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]Division of Biology and Biological Engineering,
California Institute of Technology
Pasadena, CA 91125, 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: lmfranco@caltech.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 information is stored with varying levels of stability, maintaining an adequate balance between stability and flexibility to subserve 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 the activity of dynamic 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. In the piriform cortex, by contrast, odor representations constantly drift over days unless odors are frequently experienced[12]. Other cortical regions, such as the auditory cortex, experience a monotonic decay in their spontaneous activity, despite a small fraction of cells remaining strongly correlated over multiple days[13]. 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[1–4,14]. 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[3,8,15] .

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[16]. 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[17,18], although granule cells of the dentate gyrus exhibit more stable activity[18,19]. Interestingly, recent findings suggest that in the bat hippocampus, CA1 neurons can stably map the trajectory of highly reproducible flights in familiar environments over days and weeks[20]. In the retrosplenial cortex (RSC), a region important for the association of allocentric and egocentric spatial representations[21], as well as for the association between environmental context with motor choice[22–24], a considerable fraction of cells maintain stable retinotopic tuning properties to simple visual stimuli[25]. 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[6,26]. 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 sensory and decision variables during performance of a behavioral 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[22]. 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 used a behavioral task requiring mice to associate familiar virtual contexts with a left or right joystick turn in order to obtain liquid rewards (**Fig. 1A**; for a full description of this task, see **Methods** and ref. [22]. During behavioral performance, we used two-photon calcium imaging to record the activity of 4,701 neurons in layer 2/3 of the 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 achieve an unbiased analysis of neuronal responses across all 5 days, we only consider experiments in which mice had at least 2 minimum trials in each of the 4 possible trial outcomes (Context 1 - Left, Context 1 - Right, Context 2 - Left, or Context 2 - Right) for all 5 days. This yielded 2,220 neurons tracked in 8 experiments in 4 mice (see **Fig. S1** for details). On average, these 4 mice made (mean ± s.e.m.): 20.7 ± 1.4 correct left, 7.4 ± 0.7 incorrect right, 6.6 ± 0.6 incorrect left, and 24.0 ± 1.4 correct right decisions, which translated to an overall behavioral performance (mean ± s.e.m.) = 75.7% ± 7.4% across 5 days (**Fig. 1B**).

Out of the 2,220 neurons, 1,609 neurons (72.5%) showed significant reliable responses on day 1 (see **Methods** for details), with 889 neurons (55.3%) preferring context 1, and 720 neurons (44.7%) preferring context 2 (**Fig. 1C**). To cross-validate these results, we used odd trials in day 1 as reference, and compared to even trials in days 1 to 5 (**Fig. 1C**). We noticed that responses in correct trials appear more consistent across days, as observed by a similar ridge in the population response (**Fig. 1C**, yellow diagonal). To evaluate this, we calculated the Pearson's correlation coefficient of population vectors for all of the 1,609 neurons in correct or incorrect trials (**Fig. 1D-E**). We find that activity in correct trials is indeed more stable than in incorrect trials, as shown by lower 95% CI in correct trials higher than higher 95% CI in incorrect trials (**Fig. 1F-G**). To ensure this was not due to unequal trial numbers, we sampled equally across trial types (**Fig. S2A**), and confirmed that population vectors in correct trials exhibit higher stability (**Fig. 2B-E**). In addition, we observe that responses in correct trials experience a

quick decay in stability after the first day, with a slower decay on following days (**Fig. 1C, S2A**). Despite these changes in neuronal activity, the population conserved high levels of structure and stability across the 5 days of experimentation, as revealed by high similarity of population vectors in correct trials (**Fig. 1D-G, S2B-E**). This suggests that although activity in individual neurons undergo some level of drift (**Fig. 1C, S2A**), the population as a whole preserves substantial information in correct trials across days. To assess whether the observed drift could be due to changes in response properties across days, we compared response characteristics across imaging sessions, finding similar levels of neuronal activity (**Fig. S3A**), similar fraction of neurons mapping the trial duration (**Fig. S3B**), and similar preference indices for a particular context (**Fig. S3C**). Thus, although some neurons lose their task-driven responses, other neurons become tuned to the task each day, contributing to homeostasis of task representation across the neuronal population. Further, these results indicate that population dynamics are preserved in the RSC circuit across days, despite some variability in single cell activity.

