Differential Stability of Task Variable Representations in Retrosplenial Cortex

Luis M. Franco[1,2,*], and Michael J. Goard[1,3,4,*]

[1]Neuroscience Research Institute,
University of California, Santa Barbara,
Santa Barbara, CA 93106, USA

[2]Institute of Neuroscience,
University of Oregon,
Eugene, OR 97403, USA

[3]Department of Molecular, Cellular, and Developmental Biology,
University of California, Santa Barbara,
Santa Barbara, CA 93106, USA

[4]Department of Psychological & Brain Sciences,
University of California, Santa Barbara,
Santa Barbara, CA 93106, USA

*Correspondence to: luisfran@uoregon.edu, michael.goard@lifesci.ucsb.edu

ABSTRACT

Cortical neurons store information across different timescales, from seconds to years. Although information stability is variable across regions, it can vary within a region as well. Association areas are known to multiplex behaviorally relevant variables, but the stability of their representations is not well understood. Here, we longitudinally recorded the activity of neuronal populations in the retrosplenial cortex (RSC) during the performance of a context-choice association task. We found that the activity of neurons exhibits different levels of stability across days. Using linear classifiers, we quantified the stability of three task-relevant variables. We find that RSC representations of context and trial outcome display higher stability than motor choice, both at the single cell and population levels. Together, our findings show an important characteristic of association areas, where diverse streams of information are stored with varying levels of stability, which may balance representational reliability and flexibility according to behavioral demands.

INTRODUCTION

The mammalian brain experiences plasticity throughout its lifespan due to development, learning, ongoing synaptic turnover, and homeostatic maintenance[1–4]. Nevertheless, it must somehow preserve information about known sensory features of the environment, as well as learned associations that are important for survival. How the brain maintains stable representations of the environment in dynamic cortical neural circuits is not well understood. In the visual cortex, for instance, processing of relatively simple features such as orientation tuning seems very stable across many days, and even weeks[5–8]. Similarly, representations of learned movements in the motor cortex[9,10] and somatosensory representations of repeated stimuli in the barrel cortex[11] exhibit considerable stability across multiple days of experimentation (although some neurons can be highly unstable[12]). In the piriform cortex, by contrast, odor representations constantly drift over days unless odors are frequently experienced[13]. Other cortical regions, such as the auditory cortex, experience a monotonic decay in the similarity of their spontaneous activity, despite a small fraction of cells remaining strongly correlated over multiple days[14]. These changes in single cell response properties can eventually lead to a representational drift in the population code, even in the face of stable behavioral output[1–4,15]. One possibility is that redundancy in the population code compensates for the turnover of cellular responses, maintaining stable information and animal behavior over long periods of time[3,8,16].

Although sensory and motor cortices can maintain significant levels of stability across long timescales, association areas and limbic structures appear less stable. In the posterior parietal cortex, a region important for motor planning and spatial decisions, the activity of neurons informative about task features changes considerably across days[17]. In the hippocampus, which is important for the encoding of spatial variables that depend on the integration of multimodal sensory features, CA1 pyramidal cells are highly dynamic, with many unique cells encoding place fields on each particular day[18,19], although granule cells of the dentate gyrus exhibit more stable activity[19,20]. However, recent findings indicate that in the bat hippocampus, CA1 neurons can stably map the trajectory of highly reproducible flights in familiar environments over days and weeks[21], suggesting that in some cases CA1 hippocampal neurons can exhibit very stable coding. In fact, another study recently showed that two subpopulations can coexist in CA1, one with more stable coding of environmental contexts, and one with high remapping rate, sensitive to landmark changes[22]. In the retrosplenial cortex (RSC), a region important for the association of egocentric and allocentric spatial representations[23], as well as for the association between environmental context with motor choice[24–26], a considerable fraction of cells maintain stable retinotopic tuning properties to simple visual stimuli[27]. However, it is not known how stable sensory and motor variables are represented in RSC during the performance of a cognitive task.

In the primary visual cortex, recent findings have shown that different stimuli are encoded with differing levels of stability. Representations of stimuli with simple visual statistics such as drifting gratings are highly stable across several weeks of experimentation[5,6]. However, when the same population of neurons is presented with complex stimuli such as naturalistic movies, their representation is much more dynamic, exhibiting progressive changes in response to the same movie over the course of days to weeks[6,28]. This configuration could potentially preserve fundamental tuning properties while allowing more flexible responses to complex naturalistic stimuli. In association areas, such as posterior parietal and retrosplenial cortices, a similar principle might allow highly stable representations of particular behavioral variables during performance of a learned task.

To investigate the stability of neural representations in the RSC during performance of a cognitive task, head-fixed mice were presented two different virtual T-mazes, each with a distinct wall pattern, that determined the correct left or right turn to obtain water rewards[24]. Using longitudinal two-photon microscopy, we measured calcium responses in thousands of neurons

in RSC across five consecutive days. We find that neural activity in correct trials is highly stable for both contexts, with less stability in incorrect trials. Support vector machine classifiers trained on the first day of experimentation and tested on the subsequent days revealed varying stability of task-relevant variables in the same population of RSC neurons, both at the single cell and population level. Together, our results indicate that RSC can reliably store sensory information about the environmental context and trial outcome while exhibiting flexible coding of motor choice, which we propose allows adaptability in the face of changing behavioral demands.

RESULTS

Differential stability in RSC activity during memory-guided behavior

To investigate the stability of task-relevant representations in the RSC, we developed a behavioral task requiring mice to associate virtual contexts with a left or right joystick turn in order to obtain liquid rewards (**Fig. 1A**; **Video S1**; for a full description of this task, see **Methods** and reference[24]). We then used two-photon calcium imaging to record neural activity in transgenic mice with pan-excitatory expression of GCaMP6s. We imaged a total of 4,701 neurons in layer 2/3 of the dysgranular RSC across 5 consecutive days (*n* = 18 experiments in 6 mice). The same field of view was identified in each session, and video stacks were registered using day 1 as reference (see **Methods**). To ensure that we had enough data for all subsequent analyses, we only considered sessions in which mice had ≥2 trials of each type (correct-left, incorrect-left, incorrect-right, correct-right), and also maintained a behavioral performance above chance level (**Fig. 1B-E**; more details in **Fig. S1**). On average, mice made (mean ± s.e.m.): 19.1 ± 1.0 correct left, 7.7 ± 0.7 incorrect right, 6.3 ± 0.5 incorrect left, and 23.0 ± 1.4 correct right decisions (**Fig. 1C**), which translated to an overall behavioral performance = 74.1% ± 1.0% (mean ± s.e.m.) across all experiments (**Fig. 1B**, **Fig. S1**). Their performance, response bias and response latency stabilized during training, and was consistent across imaging sessions (**Fig. S2**). In addition, we only considered neurons with significant morphological stability[29], as measured by the structural similarity index (SSI) of their corresponding ROI regions across days (**Fig. 1F-G**, **Fig. S3**), and significant reliable responses across days (see **Methods**). Together, this yielded different overlapping groups of neurons (day 2 = 1,868 neurons; day 3 = 1,908 neurons; day 4 = 1,609 neurons; day 5 = 2,020 neurons) whose responses we compared to those in day 1 (3,087 neurons with at least one pair in days 2 to 5). **Fig. 1H** shows two example neurons with preserved activity dynamics across days.

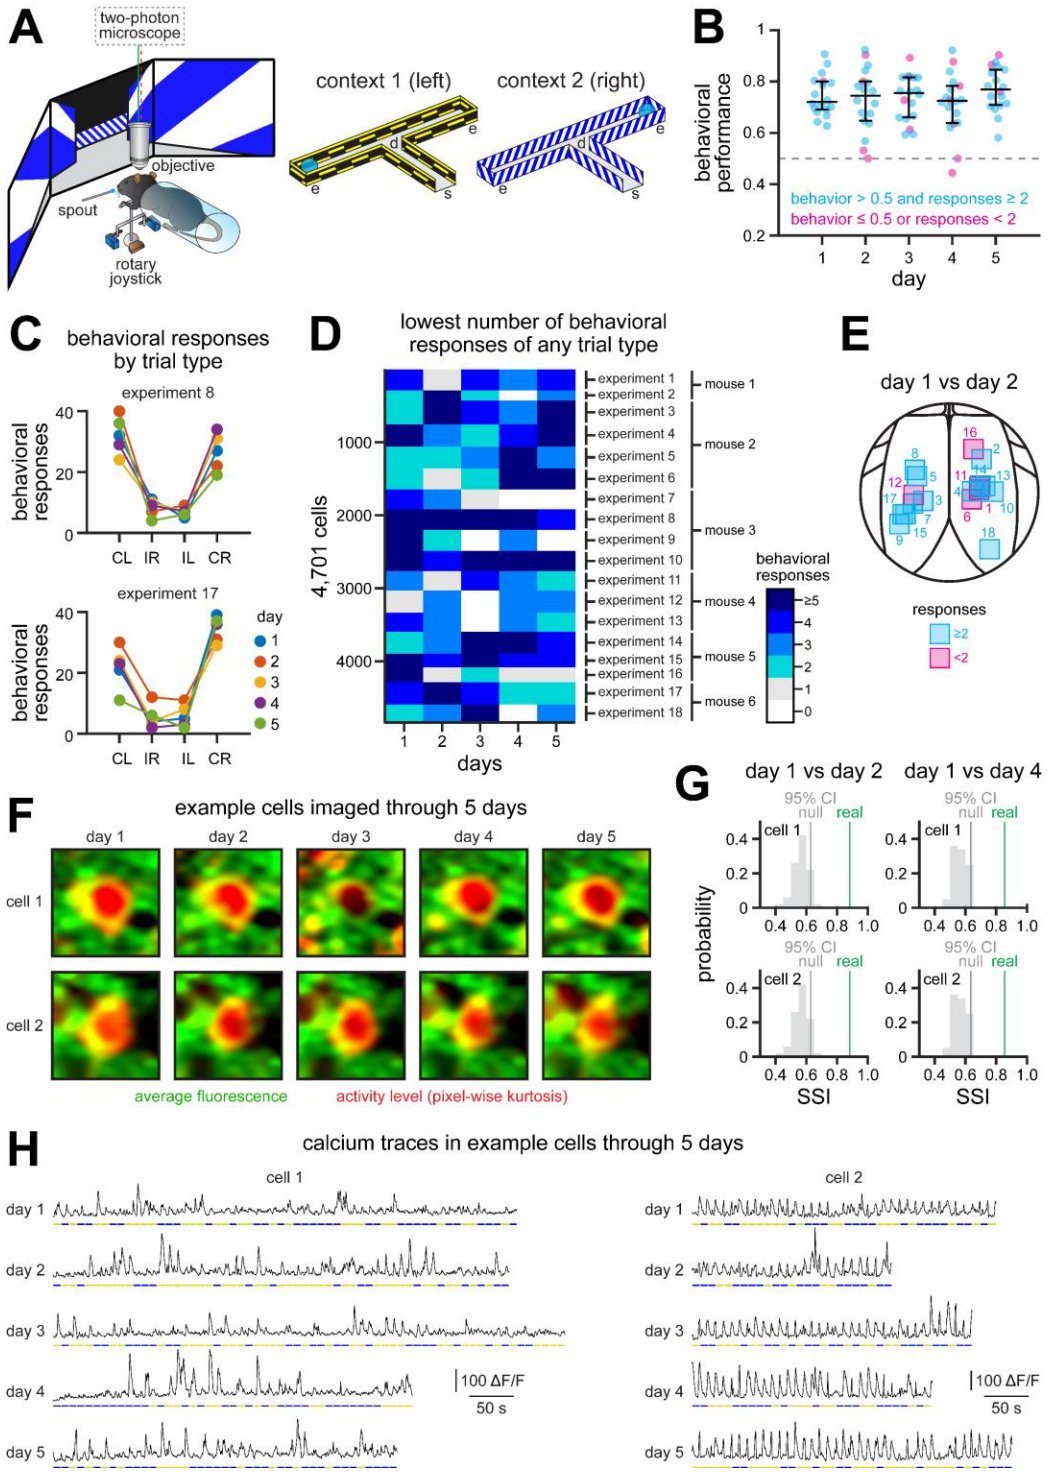

Fig. 1. Context-choice association task for studying stability of cortical representations in mice.

A. Experimental setup depicting a mouse using a rotary joystick to choose left or right trajectories depending on the visual context displayed on monitors spanning the horizontal visual field. Neural activity from populations of L2/3 RSC neurons is recorded using a two-photon microscope[23]. s = trial start; d = decision point; e = trial end.

B. Behavioral performance across 5 consecutive days (n = 18 experiments in 6 mice). Mouse performance was significantly above chance in all experiments (median): day 1 = 72.1%, day 2 = 74.5%, day 3 = 75.5%, day 4 = 72.5%, day 5 = 76.9%. Gray dots, individual sessions; black overlay, median ± 75% CIs. Only experiments with a behavioral performance above chance, and at least 2 behavioral responses of any type (see **Fig.1D**) are considered for further analyses.

C. Example experiments showing the number of behavioral responses across 5 days of experimentation. Note the higher number of correct responses in these expert mice (more examples in **Fig. S1A**). CL = correct left, IR = incorrect right, IL = incorrect left, CR = correct right.

D. Heatmap showing the lowest number of behavioral responses of any type in each experiment. For instance, on day 4 of experiment 8 (**Fig. 1C**), there are 29 correct left, 9 incorrect right, 7 incorrect left and 34 correct right responses. Thus, the lowest number of behavioral responses is 7. Only experiments with at least 2 behavioral responses from each trial type are considered for further analyses.

E. Map showing the location of included (cyan) and excluded (magenta) experiments for all comparisons between days 1 and 2. See **Fig. S1B-C** for maps of all comparisons between days.

F. Example cells imaged across 5 days of experimentation (30 x 30 μm field). Note the stability of the imaging field (green), regardless of variability in activity levels (red).

G. We estimated the stability of ROIs by calculating the structural similarity index (SSI) between 30 μm square regions surrounding each ROI across days, and compared that to a null distribution of SSIs obtained by comparing the original ROI in day 1 against random ROIs in other days (**Fig. S3**). Only ROIs with a real SSI higher than the 95th CI of their respective null SSI were considered for further analyses.

H. Activity traces for two example cells along the duration of each experimental session. Note conserved dynamics across days (cell 1 principally responds in context 1; cell 2 responds in both contexts).

At the population level, RSC neurons appear to maintain a similar structure across days. To visualize this, we plotted the activity of the neural population sorted by cross-validated latency, using odd trials in day 1 to sort neural responses by latency, and then plotting the responses in even trials across days (**Fig. 2A**). We observe that responses in correct trials are more consistent across days, with a similar ridge in population activity (yellow diagonal, **Fig. 2A**). To test this, we calculated the Pearson's correlation coefficient of population vectors in correct and incorrect trials (**Fig. 2B**). We find that activity in correct trials is indeed more stable than in incorrect trials, as revealed by higher correlations between cross-validated population vectors (**Fig. 2B**), which are significantly higher along the diagonal describing the trial duration (**Fig. 2C**; n = 1,609 neurons resampled over 100 iterations; mean correlation in correct trials > 95th CI of the correlation in incorrect trials). To ensure this was not due to unequal trial numbers, we sampled equally across trial types and confirmed that population vectors in correct trials exhibit higher stability (**Fig. S4**). In addition, we noticed that population activity in correct trials experienced a steeper decay in stability after one day, and a slower decay on the following days, but maintaining substantial levels of structure (**Fig. 2A**) and stability (**Fig. 2B-C**) across

the 5 days of experimentation. This suggests that although activity in individual neurons undergoes different levels of drift, the population preserves considerable task-related information across days, particularly in correct trials. Moreover, when we analyzed the responses of RSC neurons in each day independently, we found similar levels of activity (**Fig. S5A**), a similar fraction of neurons mapping the trial duration (**Fig. S5B**), and similar preference indices for a particular context (**Fig. S5C**). We also observed similar levels of activity in correct and incorrect trials (**Fig. S4D-E**). This indicates that while some neurons lose their task-driven responses across days (**Fig. 2A**), other neurons become tuned to the task on each day (**Fig. S5A**), contributing to the homeostasis of task-related representations across the population. Moreover, these results indicate that population dynamics are preserved in the RSC circuit across days, despite some variability in single cell activity.

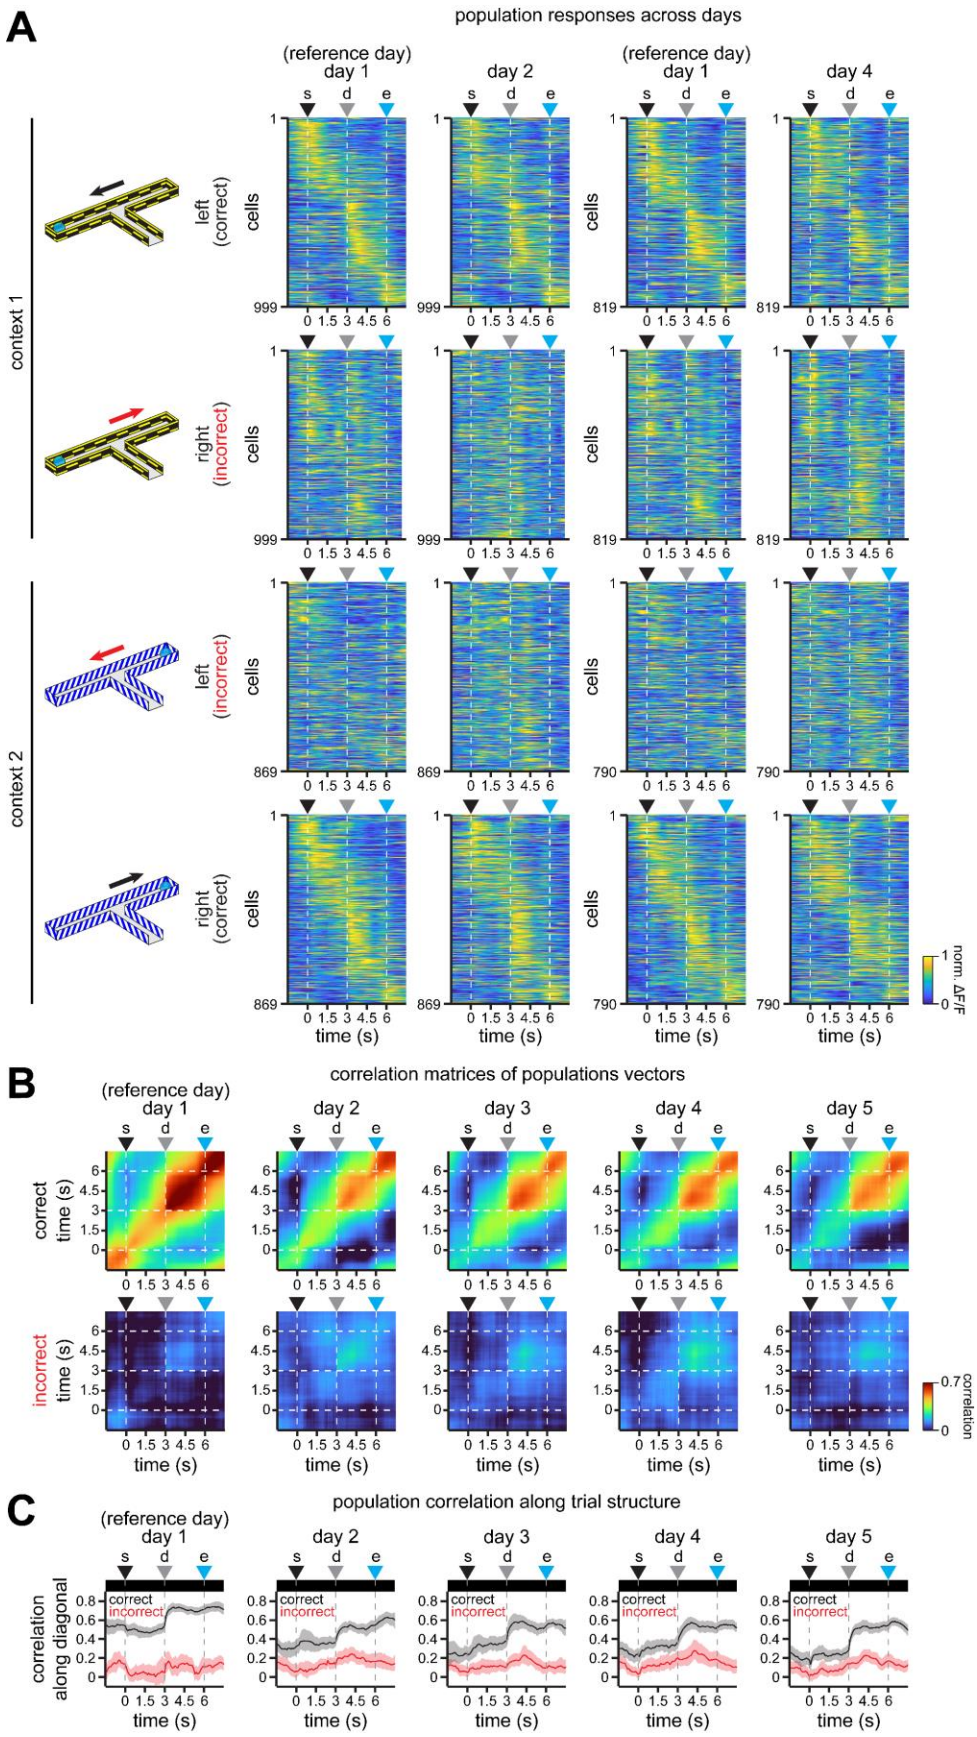
A
population responses across days
(reference day)
day 1
day 2
(reference day)
day 1
day 4
s
d
e
context 1
context 2
left
(correct)
right
(incorrect)
left
(incorrect)
right
(correct)
cells
time (s)
norm. ΔF/F
B
correlation matrices of populations vectors
day 3
day 5
correct
incorrect
correlation
C
population correlation along trial structure
correlation
along diagonal

Fig 2. Population activity in RSC exhibits different levels of stability in correct and incorrect trials.
A. Average normalized activity in the recorded population of neurons separated by their preferred context. For clarity, only comparisons between days 1 and 2, and days 1 and 4 are shown. Activity is cross-validated, sorted by latency using odd trials in day 1, and plotted using even trials for days 2 or 4.
B. Pearson's correlation coefficient matrices showing the similarity of population activity along trial duration (averaged across 100 iterations of cross-validated population vectors constructed from the preferred responses in each context, as shown in **A**). Briefly, population vectors in one half of resampled trials in day 1 are compared to population vectors in the other half of resampled non-overlapping trials in day 1. For the other days, all resampled trials in day 1 are compared to all resampled trials in days 2 to 5. In addition, to compare vectors of the same length across days, the population is resampled to 1,609 cells within each iteration (the lowest number of cells that meet criteria in all 5 days).
C. Population vector correlation coefficients along the diagonals in **B**. Note the higher stability and similarity of population activity in correct trials. Top insets indicate time bins (in black) where the mean correlation in correct trials > 95th CI of the correlation in incorrect trials throughout the trial. Solid line, mean; shaded area, 95% CIs.
For all panels: s = trial start; d = decision point; e = trial end.

Although population responses were fairly stable in correct trials, this does not necessarily mean that the responses of individual neurons to specific task epochs are stable across days. To further investigate the drift in the activity of individual cells, we compared responses exhibiting a significant difference between the two visual environments (context 1 or context 2), the two motor choices (left or right), or the two trial outcomes (rewarded or not rewarded), using balanced numbers of correct and incorrect trials to distinguish correlated task variables (see **Methods**). **Fig. 3A** shows example cells with significant activity difference between contexts 1 and 2 (cell 1), between left and right choices (cell 2), and between rewarded and not rewarded trials (cell 3, before trial end; cell 4, after trial end). **Fig. S6** contains additional example cells. Note that responses tuned to context and trial outcome appear more stable than responses tuned to motor choice in these example cells (**Fig. 3A**; **Fig. S6**). To quantify this, we calculated the Pearson's correlation coefficient between responses grouped by their preferred context, motor choice or trial outcome, comparing odd trials in day 1 against even trials in days 1 to 5 (**Fig. 3B-C**). We find that responses grouped by context tend to be slightly more stable than responses grouped by motor choice, and that responses grouped by trial outcome (either before or after trial end) show the highest stability (**Fig. 3C**; n = 3,087 [day 1], 1,868 [day 2], 1,908 [day 3]; 1,609 [day 4], and 2,020 [day 5] neurons; data were fit for each individual cell, their model parameters were rank transformed and then compared using linear mixed-effects (LME) models, with a fixed effect for task variable, and random effects for mouse and experimental session; see **Methods** for details). However, post-decision and post-trial outcome responses are difficult to disentangle through correlation analysis, since many neurons respond to both, and are investigated further in the next section. In addition, we find that responses in

correct trials are more correlated with each other than responses in incorrect trials (**Fig. 3D**; n = 1,158 to 1,527 [day 1], 601 to 841 [day 2], 763 to 1,045 [day 3], 529 to 814 [day 4], 785 to 1,040 [day 5] overlapping subpopulations of neurons; data were fit for each individual cell, the model parameters were ranked and then compared using LME models; fixed effect for trial type, random effects for mouse and experimental session), confirming our observations in the population data (**Fig. 2**). This suggests that most task-related information is conserved in correct trials across days. However, correlations between responses in incorrect trials consistently stayed above zero, suggesting a contribution in the stability of task-relevant representations (**Fig. 3D**). In terms of their response dynamics, both context and motor related responses exhibit similar profiles in correct and incorrect trials, with more prominent responses typically between trial start and trial end (cells 1 and 2 in **Fig. 3A**; cells 5 and 6 in **Fig. S6**). Interestingly, responses grouped by trial outcome can either peak after the decision point or after trial end, with some cells exhibiting mixed dynamics (cells 3 and 4 in **Fig. 3A**; cells 7 and 8 in **Fig. S6**). Thus, for all subsequent analyses we consider these two outcome-related responses separately, as they likely underlie different cognitive processes. Overall, our analyses of calcium responses indicate that activity in the RSC circuit is partially stable, with conserved population structure across days, particularly in correct trials (**Fig. 2**), which likely depends on the specific tuning properties of each cell (**Fig. 3**). However, although the response correlation is useful to draw general comparisons between activity traces in individual neurons or between population vectors, a more detailed examination that considers the structure of the trial is necessary in order to determine whether representations of specific task variables exhibit distinct levels of stability.

