Constrained plasticity can compensate for ongoing drift in neural populations

Michael E. Rule[1], Adrianna R. Loback[2], Dhruva V. Raman[1], Christopher D. Harvey[3], Timothy S. O'Leary[1]

[1]University of Cambridge, Cambridge UK [2]Co-first-author [3]Harvard Medical School, Boston, MA, USA

How does the brain achieve stable representations, despite noisy biological components and ongoing plasticity?

We examine how population codes for navigation in the mouse posterior parietal cortex (PPC) might remain stable, despite changes in single neurons.

We address two main questions:

1. How disruptive is drift in neuronal tuning in practice, considering that the brain employs a redundant population code?
2. Could plasticity compensate for this disruption, using physiologically plausible rates of synaptic weight change?

1. Single-neuron encoding of spatial navigation drifts over time

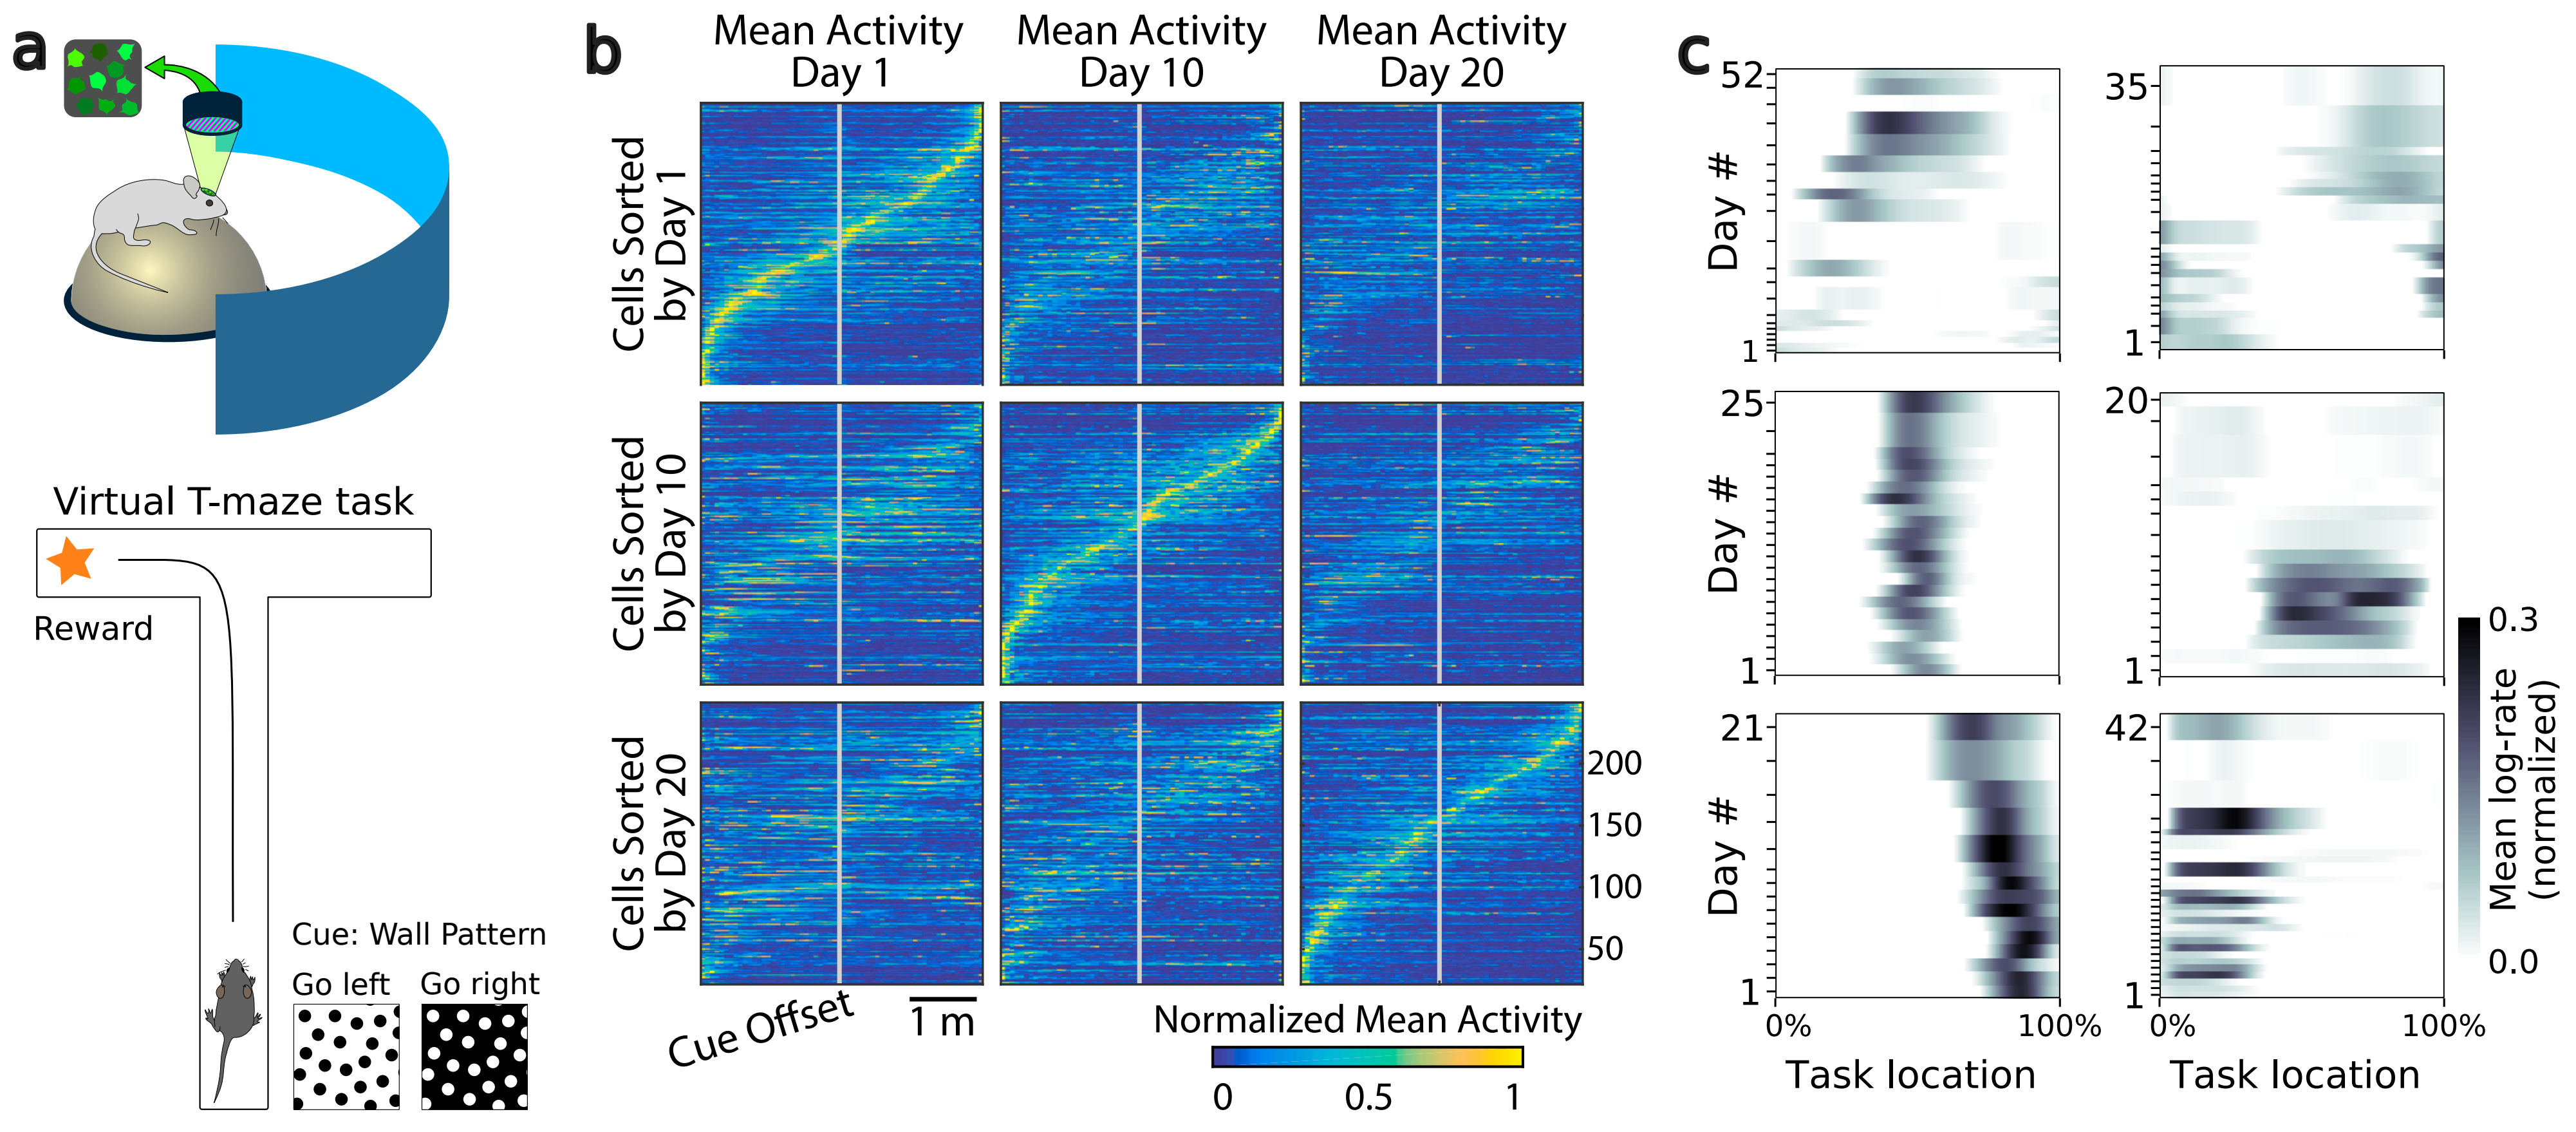

Fig. 1. **(a)** Mice were trained to use visual cues to navigate to a reward in a virtual-reality maze; neural population activity was recorded using Ca^{2+} imaging (1). **(b)** (Reprinted from 1) Neurons in PPC (y-axes) fire at various regions in the maze (x-axes). Over days to weeks, individual neurons change their tuning, re-configuring the population code. *This occurs even at steady-state behavioral performance (after learning).* **(c)** Neurons show diverse changes in tuning over days, including instability, relocation, long-term stability, gain/loss of selectivity, advancement, and intermittent responsiveness.