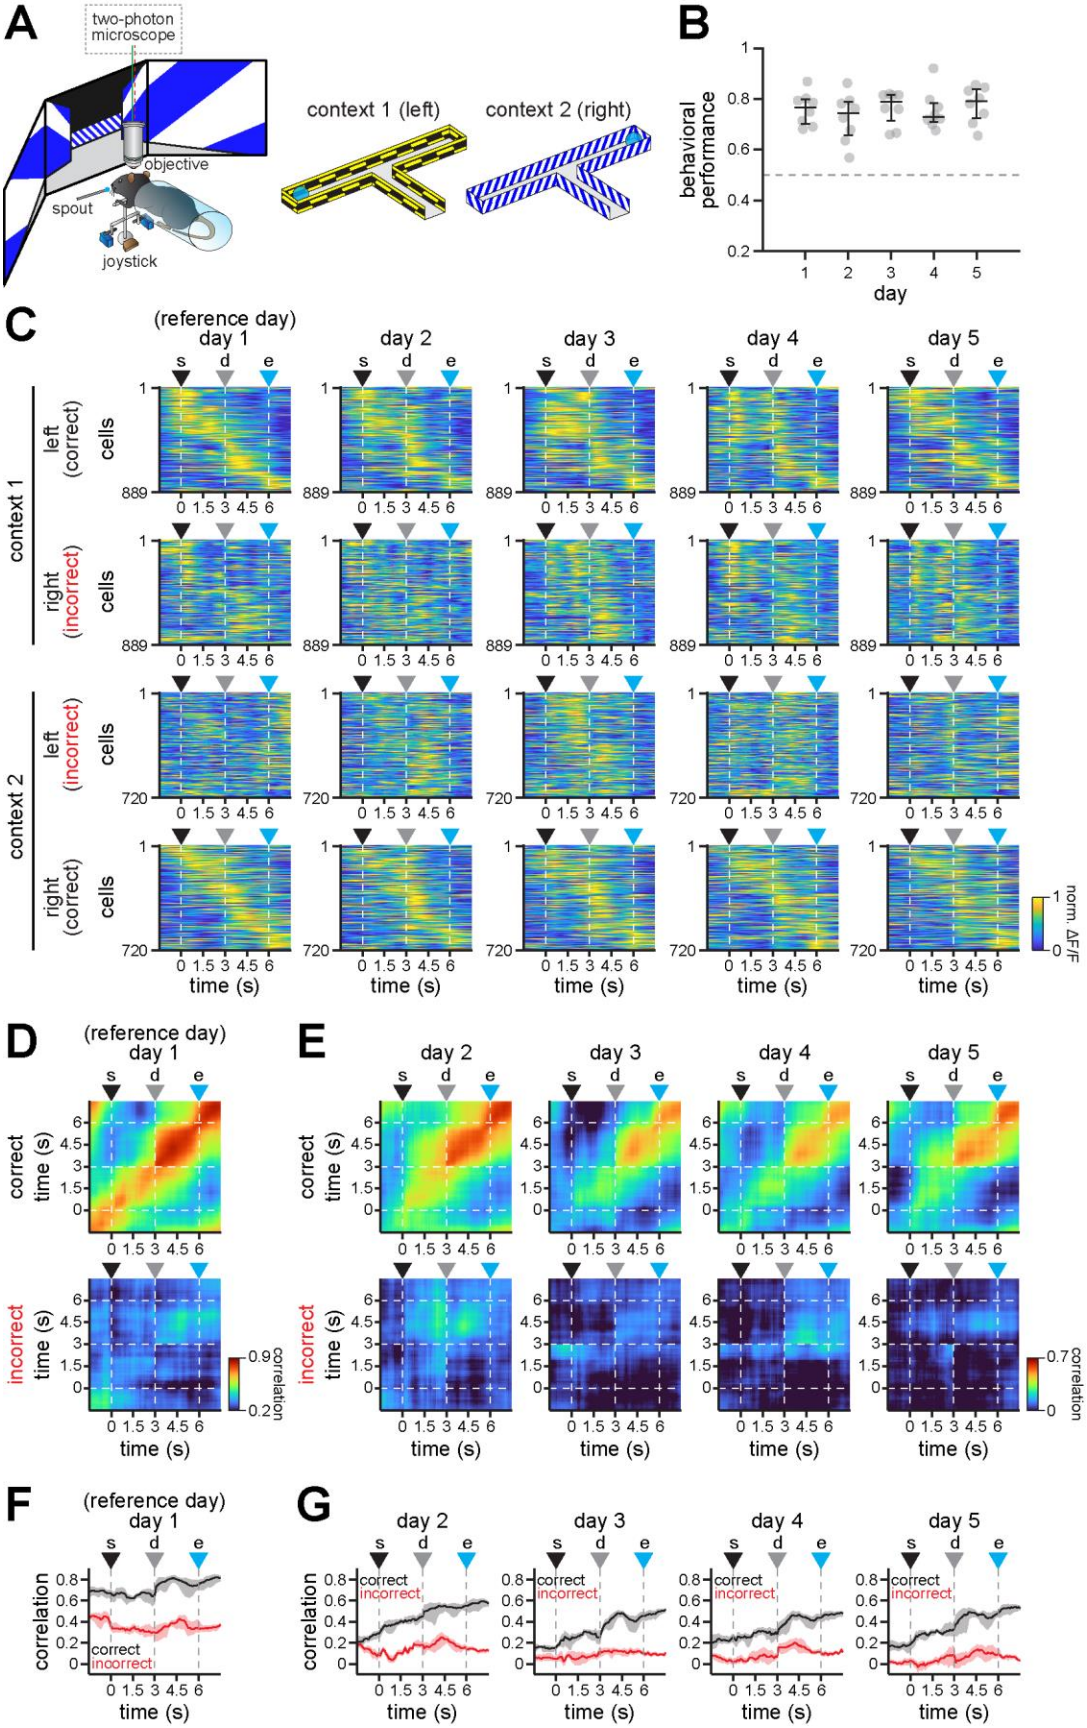
A
two-photon microscope
objective
spout
joystick
context 1 (left)
context 2 (right)
B
behavioral performance
day
C
(reference day) day 1
day 2
day 3
day 4
day 5
s
d
e
context 1
context 2
left (correct)
right (incorrect)
left (incorrect)
right (correct)
cells
time (s)
norm. ΔF/F
D
correct
incorrect
correlation
E
F
G

Fig. 1. Population activity in RSC exhibits different levels of stability in correct and incorrect trials.
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. Single cell neural activity is recorded using a two-photon microscope. For details, see **Methods** and ref. [22].
B. Behavioral performance across 5 consecutive days (n = 8 experiments in 4 mice). Mouse performance was significantly above chance in all experiments (median): day 1 = 76.8%, day 2 = 74.5%, day 3 = 79.0%, day 4 = 73.0%, day 5 = 79.3%. Gray dots, individual sessions; black overlay, median ± 75% CI.
C. Average normalized activity for the 1,609 neurons with at least 2 minimum trials in correct and incorrect trials for both contexts across the 5 days of experimentation. Neurons are grouped according to their context preference in correct trials. Activity is cross-validated, sorted by latency using odd trials on day 1, and plotted using even trials for days 1 to 5.
D-E. Pearson's correlation coefficient matrices showing the similarity of activity population vectors across trial duration. Data are cross-validated by comparing population vectors in odd trials in day 1 against population vectors in even trials from days 1 to 5. Population vectors were constructed from the preferred responses in each context, as shown in C (1,609 neurons).
F-G. Pearson's correlation coefficient along the diagonal in D and E in a window ± 1 second perpendicular to the diagonal. Note higher stability and similarity of population activity on correct trials. Solid line, mean; shaded area, 95% CI. For all panels: s = trial start; d = decision point; e = trial end.

To further investigate the observed drift in activity of individual cells, we grouped neurons with significant responses by preference to environmental context (context 1 or 2), motor choice (left or right turn), or trial outcome (reward or no reward), taking into account both correct and incorrect trials (**Fig. 2A**, **S4**). We then determine their preference in day 1 (**Fig. 2B**), and measure their stability, comparing odd trials in day 1 against even trials in days 1 to 5 for their preferred context, motor choice or trial outcome (**Fig. 2C**). We find that responses grouped by context appear more stable than responses grouped by motor choice (**Fig. 2A, S4**), as shown by significantly higher Pearson's correlation coefficients (**Fig. 2C**, left), and significantly lower change in their center of mass (**Fig. 2C**, right). 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. However, responses grouped by trial outcome, characterized by increased activity around trial end, show the highest levels of stability (**Fig. 2A,C-D, S4**). Together, these results indicate that activity in the RSC is partially stable, with conserved population structure across days, particularly in correct trials (**Fig. 1D-E**), which likely depends on the specific tuning of each cell (**Fig. 2**). However, a more detailed examination is required in order to determine whether single cell and population level representations of specific task variables exhibit distinct levels of stability.