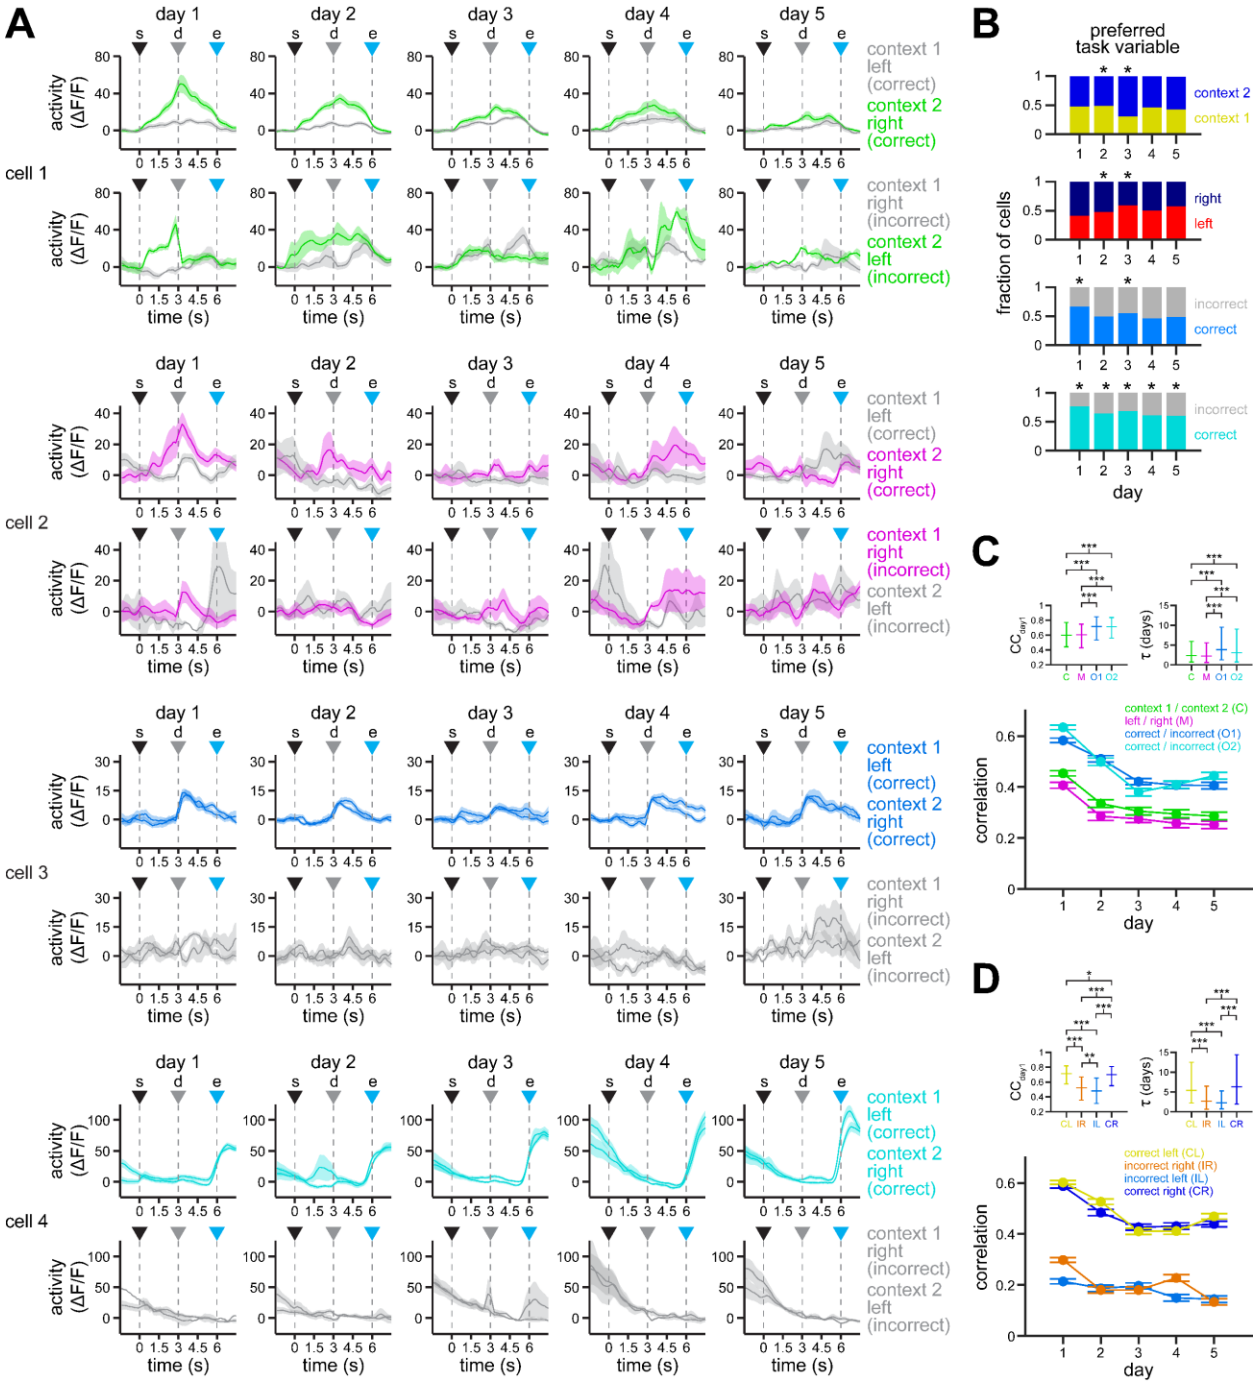

Fig. 3. Activity in individual RSC neurons tuned to visual context, motor choice or trial outcome exhibit different levels of stability.

A. Example neurons exhibiting higher responses for context 2 (cell 1), for right turns (cell 2), for correct post-decision outcome (cell 3), and for correct post-trial outcomes (cell 4), respectively. Traces are color-coded to highlight the difference between the preferred and non-preferred responses in both correct (top row) and incorrect (bottom row) trials. Solid line, mean; shaded area, s.e.m. s = trial start; d = decision point; e = trial end.

B. Fraction of cells with a preference either context, motor choice or outcome (evaluated before or after trial end). Preference is defined as higher activity in the indicated subset of trials. Briefly, trials are equally sampled within each

subset, then compared using Wilcoxon signed-rank test. If activity is significantly different over 5 consecutive time bins (0.5s), then a cell is considered to have a significant preference. Asterisks denote differential fraction of cells, assessed by non-overlap of the bootstrapped 95% CI ranges.

C-D. Average Pearson's correlation coefficients (CC) between odd trials in day 1 and even trials in days 1 to 5 (mean ± s.e.m.). **C** shows data grouped by their preferred task variable, whereas **D** shows data grouped by trial type. Top insets show significantly different comparisons for fit parameters (*$p < 0.05$, **$p < 0.01$, ***$p < 0.001$; LME models; see **Methods** for details).

Encoding of task variables exhibits distinct levels of stability in individual RSC neurons

The analysis of responses across days indicates that responses to particular task variables are unstable across days, and may show differential stability between task variables. However, response correlation might be biased for transient responses (e.g., outcome responses) and not account for changes in coding strength across the entire trial. To assess the stability in the encoding of task variables throughout the time course of the trial, we evaluated the performance of linear decoders across each individual time point[24,30]. To this end, we trained support vector machine classifiers for each specific task variable (context, choice, outcome) using activity in day 1 as reference, and tested the performance of these classifiers on the activity of the same cells in the following days (see **Methods** for details). To analyze day 1 activity without overfitting, we used half of the trials for fitting the classifier, and the other half for testing. We find that some neurons specifically and significantly encode task-related information in day 1 and, in some cases, can maintain their encoding preference across days with similar temporal dynamics as in day 1 (**Fig. 4A**; more examples in **Fig. S7**; for details about defining significant encoding see **Fig. S8**). However, it appears different task variables are encoded with different levels of stability. From the examples in **Fig. 4A** and **Fig. S7**, we observe that post-trial outcome (cells 4, 8 and 12) and context (cells 1, 5 and 9) neurons exhibited the highest coding stability. Post-decision outcome (cells 3, 7 and 11) neurons also show some level of coding stability, while motor choice (cells 2, 6 and 10) neurons display a rapid decay. Indeed, when we look at the population of neurons displaying significant encoding of each task variable in day 1, we observe a similar trend, with outcome and context neurons sustaining higher decoding performance than motor neurons (**Fig. 4B**). This is also confirmed by the higher average decoding performance of neurons that significantly encode outcome and context across days (**Fig. 4C**; n = 164 to 488 [day 1], 86 to 234 [day 2], 115 to 324 [day 3], 94 to 319 [day 4], 107 to 353 [day 5] overlapping subpopulations of neurons; data were fit for each individual cell, the model parameters were ranked and then compared by LME models; fixed effect for task variable, random effects for mouse and experimental session). Our results suggest that the

RSC constantly experiences a representational drift in the storage and encoding of task-related information, with variables exhibiting different drift rates.

One implication of the observed progressive decay in the encoding of information across days is that task information will eventually be lost by the circuit, given enough time. Yet, we observe that mice maintain behavioral performance for several weeks[24]. For this reason, we trained and tested linear classifiers in each day of experimentation independently. We find consistent subpopulations of neurons encoding context, motor choice, post-decision outcome, and post-trial outcome in each particular day (**Fig. S9**). This suggests that although some neurons lose their ability to encode features of the task, other neurons show increased task variable encoding, maintaining homeostasis in the RSC circuit.

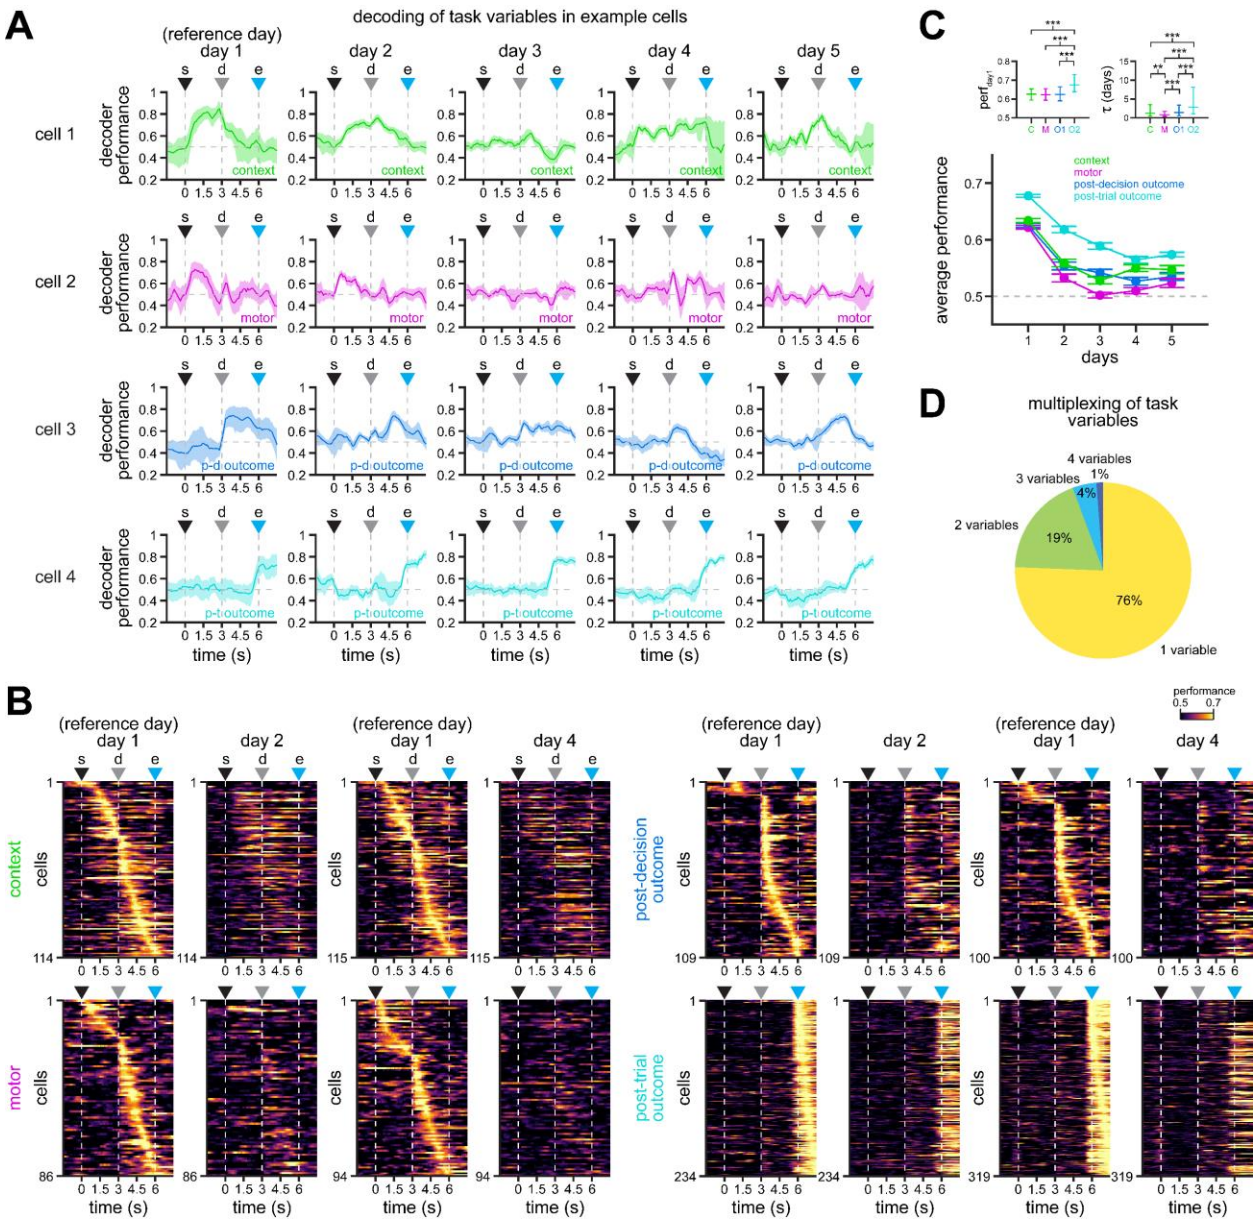

Fig. 4. RSC neurons encode task variables with differential stability.

A. Example cells displaying significant decoding of context (cell 1, top), motor choice (cell 2, top-middle), post-decision outcome (cell 3, bottom-middle), and post-trial outcome (cell 4, bottom) in day 1. Solid line, mean; shaded area, 95% CIs. Decoding performance obtained using support vector machine classifiers.

B. Average decoding of the indicated task variables in single cells. All cells exhibit significant decoding in day 1. Only comparisons to days 2 and 4 are shown for clarity. Note that decoding of outcome and context is more stable than decoding of motor choice.

C. Average decoding performance of cells that show significant decoding in day 1 (mean ± s.e.m.). Top inset shows significantly different comparisons for fit parameters (*$p < 0.05$, **$p < 0.01$, ***$p < 0.001$; LME models; see **Methods** for details).

D. Percentage of cells with significant decoding of 1, 2, 3 or 4 task variables in day 1.

For all panels: s = trial start; d = decision point; e = trial end.

Our analysis of encoding of task information in individual neurons revealed another important finding. We observe that although most neurons specifically encode information about a single task variable, some neurons can multiplex different sets of variables (**Fig. 4D**, **Fig. S10**). For this reason, we investigated how different types of information within individual neurons are preserved over time. To evaluate this, we compared the performance of support vector machine classifiers in neurons with significant encoding of two task variables in day 1 (**Fig. 5A-B**, **Fig. S11A-B**). By examining the population of neurons, we find a larger fraction of cells that maintain their encoding of context and outcome information than their encoding of motor choice (**Fig. 5A**, **S11A**). This trend can also be observed in individual examples (**Fig. 5B**, **Fig. S11B**). We confirmed these findings by calculating the decoding performance across days (**Fig. 5C**; $p < 0.02$; n = 43 to 132 [day 1], 20 to 71 [day 2], 30 to 88 [day 3], 24 to 81 [day 4], 29 to 93 [day 5] overlapping subpopulation of neurons multiplexing two task variables; data were fit for each individual cell, the model parameters were ranked and then compared by LME models; fixed effect for task variable, random effects for mouse and experimental session). Moreover, of the two different types of outcome encoding, post-trial outcome showed higher stability (**Fig. S12**), which might be related to the presence of a reward and/or increased licking in correct trials. To further characterize neurons multiplexing different task variables, we compared the density of the area under the curve of the classifier performances (**Fig. S13A**) for the subpopulations shown in **Fig. 5A** and **Fig. S11A**. We observe higher stability in the encoding of context in context-motor cells, particularly in the first half of the trial, before the decision point (top row, **Fig. S13**). Similarly, we observe higher stability in the encoding of trial outcome in motor-outcome cells, specifically after the decision point during the second half of the trial (two bottom rows, **Fig. S13**). In context-outcome cells, however, stability seems bimodal, with context showing higher stability before the decision point, and trial outcome higher stability by the end of the trial (second and third rows, **Fig. S13**). These results highlight an important feature of RSC neurons that multiplex task information, which favor the maintenance of environmental context and trial outcome over the encoding of motor choice.

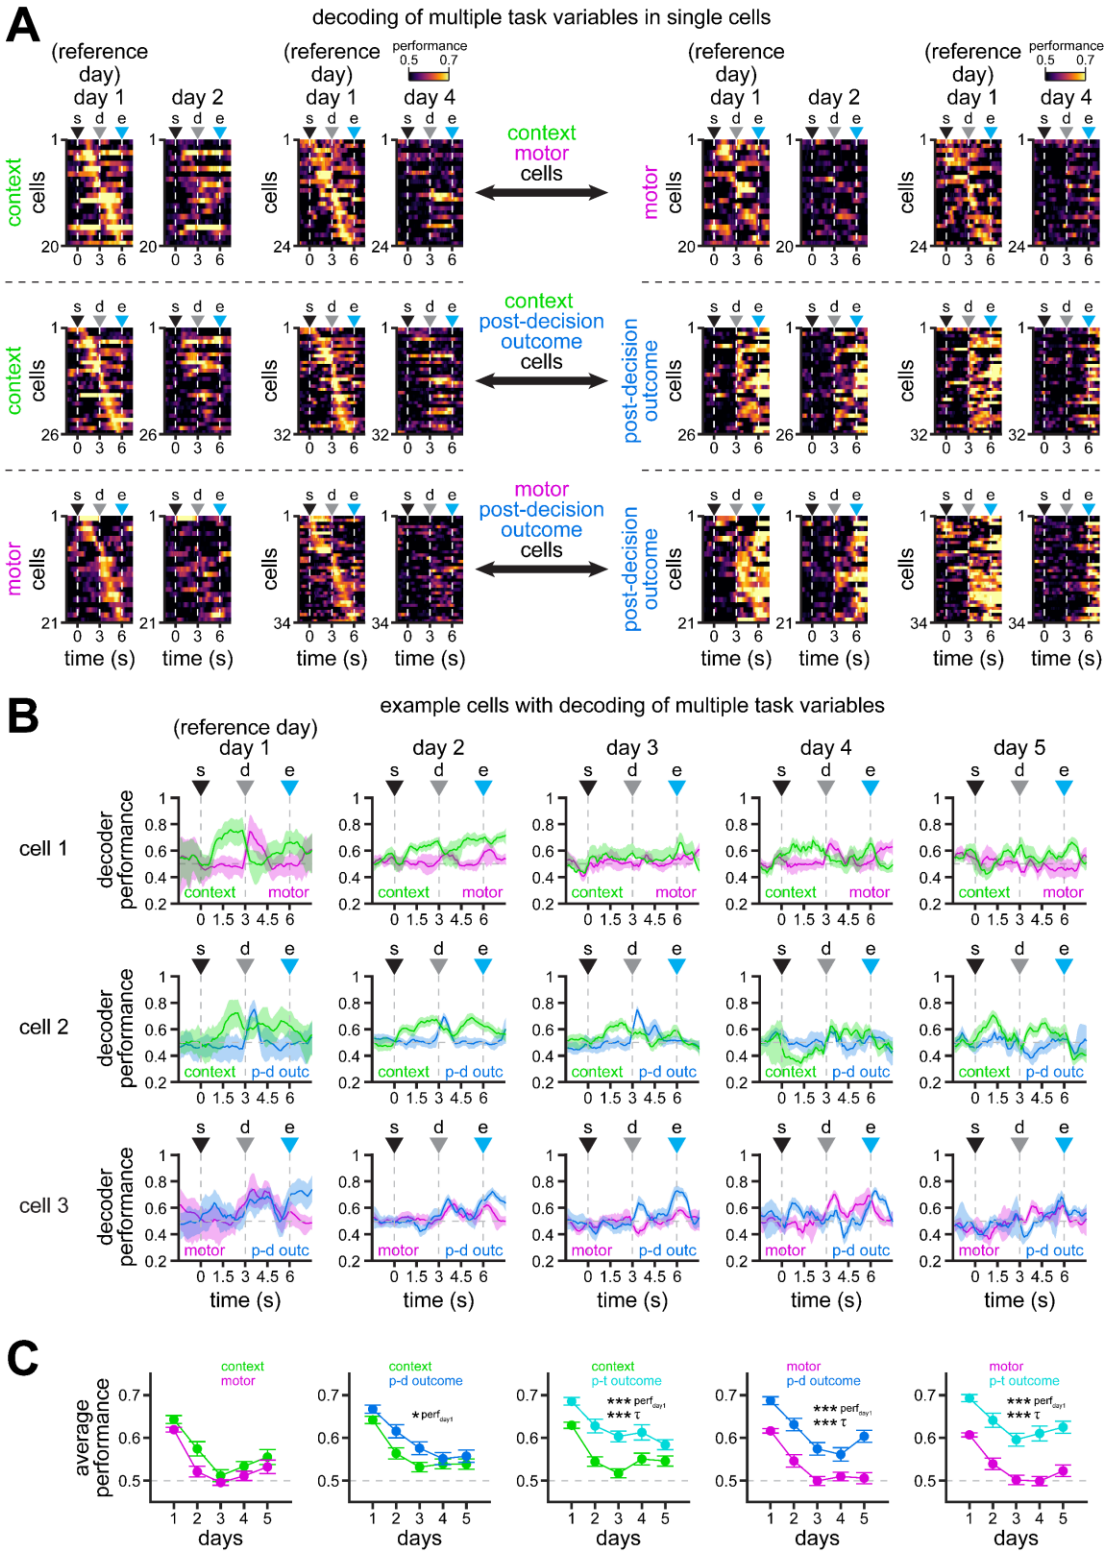

Fig. 5. Multiplexing of task information with varying stability within single RSC neurons.

A. Average performance of support vector machine classifiers on the activity of single cells that significantly encode two task variables: context and motor choice (top row), context and post-decision outcome (middle row), and motor choice and post-decision outcome (bottom row) in day 1. Additional cells with multiplexing decoding of different pairs

of variables are shown in **Fig. S11** and **Fig. S12**. Note the better decoding performance of models trained in day 1 and tested in days 2 and 4 for context and outcome variables.

B. Example cells with significant encoding of context and motor choice (cell 1, top), context and post-decision outcome (cell 2, middle), and motor choice and post-decision outcome (cell 3, bottom) in day 1. Solid line, mean; shaded area, 95% CIs.

C. Average decoding performance of the two task variables indicated in each plot within individual cells (mean ± s.e.m.). Asterisks denote significantly different decoding performance, as measured by different fit (**p* < 0.05, ***p* < 0.01, ****p* < 0.001; LME models; see **Methods** for details). Note the higher decoding performance of the classifiers decoding context and outcome when compared to motor choice.

For all panels: s = trial start; d = decision point; e = trial end.

Differential stability in population coding of task variables in the RSC

Our findings on the stability of population activity (**Fig. 2B-C**) and, in particular, on the differential stability of information encoding in individual neurons (**Fig. 4B-C**) suggest there may be similar dynamics in the population-level encoding of task variables in the RSC. To test this hypothesis, we trained support vector machine classifiers on population vectors in day 1, and tested those models on population vectors in the following days for each individual 5-day experiment. Inspection of an example experiment (**Fig. 6A**) as well as the entire group of experiments (**Fig. 6B**) further support our previous findings, in which encoding of context and trial outcome exhibit the highest stability while encoding of motor choice shows lower stability. This is also confirmed by fitting an exponential decay function to the performance of classifiers in the group of experiments shown in **Fig. 6B**. Briefly, we calculate the area between the classifier performance curve and the chance level (0.5) at relevant windows during task performance for each day: 1.5-3s for environmental context (immediately prior to the decision point), 3-4.5s for motor choice and post-decision outcome (immediately after the decision point), and 6-7.5s for post-trial outcome (during the reward period). We then normalize the area in day 1 (when the classifier is trained), and fit an exponential decay function ($\mathrm{NP} = e^{-1/\tau}$) to estimate the time at which decoding performance decays to 1/e (~37%) of its original value in day 1. The normalized performance (NP) allows us to measure the encoding stability of task variables, where a value = 1 signifies same decoding performance as in day 1 and a value = 0 indicates that decoding performance has decreased to chance level. Consistent with our observations in individual neurons, we find that context (τ = 4.3 days) and post-trial outcome (τ = 5.1 days) exhibit the highest stability, whereas post-decision outcome (τ = 3.1 days) and motor choice (τ = 1.6 days) display lower stability (**Fig. 6C**). Moreover, fitting exponential decays to each individual experiment revealed similar results (**Fig. 6D**, **Fig. S14A-B**; n = 8 [context], 11 [motor], 11 [post-decision outcome], 11 [post-trial outcome]). Pairwise variable comparisons using LME

models found significant differences between context and motor (p = 0.02), motor and post-decision outcome (p = 0.01), motor and post-trial outcome (p < 0.0001), and between post-decision and post-trial outcome (p = 0.01), further supporting differential stability of task-related representations in the RSC, and its persistence within single experiments and across the population of experiments. Data grouped by mouse show qualitatively similar results (**Fig. S14C-D**). Importantly, this representational drift might indicate that individual experiments in some mice no longer encode task variables after a few days. However, models trained and tested on each day independently showed that there is sufficient coding power in the population throughout the duration of the experiments (**Fig. S15**).

