.5
0
50
100
150
200
250
300
Time (minutes)
LMS
Past days
Concatenated

… (modest) plasticity is required

… (modest) plasticity is required

Good decoding performance with small weight changes is possible

… (modest) plasticity is required

Good decoding performance with small weight changes is possible

Continual learning could achieve this

... (modest) plasticity is required

Good decoding performance with small weight changes is possible

Continual learning could achieve this

For 100-200 neurons, $\mathcal{O}(10\%)$ weight-change/day

… (modest) plasticity is required

Good decoding performance with small weight changes is possible

Continual learning could achieve this

For 100-200 neurons, $\mathcal{O}(10\%)$ weight-change/day

Could be less for larger populations (depends on correlations)

… (modest) plasticity is required

Good decoding performance with small weight changes is possible

Continual learning could achieve this

For 100-200 neurons, $\mathcal{O}(10\%)$ weight-change/day

Could be less for larger populations (depends on correlations)

Drift is compatible with stable performance:

Drift is compatible with stable performance:

Neural codes in mouse PPC change w/o disrupting behavior

Drift is compatible with stable performance:

Neural codes in mouse PPC change w/o disrupting behavior

- Redundancy: *Many neurons, low-D task*

Drift is compatible with stable performance:

Neural codes in mouse PPC change w/o disrupting behavior

- Redundancy: *Many neurons, low-D task*
    - Ignore the noise (much drift resembles noise)

Drift is compatible with stable performance:

Neural codes in mouse PPC change w/o disrupting behavior

- Redundancy: *Many neurons, low-D task*
  - Ignore the noise (much drift resembles noise)
- Space of possible representations remains ‘sloppy’

Drift is compatible with stable performance:

Neural codes in mouse PPC change w/o disrupting behavior

- Redundancy: *Many neurons, low-D task*
  - Ignore the noise (much drift resembles noise)
- Space of possible representations remains 'sloppy'
  - There is a stable subspace

Drift is compatible with stable performance:

Neural codes in mouse PPC change w/o disrupting behavior

- Redundancy: *Many neurons, low-D task*
  - Ignore the noise (much drift resembles noise)
- Space of possible representations remains 'sloppy'
  - There is a stable subspace
- Plasticity

Drift is compatible with stable performance:

Neural codes in mouse PPC change w/o disrupting behavior

- Redundancy: *Many neurons, low-D task*
  - Ignore the noise (much drift resembles noise)
- Space of possible representations remains 'sloppy'
  - There is a stable subspace
- Plasticity
  - Slow change may sill occur, but could be tracked

Drift is compatible with stable performance:

Neural codes in mouse PPC change w/o disrupting behavior

- Redundancy: *Many neurons, low-D task*
  - Ignore the noise (much drift resembles noise)
- Space of possible representations remains ‘sloppy’
  - There is a stable subspace
- Plasticity
  - Slow change may sill occur, but could be tracked
  - (internal prediction error feedback?)

Drift is compatible with stable performance:

Neural codes in mouse PPC change w/o disrupting behavior

- Redundancy: *Many neurons, low-D task*
  - Ignore the noise (much drift resembles noise)
- Space of possible representations remains 'sloppy'
  - There is a stable subspace
- Plasticity
  - Slow change may sill occur, but could be tracked
  - (internal prediction error feedback?)

Observed changes in PCC:

Drift is compatible with stable performance:

Neural codes in mouse PPC change w/o disrupting behavior

- Redundancy: *Many neurons, low-D task*
  - Ignore the noise (much drift resembles noise)
- Space of possible representations remains 'sloppy'
  - There is a stable subspace
- Plasticity
  - Slow change may sill occur, but could be tracked
  - (internal prediction error feedback?)

Observed changes in PCC:

- Learning of other tasks (or compensation for learning elsewhere?)

Drift is compatible with stable performance:

Neural codes in mouse PPC change w/o disrupting behavior

- Redundancy: *Many neurons, low-D task*
  - Ignore the noise (much drift resembles noise)
- Space of possible representations remains ‘sloppy’
  - There is a stable subspace
- Plasticity
  - Slow change may sill occur, but could be tracked
  - (internal prediction error feedback?)