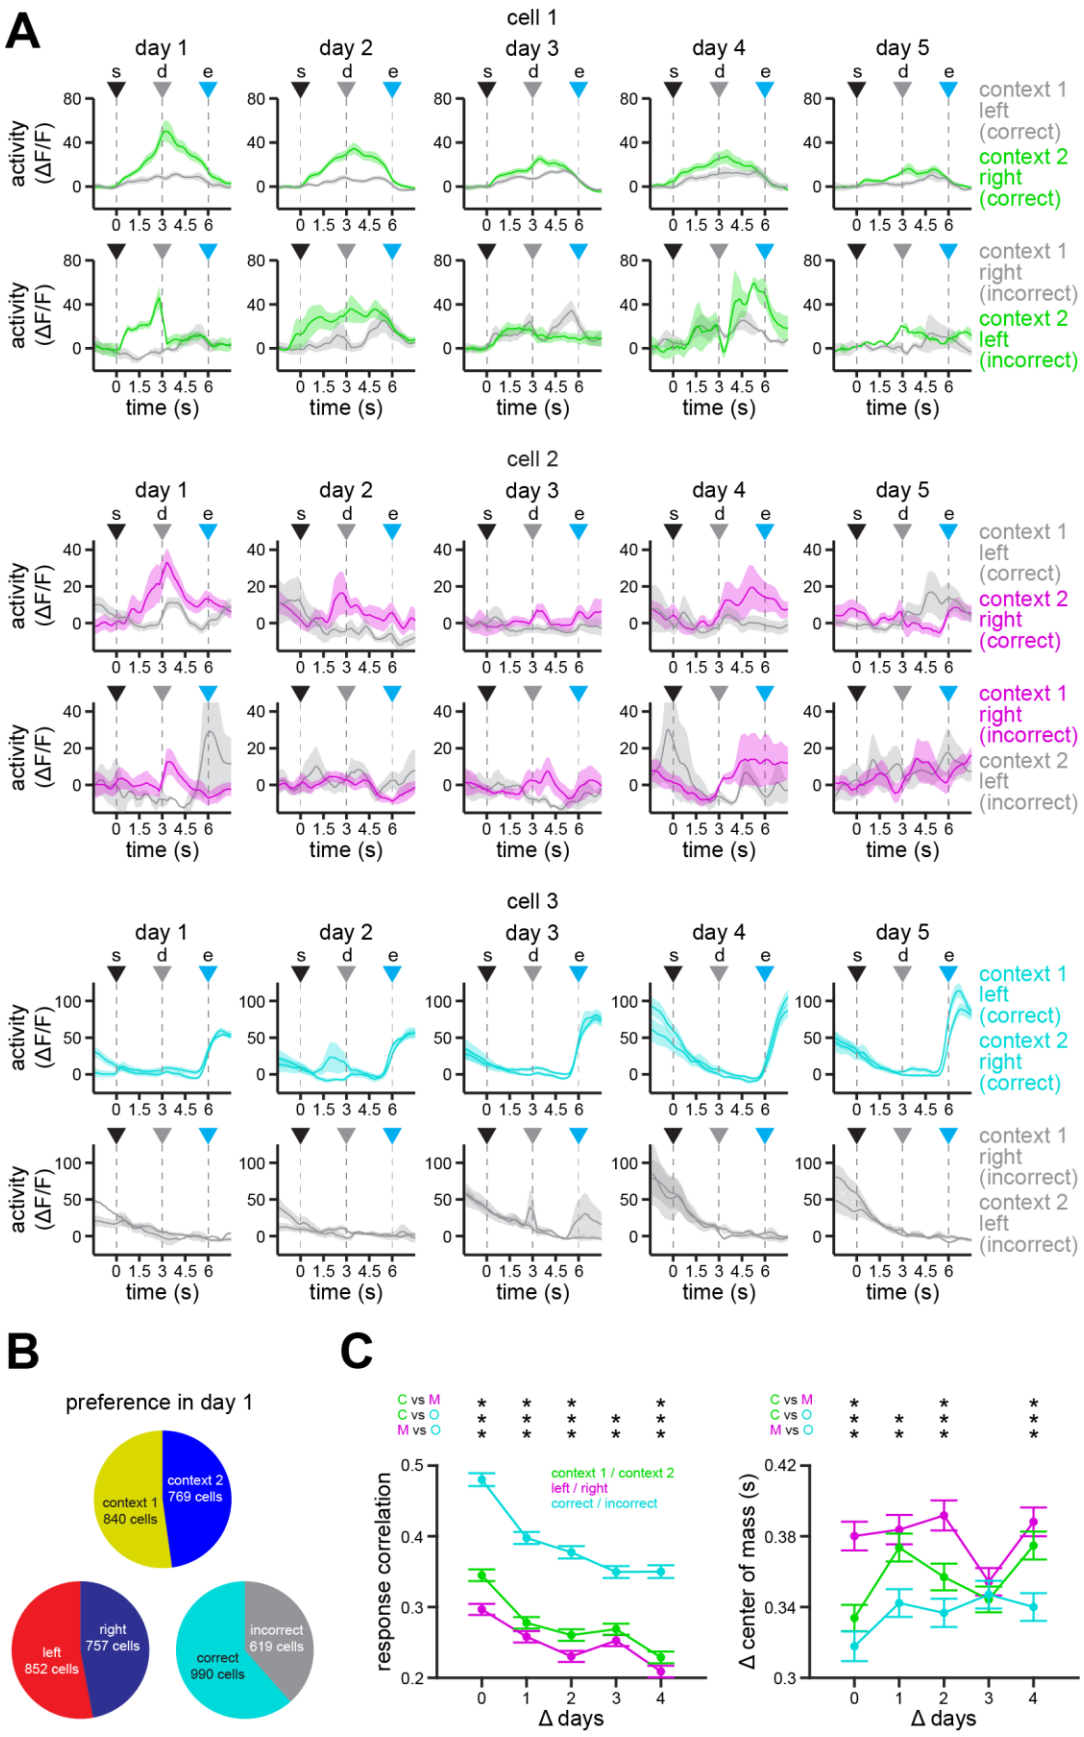
A
cell 1
day 1
day 2
day 3
day 4
day 5
s
d
e
activity (ΔF/F)
time (s)
context 1 left (correct)
context 2 right (correct)
context 1 right (incorrect)
context 2 left (incorrect)
cell 2
context 1 left (correct)
context 2 right (correct)
context 1 right (incorrect)
context 2 left (incorrect)
cell 3
context 1 left (correct)
context 2 right (correct)
context 1 right (incorrect)
context 2 left (incorrect)
B
preference in day 1
context 1 840 cells
context 2 769 cells
left 852 cells
right 757 cells
correct 990 cells
incorrect 619 cells
C
C vs M
C vs O
M vs O
context 1 / context 2
left / right
correct / incorrect
response correlation
Δ center of mass (s)
Δ days

Fig 2. Activity in individual RSC neurons tuned to context, choice or outcome exhibit different levels of stability.

A. Example neurons exhibiting higher responses for context 1 (cell 1), for right turns (cell 2), and for correct outcomes (cell 3), 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. Number of cells with a preference for environmental context, motor choice or trial outcome in day 1. Preference is defined as higher activity in the indicated subset of trials (trials are equally weighted within each subset, then their corresponding averages are compared).

C. Average Pearson's correlation coefficient (left) and change in center of mass (right) between odd trials in day 1 and even trials in days 1 to 5 (mean ± s.e.m.). Trials are equally weighted, then their average values are compared (n = 1,609 cells, Wilcoxon signed-rank test, $p < 0.05$). Asterisks indicate significant differences between context and motor choice (top), context and trial outcome (middle), or motor choice and trial outcome (bottom).