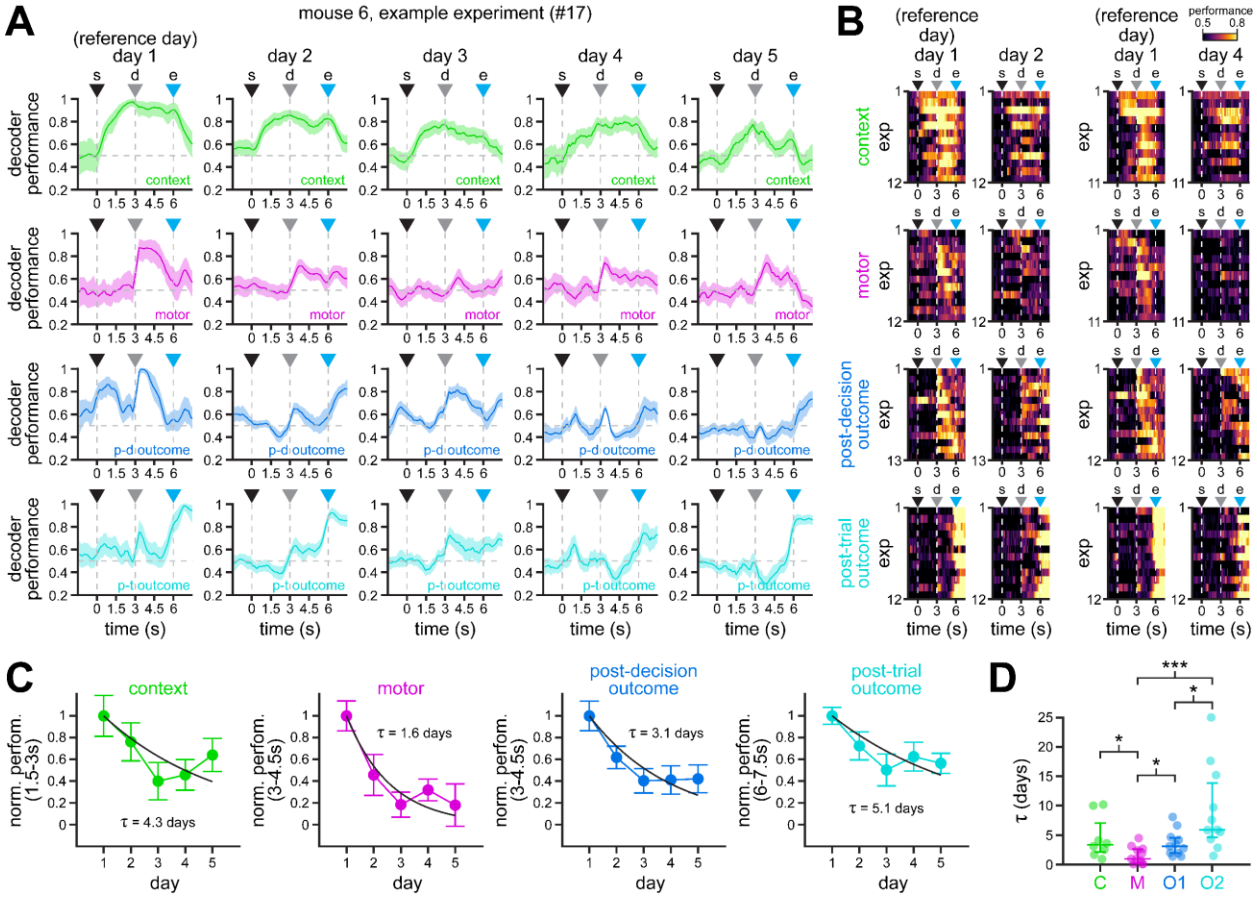

Fig. 6. Subpopulations of RSC neurons recorded in individual experiments exhibit different stability of task-related representations.

A. Encoding of context (top row), motor choice (top-middle row), post-decision outcome (bottom-middle row), and post-trial outcome (bottom row) by populations of neurons recorded in individual experiments. Solid line, mean; shaded area, 95% CIs. Note the gradient in population encoding, where post-trial outcome and context are more stable than motor choice and post-decision outcome.

B. Heat maps of context (top), motor choice (top-middle), post-decision outcome (bottom-middle), and post-trial outcome (bottom) coding performance along trial duration for the different experiments in this study. For clarity, only

comparisons between day 1 and days 2 and 4 are shown. Note the higher stability of context (top row) and post-trial outcome (bottom row) encoding in days 2 and 4 between trial start and trial end, and after trial end, respectively. By contrast, motor choice and post-decision outcome exhibit lower stability.

C. Normalized decoding performance (averaged across experiments) integrated over the indicated 1.5 s windows relevant for each task variable (mean ± s.e.m.). Exponential decay functions were fit to estimate the decay in encoding stability across days (black lines). Note the faster decay in motor choice and post-decision outcome encoding, and the slower decay in context and post-trial outcome encoding.

D. Time constants (τ) calculated after fitting exponential decay functions to the decoding performance of each task variable in each individual experiment (see **Fig. S14A-B**). Dots, individual sessions; overlaid lines, median ± 75% CIs. For clarity, one outlier τ value for post-trial outcome (33.5 days) is shown at saturation (25 days). Asterisks denote significantly different decay constants (*p < 0.05, **p < 0.01, ***p < 0.001; LME models; see **Methods** for details). For all panels: s = trial start; d = decision point; e = trial end.

In order to reduce sampling variability in individual experiments and to evaluate whether the encoding of task variables exhibit similar dynamics across the RSC circuit, we used pseudo-populations of neurons pooled across experimental sessions to analyze the data across all experiments and all mice. In accordance with our previous findings[24], classifiers trained on population vectors show high decoding performance of environmental context between trial start and trial end, which corresponds to the presence of visual stimuli (upper left panel, **Fig. 7A**). Decoding of motor choice is highest around the decision point, when mice rotate the joystick to select a trajectory through the maze (top-middle left panel, **Fig. 7A**). Decoding of post-decision outcome exhibits a prominent peak around the decision point (bottom-middle left panel, **Fig. 7A**), whereas decoding of post-trial outcome displays ceiling level performance during the reward period, after trial end (bottom left panel, **Fig. 7A**). In terms of their stability, decoding of environmental context shows a clear and smooth decay across the 5 days of experimentation, with performance above chance level (0.5, dashed horizontal line) even on day 5 (top row, **Fig. 7A-B**). By contrast, population decoding of motor choice undergoes a rapid decay from day 1 to day 2, and declines to near chance levels by day 3 (top-middle row, **Fig. 7A-B**). In addition, post-trial outcome exhibits very stable population decoding at the end of the trial (6-7.5s) across the 5 days of experimentation (bottom row, **Fig. 7A-B**), whereas post-decision outcome (3-4.5s) displays a more rapid decay (bottom-middle row, **Fig. 7A-B**). To quantify this, we fit exponential decay functions to the performance of classifiers trained on population data (**Fig. 7C**). Consistent with our observations in single cells (**Fig. 4**, **Fig. 5**) and individual experiments (**Fig. 6**), we find that context (**τ** = 7.3 days; 95% CIs = 5.9 to 9.7 days) and post-trial outcome (**τ** = 35.7 days; 95% CIs = 25.9 to 50.0 days) are considerably more stable than motor choice (**τ** =

1.9 days; 95% CIs = 1.4 to 2.5 days) and post-decision outcome (τ = 2.2 days; 95% CIs = 1.9 to 2.5 days), as shown by the non-overlap between their 95% CIs (**Fig. 7C**).

To ensure our findings can be generalized regardless of the chosen reference day, we also trained support vector machine classifiers on population data from days 2 to 5 and tested them on the reminder of days (**Fig. S16**). We find a similar trend, with representation of context and post-trial outcome being considerably more stable than representations of motor choice and post-decision outcome, even when evaluating models in future or past days (**Fig. S16C**). In addition, since we maximized the usage of available data throughout this study based on behavioral parameters (**Fig. 1B-E**, **Fig. S1**), morphological stability of ROIs (**Fig. 1F-G**, **Fig. S3**), and reliability of task-related neuronal activity, we then analyzed only neurons that pass all these tests consecutively across all 5 days of experimentation. We confirmed that this smaller subpopulation of 601 neurons also exhibits a similar differential stability profile in task-related representations of the RSC (**Fig. S17**). Training and testing support vector machine classifiers on the inferred spikes obtained by deconvolution of calcium signals (**Fig. S18A-B**), also reveals similar dynamics in RSC representations, with context and post-trial outcome displaying higher stability across days (**Fig. S18C-D**). Moreover, we confirmed that our stability analyses are not biased by the normalization of decoding performance, which allowed us to estimate decay constants for the decline in decoding performance across days. The average performance of classifiers trained on population data also shows higher stability for the context and outcome variables, and lower stability for motor choice at behaviorally relevant windows (**Fig. S19A**). In addition, since the size of the subpopulations of neurons exhibiting significant decoding in day 1 can vary (**Fig. 4B**), we trained classifiers with equally sampled numbers of neurons. We again find higher stability in the encoding of context and post-trial outcome (**Fig. S19B**). To ensure that greater stability is not solely a feture of higher day 1 performance, we balanced day 1 performance by determining the number of randomly sampled neurons necessary for 90% decoding of each task variables on day 1 (**Fig. S19C**). Using these subpopulations, we once again found higher stability for context and post-trial outcome, and lower stability for motor choice and post-decision outcome (**Fig. S19D**).

Together, our results show a gradient in the stability of information encoding and storage in the RSC, in which environmental context is more stable than motor choice. This, in turn, suggests a mechanism where the RSC reliably stores sensory information from the environment while, at the same time, allows flexibility in the mapping of motor choice, which likely depends

on current sensory demands and reward history (**Fig. 7D**). Moreover, it is likely that association areas, which multiplex different streams of information, encode sensory and motor signals with lower levels of stability than the dedicated sensory and motor cortices they connect to (**Fig. 7E**). This ability of association areas to process information in a more flexible manner than dedicated sensory and motor cortices might help to associate incoming sensory signals to appropriate choices in the continuously changing situations that animals encounter in the environment.

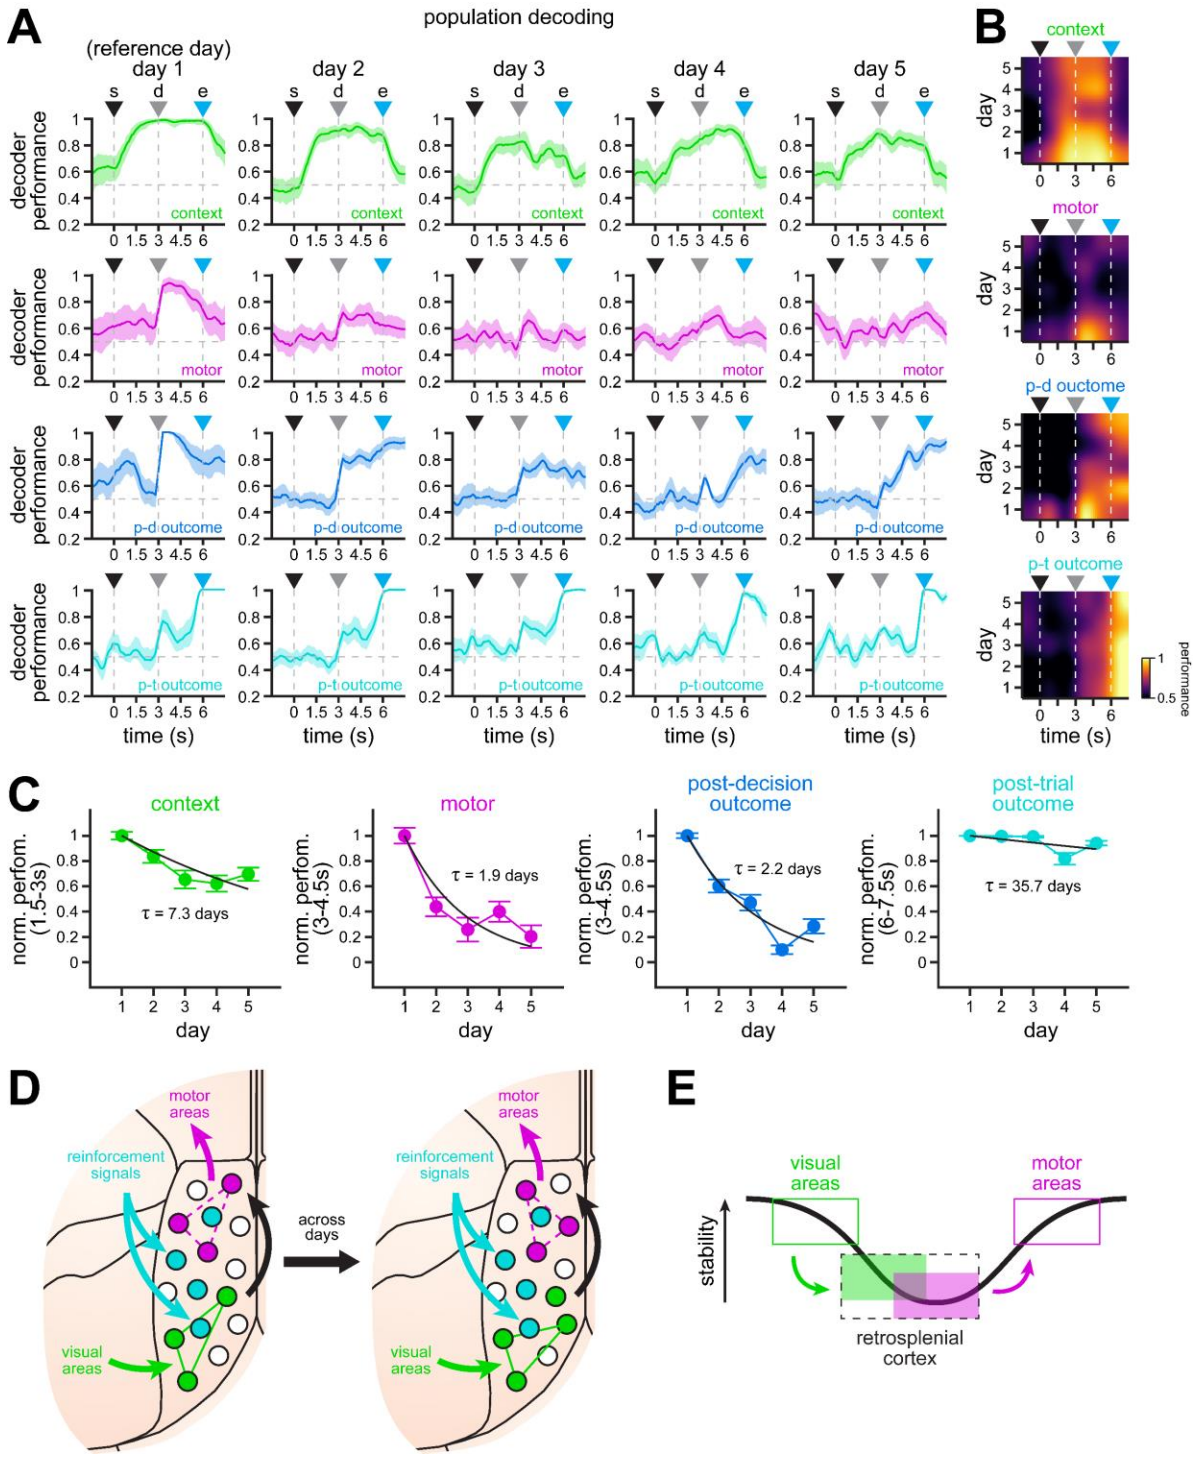

Fig. 7. Population-level encoding of task variables exhibits differential stability in RSC.

A. Encoding of context (top row), motor choice (top-middle row), post-decision outcome (bottom-middle row) and post-trial outcome (bottom row) by pseudo-populations of RSC neurons pooled across all mice. Solid line, mean; shaded area, 95% CIs. Note the gradient in population encoding, where post-trial outcome and context are more stable than post-decision outcome and motor choice.

B. Heat maps of context (top), motor choice (top-middle), post-decision outcome (bottom-middle), and post-trial outcome (bottom) coding performance along trial duration and across all 5 days of experimentation. Note the stability of context encoding between trial start and trial end. Also, note the stability of post-trial outcome encoding after trial end. By contrast, the encoding of motor choice and post-decision outcome around the decision point decays more rapidly after day 2.

C. Normalized decoding performance integrated over the indicated 1.5 s windows relevant for each task variable (mean ± bootstrapped s.e.m.). Exponential decay functions were fit to estimate the decay in encoding stability across days (black lines denote fits to the average normalized performance). Note the faster decay in motor choice and post-decision outcome encoding, and the slower decay in context and post-trial outcome encoding. To obtain CIs for the decay constants, we also fit decay functions to each iteration of the decoder performance (95% CIs, context = 5.9 – 9.7 days, motor = 1.4 – 2.5 days, post-decision outcome = 1.9 – 2.5 days, post-trial outcome = 25.9 – 50.0 days).

D. Schematized working model of RSC processing and representational stability. We hypothesize that current visual sensory features are matched to stored reference features (green arrow), which then triggers a motor output (magenta arrow) from a collection of possible choices during the association process (black arrow). Our model also postulates that sensory features are more stably represented in the RSC (solid green lines), which is supported by strong connectivity with primary and secondary visual areas[24,27,31], and that motor choice is less stable (dashed magenta lines), possibly favoring flexibility during the decision process. In addition, reinforcement signals (cyan arrows) help maintain stable responses of RSC neurons tuned to task features, particularly those related to trial outcome and environmental context.

E. Previous studies have reported higher levels of stability in sensory and motor cortices [5,6,9,32]. Our results in retrosplenial cortex suggest that stability in association areas, which typically comprise neurons multiplexing different streams of information, follows a different function. First, sensory inputs are represented with lower stability than in the region they originate from. Motor coding, although stable within a session, changes considerably between days, which may lead to more flexible remapping to new contingencies. We expect that motor representations will become very stable again in more peripheral motor cortices.

For all panels: s = trial start; d = decision point; e = trial end.

DISCUSSION

It has been proposed that a principal function of the mammalian cortex is to preserve memories through stable representations of sensory, motor, and cognitive variables[33–35]. However, recent technological developments for longitudinal recording of neural activity have shown that neuronal response tuning can be dynamic. Often, different subsets of neurons encode information about the same cognitive variable on different days. Yet, these representations must remain stable in order to generate reproducible behaviors over time, particularly in cortical circuits critical for performing cognitive tasks. In this study, we set out to investigate how neurons in the RSC, important for the association of contextual sensory information and motor choice[24,36], maintain their coding properties over consecutive days. We find that different task variables are represented with distinct levels of stability in the RSC, both at the population and at the single cell levels. Environmental context, which likely reflects visual sensory features, as

well as post-trial outcome are the two variables represented with the highest stability. By contrast, motor choice is less stable, likely favoring flexibility during the decision process. We argue that the representation of cognitive variables with different stability might be a common feature of association areas.

One key finding in our study is the differential stability of correct and incorrect responses during task performance. A possible confound is that activity levels during correct trials might be higher than in incorrect trials, thus leading to lower stability in incorrect trials consisting mainly of noise-level fluctuations with no noticeable responses. However, we did not find differences in averaged activity levels between correct and incorrect responses, except at the end of the trial, when mice receive rewards only on correct trials. Hence, the observed differential stability is likely due to more consistent and reliable responses in correct trials, both within and across days. Moreover, reward may contribute to the stability of responses tuned to correct trials by enhancing memory consolidation[37,38]. Behaviorally, reduced attention and engagement in task performance might explain the higher variability and lower stability of incorrect responses across days, which might be dominated by task-independent variables in mice[39–45]. In addition, expectation, particularly after a streak of correct decisions in a single context, might modulate behavior and its underlying neural activity, biasing representations to encode previously rewarded choices[46,47], providing another substrate for increased variability in incorrect trials.

Several studies have reported that dedicated sensory or motor cortices typically encode information in a stable manner over days and weeks[5,6,9,10,14]. However, the stability of sensory and motor representations in association areas, which comprise much of human neocortex, is less understood. Our results indicate that RSC neurons exhibit more stable responses to the environmental context than to motor choices. We postulate that the functional connectivity between context-encoding neurons and motor-encoding neurons within the RSC, which likely mediate the association process, are weak and flexible across days, while staying robust within a single day. Indeed, on individual days both context and motor choice are accurately encoded with distinct dynamics, reflecting the duration of the sensory stimulus and the window for decision, respectively, as we show here and in prior work[24].

It is possible that RSC neurons store reference sensory features that are rapidly matched to incoming sensory inputs, eliciting relatively strong context-dependent responses. However, motor planning responses might be shaped associatively on a daily basis, taking

advantage of synaptic plasticity to generate a stable motor output across days, as has been proposed for posterior parietal cortex in similar tasks[48]. In addition, motor areas might help consolidate the encoding of task features via a feedback mechanism, involving other brain regions such as the dorsal striatum[49]. Together, this evidence suggests that stable representations of sensory features in the RSC may improve the comparison of incoming inputs to stored reference contexts. At the same time, this circuit encodes higher order motor choice outputs more dynamically, which allows updating according to task demands and provides flexibility in decision making. This is particularly relevant during navigation, which is dependent on RSC[23,24,50–58], where motor decisions must be flexibly mapped, even within a particular context. For example, the same intersection may prompt a left turn or a right turn depending on the goal location. In fact, RSC populations were recently shown to play a role in context-value associations and updating, increasing in proportion with task demands, and thus supporting cognitive flexibility across learning[25].

The implications of stable encoding of trial outcome require further consideration. One possible explanation is that reinforcement signals might help to consolidate neural activity and maintain neuronal tuning to task variables[37,38]. In this regard, a recent study found that RSC neurons have been shown to encode contexts more efficiently when they are paired with a reward[25]. Moreover, the observed stability of outcome representations, combined with stable sensory representations, might help shape motor coding on specific days, given the absence of stable motor coding across days. However, it is important to consider that the reward obtained at the end of the trial might cause lick-induced neural activity in correct trials in the RSC. This makes the stable coding at the end of the trial difficult to interpret. Nonetheless, trial outcome is also encoded right after a decision has been made, and well prior to a reward, which might reflect higher order signals related to the expectation (or the lack thereof) of a reward. Alternatively, the encoding of trial outcome during the post-decision period could represent sensory signals (such as visual or auditory) associated with the route leading to a particular outcome.

Recent findings indicate that encoding of learned cognitive variables is dynamic, exhibiting different levels of representational stability[6,17,19]. Yet, neural circuits must perform computations reliably over long periods of time to provide a substrate for robust behavior. In this regard, it has become clear that multiplexed coding of task variables within neural circuits underlies many task-relevant transformations[22,27,40,53–56]. By balancing differential

representational stability of different cognitive variables, animals can maintain stable representations of particular variables, such as the current environmental context, while maintaining flexible associations, such as the appropriate action within that environment.

METHODS

Animals

To achieve widespread calcium indicator expression across the dorsal cortex, we bred Emx1-Cre (Jax Stock #005628) × ROSA-LNL-tTA (Jax Stock #011008) × TITL-GCaMP6s (Jax Stock #024104) triple transgenic mice to express GCaMP6s in cortical excitatory neurons. 12 to 16 week old mice of both sexes were implanted with a head plate and cranial window (n = 6). Water restriction started 7 days after recovery from surgical procedures, and behavioral training started after 7 days of water restriction (14 days after surgery). The animals were housed on a 12 hour light/12 hour dark cycle in cages of up to five animals before the implants, and individually after the implants. All animal procedures were approved by the Institutional Animal Care and Use Committee at the University of California, Santa Barbara.

Surgical procedures

Surgeries were conducted under isoflurane anesthesia (3.5% induction, 1.5 to 2.5% maintenance). Deep anesthesia was monitored throughout the procedure every 15 min by confirming that respiratory rate was kept within 54-80 beats/min, body temperature within 36 ± 1 °C, and there was no response to a toe pinch. Before incision, the scalp was infiltrated with lidocaine (5 mg/kg, subcutaneous) for analgesia. Meloxicam (1 mg/kg, intraperitoneal) was administered preoperatively to reduce inflammation. Once anesthetized, the scalp overlying the dorsal skull was sanitized and removed, and the periosteum was removed with a scalpel. The skull was abraded with a drill burr to improve adhesion of dental acrylic. Then, a 4 to 5 mm diameter craniotomy was made over the midline (centered at 2.5 to 3.0 mm posterior to bregma), leaving the dura intact. A cranial window was implanted over the craniotomy and sealed first with silicon elastomer (Kwik-Sil, World Precision Instruments) and then with dental acrylic (C&B Metabond, Parkell) mixed with black ink to reduce light transmission. The cranial windows were made of two rounded pieces of coverglass (Warner Instruments) bonded with an ultraviolet-cured optical adhesive (Norland, NOA61). The bottom coverglass (4 mm) fit tightly inside the craniotomy, while the top coverglass (5 mm) was bonded to the skull using dental acrylic. A custom-designed stainless steel head plate (eMachineShop.com) was then affixed

using dental acrylic. After surgery, mice were administered carprofen (5 to 10 mg/kg, oral) every 24 hours for 3 days to reduce inflammation.