Observed changes in PCC:

- Learning of other tasks (or compensation for learning elsewhere?)
- Ongoing plasticity at steady-state?

Drift is compatible with stable performance:

Neural codes in mouse PPC change w/o disrupting behavior

- Redundancy: *Many neurons, low-D task*
  - Ignore the noise (much drift resembles noise)
- Space of possible representations remains ‘sloppy’
  - There is a stable subspace
- Plasticity
  - Slow change may sill occur, but could be tracked
  - (internal prediction error feedback?)

Observed changes in PCC:

- Learning of other tasks (or compensation for learning elsewhere?)
- Ongoing plasticity at steady-state?
- ‘Time-stamping’?

Thanks!

Timothy O'Leary

Adrianna Loback

Dhruva Raman

Fulvio Forni

Laura Driscoll

Chris Harvey

Alon Rubin

Yaniv Ziv

This work was supported by the Human Frontier Science Program (RGY0069), ERC Starting Grant (StG FLEXNEURO 716643) and grants from the NIH (NS089521, MH107620, NS108410)

end

Mongillo, Rumpel, Loewenstein (2017), review:

Mongillo, Rumpel, Loewenstein (2017), review:

- Synapses are unstable; Preserve **effective** connectivity:

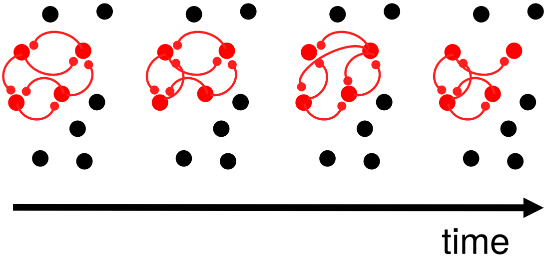

Mongillo, Rumpel, Loewenstein (2017)

Mongillo, Rumpel, Loewenstein (2017), review:

- Synapses are unstable; Preserve **effective** connectivity:
- Multiple configurations → same computation

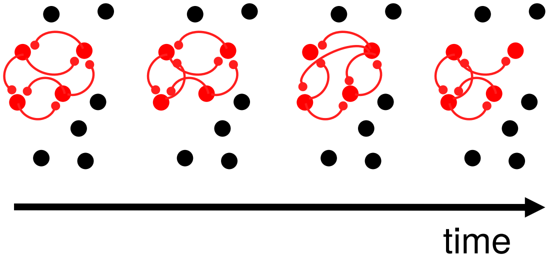

Mongillo, Rumpel, Loewenstein (2017)

Mongillo, Rumpel, Loewenstein (2017), review:

- Synapses are unstable; Preserve **effective** connectivity:
- Multiple configurations → same computation
- E.g. Brinkman & al. (2018)

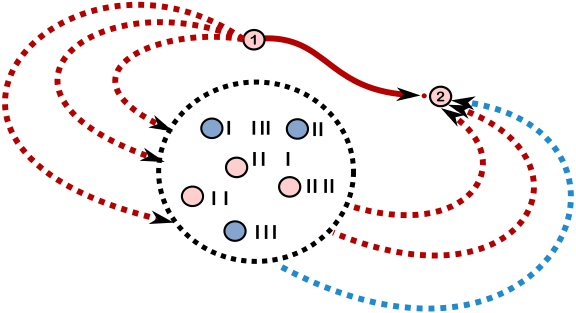

Brinkman & al. (2018)

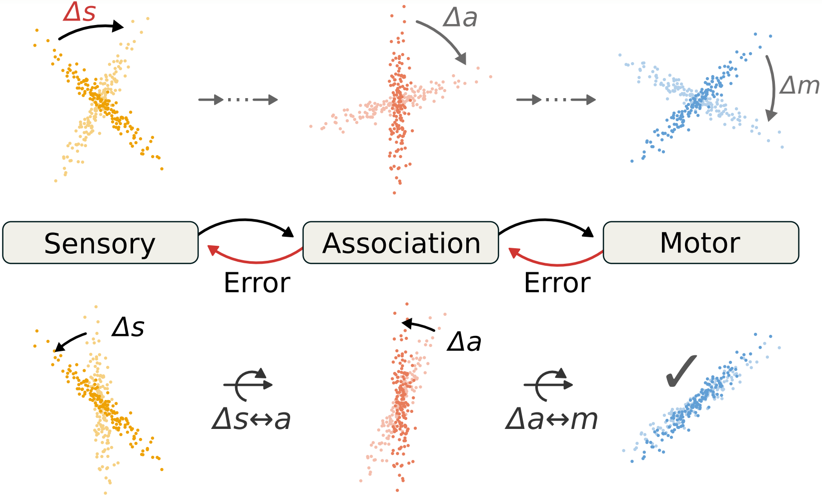
Δs
Δa
Δm
Sensory
Association
Motor
Error
Error
Δs
Δa
Δs↔a
Δa↔m
✓