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

To assess the stability of task variables throughout the time course of the trial, we next evaluated the performance of linear decoders across each individual time point[22,27]. To this end, we trained support vector machine classifiers for each specific task variable (context, motor choice, outcome) using activity in reference day 1, and tested the performance of these classifiers on the activity of the same cells in the following days (see **Methods**). To analyze day 1 activity without overfitting, we divided the trials in half for fitting the classifier and half for testing, respectively. Out of the 1,609 cells with significant reliable responses, we find slightly overlapping subpopulations consisting of 125, 111 and 406 cells (533 cells in total) significantly encoding context, motor and outcome information, respectively, in day 1 (**Fig. 3A**). Moreover, we observe that some cells maintain their decoding performance across days with similar temporal dynamics as in reference day 1 (**Fig. 3B**). However, these subpopulations exhibit different levels of stability (**Fig. 3A-D**, see also **Fig. S5**). For instance, we observe a higher fraction of cells consistently encoding trial outcome across the 5 days of experimentation (**Fig. 3A-B, E**). In addition, cells encoding environmental context exhibit significantly higher stability than cells encoding motor choice from days 3 to 5 (**Fig. 3A-B, E**). These changes in the encoding stability of single cells likely reflect a drift in population activity (**Fig. 1C**), yet some task variables appear to drift more slowly than others. One possibility is that the differential stability we observe across days might simply reflect better overall decoding of those variables in RSC neurons. To untangle this, we grouped cells according to their decoding performance in day 1,

and quantified the fractions of cells that maintain significant decoding (**Fig. S6**). We observe a similar gradient, where trial outcome is more stable than environmental context, and in turn more stable than motor choice, even across cells with similar decoding performance in day 1 (**Fig. S6**). One implication of progressive decay in information storage and encoding across days is that task information will be eventually lost by the circuit, given enough time. However, we find that models trained and tested within each independent day show consistent subpopulations of cells encoding context, motor and trial outcome (**Fig. S7**). This suggests that although some neurons lose their ability to encode features of the task, other neurons gain coding computational power, maintaining the homeostasis of the RSC circuit, and contributing to the storage of information across days.

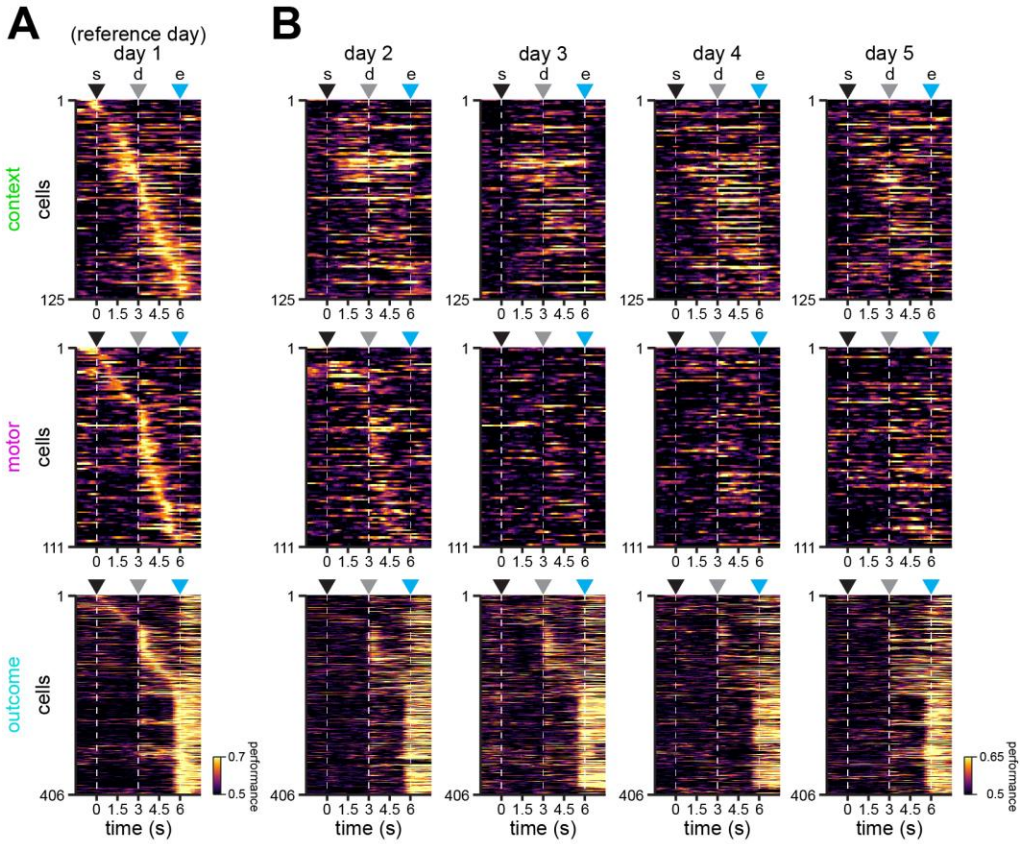
A
(reference day)
day 1
B
day 2
day 3
day 4
day 5
s
d
e
context
motor
outcome
cells
1
125
111
406
0 1.5 3 4.5 6
time (s)
0.7
0.5
0.65
performance

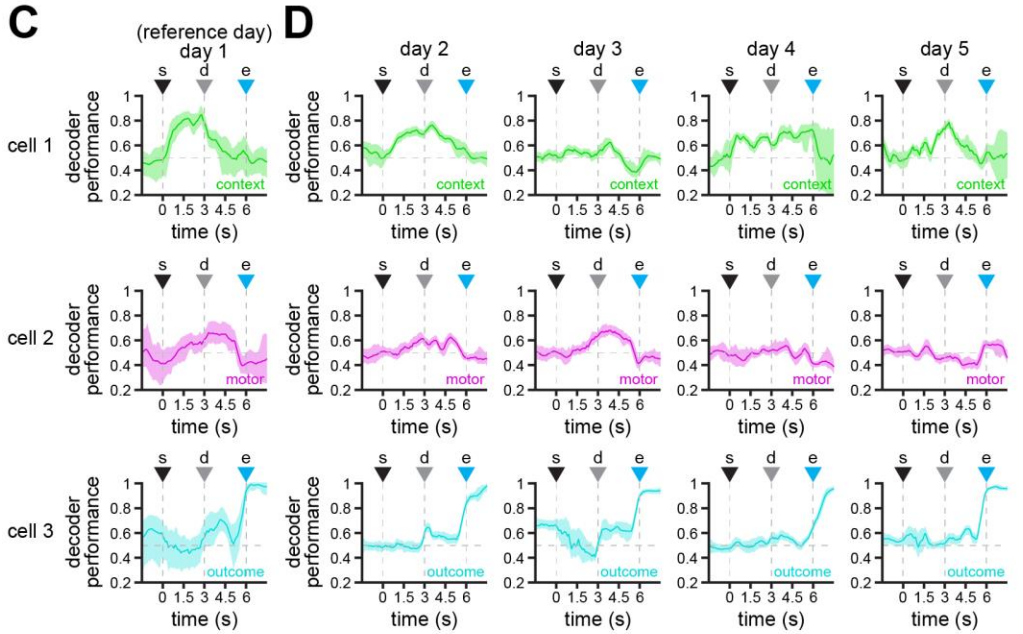
C
(reference day)
day 1
D
day 2
day 3
day 4
day 5
cell 1
cell 2
cell 3
decoder performance
context
motor
outcome
s
d
e
0 1.5 3 4.5 6
time (s)

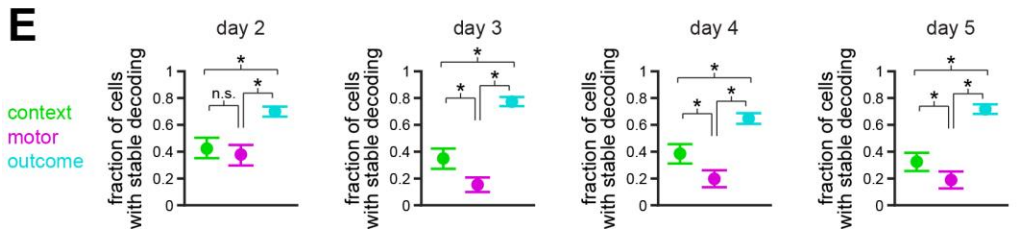
E
day 2
day 3
day 4
day 5
context
motor
outcome
fraction of cells with stable decoding
n.s.
*

Fig. 3. RSC neurons encode task variables with differential stability.