Virtual T-maze design

Mice were head-fixed using custom restraint hardware (https://goard.mcdb.ucsb.edu/resources) and placed in polypropylene tubes to limit movement. A custom rotatory joystick, located within reach of mouse forelimbs, was mounted on optical hardware (Thorlabs) and attached to an optical encoder (Digi-key Electronics). Two servomotors (Adafruit Industries) were used to constrict the position of the joystick throughout the trial, allowing movement only during the decision window. A spout was placed near to the mouth of the mouse for delivery of water rewards (10 to 12 μl), which were controlled by a solenoid valve (Parker Hannifin). Licks were detected through a capacitive touch sensor connected to the metallic spout and to a metallic mesh inside the polypropylene tube. A virtual T-maze was displayed across three screens arranged in an arc to subtend 180° of the mouse visual field. All electronic components were controlled by custom code written in MATLAB (MathWorks) through Arduino Uno (Arduino).

Virtual mazes were built using the ViRMEn package[59], with modifications to control progression through the maze. We created two different virtual T-mazes with unique wall patterns, one consisting of horizontal yellow dashes on a black background (context 1; luminance: 10.6 ± 0.5 lux, mean ± s.d.; root mean square contrast: 0.496) and one with oblique blue stripes on a white background (context 2; luminance: 13.4 ± 1.0 lux, mean ± s.d.; root mean square contrast: 0.414). The mazes were presented in pseudorandom order across an experimental session. Progression through the virtual environment proceeded at a fixed speed throughout the trial, such that both traversal from the maze start to the decision point and traversal from the decision point to the end of the maze lasted 3 s each. During the decision window, two servomotors opened to allow rotation of the joystick with the forelimbs. All mice used both forelimbs to rotate the joystick. Rotation of the joystick was coupled to rotation of the field of view in the virtual environment. If the field of view rotation exceeded 45° to the left or to the right from the central position, a decision was registered, and the joystick was returned to the central position using the servomotors. If a decision was registered, then a tone was played, indicating the trial outcome (correct: 5 kHz, 70 dB tone for 0.25 s; incorrect: broadband white noise, 80 dB for 0.25 s), and the mouse was automatically rotated 90° in the chosen direction and then progressed to the end of the selected T-maze arm. The end of the rewarded arm contained a reward cue (a blue sphere) that was reached at the same time as water was

administered, whereas the unrewarded arm did not contain the reward cue. As the mouse reached the end of the maze, the outcome tone was played again at the same frequency and volume but with longer duration (correct: 5 kHz, 70 dB tone for 1 s; incorrect: broadband white noise, 80 dB for 2 s). If no decision was detected within 3 s, then the joystick was returned to the center position using the servos, and the trial was aborted and not included in later analyses. At the end of each trial, there was a 3 s interval with a uniform blue (correct trial) or black (incorrect trial) screen before teleporting back to the start of the T-maze for the next trial. To avoid repeated presentations of the same maze, the probability for consecutively displaying the same context progressively decreased after a particular context was displayed, from 50% to 0% over 10 trials. To prevent bias, if the number responses to the left or to the right exceeded double the number of responses in the nonpreferred direction over a 10 trial window, then the context rewarded for the nonpreferred direction was consecutively displayed until bias was reduced below threshold.

Behavioral training

Before training, mice were gradually water restricted beginning 7 days after surgery. HydroGel (99% H2O, ClearH20) was provided in decreasing amounts each day for 7 days (2.0 to 1.2 g). During training, mice were supplemented with HydroGel depending on the amount of water obtained in each particular session to keep their body weight at or above 85% of their initial weight, typically 1.2 mL of water (or HydroGel equivalent) per day. In addition, once per week (before nontraining days), mice received 2.0 g of HydroGel.

Starting 7 days after water restriction, mice spent 5 days of training on a habituation and shaping protocol. During this time, we unilaterally opened the servo to only allow correct decisions (forced choice trials). This stage was important for training mice to move the joystick in both directions without bias while also instructing them on the correct context-reward associations. Following shaping, the percentage of trials with free choices was progressively increased from 10 to 80%. Only free choices were used for all further analyses. Typically, it took mice 25 to 40 sessions to achieve plateau levels of performance. Mice were trained once a day, 5-6 days a week. A single training session lasted 30 min, divided in 2 blocks of 15 min each.

Two-photon imaging

GCaMP6s fluorescence was imaged using a Prairie Investigator two-photon microscopy system with a resonant galvo scanning module (Bruker). For fluorescence excitation, we used a

Ti:Sapphire laser (Mai-Tai eHP, Newport) with dispersion compensation (Deep See, Newport) tuned to λ = 920 nm. For collection, we used GaAsP photomultiplier tubes (Hamamatsu). To achieve a wide field of view, we used a 16×/0.8–numerical aperture microscope objective (Nikon) at an optical zoom of 2× (425 × 425 μm field). Imaging planes at a depth of 90 to 150 μm were imaged at a frame rate of 10 Hz. Laser power ranged from 40 to 75 mW at the sample depending on GCaMP6s expression levels. Photobleaching was minimal (<1%/min) for all laser powers used. A custom stainless steel light blocker (eMachineShop) was mounted to the head plate and interlocked with a tube around the objective to prevent light from the visual stimulus monitor from reaching the photomultiplier tubes.

To achieve longitudinal imaging across days, we identified the same field of view using the blood vessel pattern as a guide, and by measuring the depth from the surface of the cortex with micrometric precision. Landmark cells and/or small blood vessels were used for fine adjustments to match the field of view.

Two-photon postprocessing

Images were acquired using PrairieView acquisition software and converted into TIF files. All subsequent analyses were performed in MATLAB (MathWorks) using custom code (https://goard.mcdb.ucsb.edu/resources). First, images from a single 15 min recording block (2 blocks per 30 min experimental session) were corrected for X-Y movement by rigid registration to a reference image (the pixel-wise mean of all frames in the middle third of the recording block) using two dimensional cross-correlation. Then, each block (2 per day) was independently registered to the first block of the first day of experimentation using the two dimensional cross-correlation of their average projections. This yielded a 150 min long stack containing images across a 5-day experiment. To identify responsive neural somata, a pixel-wise activity map was calculated over each 150 min long stack using a modified kurtosis measure. Neuron cell bodies were identified using local adaptive threshold and iterative segmentation. Automatically defined regions of interest (ROIs) were then manually checked for proper segmentation in a graphical user interface (GUI) that allows comparison between raw fluorescence and activity map images. The corresponding GUI for registration of video stacks and refinement of ROIs can be found at: https://github.com/ucsb-goard-lab/defineCellROIs.

To ensure that the response of individual neurons was not due to local neuropil contamination of somatic signals, a corrected fluorescence measure was estimated according to

$$F_{corrected}(n) = F_{soma}(n) - \alpha * \left(F_{neuropil}(n) - \underline{F_{neuropil}}\right)$$

where $F_{neuropil}$ was defined as the fluorescence in the region within 30 μm from the ROI border (excluding other ROIs) for frame n, $F_{neuropil}$ is the average neuropil fluorescence across the entire time series, and α was chosen from [0 1] for each neuron to minimize the Pearson's correlation coefficient between $F_{corrected}$ and $F_{neuropil}$. The $\Delta F/F(\%)$ for each neuron was then calculated as:

$$\Delta F/F(\%) = 100 * (F_n - F_0)/F_0$$

where F_n is the corrected fluorescence ($F_{corrected}$) for frame n, and F_0 is defined as the mode of the corrected fluorescence density distribution across the entire time series.

For some analyses (**Fig. S18**), we estimated neuronal spikes from calcium signals using a MATLAB implementation of a sparse, nonnegative convolution algorithm (OASIS) on $\Delta F/F$ traces with an autoregressive model of order 1 for the convolution kernel[60].

Analysis of two-photon imaging data

To determine whether a neuron exhibited consistent task-related activity, we calculated the Pearson's correlation coefficient between odd and even correct trials and compared it to the correlation coefficient of trials in which activity was circularly permuted at a random point on each trial (time shuffled). If the original correlation coefficient was greater than the 95th percentile of the correlation coefficients from the time shuffled activity (100 iterations), then a cell was considered to exhibit reliable task-related activity for that specific day. Only cells with reliable activity in day 1 and in at least one of days 2 to 5 were used for specific day-to-day comparisons. Furthermore, we evaluated whether each cell ROI was morphologically stable across experimental sessions (**Fig S3A**). To achieve this, we centered each ROI within a 30 x 30 μm window, and calculated the structural similarity index (SSI; Matlab `ssim` function) between the windows in day 1 and days 2 to 5. We then compared each SSI to a null distribution of SSIs between the corresponding window in day 1 and random windows in other days. If the true SSI value was higher than the 95th CI of the SSI null distribution, then an ROI was considered stable for that specific day-to-day comparison (**Fig. S3B**). We also checked

using the 2-D correlation coefficient as a metric of similarity between windows containing ROIs, obtaining similar results (**Fig. S3C**). In addition to activity reliability and ROI morphological stability, each day-to-day comparison had to meet behavioral criteria: performance above chance level (>0.5), and ≥2 trials of each trial type (correct-left, incorrect-left, incorrect-right, correct-right; **Fig. S1**).

Trained mice typically make more correct than incorrect behavioral responses. Thus, in order to achieve an unbiased comparison of the neuronal activity in correct and incorrect trials, we resampled trials without replacement (100 iterations) to the lowest number of trials in each context for each neuron. In most cases, this resulted in neural responses resampled to the number of incorrect trials, although this was not always the case (see **Fig. S1A**).

Preference for a task variable was determined by resampling trials with replacement (100 iterations) and grouping accordingly. For instance, to evaluate context preference, correct left and incorrect right trials (context 1) are resampled and grouped together, and compared to resampled incorrect left and correct right trials (context 2). Then, if activity is significantly different (Wilcoxon signed-rank test) for 5 consecutive bins between the two populations of resampled trials, a cell is considered to exhibit preference for the trials with higher levels of activity. In addition, to evaluate differences in task variable preference, the subpopulations of cells with preference for a particular task variable were resampled for each day (100 iterations). Preference was considered different when the 95% CIs of their bootstrapped distribution did not overlap (**Fig. 3B**).

To measure the stability of responses in individual cells, we calculated the Pearson's correlation coefficient (CC) between odd trials in day 1 and even trials in days 1 to 5. This was done either directly, for each trial type separately (correct left, incorrect right, incorrect left, and correct right; **Fig. 3D**), or by grouping activity by their preferred task variable (context 1/context 2, left/right choice, or correct/incorrect outcome; **Fig. 3C**). Next, to quantify the stability of the CC across days, we fit exponential decay functions to each individual cell:

$$CC = CC_{day1} * e^{\frac{1-day}{\tau}}$$

to obtain their fit parameters CC_{day1} and τ. However, since fit parameters were not normally distributed, we first calculated their rankit:

$$\text{fit parameter} = invnormal\left(\frac{rank\ of\ \text{fit parameter} - 0.5}{n}\right)$$

and then used LME models to determine significant differences among groups.

The coding performance of task variables was determined by support vector machine classifiers[24]. To achieve this, trials were homogenized by sampling an equal number of each trial type with replacement (correct left, incorrect right, incorrect left, and correct right), and then splitting them into two nonoverlapping groups containing 50% trials of each type. This process ensures that all conditions comprise the same number of trials. For analyses within a single day, a support vector machine classifier was trained using 50% of trials and tested on the remaining 50% of trials per time bin. For analyses across days, and to maintain similar sampling, 50% of randomly sampled trials were used for training the classifiers in reference day 1, and 50% of randomly sampled trials were used for testing these classifiers in days 2 to 5. The performance confidence intervals (CI) were obtained by calculating the 5th and 95th percentiles of the bootstrapped data (100 iterations). The encoding of the different variables was determined by comparing conditions as follows: (1) context, correct left and incorrect right versus incorrect left and correct right; (2) motor, correct left and incorrect left versus incorrect right and correct right; and (3) outcome, correct left and correct right versus incorrect right and incorrect left. A cell was considered to encode any of the task variables if the performance of the decoder was higher than chance for at least 5 consecutive time bins (lower 5th CI of the bootstrapped distribution > 0.5; see **Fig. S8**).

To measure the stability of decoding performance (P) in individual cells (**Fig. 4C, Fig. 5C, Fig. S12C**), we fit exponential decay functions to each cell:

$$P = P_{day1} * e^{\frac{1-day}{\tau}} + 0.5$$

This allowed us to compare the fit parameters P_{day1} and τ among groups using LME models on their corresponding rankit values. The constant 0.5 is added to account for the baseline where decoding performance reaches chance level. In **Fig. 4C**, the top left inset for P_{day1} takes this 0.5 into account.

For some analyses of decoding stability, we integrated between chance level (0.5) and performance level for time windows relevant for each task variable (1.5 s), and normalized by the integrated performance in day 1 (when the classifier is trained). We then fit an exponential decay function to this normalized performance (NP):

$$NP = e^{\frac{1-day}{\tau}}$$

to estimate the time at which decoding over chance level decays to 1/e (~37%) of its original value in day 1. Note that NP values = 1 indicate a decoding performance equal to that in day 1, and that NP values = 0 signify decoding performance equal to chance level. Thus, NP is a metric that allows us to evaluate the stability in the decoding of task variables across days. In addition, to obtain CIs for the decay constants, we also fit decay functions to each iteration of the classifier performance. This exponential decay function generally fit the decay in decoding performance well (**Fig. 6C, Fig. 7C, Fig. S14, Fig. S16C, Fig. S17C, Fig. S18D**), particularly when enough data was available to calculate the decoding performance of task variables. However, for one analysis, we estimated decay constants for each individual experiment or each individual mouse (**Fig. 6D, S14**), which, in some cases, resulted in poor fits. Thus, we only considered fits with low sum of squares (SSE) values (≤ 1.6, preserving 87.9% of fits; based on the distribution of SSE values), and positive R-squared values (fits are better than a horizontal line through the mean value). Then, the rankit of their corresponding decay constants were compared using LME models. In addition, we used a slightly different model for the fits on not normalized data (**Fig. S19**), adding one term for the baseline = 0.5, and one term for the performance in day 1:

$$P = P_{day1} * e^{\frac{1-day}{\tau}} + 0.5$$

In **Fig. S19**, the value for P_{day1} takes this 0.5 into account.

Statistical information

When possible, data groups are compared using linear mixed-effects (LME) models using task variables or trial type for the fixed effect term, and the experiment number and the mouse number as nominal values for the random-effects terms. All three terms are modelled as predictor variables of the particular data being compared (e.g., correlation coefficient, classifier performance). We fit the parameters of mixed-effect models using the restricted maximum likelihood (REML) estimation, which estimates the variance parameters independently of the estimates for the fixed effects. When independent sampling was not possible (e.g., population vector correlation coefficients, population decoder performance), data were compared using bootstrap estimates via sampling with replacement (100 iterations), and differences were

determined by non-overlap of the 95% CIs. Bootstrap estimates of s.e.m. were calculated as the s.d. of values evaluated in shuffled iterations.

DATA AVAILABILITY

All processed imaging data used for making Figures 1-7 is available in the following repository: (*url to be posted upon acceptance*)**.**

ACKNOWLEDGEMENTS

We would like to thank Ralf Wessel, Tina Xia, and Jhoseph Shin for their comments on the manuscript and code review.

Funding: This work was supported by the following: Harvey Karp Discovery Award (L.M.F.), UC MEXUS-CONACYT Postdoctoral Fellowship (L.M.F.), NIH R01 RF1NS121919 (M.J.G.), NSF 1707287 (M.J.G.), Larry L. Hillblom Foundation (M.J.G.), and Whitehall Foundation (M.J.G.).

AUTHOR CONTRIBUTIONS

L.M.F. and M.J.G. designed the experiments; L.M.F. performed the surgical implants; L.M.F. conducted the experiments and analyzed the data; L.M.F. and M.J.G. wrote the manuscript.

COMPETING INTERESTS

The authors declare no competing financial interests.

REFERENCES

1. Rule, M. E., O'Leary, T. & Harvey, C. D. Causes and consequences of representational drift. *Curr Opin Neurobiol* **58**, 141–147 (2019).
2. Mau, W., Hasselmo, M. E. & Cai, D. J. The brain in motion: How ensemble fluidity drives memory-updating and flexibility. *Elife* **9**, 1–24 (2020).
3. Raman, D. V. & O'leary, T. Optimal plasticity for memory maintenance during ongoing synaptic change. *Elife* **10**, (2021).

4. Clopath, C., Bonhoeffer, T., Hübener, M. & Rose, T. Variance and invariance of neuronal long-term representations. *Philosophical Transactions of the Royal Society B: Biological Sciences* **372**, (2017).
5. Jeon, B. B., Swain, A. D., Good, J. T., Chase, S. M. & Kuhlman, S. J. Feature selectivity is stable in primary visual cortex across a range of spatial frequencies. *Sci Rep* **8**, (2018).
6. Marks, T. D. & Goard, M. J. Stimulus-dependent representational drift in primary visual cortex. *Nature Communications 2021 12:1* **12**, 1–16 (2021).
7. Pérez-Ortega, J., Alejandre-García, T. & Yuste, R. Long-term stability of cortical ensembles. *Elife* **10**, (2021).
8. Montijn, J. S., Meijer, G. T., Lansink, C. S. & Pennartz, C. M. A. Population-Level Neural Codes Are Robust to Single-Neuron Variability from a Multidimensional Coding Perspective. *Cell Rep* **16**, 2486–2498 (2016).
9. Huber, D. *et al.* Multiple dynamic representations in the motor cortex during sensorimotor learning. *Nature* **484**, 473–478 (2012).
10. Jensen, K. T., Kadmon Harpaz, N., Dhawale, A. K., Wolff, S. B. E. & Ölveczky, B. P. Long-term stability of single neuron activity in the motor system. *Nature Neuroscience 2022 25:12* **25**, 1664–1674 (2022).
11. Mayrhofer, J. M., Haiss, F., Helmchen, F. & Weber, B. Sparse, reliable, and long-term stable representation of periodic whisker deflections in the mouse barrel cortex. *Neuroimage* **115**, 52–63 (2015).
12. Wang, H. C., LeMessurier, A. M. & Feldman, D. E. Tuning instability of non-columnar neurons in the salt-and-pepper whisker map in somatosensory cortex. *Nature Communications 2022 13:1* **13**, 1–18 (2022).
13. Schoonover, C. E., Ohashi, S. N., Axel, R. & Fink, A. J. P. Representational drift in primary olfactory cortex. *Nature 2021 594:7864* **594**, 541–546 (2021).
14. Betzel, R. F., Wood, K. C., Angeloni, C., Geffen, M. N. & Bassett, D. S. Stability of spontaneous, correlated activity in mouse auditory cortex. *PLoS Comput Biol* **15**, e1007360 (2019).
15. Gallego, J. A., Perich, M. G., Chowdhury, R. H., Solla, S. A. & Miller, L. E. Long-term stability of cortical population dynamics underlying consistent behavior. *Nature Neuroscience 2020 23:2* **23**, 260–270 (2020).
16. Xia, J., Marks, T. D., Goard, M. J. & Wessel, R. Stable representation of a naturalistic movie emerges from episodic activity with gain variability. *Nature Communications 2021 12:1* **12**, 1–15 (2021).
17. Driscoll, L. N., Pettit, N. L., Minderer, M., Chettih, S. N. & Harvey, C. D. Dynamic Reorganization of Neuronal Activity Patterns in Parietal Cortex. *Cell* **170**, 986-999.e16 (2017).
18. Ziv, Y. *et al.* Long-term dynamics of CA1 hippocampal place codes. *Nat Neurosci* **16**, 264–266 (2013).
19. Hainmueller, T. & Bartos, M. Parallel emergence of stable and dynamic memory engrams in the hippocampus. *Nature* **558**, 292–296 (2018).
20. Kim, S., Jung, D. & Royer, S. Place cell maps slowly develop via competitive learning and conjunctive coding in the dentate gyrus. *Nature Communications 2020 11:1* **11**, 1–15 (2020).

21. Liberti, W. A., Schmid, T. A., Forli, A., Snyder, M. & Yartsev, M. M. A stable hippocampal code in freely flying bats. *Nature 2022 604:7904* **604**, 98–103 (2022).
22. Shin, J., Lee, H. W., Jin, S. W. & Lee, I. Subtle visual change in a virtual environment induces heterogeneous remapping systematically in CA1, but not CA3. *Cell Rep* **41**, 111823 (2022).
23. Alexander, A. S. & Nitz, D. A. Retrosplenial cortex maps the conjunction of internal and external spaces. *Nat Neurosci* **18**, 1143–1151 (2015).
24. Franco, L. M. & Goard, M. J. A distributed circuit for associating environmental context with motor choice in retrosplenial cortex. *Sci Adv* **7**, eabf9815 (2021).
25. Sun, W. *et al.* Context value updating and multidimensional neuronal encoding in the retrosplenial cortex. *Nature Communications 2021 12:1* **12**, 1–17 (2021).
26. Tseng, S. Y., Chettih, S. N., Arlt, C., Barroso-Luque, R. & Harvey, C. D. Shared and specialized coding across posterior cortical areas for dynamic navigation decisions. *Neuron* **110**, 2484-2502.e16 (2022).
27. Powell, A. *et al.* Stable Encoding of Visual Cues in the Mouse Retrosplenial Cortex. *Cerebral Cortex* **30**, 4424–4437 (2020).
28. Deitch, D., Rubin, A. & Ziv, Y. Representational drift in the mouse visual cortex. *Current Biology* **31**, 4327-4339.e6 (2021).
29. Sheintuch, L. *et al.* Tracking the Same Neurons across Multiple Days in Ca2+ Imaging Data. *Cell Rep* **21**, 1102–1115 (2017).
30. Chen, T. W., Li, N., Daie, K. & Svoboda, K. A Map of Anticipatory Activity in Mouse Motor Cortex. *Neuron* **94**, 866-879.e4 (2017).
31. Oh, S. W. *et al.* A mesoscale connectome of the mouse brain. *Nature 2014 508:7495* **508**, 207–214 (2014).
32. Jensen, K. T., Harpaz, N. K., Dhawale, A. K., Wolff, S. B. E. & Ölveczky, B. P. Long-term stability of single neuron activity in the motor system. *bioRxiv* 2021.10.27.465945 (2022) doi:10.1101/2021.10.27.465945.
33. Frankland, P. W. & Bontempi, B. The organization of recent and remote memories. *Nat Rev Neurosci* **6**, 119–130 (2005).
34. Ranganath, C. & Ritchey, M. Two cortical systems for memory-guided behaviour. *Nature Reviews Neuroscience 2012 13:10* **13**, 713–726 (2012).
35. Tonegawa, S., Morrissey, M. D. & Kitamura, T. The role of engram cells in the systems consolidation of memory. *Nature Reviews Neuroscience 2018 19:8* **19**, 485–498 (2018).
36. Kira, S., Safaai, H., Morcos, A. S., Panzeri, S. & Harvey, C. D. A distributed and efficient population code of mixed selectivity neurons for flexible navigation decisions. *bioRxiv* 2022.04.10.487349 (2022) doi:10.1101/2022.04.10.487349.
37. Singer, A. C. & Frank, L. M. Rewarded Outcomes Enhance Reactivation of Experience in the Hippocampus. *Neuron* **64**, 910–921 (2009).
38. Patil, A., Murty, V. P., Dunsmoor, J. E., Phelps, E. A. & Davachi, L. Reward retroactively enhances memory consolidation for related items. *Learning & Memory* **24**, 65–69 (2017).
39. Goard, M. & Dan, Y. Basal forebrain activation enhances cortical coding of natural scenes. *Nat Neurosci* **12**, 1444–1449 (2009).
40. Pinto, L. *et al.* Fast modulation of visual perception by basal forebrain cholinergic neurons. *Nat Neurosci* **16**, 1857–1863 (2013).