2. Linear decoding recovers behavior from population activity

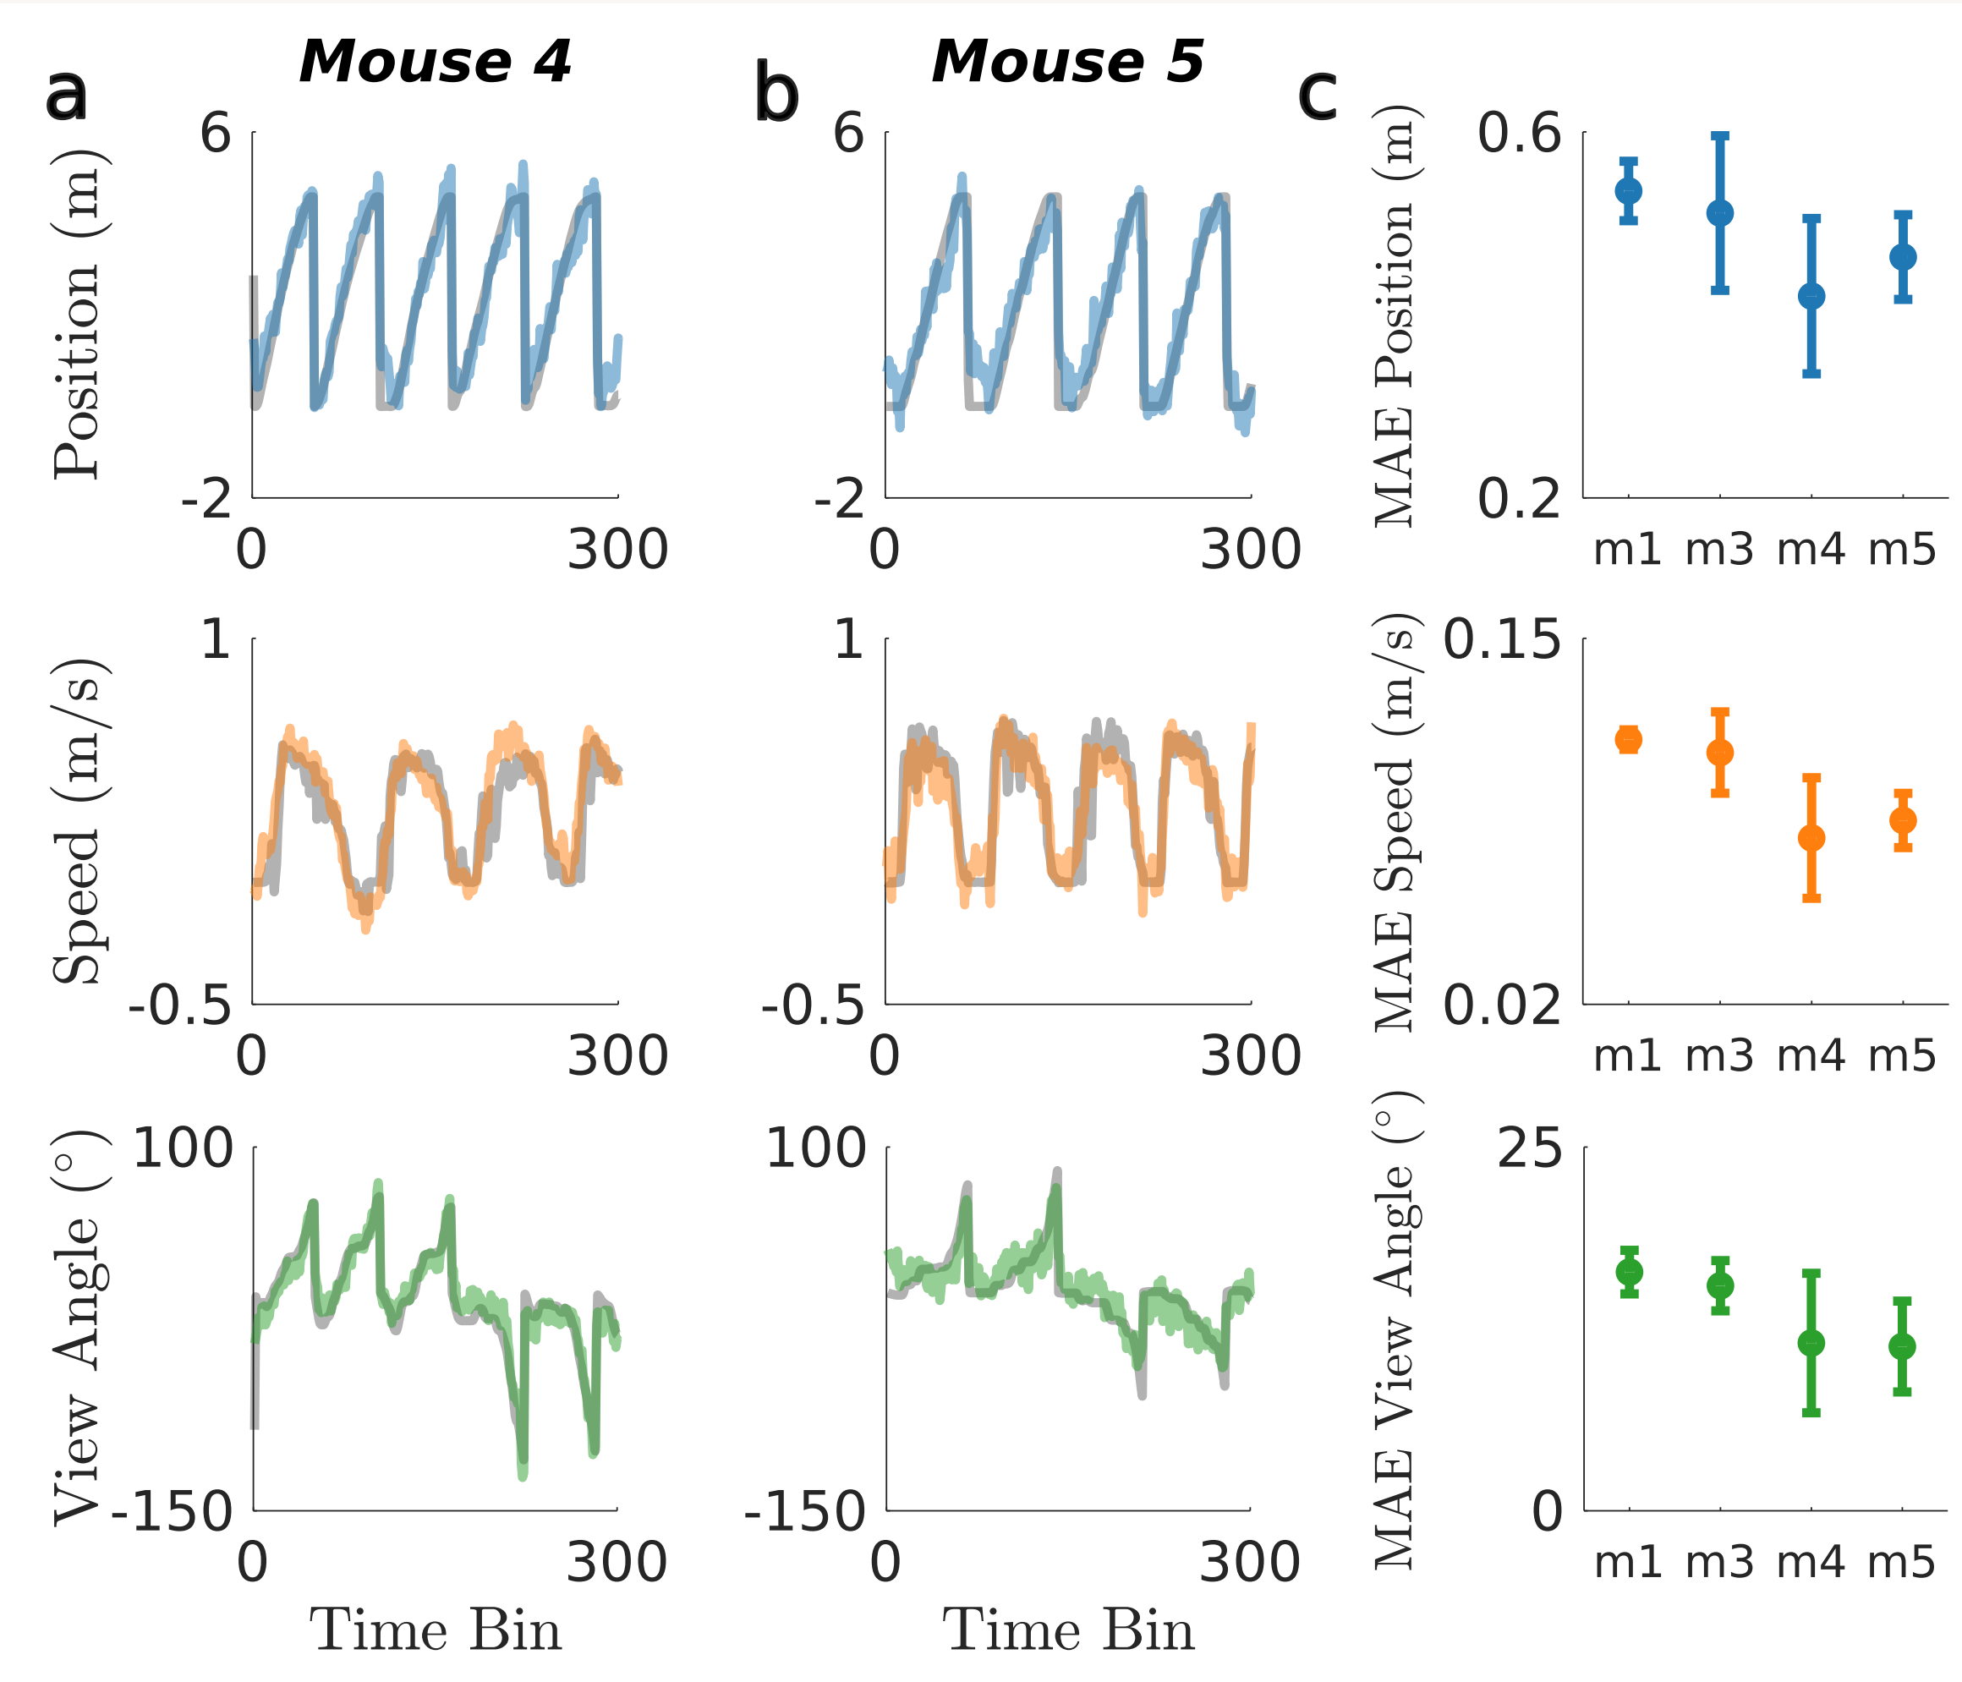

Fig. 2. **(a)** Decoding performance for a single session (mouse 4). Shaded colors: prediction under cross-validation; Solid lines: ground truth. **(b)** Example (mouse 5), same as (a). **(c)** Linear decoding performance summary; each point denotes one mouse. Error bars: 1σ over all sessions.

Can we understand the impact of population drift on encoding by studying how linear coding generalizes across days?