A-B. Average performance of a support vector machine classifier on the activity of cells exhibiting significant decoding of environmental context (top row), motor choice (middle row), or trial outcome (bottom row) in day 1. For the cells shown in A, classifiers were computed and tested using activity within day 1. In B, for those same cells, models were computed using activity in day 1, and then tested using activity in days 2 to 5. Note the gradient in the stability of information decoding, where trial outcome exhibits higher stability than context, which in turn is more stable than motor choice.

C-D. Example cells with significant decoding of context (cell 1, top), motor choice (cell 2, middle), and trial outcome (cell 3, bottom) in day 1. Solid line, mean; shaded area, 95% CI.

E. Fraction of cells that maintain significant decoding in days 2 to 5 (mean ± bootstrapped 95% CI). Stability was defined as cells with at least 5 bins with significant decoding that were also significant in day 1. Asterisks denote significantly higher stability. Significance was defined as lower 5% CI higher than higher 95% CI among task variables. For all panels: s = trial start; d = decision point; e = trial end.

Our findings on the encoding of task variables by individual cells suggest that differential stability might also be present within individual RSC neurons multiplexing information of at least two variables[22]. We evaluated this by comparing the performance of support vector machine classifiers in cells with significant encoding of two task variables in reference day 1 (**Fig. 4A-C**). Indeed, we find a similar gradient in which both context and trial outcome show higher stability than motor choice, as quantified by a consistently higher number of time bins with significant decoding (**Fig. 4D**). In order to further characterize this, we compare the density of the area under the curve of the classifier performances (**Fig. S8A**) for the subpopulations shown in **Fig. 4A-B**. We observe higher stability in the encoding of environmental context in context-motor cells, particularly in the first half of the trial, before the decision point (top row, **Fig. S8**). 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 (bottom row, **Fig. S8**). 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 (middle row, **Fig. S8**). These results highlight an important feature of RSC neurons that multiplex task information, which favor the maintenance of environmental context and trial outcome over a more flexible encoding of motor choice.

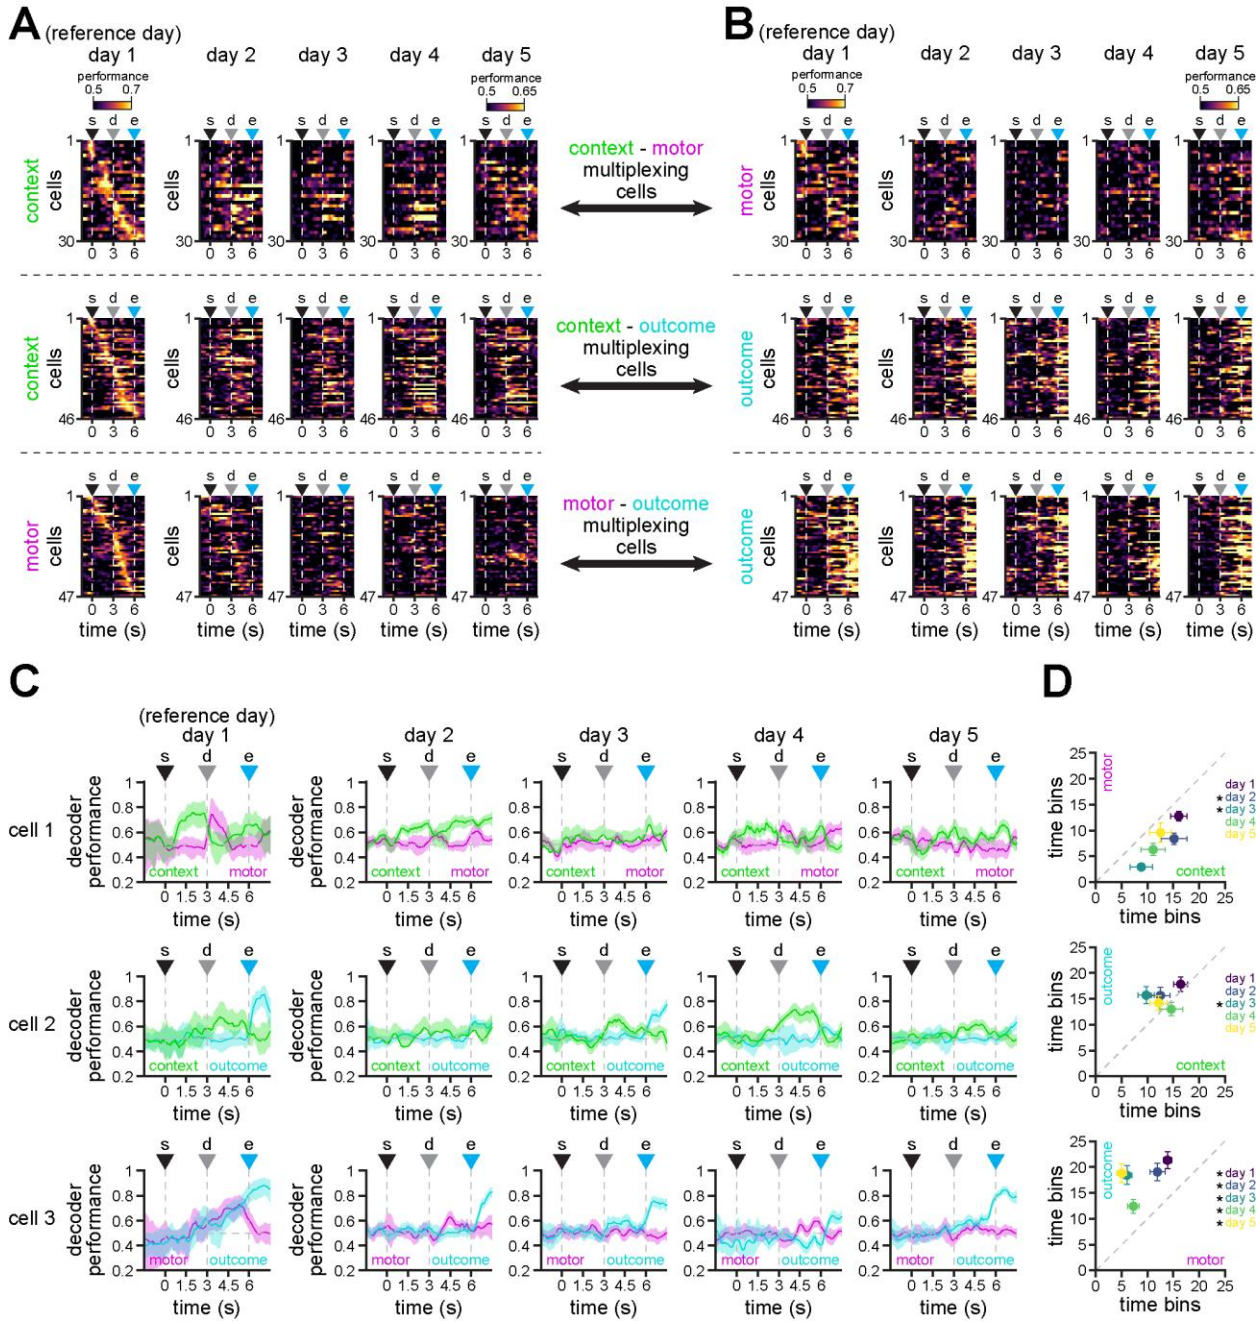

Fig. 4. RSC neurons multiplex task information with differential stability.