41. McGinley, M. J., David, S. V. & McCormick, D. A. Cortical Membrane Potential Signature of Optimal States for Sensory Signal Detection. *Neuron* **87**, 179–192 (2015).
42. McGinley, M. J. J. *et al.* Waking State: Rapid Variations Modulate Neural and Behavioral Responses. *Neuron* **87**, 1143–1161 (2015).
43. Kuchibhotla, K. V. *et al.* Parallel processing by cortical inhibition enables context-dependent behavior. *Nat Neurosci* **20**, 62–71 (2016).
44. Musall, S., Kaufman, M. T., Juavinett, A. L., Gluf, S. & Churchland, A. K. Single-trial neural dynamics are dominated by richly varied movements. *Nat Neurosci* **22**, 1677–1686 (2019).
45. Stringer, C. *et al.* Spontaneous behaviors drive multidimensional, brainwide activity. *Science (1979)* **364**, (2019).
46. Morcos, A. S. & Harvey, C. D. History-dependent variability in population dynamics during evidence accumulation in cortex. *Nat Neurosci* **19**, 1672–1681 (2016).
47. Hwang, E. J., Dahlen, J. E., Mukundan, M. & Komiyama, T. History-based action selection bias in posterior parietal cortex. *Nat Commun* **8**, 1242 (2017).
48. Rule, M. E. *et al.* Stable task information from an unstable neural population. *Elife* **9**, 1–16 (2020).
49. Balleine, B. W., Delgado, M. R. & Hikosaka, O. The Role of the Dorsal Striatum in Reward and Decision-Making. *Journal of Neuroscience* **27**, 8161–8165 (2007).
50. Mao, D., Kandler, S., McNaughton, B. L. & Bonin, V. Sparse orthogonal population representation of spatial context in the retrosplenial cortex. *Nat Commun* **8**, 243 (2017).
51. Jacob, P.-Y. *et al.* An independent, landmark-dominated head-direction signal in dysgranular retrosplenial cortex. *Nat Neurosci* **20**, 173–175 (2017).
52. Miller, A. M. P., Mau, W. & Smith, D. M. Retrosplenial Cortical Representations of Space and Future Goal Locations Develop with Learning. *Curr Biol* **29**, 2083-2090.e4 (2019).
53. Fischer, L. F., Mojica Soto-Albors, R., Buck, F. & Harnett, M. T. Representation of visual landmarks in retrosplenial cortex. *Elife* **9**, (2020).
54. Vann, S. D., Aggleton, J. P. & Maguire, E. A. What does the retrosplenial cortex do? *Nat Rev Neurosci* **10**, 792–802 (2009).
55. Miller, A. M. P., Vedder, L. C., Law, L. M. & Smith, D. M. Cues, context, and long-term memory: the role of the retrosplenial cortex in spatial cognition. *Front Hum Neurosci* **8**, 586 (2014).
56. Vedder, L. C., Miller, A. M. P., Harrison, M. B. & Smith, D. M. Retrosplenial Cortical Neurons Encode Navigational Cues, Trajectories and Reward Locations During Goal Directed Navigation. *Cerebral Cortex* **27**, 3713–3723 (2017).
57. Hennestad, E., Witoelar, A., Chambers, A. R. & Vervaeke, K. Mapping vestibular and visual contributions to angular head velocity tuning in the cortex. *Cell Rep* **37**, 110134 (2021).
58. Brennan, E. K. W. *et al.* Thalamus and claustrum control parallel layer 1 circuits in retrosplenial cortex. *Elife* **10**, (2021).
59. Aronov, D. & Tank, D. W. Engagement of Neural Circuits Underlying 2D Spatial Navigation in a Rodent Virtual Reality System. *Neuron* **84**, 442–456 (2014).
60. Friedrich, J., Zhou, P. & Paninski, L. Fast online deconvolution of calcium imaging data. *PLoS Comput Biol* **13**, e1005423 (2017).

SUPPLEMENTAL INFORMATION

Fig. S1. Related to Fig. 1.

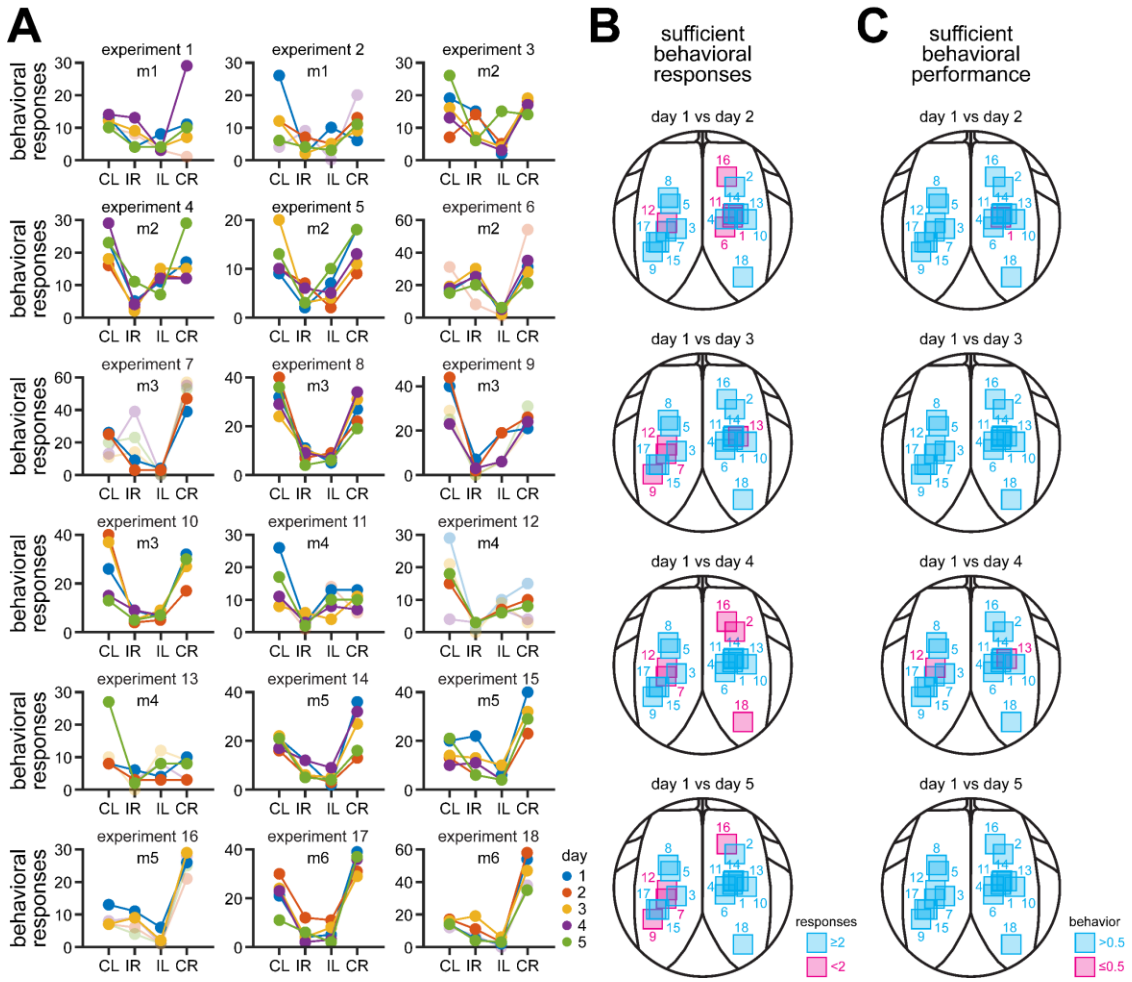

Fig. S1. Behavioral responses and location of imaging fields within RSC. Related to Fig. 1.

A. Combinations of correct (C) / incorrect (I) and left (L) / right (R) responses in both contexts across 5 days. Each day is shown in a different color. Blue, day 1; red, day 2; yellow, day 3; purple, day 4; green, day 5. Solid lines show experiments considered for further analyses, whereas faded lines indicate experiments without sufficient behavioral responses or behavioral performance, or both (**B** and **C**).

B-C. Location of imaging fields for all experiments. Fields are color-coded to show experiments included (cyan) and not included (magenta) for analyses according to: (**B**) enough behavioral responses (minimum ≥2 trials of each trial type, see **A**), or (**C**) behavioral performance above chance (0.5).

Fig. S2. Related to Fig. 1.

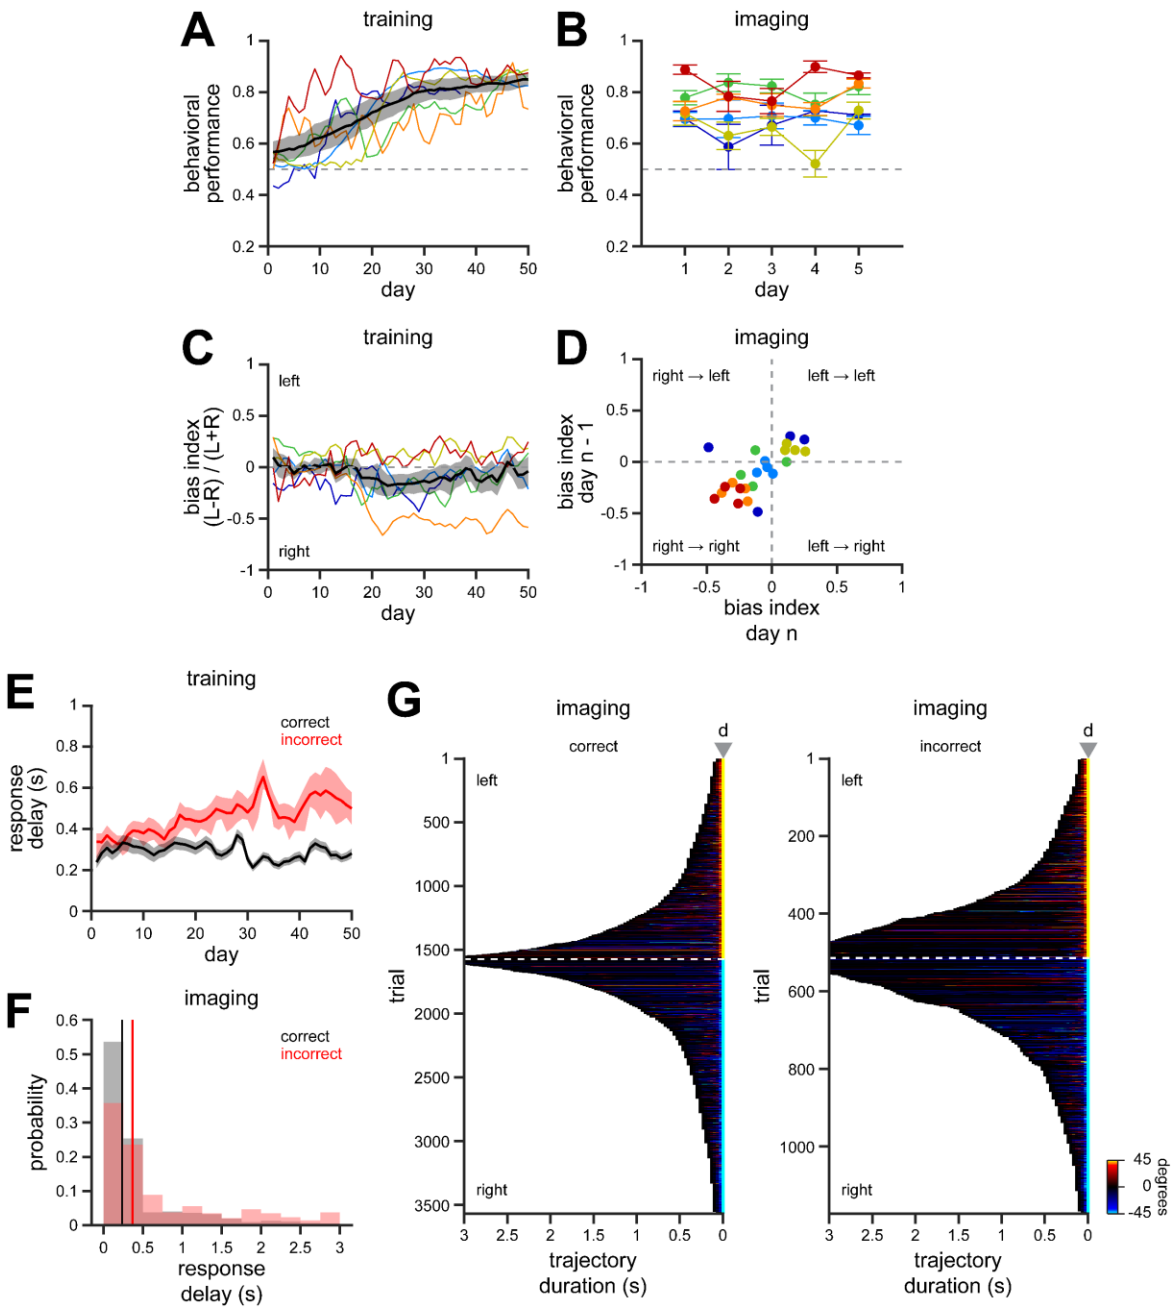

Fig. S2. Behavioral parameters during training and imaging sessions. Related to Fig. 1.

A. Behavioral performance during training (n = 6 mice). Performance progressively increased across 50 training sessions until reaching a plateau. Each mouse is shown in a different color. Black solid line, mean across all mice; shaded area, s.e.m.

B. Behavioral performance across 5 consecutive days during imaging sessions. Performance is averaged per mouse (mean ± s.e.m.; n = 18 experiments in 6 mice). Each mouse is shown in a different color.

C. Bias index during training sessions (n = 6 mice). With the exception of one mouse, mice generally perform this task in an unbiased manner, with a similar number of decisions to the left or to the right. Each mouse is shown in a different color. Black solid line, mean across all mice; shaded area, s.e.m.

D. Bias index during imaging sessions. Bias index is averaged per mouse and plotted against its bias index in the previous day (n = 18 experiments in 6 mice). Each mouse is shown in a different color. Response biases during the imaging sessions were small, but consistent across successive days.

E. Decision delay in correct and incorrect trials across training. Solid line, mean; shaded area, s.e.m. Mice tended to make faster decisions in correct trials by the end of training.

F. Histogram showing the decision delay for correct and incorrect trials during imaging sessions. The median is indicated by a solid line. Mice maintained their faster decisions for correct trials during imaging sessions.

G. Joystick trajectories during decisions in correct and incorrect trials. The angle rotation is color-coded as indicated.

Fig. S3. Related to Fig. 1.

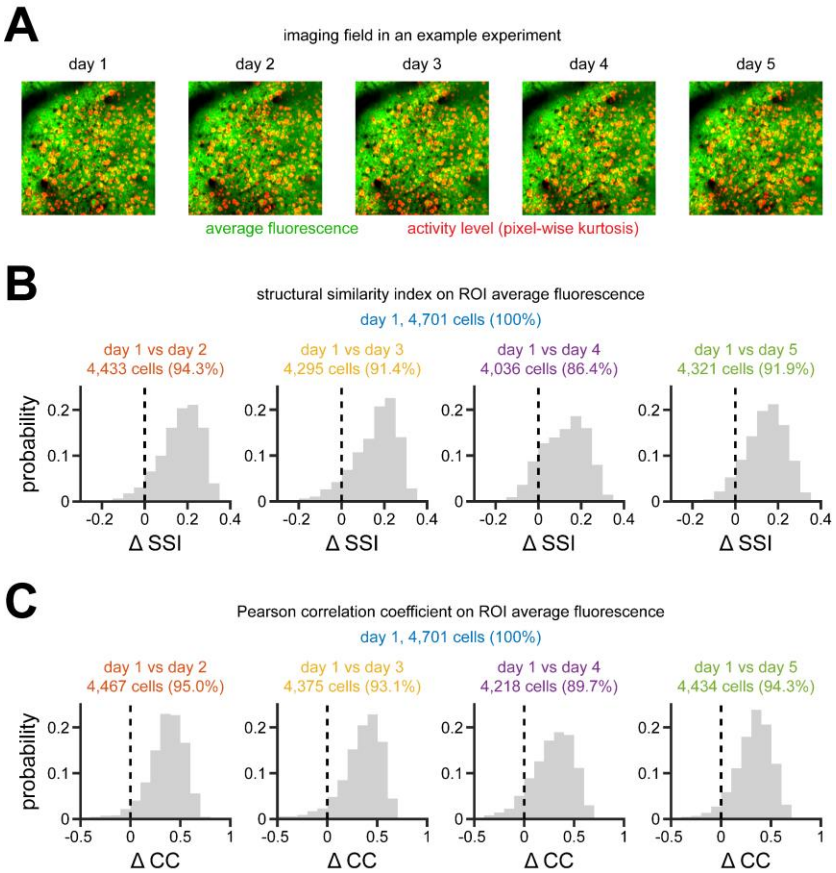

Fig. S3. ROI morphological stability across imaging sessions. Related to Fig. 1.

A. Example imaging field (425 × 425 μm) recorded across 5 days. See Figure 1F for examples of individual ROIs.

B. Difference between the structural similarity index (SSI) of each ROI and the 95th CI of their corresponding SSI null distribution (ΔSSI; see **Methods** for details). Only neurons with ΔSSI > 0 were used for subsequent analyses.

C. Similar to **B**, only in this case the 2-D correlation coefficient is used to compare the similarity between ROIs. Note that although the 2-D correlation coefficient provides similar results, it is slightly more permissive than the SSI. Thus, we used the SSI as a measure of morphological stability.

Fig. S4. Related to Fig. 2.

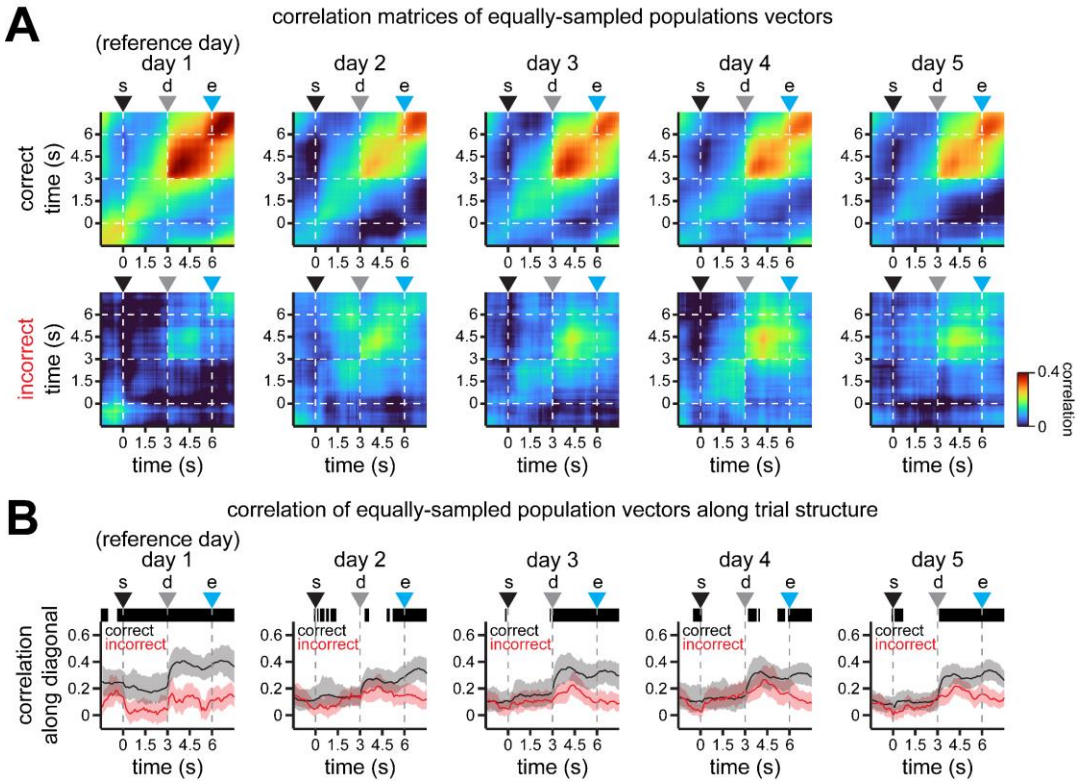

Fig. S4. Stability of population activity in equally sampled correct and incorrect trials. Related to Fig. 2.

A. Pearson's correlation coefficient matrices showing the similarity of population activity along trial duration (averaged across 100 iterations). Briefly, population vectors in one half of resampled trials in day 1 are compared to population vectors in the other half of resampled non-overlapping trials in day 1. For other days, all resampled trials in day 1 are compared to all resampled trials in days 2 to 5. In addition, trials are randomly sampled to the minimum number of trials in either correct or incorrect responses for each cell. Moreover, to compare vectors of the same length across days, the population is resampled to 1,609 cells within each iteration (the lowest number of cells that meet criteria in all 5 days). Note the higher similarity of population vectors in correct trials across days, consistent with our findings in the original dataset (**Fig. 2B-C**).

B. Population vector correlation coefficients along the diagonals in **B**. Note the higher stability and similarity of population activity in correct trials. Top insets indicate time bins (in black) where the mean correlation in correct trials > 95th CI of the correlation in incorrect trials along trial structure. Solid line, mean; shaded area, 95% CIs.

For all panels: s = trial start; d = decision point; e = trial end.

Fig. S5. Related to Fig. 2.

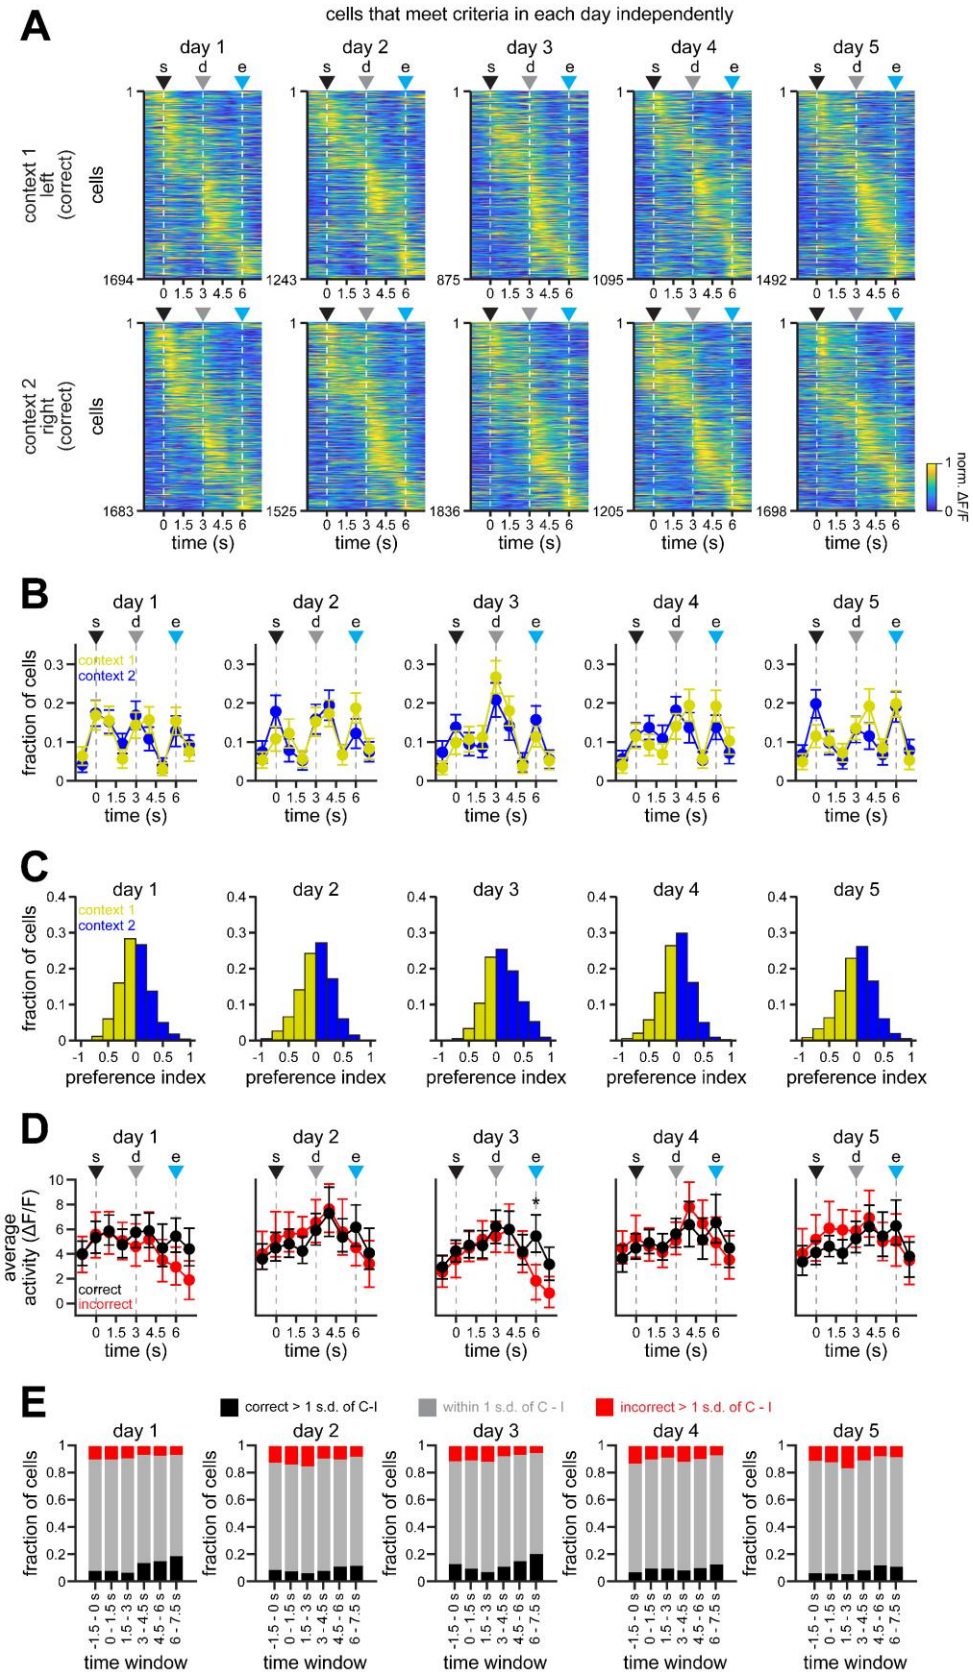

Fig. S5. Similar statistics in neural activity on independent days. Related to Fig. 2.

A. Average normalized activity in correct trials for both contexts across the 5 days of experimentation. Responses are cross-validated, sorted in odd trials, and plotted in even trials for each day independently. Note a similar mapping along trial duration by each subpopulation of cells determined in each day independently.

B. Fraction of neurons with peak activity along trial duration for context 1 (yellow) and for context 2 (blue). To estimate population error, the population was randomly sampled (n = 100 neurons per iteration), and the mean ± boostrapped s.e.m. of the response fraction was plotted.

C. Histograms of the preference index for the cells in **A**. Note a similar preference for either context in each population of cells.

D. Average activity in individual neurons for correct and incorrect responses along trial duration. To estimate population error, the population was randomly sampled (n = 100 neurons per iteration), and the mean ± boostrapped s.e.m. of the average activity was plotted.

E. Fraction of neurons with activity levels within 1 standard deviation of the difference between correct and incorrect (C - I) in gray, and higher that 1 standard deviation in correct (black) and incorrect (red) responses.