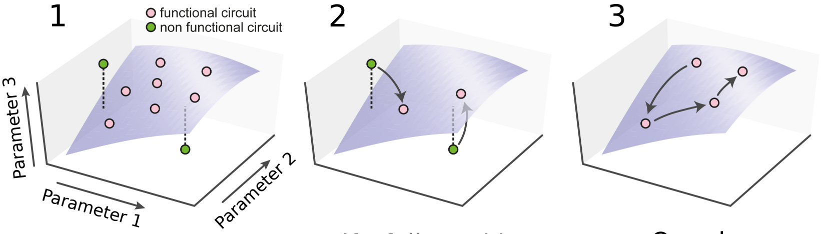

Multiple solutions

Drift, followed by error feedback

Ongoing reconfiguration

Premise

Premise

- The Posterior Parietal Cortex (PPC) in mice contains neural correlates of spatial navigation

Premise

- The Posterior Parietal Cortex (PPC) in mice contains neural correlates of spatial navigation
- Single neuron tunings change over days

Premise

- The Posterior Parietal Cortex (PPC) in mice contains neural correlates of spatial navigation
- Single neuron tunings change over days
- The overall population continues to tile the task

Premise

- The Posterior Parietal Cortex (PPC) in mice contains neural correlates of spatial navigation
- Single neuron tunings change over days
- The overall population continues to tile the task
- Similar latent low-dimensional manifolds can be extracted

Premise

- The Posterior Parietal Cortex (PPC) in mice contains neural correlates of spatial navigation
- Single neuron tunings change over days
- The overall population continues to tile the task
- Similar latent low-dimensional manifolds can be extracted

Hypotheses

Premise

- The Posterior Parietal Cortex (PPC) in mice contains neural correlates of spatial navigation
- Single neuron tunings change over days
- The overall population continues to tile the task
- Similar latent low-dimensional manifolds can be extracted

Hypotheses

- Reconfiguration occurs within a ‘null’ space

Premise

- The Posterior Parietal Cortex (PPC) in mice contains neural correlates of spatial navigation
- Single neuron tunings change over days
- The overall population continues to tile the task
- Similar latent low-dimensional manifolds can be extracted

Hypotheses

- Reconfiguration occurs within a ‘null’ space
- Concurrent changes in other brain areas preserve sensorimotor transformations

Premise

- The Posterior Parietal Cortex (PPC) in mice contains neural correlates of spatial navigation
- Single neuron tunings change over days
- The overall population continues to tile the task
- Similar latent low-dimensional manifolds can be extracted

Hypotheses

- Reconfiguration occurs within a ‘null’ space
- Concurrent changes in other brain areas preserve sensorimotor transformations

Approach

Premise

- The Posterior Parietal Cortex (PPC) in mice contains neural correlates of spatial navigation
- Single neuron tunings change over days
- The overall population continues to tile the task
- Similar latent low-dimensional manifolds can be extracted

Hypotheses

- Reconfiguration occurs within a ‘null’ space
- Concurrent changes in other brain areas preserve sensorimotor transformations

Approach

- Can we find a stable encoding subspace?

Premise

- The Posterior Parietal Cortex (PPC) in mice contains neural correlates of spatial navigation
- Single neuron tunings change over days
- The overall population continues to tile the task
- Similar latent low-dimensional manifolds can be extracted

Hypotheses

- Reconfiguration occurs within a ‘null’ space
- Concurrent changes in other brain areas preserve sensorimotor transformations

Approach

- Can we find a stable encoding subspace?
- How much plasticity is needed to track an evolving code?