A-B. Average performance of a support vector machine classifier on the activity of cells significantly multiplexing the encoding of at least 2 task variables: context and motor choice (top row), context and trial outcome (middle row), and motor choice and trial outcome (bottom row) in day 1. For the reference day (most left columns), models were computed and tested using activity within day 1. For the other days, models were computed using activity in day 1, and then tested using activity in days 2 to 5.

C. Example cells with significant multiplexing of context and motor choice (cell 1, top), context and trial outcome (cell 2, middle), and motor choice and trial outcome (cell 3, bottom) in day 1. Solid line, mean; shaded area, 95% CI.

D. Number of time bins with significant decoding of the 2 task variables indicated in each plot. Asterisks denote a significantly different number of bins with significant decoding (Wilcoxon rank-sum test, $p < 0.05$). Note the higher performance of the classifiers decoding context and trial 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. 1C**) as well as on the stability of information encoding in single cells (**Fig. 3A-B**) suggest that a similar tendency is present in population encoding of task variables in the RSC. To test this hypothesis, we used support vector machine classifiers trained on population vectors on day 1 and tested on population vectors in the following days. For these analyses, we only considered the subpopulations of cells shown in **Fig. 3**, although similar results can be observed when evaluating the whole population of 1,609 cells (**Fig. S9**). In accordance with our previous findings[22], 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 (**Fig. 5A**, upper left panel). Decoding of motor choice is highest around the decision point, when mice rotate the joystick to select a trajectory through the maze (**Fig. 5A**, middle left panel). Decoding of trial outcome exhibits two distinct peaks, with a small peak around the decision point, and ceiling level performance during the reward period after the trial end (**Fig. 5A**, bottom left panel). In terms of the stability of these representations, 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. 5A-B**). By contrast, population decoding of motor choice undergoes a fast and abrupt decay from day 1 to day 2, and is practically extinguished by day 3 (middle row, **Fig. 5A-B**). Trial outcome, on the other hand, exhibits very stable population decoding at the end of the trial (6-7.5s) across the 5 days of experimentation, although the signal appearing right after the decision point (3-4.5s) decays more quickly (bottom row, **Fig. 5A-B**). To further characterize the stability of population encoding, we calculated the area under the curve at relevant windows during task performance. Namely, 1.5-3s for environmental context (right before the decision point), 3-4.5s for motor choice and for trial outcome (right after the decision point), and 6-7.5s for trial outcome (when mice either receive or do not receive a reward). Then, we fit an exponential decay function ($\mathrm{AUC} = \mathrm{c} * e^{-\mathrm{day}/\tau}$) to estimate the time at which decoding decays to 1/e (~37%) of its original

value in day 1 (**Fig. 5C**). Consistent with our observations in single cells, we find that context (**τ** = 6.4 days) is considerably more stable than motor choice (**τ** = 1.6 days). Also, trial outcome exhibits bimodal stability along the trial time course, with lower stability for the signal following the decision point (**τ** = 4.0 days), and very high stability for the end of the trial, when mice either receive or do not receive a reward (**τ** = 179.1 days). Note that responses to uninstructed movements, such as licking of the water reward, may contribute to the high encoding and stability of trial outcome at the end of the trial[28,29]. 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. 5D**). 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. 5E**).