For all panels: s = trial start; d = decision point; e = trial end.

Fig. S6. Related to Fig. 3.

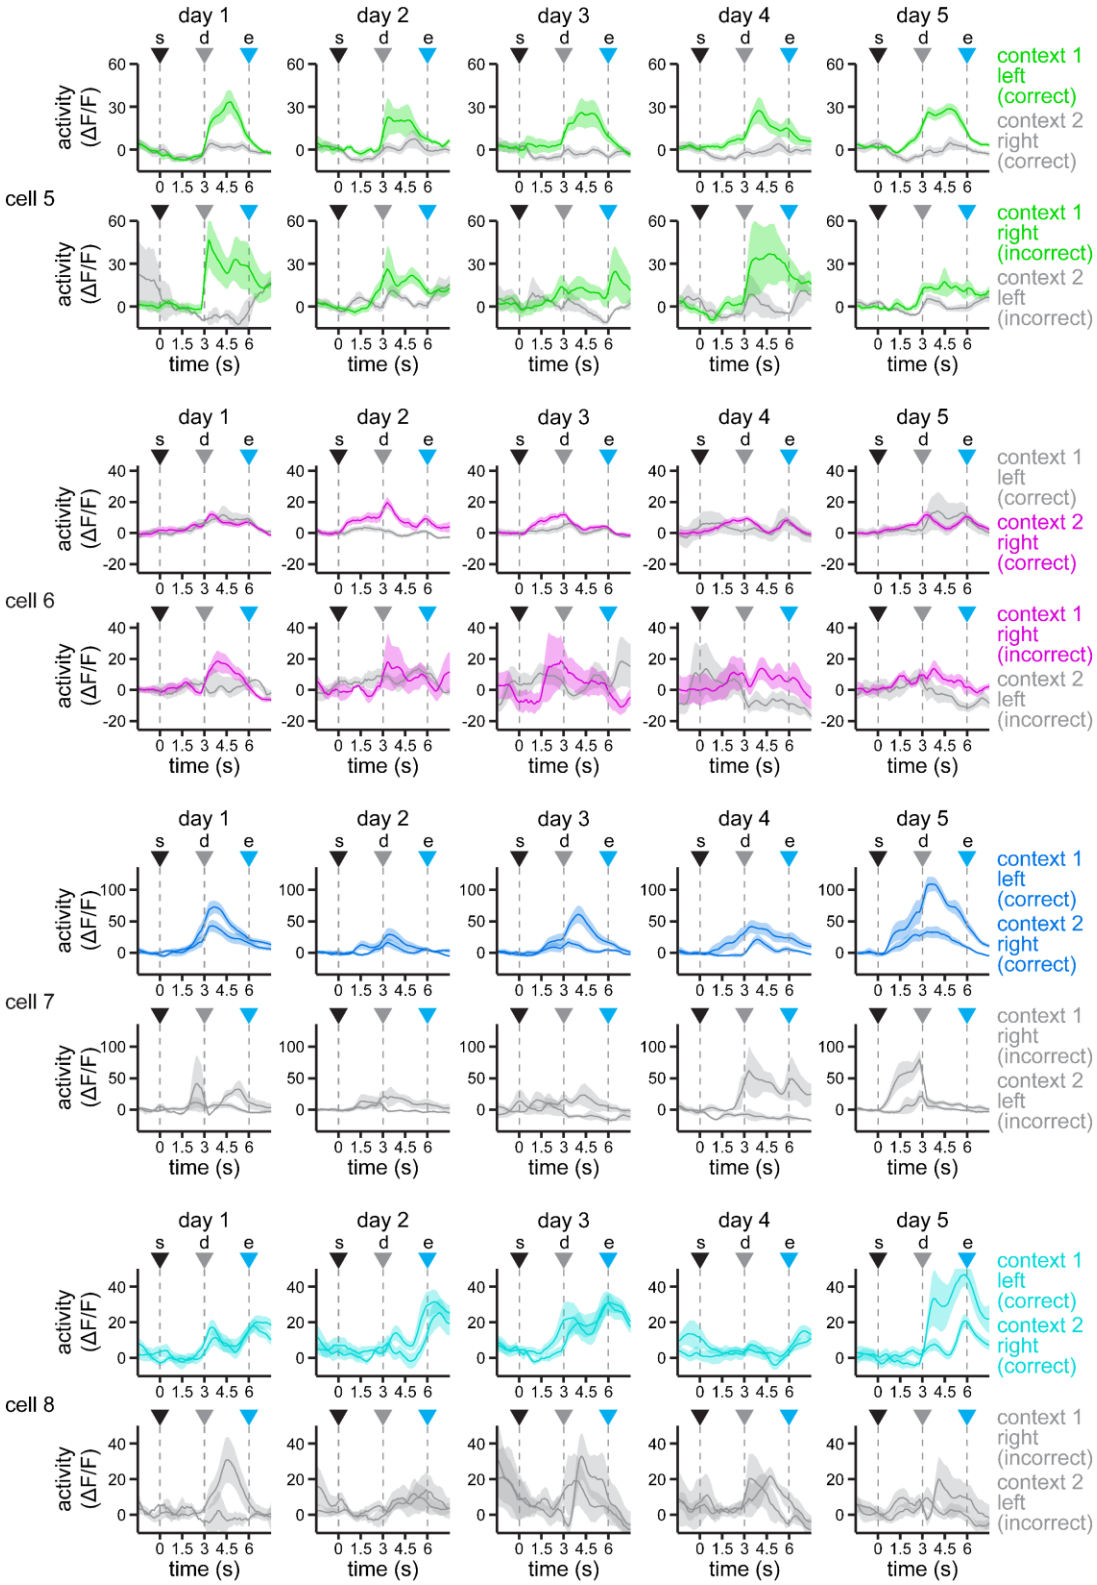

Fig. S6. Additional examples of RSC neuron responses during task performance. Related to Fig. 3.

Example neurons exhibiting higher responses for context 1 (cell 5), for right turns (cell 6), and for correct post-decision (cell 7) and post-trial outcomes (cell 8), respectively. Traces are color-coded to highlight the difference between the preferred and non-preferred responses in both correct (top row) and incorrect (bottom row) trials. Solid line, mean; shaded area, s.e.m. s = trial start; d = decision point; e = trial end.

Fig. S7. Related to Fig. 4.

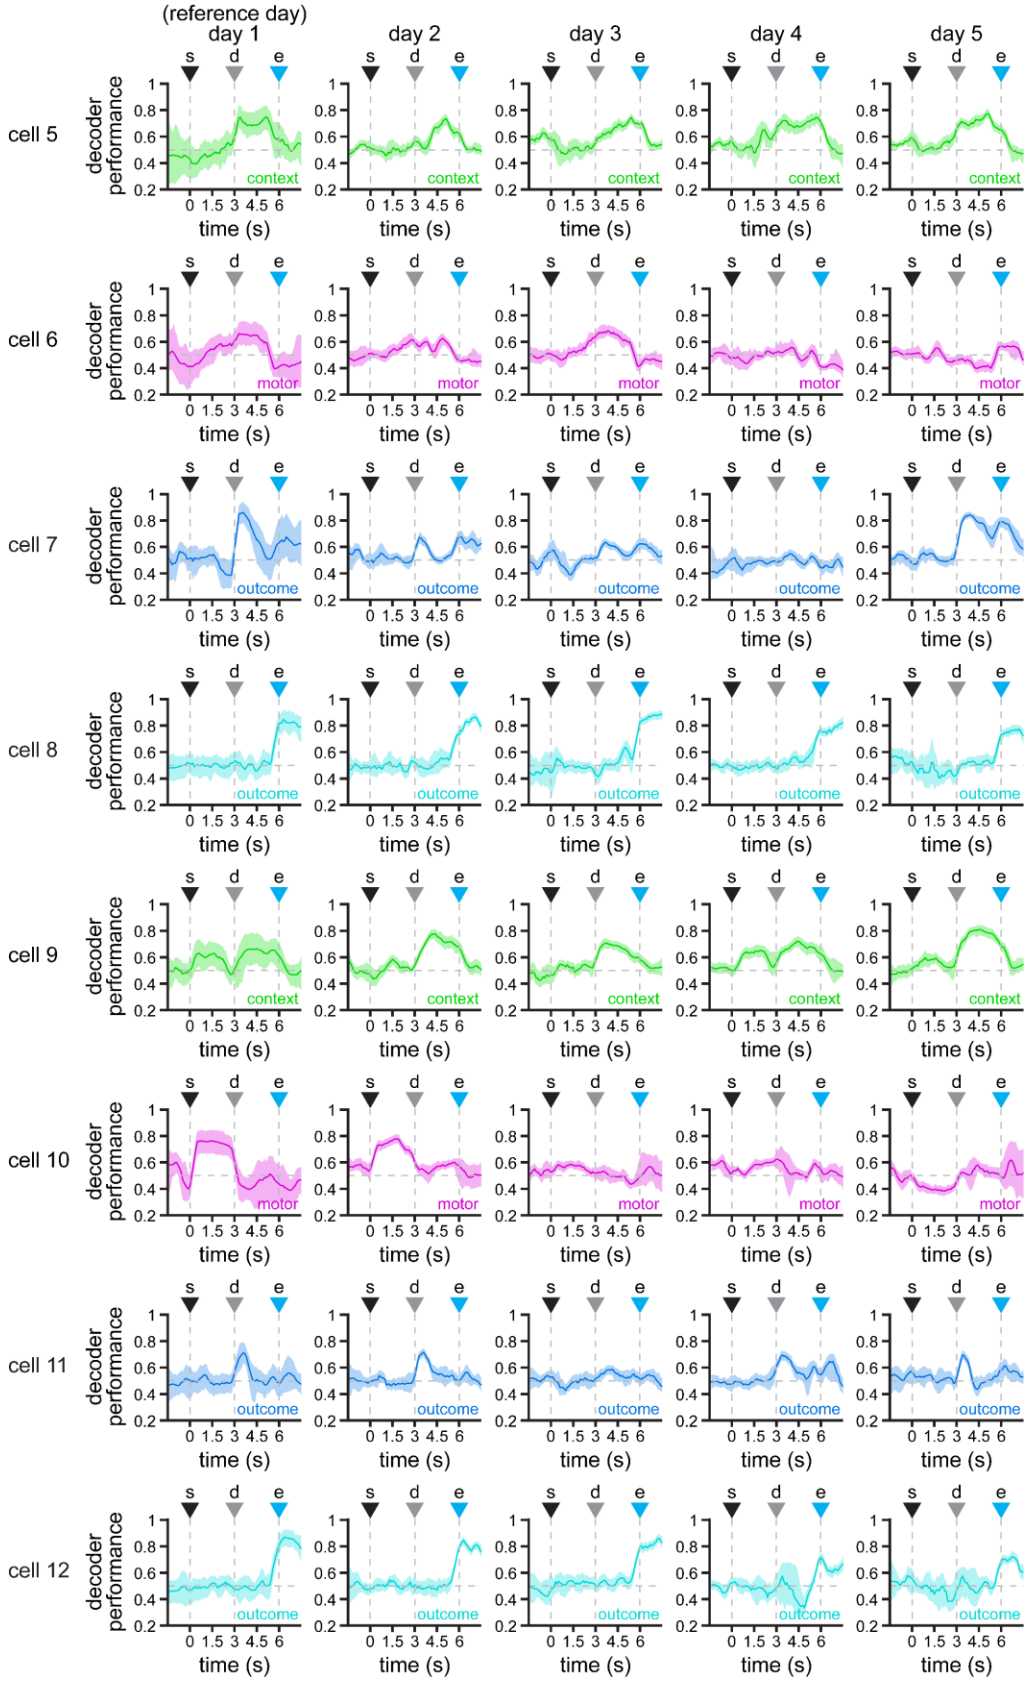

Fig. S7. Additional examples of RSC neurons with significant coding of task variables. Related to Fig. 4. Example cells with significant decoding of context (cells 5 and 9), motor choice (cells 6 and 10), post-decision outcome (cells 7 and 11), and post-trial outcome (cells 8 and 12) in day 1. Solid line, mean; shaded area, 95% CIs. s = trial start; d = decision point; e = trial end.

Fig. S8. Related to Fig. 4.

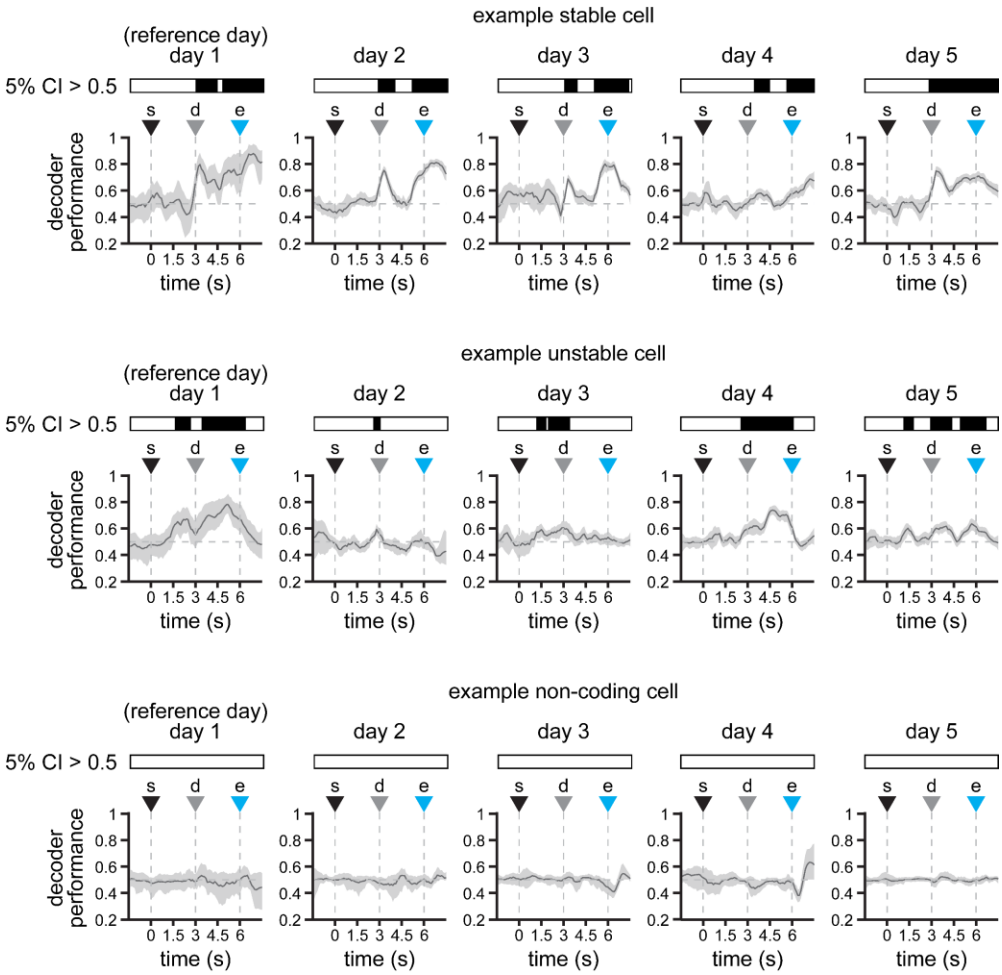

Fig. S8. Determination of cells with significant encoding of task variables. Related to Fig. 4.

The decoding performance of individual cells is calculated for each time bin by using half of the data for training a support vector machine classifier, and the other half for testing it (see **Methods** for details). We repeat this process multiple times, randomizing the trials used for model training and testing in each iteration. Since we cannot use conventional hypothesis tests with bootstrapped data, the confidence intervals (CI) of decoding performance were used to determine whether a cell exhibited consistent decoding. If the lower 5th CI of decoding performance stayed above chance level (0.5) for 5 consecutive bins, a cell was considered to significantly encode a specific task variable.

Fig S9. Related to Fig. 4.

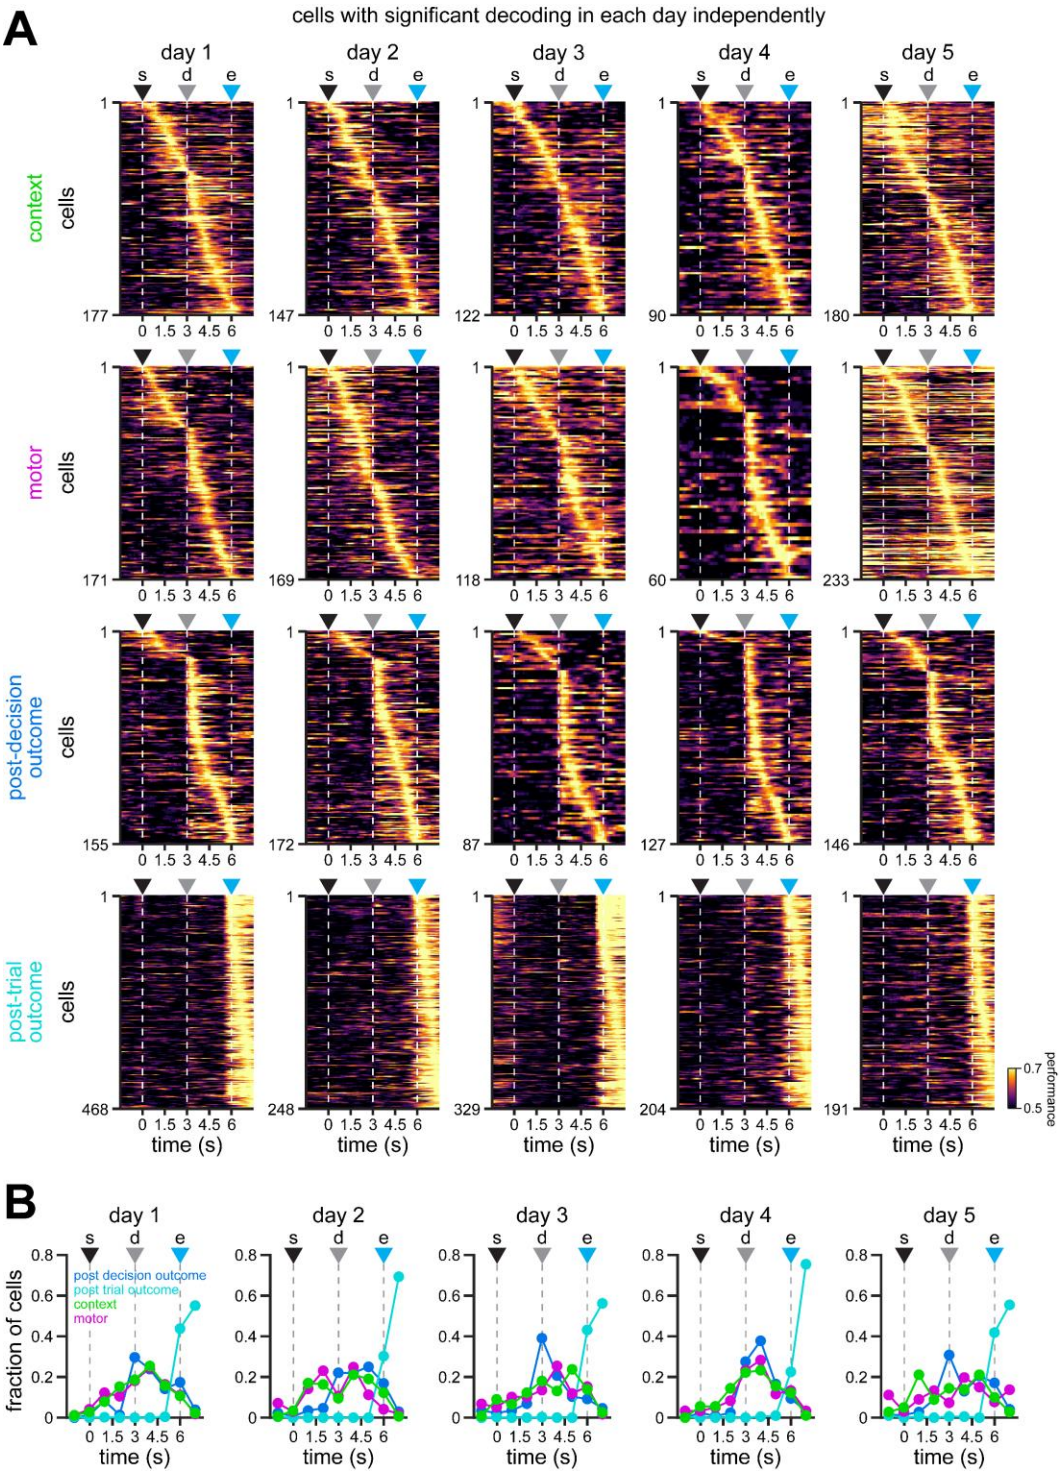

Fig. S9. Similar dynamics in the encoding of task variables by individual RSC neurons in each day independently. Related to Fig. 4.

A. Average performance of support vector machine decoders on the activity of cells exhibiting significant decoding of environmental context (top row), motor choice (top-middle row), post-decision outcome (bottom-middle row), and

post-trial outcome (bottom row). Models were trained and tested using activity within each day independently. Note similar decoding performance for each task variable in each of the 5 days of experimentation. s = trial start; d = decision point; e = trial end.

B. Fraction of cells exhibiting significant decoding of environmental context, motor choice, post-decision outcome, and post-trial outcome along trial duration. Note similar dynamics for each variable on each independent day.

Fig. S10. Related to Fig. 4

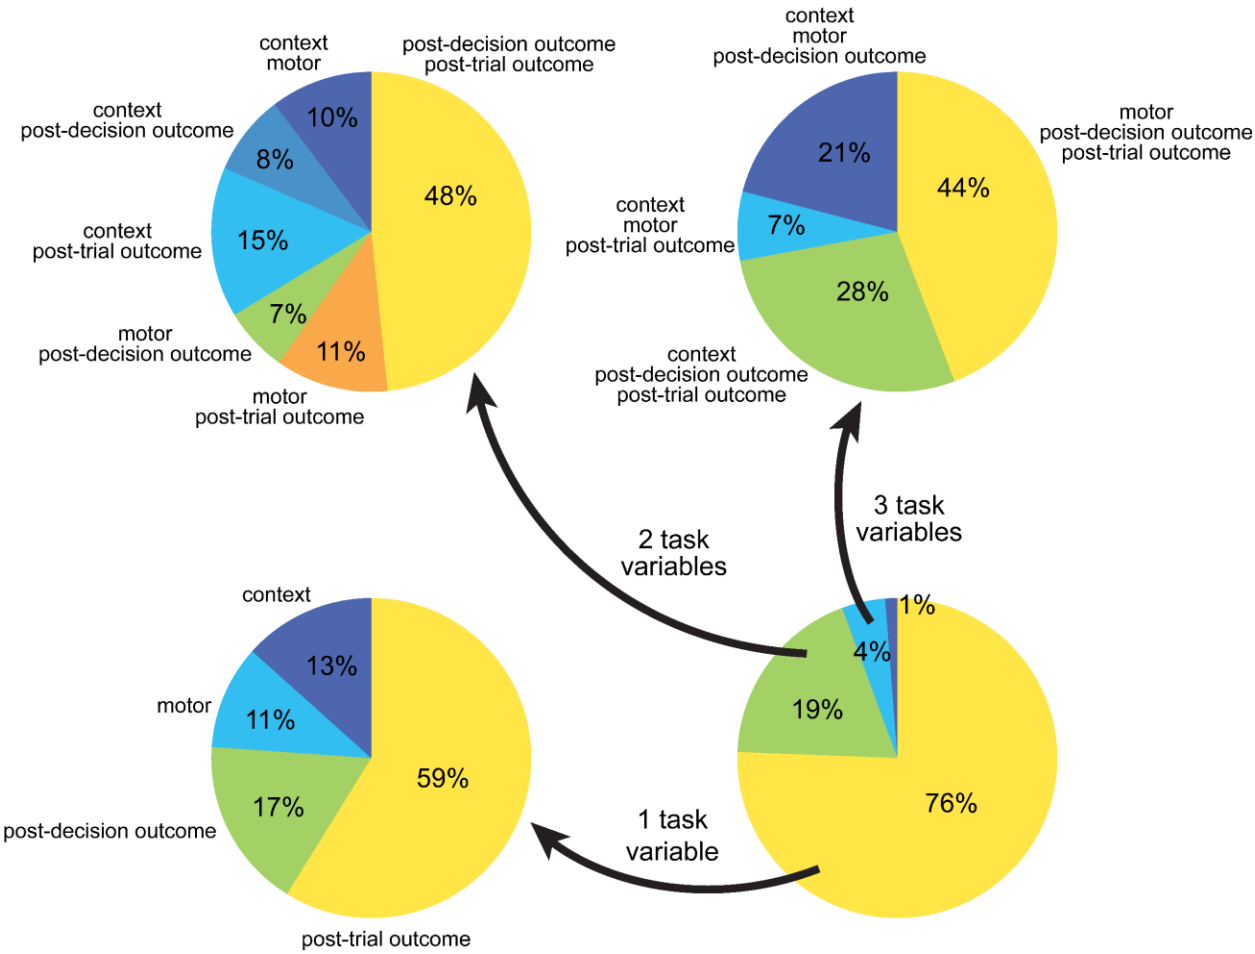

Fig. S10. Fraction of RSC neurons with significant decoding of task-variables. Related to Fig. 4. Fraction of cells with significant encoding of 1, 2, 3, or all 4 task variables (bottom right), broken down by their corresponding fractions when encoding 1 variable (bottom left), 2 variables (top left), and 3 variables (top right).

Fig. S11. Related to Fig. 5.