3. Redundancy allows a fixed linear readout to work over 5-7 days despite drift

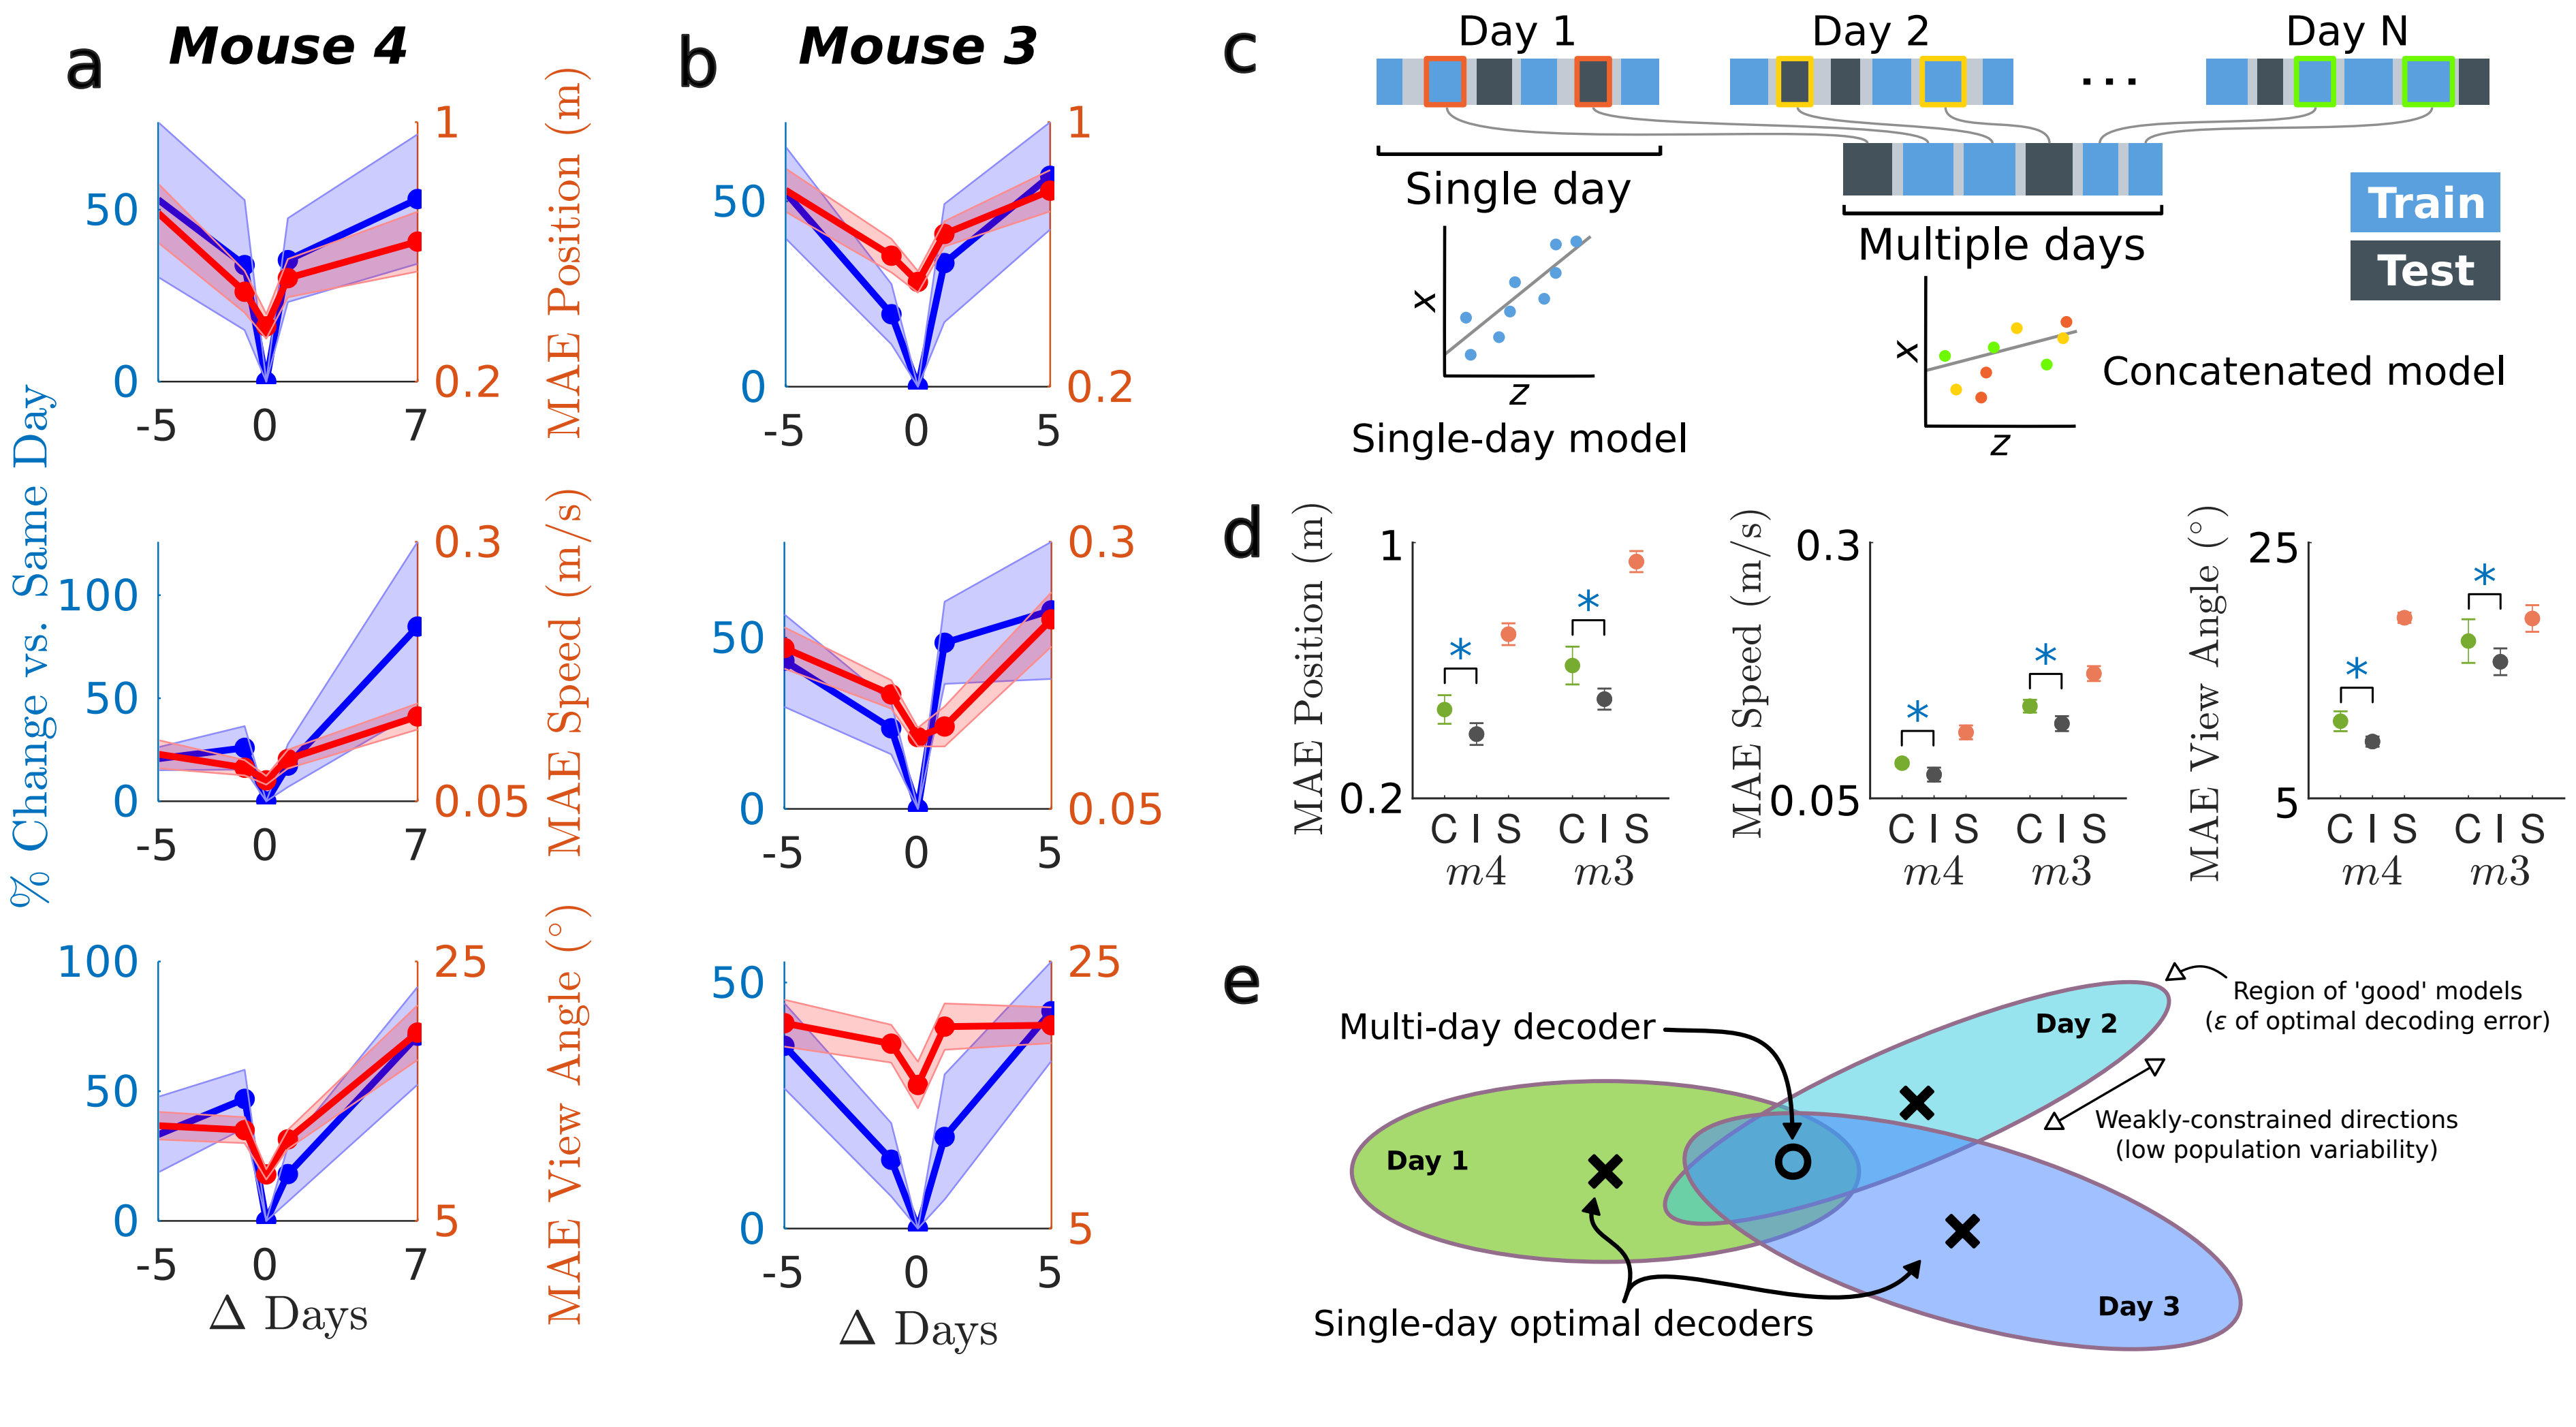

Fig. 3. **(a,b)** Red: Generalization error (mean absolute error, 'MAE') for decoders trained on a single day ('0') & tested on past/future days. Blue: % increase in error over the optimal decoder for the testing day.

(c) Is there a fixed linear decoder that could decode well over several days?