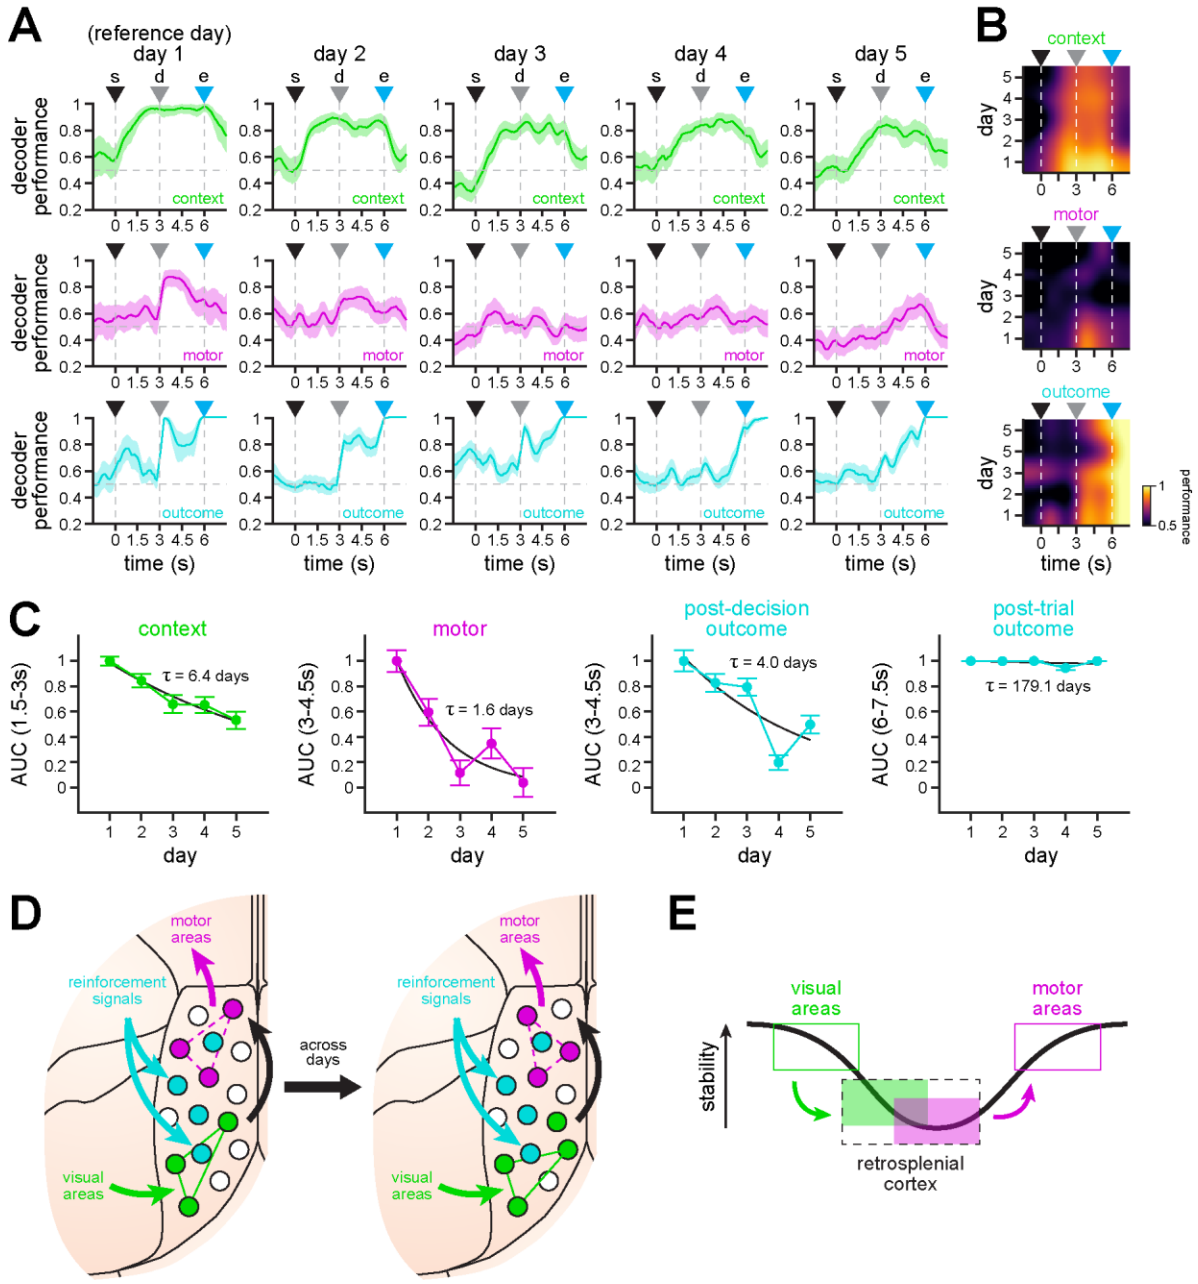

Fig. 5. Population-level coding of task variables exhibits differential stability in RSC.

A. Encoding of environmental context (top row), motor choice (middle row), and trial outcome (bottom row) by the subpopulations of neurons shown in **Fig. 3A-B** (context, 125 cells; motor, 111 cells; outcome, 406 cells) across 5 consecutive days of experimentation. Solid line, mean; shaded area, 95% CI. Note the gradient in population encoding, where outcome is more stable than context, which in turn is more stable than motor choice. s = trial start; d = decision point; e = trial end.

B. Heat maps of context (top), motor choice (middle), and trial outcome (bottom) coding performance along the trial duration, and across the 5 days of experimentation. Note the stability of context encoding

between trial start and trial end. Also, note the stability of outcome encoding after trial end. By contrast, the encoding of motor choice around the decision point decays abruptly after day 2.

C. Area under the curve (AUC) for the indicated 1.5 s time windows relevant for each task variable (mean ± s.d.). An exponential decay function was fit to estimate the decay in population encoding stability across days. Note the fast decay in motor choice encoding, and the slower decay in context and in trial outcome encoding around the decision point. Also note the very stable encoding of trial outcome after trial end.

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[22,25,30], 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 [5,6,9,10]. 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. During the association process, motor choice signals are represented with even lower stability. Once this information has been processed, it is then conveyed to motor areas, where the representation of motor plans and descending motor commands is much more stable.

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[31–33]. However, recent technological developments for longitudinal recording of neural activity have shown that tuning to specific variables 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[22,34], 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 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[35,36]. 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[28,29,37–41]. 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[42,43], providing another substrate for increased variability in incorrect trials.

Dedicated sensory or motor cortices were recently shown to stably encode information over days and weeks[5,6,9,10,13]. However, the stability of sensory and motor representations in association areas, which comprise the majority of the human cortex, is less understood. Our results indicate that RSC neurons exhibit more stable responses to sensory environmental contexts 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 ref. [22].

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[44]. In addition, motor areas might help consolidate the encoding of task features by a feedback mechanism, involving other brain regions such as the dorsal striatum[45]. Together, this evidence suggests that stable representations of sensory features in the RSC 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[21,22,46–54], 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.

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[35,36]. Moreover, the observed stability of outcome representations, combined with stable sensory representations, might help shape motor coding on specific days, even in the absence of stable motor coding. A possible confound, however, is that the reward obtained at the end of the trial might drive activity in correct trials in the RSC. This makes the stable coding at the end of the trial difficult to interpret. Nonetheless, the encoding of trial outcome after the decision (and well prior to reward) indicates that trial outcome is encoded more robustly than motor choice.

Recent findings indicate that encoding of learned cognitive variables is dynamic, exhibiting different levels of representational stability[6,16,18]. 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[22,27,29,53–56]. 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 54-66 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[57], 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 (allowing comparison to raw fluorescence and activity map images). 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 ΔF/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.

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 97.5th percentile of the correlation coefficients from the time shuffled activity (1000 iterations), then the neuron was considered to exhibit reliable task-related activity and used for all subsequent analyses.

Task proficient 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 repetition to the lowest number of trials in each context for each neuron (100 iterations). In most cases, this resulted in neural responses resampled to the number of incorrect trials, although this was not always the case (see **Fig. 1SA**).

To measure stability in single-cell responses, we used 2 metrics: 1) the Pearson's correlation coefficient between odd trials in day 1 and even trials in days 1 to 5, and 2) the absolute change in the center of mass of responses in odd trials in day 1 and even trials in days 1 to 5. We calculated the center of mass as:

$$CoM = \frac{\sum_{i=1}^{n} x_i * i}{\sum_{i=1}^{n} x_i}$$

where the center of mass (CoM) is defined as the sum of neural activity (x) in the i-th time bin multiplied by the i-th time bin along a trial with length n, normalized by the total activity in that trial defined as the sum of (x).

The coding performance of task variables was determined by support vector machine classifiers [22]. 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) or s.e.m. were obtained by calculating the 5th and 95th percentiles, or the s.d. 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 significantly higher than chance for at least 5 consecutive bins (lower 5% CI of the bootstrapped distribution > 0.5).

Statistical information

All data groups were compared using nonparametric tests (Wilcoxon rank sum, or Wilcoxon signed-rank tests). Bootstrap estimates of s.e.m. were calculated as the s.d. of values evaluated in 100 to 1000 shuffled iterations, obtained by randomly sampling with replacement from the original values. Differences were considered significant if they were outside the 95th percentile of the bootstrapped distribution.

DATA AVAILABILITY

All processed imaging data used for making Figures 1-5 is available in the following repository: (*url to be posted upon acceptance*).

ACKNOWLEDGEMENTS

We would like to thank Ralf Wessel and Tina Xia for comments on manuscript.

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., 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.
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. 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).
13. 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).
14. 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).