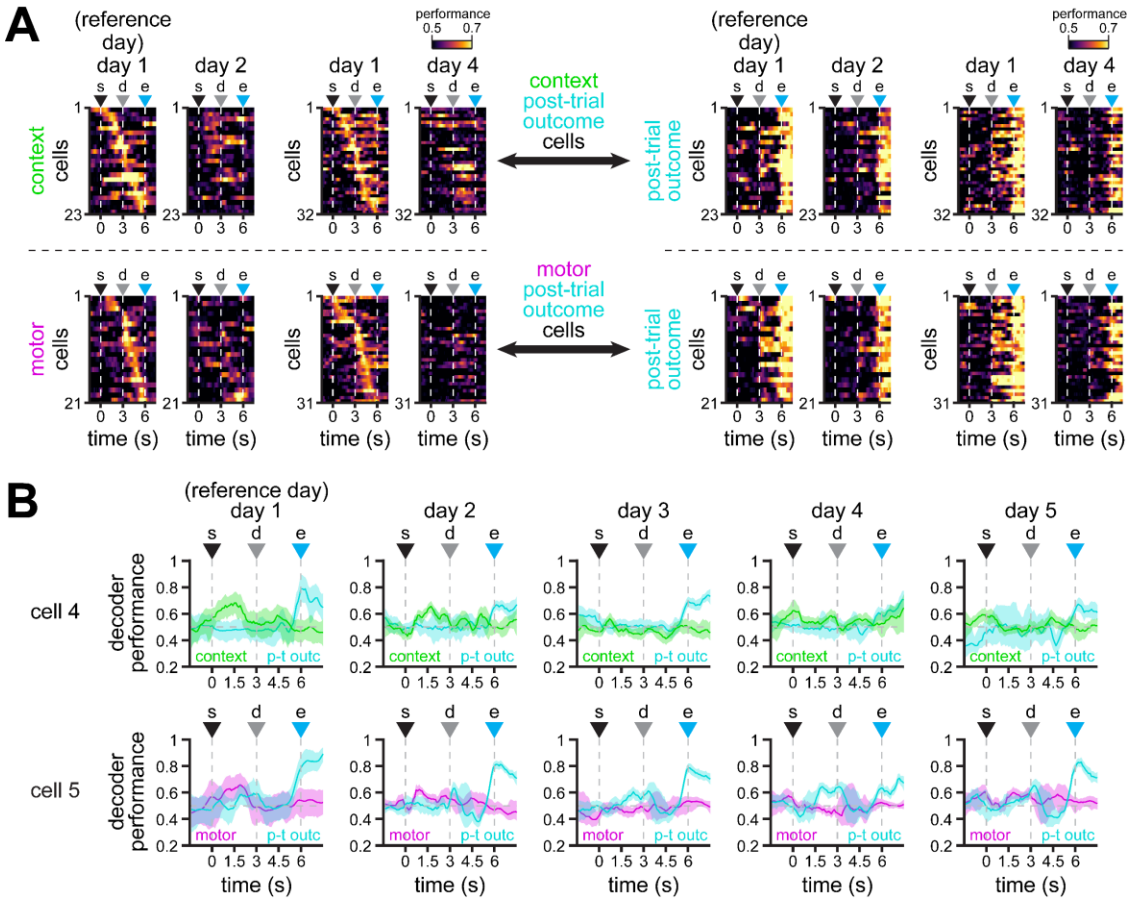

Fig. S11. Multiplexing of task information with varying stability within single RSC neurons (additional examples). Related to Fig. 5.

A. Average performance of support vector machine classifiers on the activity of single cells that significantly encode two task variables: context and post-trial outcome (top row), and motor choice and post-trial outcome (bottom row) in day 1. Note the better decoding performance of models trained in day 1 and tested in days 2 and 4 for post-trial outcome.

B. Example cells with significant encoding of context and post-trial outcome (cell 4, top), and motor choice and post-trial outcome (cell 5, bottom) in day 1. Solid line, mean; shaded area, 95% CIs.

For all panels: s = trial start; d = decision point; e = trial end.

Fig. S12. Related to Fig. 5.

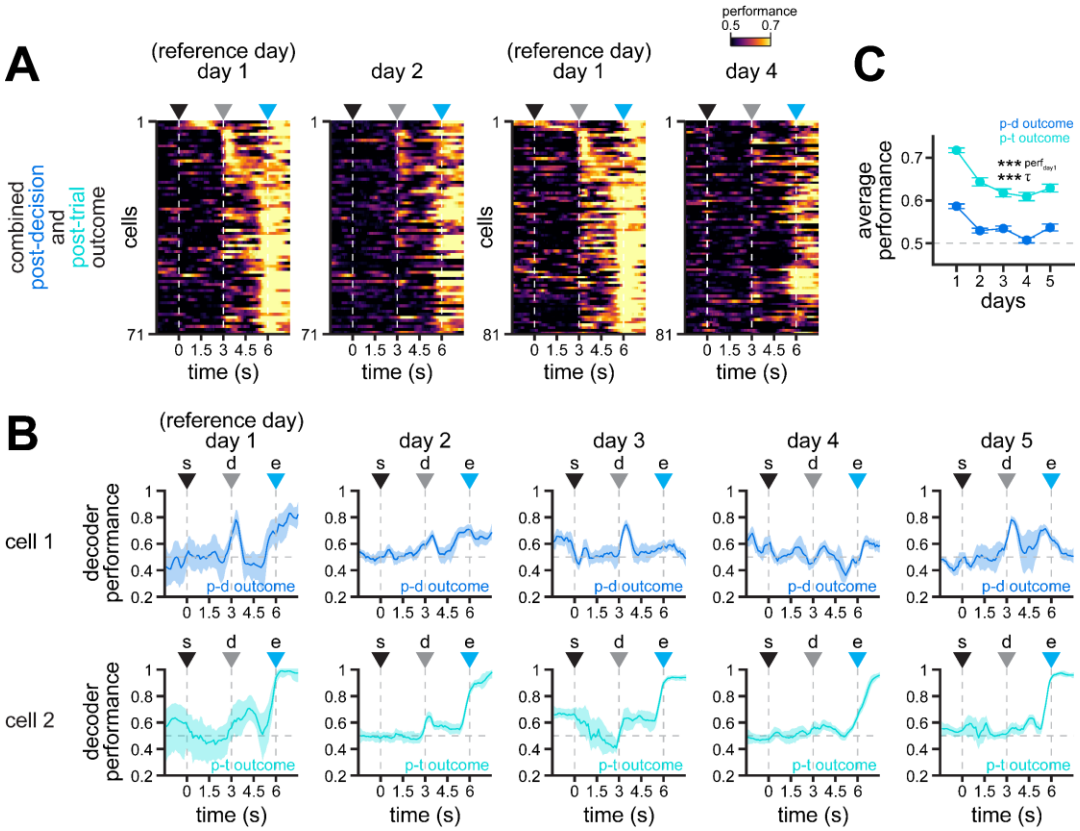

Fig. S12. Multiplexing of post-decision and post-trial outcome in individual RSC neurons. Related to Fig. 5.

A. Average performance of support vector machine classifiers on the activity of single cells that significantly encode both post-decision and post-trial outcome in day 1. Note the higher stability of post-trial outcome after trial end.

B. Examples cells with significant encoding of post-decision and post-trial outcome in day 1. Cell 1 (top) exhibits higher stability for post-decision outcome, whereas cell 2 (bottom) exhibits higher stability for post-trial outcome. Solid line, mean; shaded area, 95% CIs.

C. Average decoding performance of post-decision and post-trial outcome within individual cells (mean ± s.e.m.). Asterisks denote significantly different decoding performance, as measured by different fit parameters (*$p < 0.05$, **$p < 0.01$, ***$p < 0.001$; LME models; see **Methods** for details). Note the better decoding of the outcome variable after trial end than after the decision point.

For all panels: s = trial start; d = decision point; e = trial end.

Fig. S13. Related to Fig. 5.

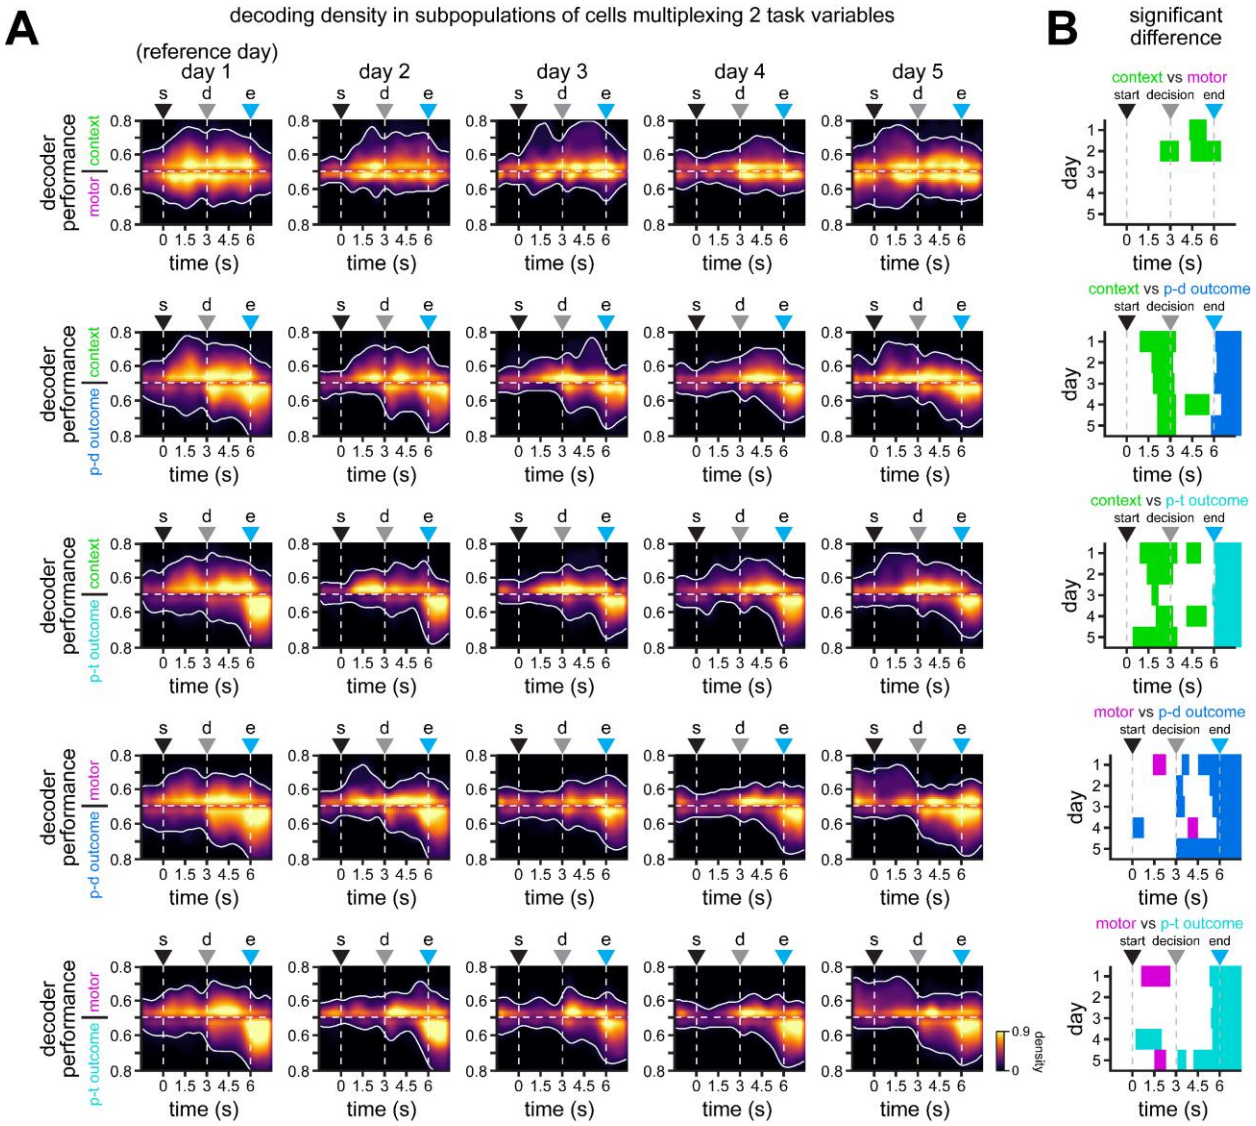

Fig. S13. Differential stability in cells with multiplexing encoding of task variables as a function of trial duration. Related to Fig. 5.

A. Heatmaps showing the density of the area under the curve of decoding performance for the subpopulations shown in **Fig. 5A** and **Fig. S11**. Task variables are plotted either on top or at the bottom of chance level (0.5), as indicated in the mirrored Y axis. Note the gradient in the stability of single cells multiplexing the indicated task variables, where context and outcome are more stable than motor choice. White lines enclose >10% density.

B. Colormaps showing significant differences between the encoding of task variables indicated in each plot for the subpopulations of cells shown in **Fig. 5A** and **Fig. S11** ($p < 0.05$ for at least 5 consecutive time bins; n = 43 to 132 [day 1], 20 to 71 [day 2], 30 to 88 [day 3], 24 to 81 [day 4], 29 to 93 [day 5] overlapping subpopulation of neurons multiplexing two task variables; LME models). Everything shown in color (not white) is significantly different. Colors indicate the task variable decoded with better performance (green = context; magenta = motor; blue = post-decision outcome, cyan = post-trial outcome).

Fig. S14. Related to Fig. 6.

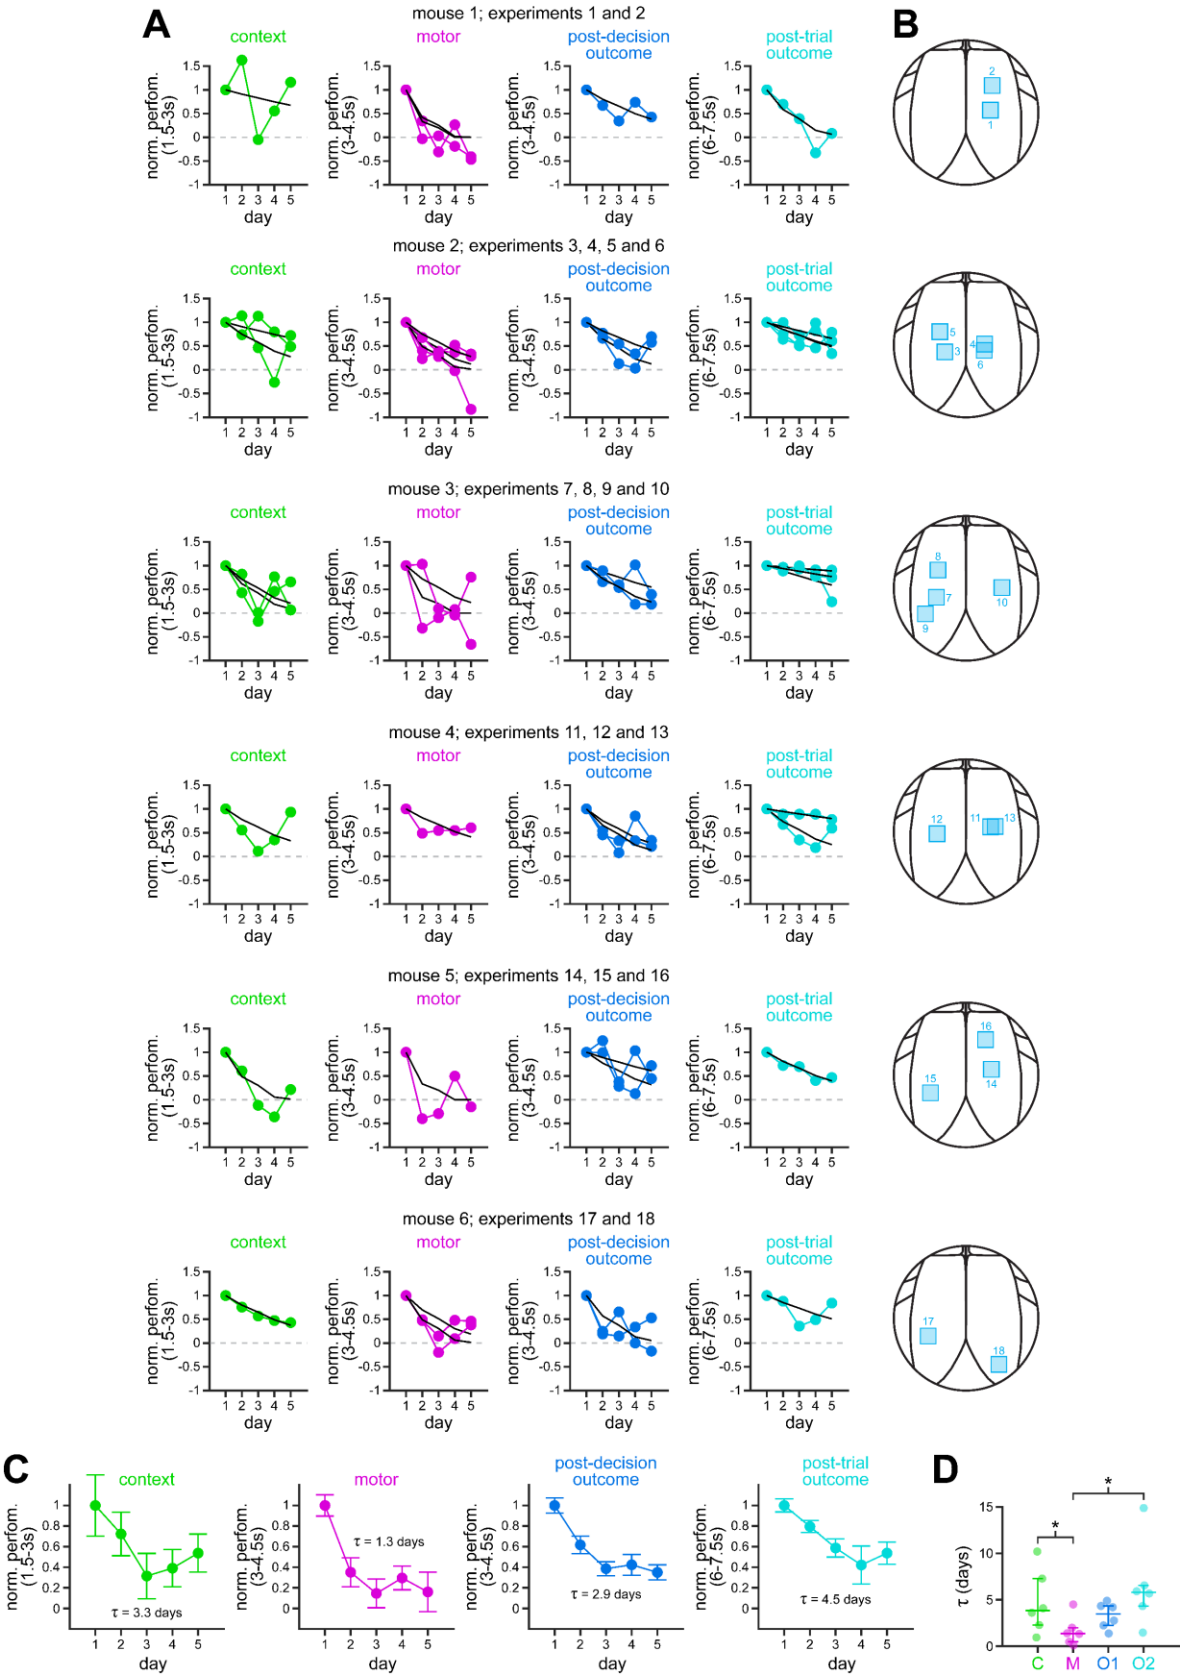

Fig. S14. Stability of task variables in individual experiments. Related to Fig. 6.

A. Normalized decoding performance integrated over the indicated 1.5 s windows relevant for each task variable. Exponential decay functions were fit to estimate the decay in encoding stability across days for each individual experiment (black lines). Their corresponding decay constants are plotted in **Fig. 6D**.

B. Maps showing the location of imaging fields within RSC for each mouse.

C. Normalized decoding performance (averaged across mice) integrated over the indicated 1.5 s windows relevant for each task variable (mean ± s.e.m.). Exponential decay functions were fit to estimate the decay in encoding stability across days (black lines). Note the faster decay in motor choice and post-decision outcome encoding, and the slower decay in context and post-trial outcome encoding.

D. Time constants (τ) calculated after fitting exponential decay functions to the decoding performance of each task variable in each individual mouse (see **Fig. S14A-B**). Dots, individual mice; overlaid lines, median ± 75% CIs. Asterisks denote significantly different decay constants (*$p < 0.05$, **$p < 0.01$, ***$p < 0.001$; LME models; see **Methods** for details).

Fig. S15. Related to Fig. 6.

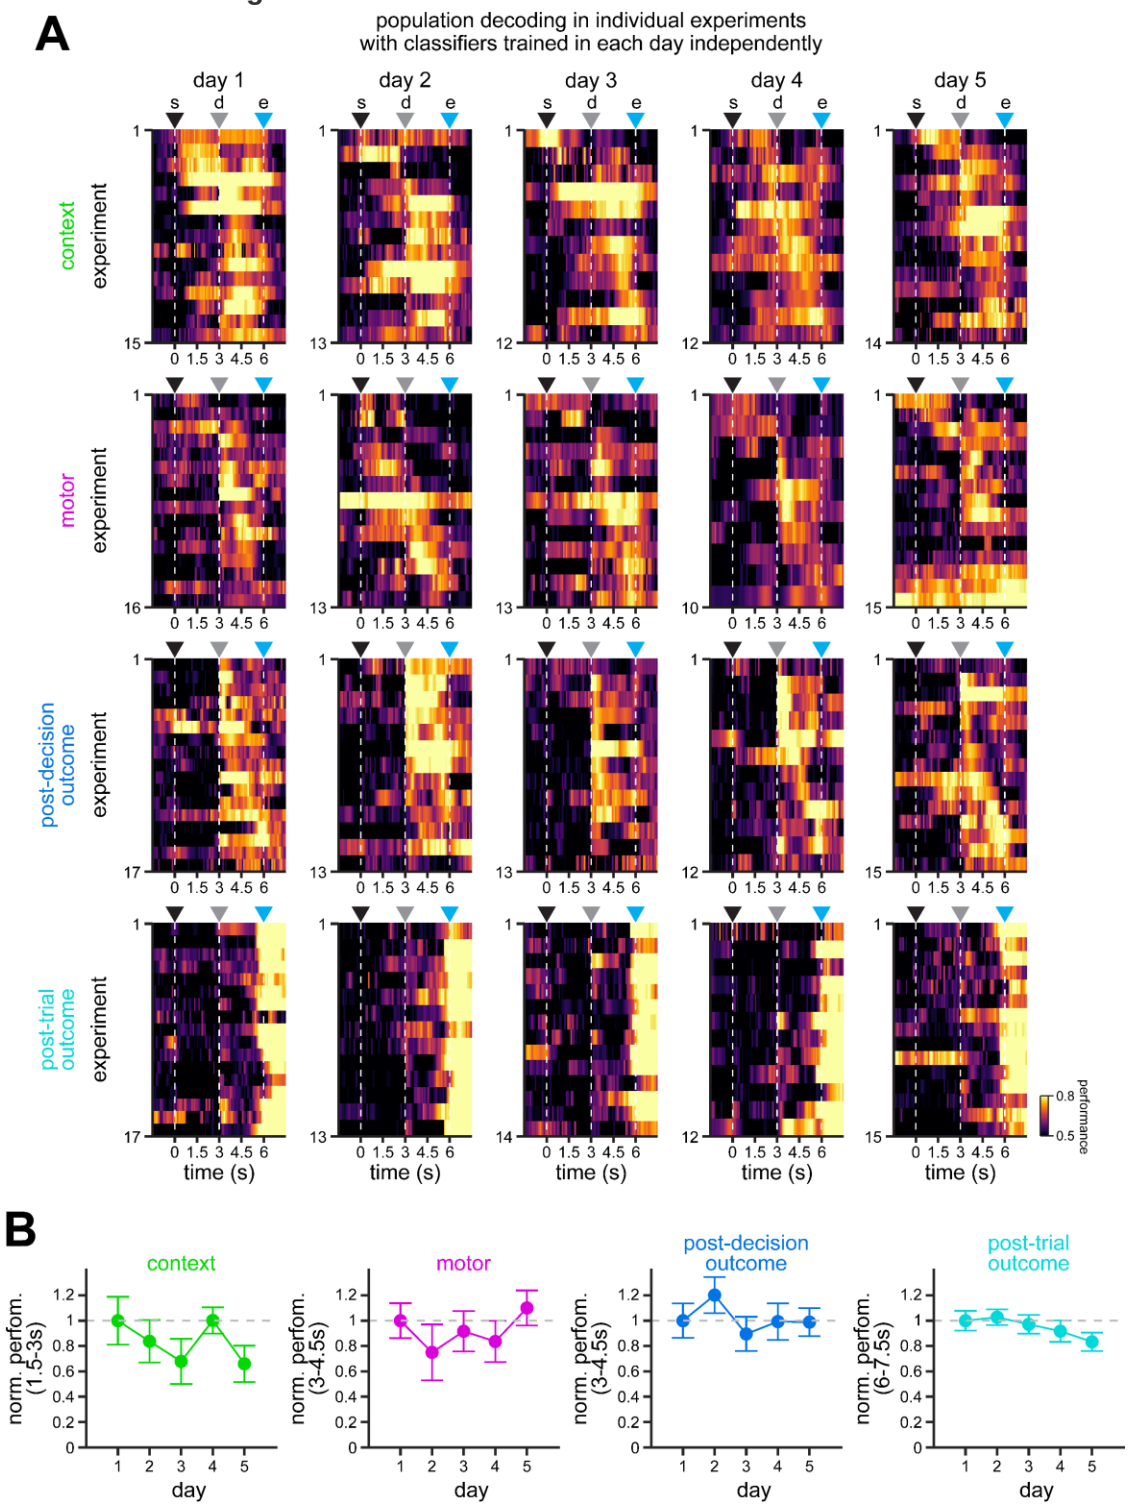

Fig. S15. Encoding of task variables in individual experiments. Related to Fig. 6.