(d) Decoding error on concatenated data (green) is only slightly worse than the mean best single-day linear decoding error (grey). (e.g. mouse 4 and 3; red: shuffle chance-level reflects extent that redundancy allows prediction across unrelated neural codes)

(e) *The relevant measure of drift is not neural activity, but decoding error* (ϵ). Redundant codes allow for many valid decoders. **Multiple days can share an encoding subspace that works well, but differs from optimal single-day decoders.**

4. Adaptive decoders can track coding-relevant drift with modest plasticity

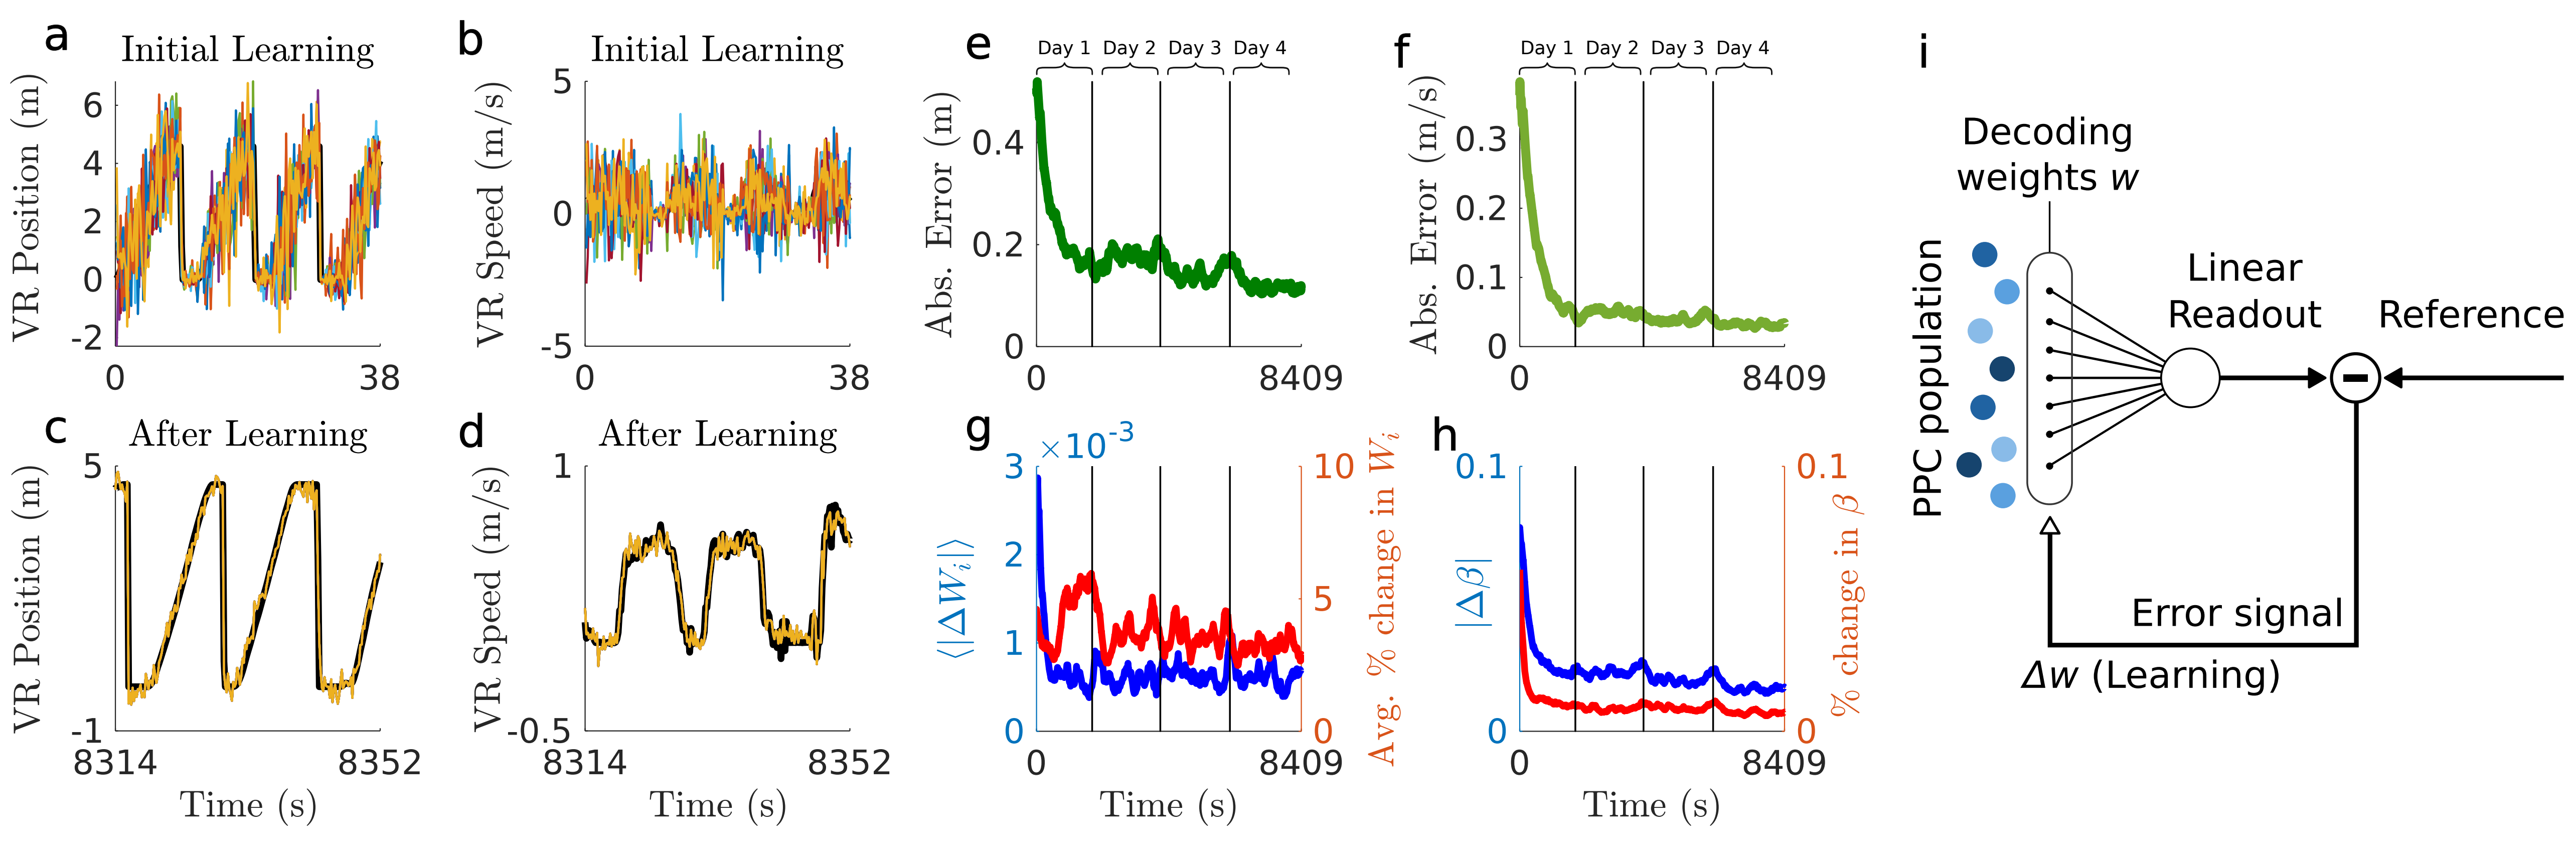

Fig. 4. **Decoding results for the online Least-Mean-Squares (LMS) algorithm over multiple days.** Example shown for mouse 3. **(a)** Position and **(b)** speed predictions during initial latent learning (various colors: 10 random parameter initializations; black: ground truth). **(c,d)** Predictions of the LMS algorithm after learning. **(e,f)** Smoothed absolute decoding error over the first four days for position and speed, respectively. **(g)** (blue) Average absolute change of decoding weights W_i over time, (red) average % change from the previous weight value (position decoder). **(h)** Similar to (g), but for the linear bias/offset parameter (speed decoder). **(i)** **An adaptive linear readout could use error feedback from downstream areas to adjust its interpretation of PPC firing and compensate for drift.**