A. Average performance of support vector machine classifiers on the activity of different subsets of cells separated by experiment: context (top row), motor choice (top-middle row), post-decision outcome (bottom-middle row), and post-trial outcome (bottom row). Models were trained and tested using activity within each day independently. Note similar

decoding performance for each task variable in each of the 5 days of experimentation. s = trial start; d = decision point; e = trial end.

B. Normalized decoding performance integrated over the indicated 1.5 s windows relevant for each task variable (mean ± s.e.m. across experiments). Note how decoding performance remains largely unchanged across all days.

Fig. S16. Related to Fig. 7.

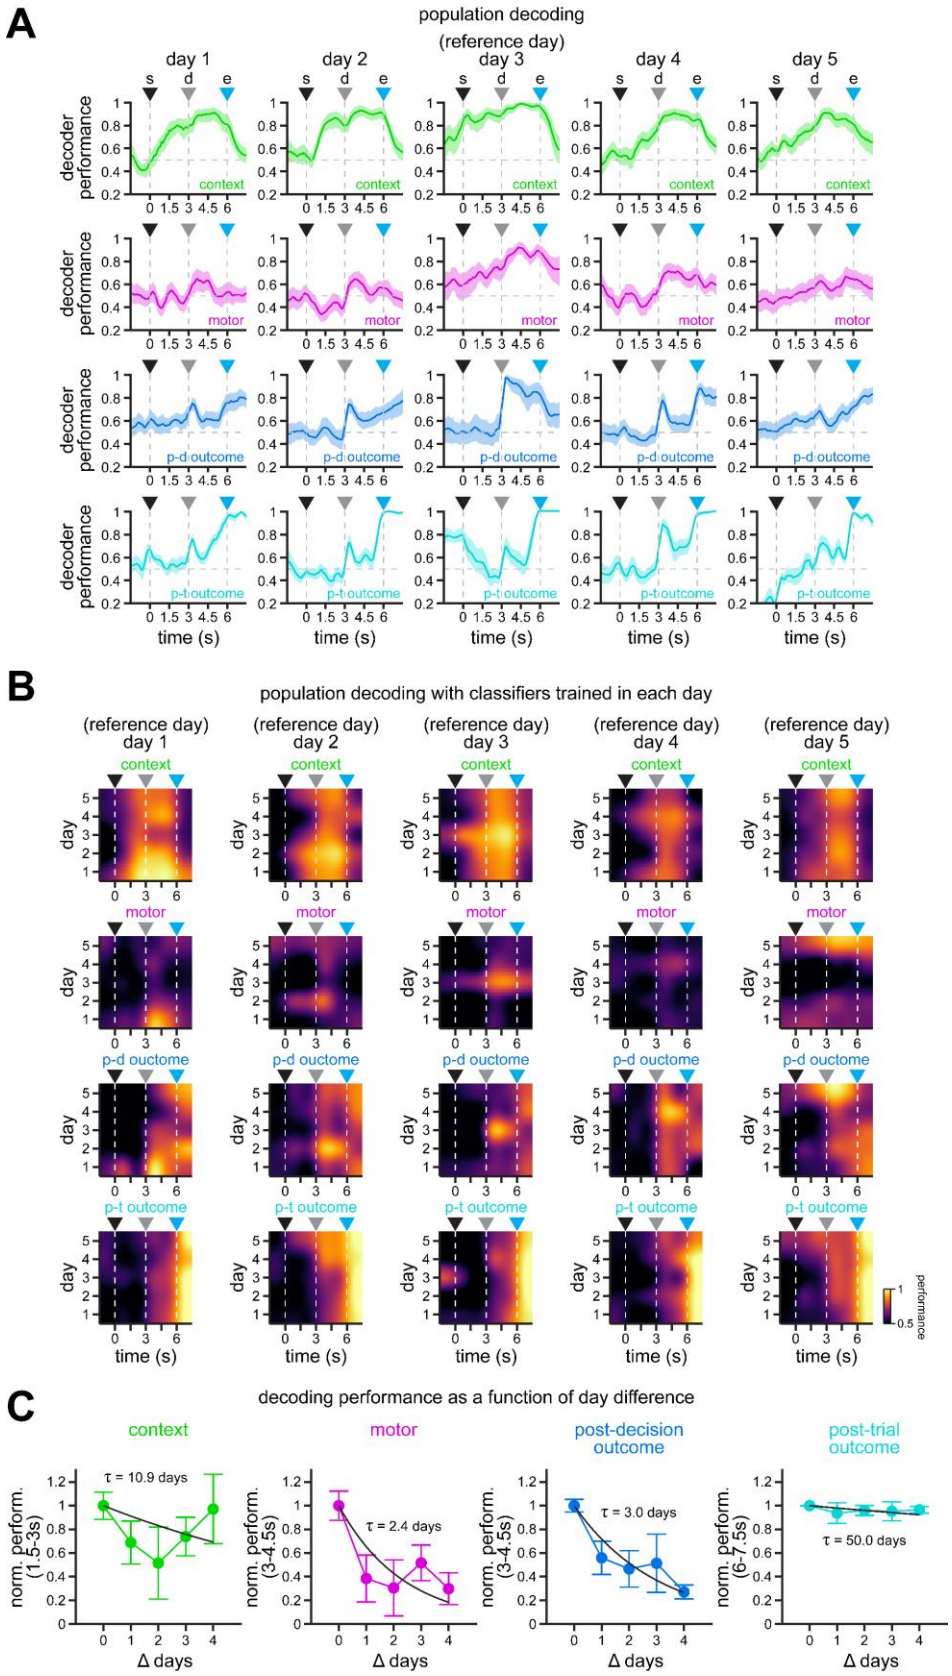

Fig. S16. Differential stability in population encoding of task variables in the RSC (classifiers trained in each day independently). Related to Fig. 7.

A. Encoding of context (top row), motor choice (top-middle row), post-decision outcome (bottom-middle row) and post-trial outcome (bottom row) by all recorded RSC neurons across all mice using day 3 as reference. Solid line, mean; shaded area, 95% CIs. Note the gradient in population encoding, where post-trial outcome and context are more stable than post-decision outcome and motor choice.

B. Heat maps of context (top row), motor choice (top-middle row), post-decision outcome (bottom-middle row), and post-trial outcome (bottom row) coding performance along trial duration and across all 5 days of experimentation when different days are used as reference. Note that the results are consistent regardless of the reference day used for training our support vector machine classifiers. Namely, higher stability of context and outcome variables, and lower stability of motor choice.

C. Normalized decoding performance integrated over the indicated 1.5 s windows relevant for each task variable (mean ± bootstrapped s.e.m.) as a function of day difference between the training of the classifier and testing on past and future days. Exponential decay functions were fit to estimate the decay in encoding stability across days (black lines denote fits to the average normalized performance). Note the faster decay in motor choice and post-decision outcome encoding, and the slower decay in context and post-trial outcome encoding. To obtain CIs for the decay constants, we also fit decay functions to each iteration of the decoder performance (95% CIs, context = 4.6 – 50.0 days, motor = 0.7 – 3.9 days, post-decision outcome = 1.7 – 4.3 days, post-trial outcome = 26.7 – 50.0 days).

Fig. S17. Related to Fig. 7.

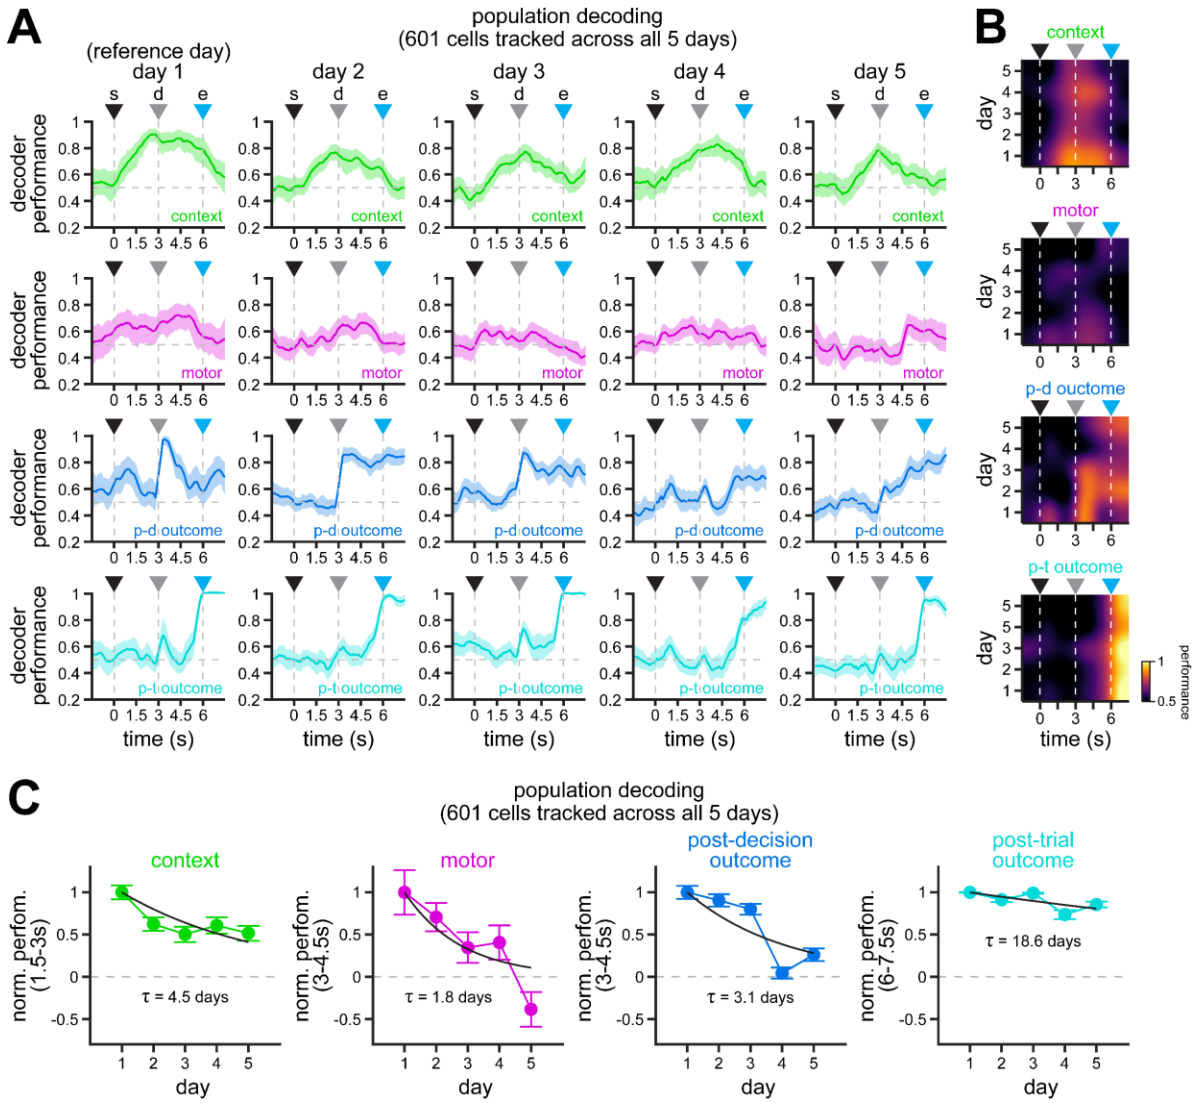

Fig. S17. Population encoding of task variables by a subpopulation of RSC neurons that can be tracked consecutively across 5 days. Related to Fig. 7.

A. Encoding of context (top row), motor choice (top-middle row), post-decision outcome (bottom-middle row) and post-trial outcome (bottom row) by the 601 cells that pass all tests in all days for inclusion in our analyses (ROI morphological stability, behavior criteria, and activity reliability). Solid line, mean; shaded area, 95% CIs. Note a similar gradient in population encoding as the one observed with the larger population of cells (**Fig. 7A**).

B. Heat maps of context (top), motor choice (top-middle), post-decision outcome (bottom-middle), and post-trial outcome (bottom) coding performance along trial duration and across all 5 days of experimentation. Note the similar dynamics in the encoding and stability of the different task variables (**Fig. 7B**).

C. Normalized decoding performance integrated over the indicated 1.5 s windows relevant for each task variable (mean ± bootstrapped s.e.m.). Exponential decay functions were fit to estimate the decay in encoding stability across days (black lines denote fits to the average normalized performance). Again, note similar dynamics (**Fig. 7C**), with a faster decay in motor choice and post-decision outcome encoding, and a slower decay in context and post-trial

outcome encoding. To obtain CIs for the decay constants, we also fit decay functions to each iteration of the decoder performance (95% CIs, context = 3.4 – 5.9 days, motor = 1.2 – 2.5 days, post-decision outcome = 2.7 – 3.7 days, post-trial outcome = 15.0 – 23.9 days).

Fig. S18. Related to Fig. 7.

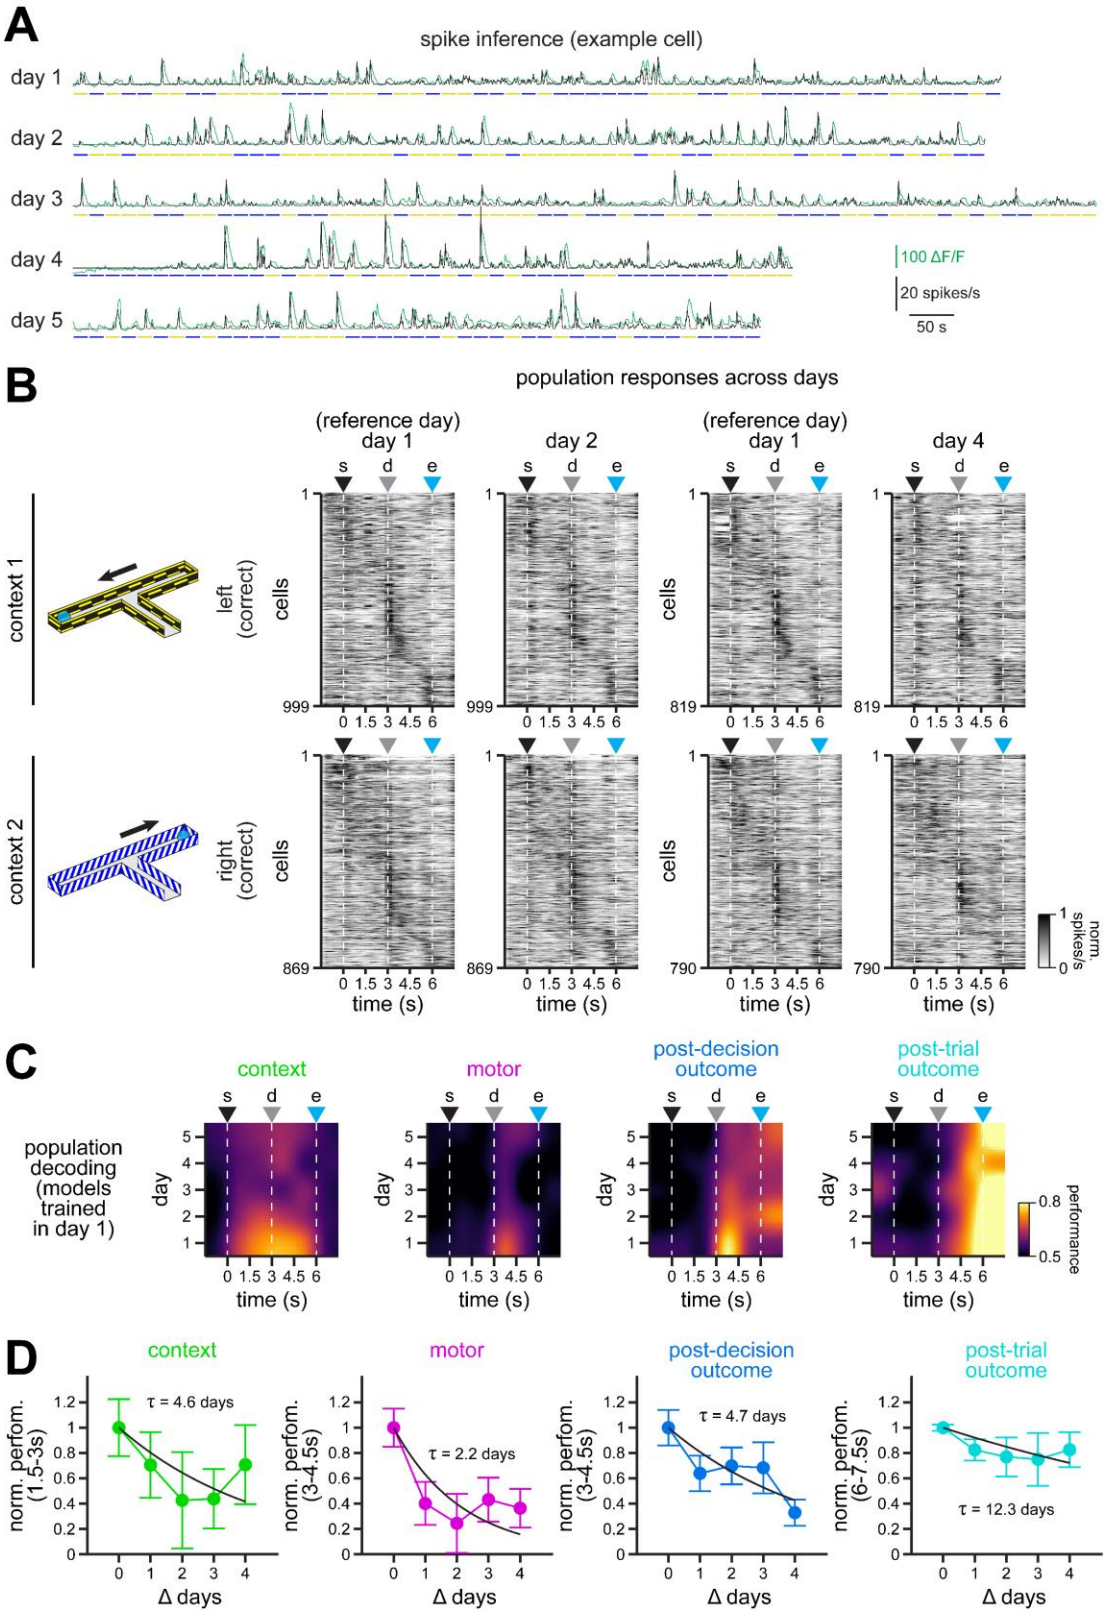

Fig. S18. Classifiers trained on inferred spikes reveal similar dynamics in the stability of task-related representations in the RSC. Related to Fig. 7.

A. Overlay of calcium signals (green) and inferred spike traces (black) for an example cell. Spikes were estimated by a sparse, nonnegative convolution algorithm (OASIS) with an autoregressive model of order 1 for the convolution kernel[60].

B. Averaged normalized spikes in the population of neurons separated by their preferred context. Only comparisons between days 1 and 2, and days 1 and 4 are shown for clarity. Spikes are cross-validated, sorted by latency using odd trials in day 1, and plotted using even trials for days 2 or 4.

C. Heat maps of context (left), motor choice (middle-left), post-decision outcome (middle-right), and post-trial outcome (right) coding performance along trial duration and across all 5 days of experimentation. Note the similar dynamics in the encoding and stability of task variables across days (**Fig. 7B**), even when estimated from inferred spikes.

D. Normalized decoding performance integrated over the indicated 1.5 s windows relevant for each task variable (mean ± bootstrapped s.e.m.) as a function of day difference between the training of the classifier and testing in past and future days. Exponential decay functions were fit to estimate the decay in encoding stability across days (black lines denote fits to the average normalized performance). Note similar dynamics, with a faster decay in motor choice and post-decision outcome encoding, and a slower decay in context and post-trial outcome encoding. To obtain CIs for the decay constants, we also fit decay functions to each iteration of the decoder performance (95% CIs, context = 1.3 – 13.1 days, motor = 0.6 – 3.8 days, post-decision outcome = 2.8 – 7.0 days, post-trial outcome = 6.4 – 40.5 days).

Fig. S19. Related to Fig. 7.

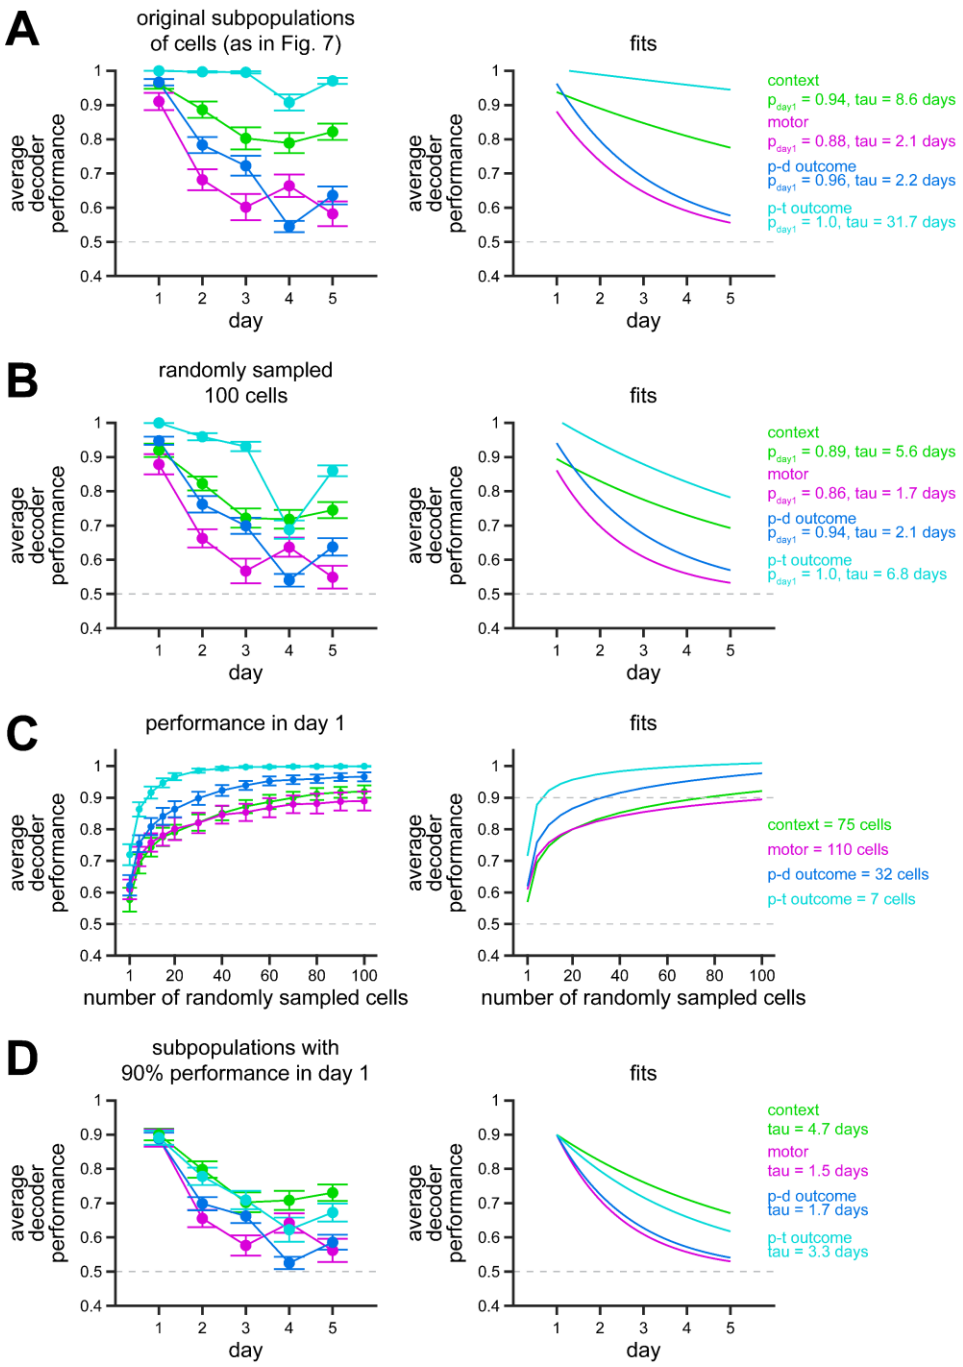

Fig. S19. Average decoding performance of classifiers trained on each day independently. Related to Fig. 7.

A. Average decoding performance over the indicated 1.5 s windows relevant for each task variable across days (mean ± bootstrapped s.e.m.). Note the faster decay in decoder performance for motor and post-decision outcome. By contrast, context and post-trial outcome show higher stability. Fits are shown on the right, with their corresponding values for performance in day 1 as well as their decay constants.

B. Similar to A, only in this case, subpopulations were resampled to 100 neurons in each iteration during training and testing of classifiers (mean ± bootstrapped s.e.m.). This ensures all task variables are evaluated with an identical

population size. Note a similar trend, where context and post-trial outcome exhibit higher stability. Fits are shown on the right, with their corresponding values for performance in day 1 as well as their decay constants.

C. Average decoding performance for randomly sampled neurons over 100 iterations in day 1 (mean ± bootstrapped s.e.m.). The decoder performance (P) as a function of the number of neurons (n) was fit to a power function ($\mathrm{P} = a * n^b + c$), and the number of neurons corresponding to 90% decoding performance in day 1 was estimated, as indicated.

D. Similar to A, but with different subpopulations of sampled neurons corresponding to 90% decoding performance in day 1 (as in **C**). Note the higher stability of context and post-trial outcome, and the lower stability of motor and post-decision outcome.

Video S1. Calcium imaging and behavior videos from the same mouse over several sessions (Day 1: upper left; Day 3: upper right; Day 4: lower left; Day 5: lower right). Each trial type is annotated.