5. Most drift is in irrelevant directions

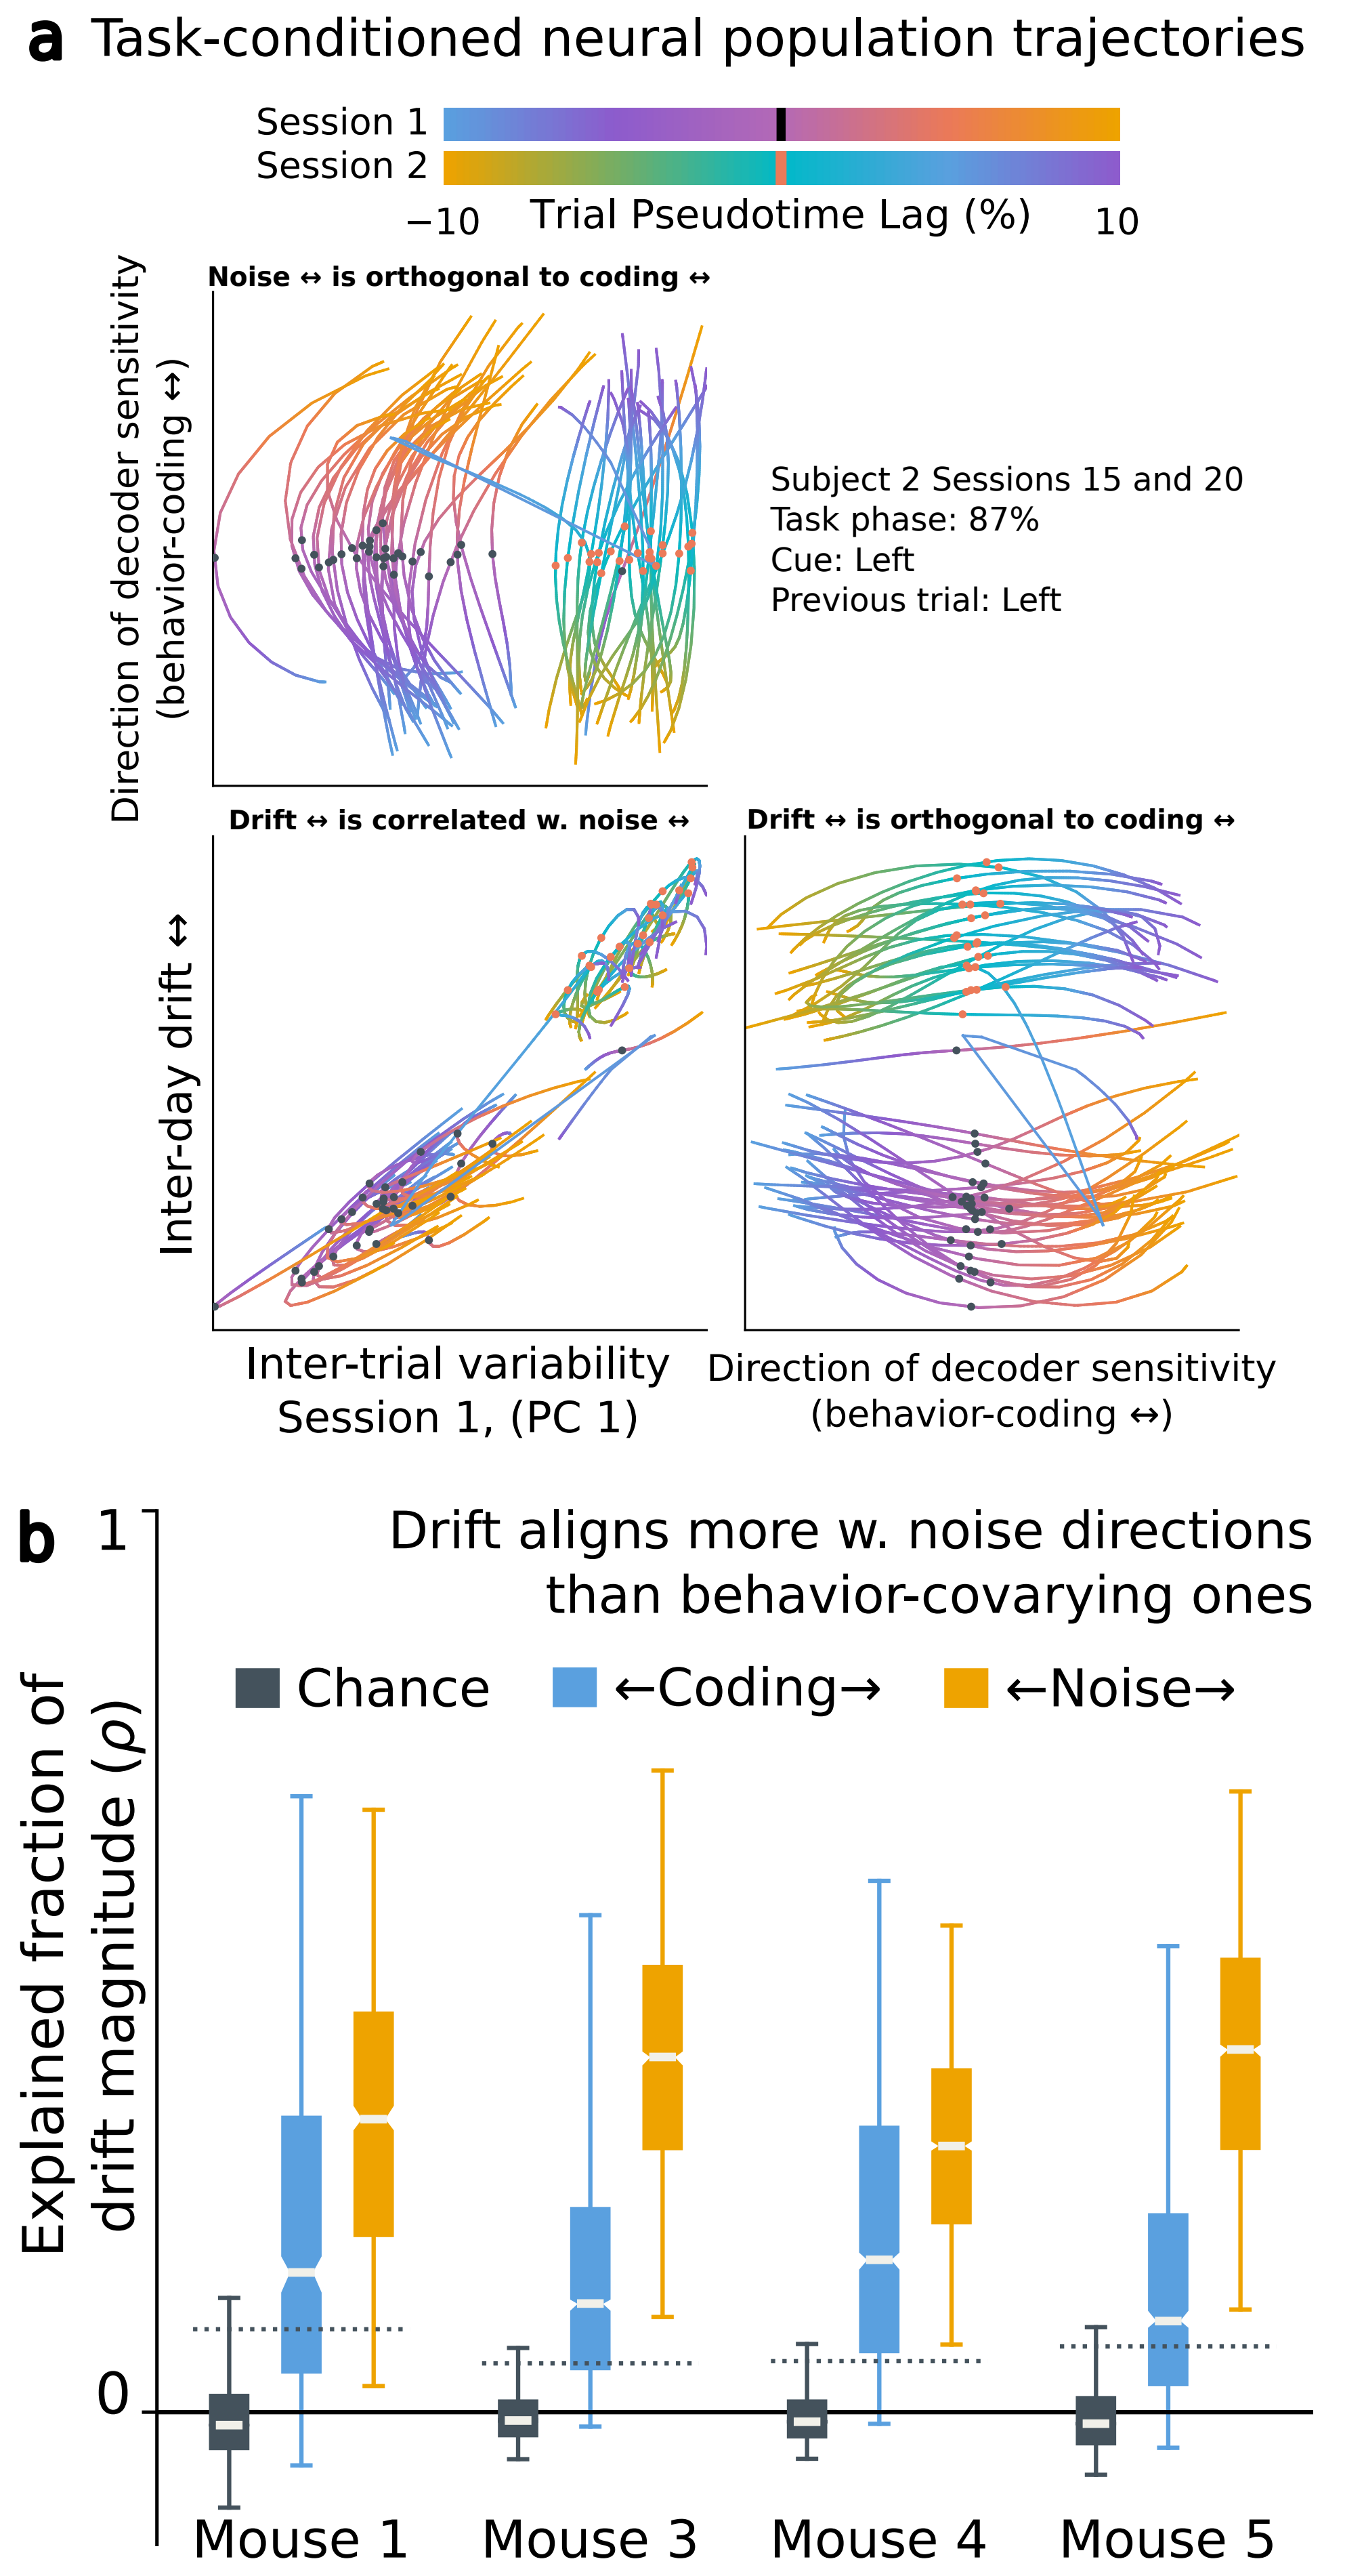

Subject 2 Sessions 15 and 20
Task phase: 87%
Cue: Left
Previous trial: Left

Fig. 5. **(a)** Projection of neural population trajectories (over a short time window) onto behavior-coding, noise (trial-to-trial variability), and inter-day drift axes. Noise and inter-day drift are often correlated and orthogonal to the behavior-coding axis. **(b)** **A larger fraction (ρ) of drift is explained by population directions exhibiting trial-to-trial variability (yellow), compared to directions that covary with behavior (blue).**

We find that

1. Distributed and redundant population codes could allow for a surprisingly stable encoding of behavior, despite population reconfiguration
2. Changes in neural population codes align more with "noisy" dimensions of neural activity than with behavior-coding ones
3. Re-arrangement in population codes could be tracked using weak error feedback and synaptic plasticity (2%/day)

---

[1] Driscoll LN, et al. Dynamic reorganization of neuronal activity patterns in parietal cortex. *Cell*, 170(5):986–999, 2017.

This work was supported by HFSP grant number RGY 0069/2017 "Building a theory of shifting representations in the mammalian brain"; principal investigators Timothy O'Leary, Christopher Harvey, and Yaniv Ziv.