15. 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).
16. 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).
17. Ziv, Y. *et al.* Long-term dynamics of CA1 hippocampal place codes. *Nat Neurosci* **16**, 264–266 (2013).
18. Hainmueller, T. & Bartos, M. Parallel emergence of stable and dynamic memory engrams in the hippocampus. *Nature* **558**, 292–296 (2018).
19. 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).
20. 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).
21. Alexander, A. S. & Nitz, D. A. Retrosplenial cortex maps the conjunction of internal and external spaces. *Nat Neurosci* **18**, 1143–1151 (2015).
22. Franco, L. M. & Goard, M. J. A distributed circuit for associating environmental context with motor choice in retrosplenial cortex. *Sci Adv* **7**, eabf9815 (2021).
23. Sun, W. *et al.* Context value updating and multidimensional neuronal encoding in the retrosplenial cortex. *Nature Communications 2021 12:1* **12**, 1–17 (2021).
24. 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).
25. Powell, A. *et al.* Stable Encoding of Visual Cues in the Mouse Retrosplenial Cortex. *Cerebral Cortex* **30**, 4424–4437 (2020).
26. Deitch, D., Rubin, A. & Ziv, Y. Representational drift in the mouse visual cortex. *Current Biology* **31**, 4327-4339.e6 (2021).
27. Chen, T. W., Li, N., Daie, K. & Svoboda, K. A Map of Anticipatory Activity in Mouse Motor Cortex. *Neuron* **94**, 866-879.e4 (2017).
28. Stringer, C. *et al.* Spontaneous behaviors drive multidimensional, brainwide activity. *Science (1979)* **364**, (2019).
29. 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).
30. Oh, S. W. *et al.* A mesoscale connectome of the mouse brain. *Nature 2014 508:7495* **508**, 207–214 (2014).
31. Frankland, P. W. & Bontempi, B. The organization of recent and remote memories. *Nat Rev Neurosci* **6**, 119–130 (2005).
32. Ranganath, C. & Ritchey, M. Two cortical systems for memory-guided behaviour. *Nature Reviews Neuroscience 2012 13:10* **13**, 713–726 (2012).
33. 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).

34. 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.
35. Singer, A. C. & Frank, L. M. Rewarded Outcomes Enhance Reactivation of Experience in the Hippocampus. *Neuron* **64**, 910–921 (2009).
36. 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).
37. Goard, M. & Dan, Y. Basal forebrain activation enhances cortical coding of natural scenes. *Nat Neurosci* **12**, 1444–1449 (2009).
38. Pinto, L. *et al.* Fast modulation of visual perception by basal forebrain cholinergic neurons. *Nat Neurosci* **16**, 1857–1863 (2013).
39. 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).
40. McGinley, M. J. J. *et al.* Waking State: Rapid Variations Modulate Neural and Behavioral Responses. *Neuron* **87**, 1143–1161 (2015).
41. Kuchibhotla, K. V. *et al.* Parallel processing by cortical inhibition enables context-dependent behavior. *Nat Neurosci* **20**, 62–71 (2016).
42. Morcos, A. S. & Harvey, C. D. History-dependent variability in population dynamics during evidence accumulation in cortex. *Nat Neurosci* **19**, 1672–1681 (2016).
43. Hwang, E. J., Dahlen, J. E., Mukundan, M. & Komiyama, T. History-based action selection bias in posterior parietal cortex. *Nat Commun* **8**, 1242 (2017).
44. Rule, M. E. *et al.* Stable task information from an unstable neural population. *Elife* **9**, 1–16 (2020).
45. 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).
46. 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).
47. Jacob, P.-Y. *et al.* An independent, landmark-dominated head-direction signal in dysgranular retrosplenial cortex. *Nat Neurosci* **20**, 173–175 (2017).
48. 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).
49. Fischer, L. F., Mojica Soto-Albors, R., Buck, F. & Harnett, M. T. Representation of visual landmarks in retrosplenial cortex. *Elife* **9**, (2020).
50. Vann, S. D., Aggleton, J. P. & Maguire, E. A. What does the retrosplenial cortex do? *Nat Rev Neurosci* **10**, 792–802 (2009).
51. 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).
52. 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).
53. 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).

54. Brennan, E. K. W. *et al.* Thalamus and claustrum control parallel layer 1 circuits in retrosplenial cortex. *Elife* **10**, (2021).
55. Goard, M. J., Pho, G. N., Woodson, J. & Sur, M. Distinct roles of visual, parietal, and frontal motor cortices in memory-guided sensorimotor decisions. *Elife* **5**, e13764 (2016).
56. Ramakrishnana, A. *et al.* Cortical neurons multiplex reward-related signals along with sensory and motor information. *Proc Natl Acad Sci U S A* **114**, E4841–E4850 (2017).
57. Aronov, D. & Tank, D. W. Engagement of Neural Circuits Underlying 2D Spatial Navigation in a Rodent Virtual Reality System. *Neuron* **84**, 442–456 (2014).

SUPPLEMENTAL INFORMATION

Fig. S1. Related to Fig. 1.

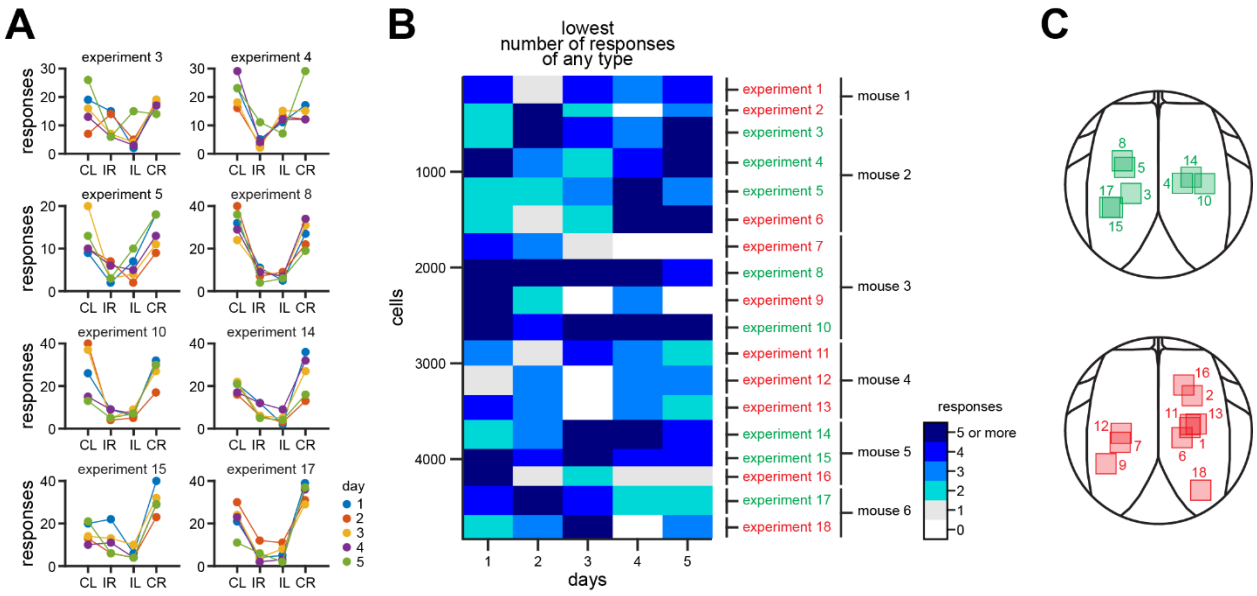

Fig. S1. Neurons with sufficient trials across 5 consecutive days. Related to Fig. 1.

A. Combinations of correct (C), incorrect (I), left (L), and right (R) responses in both contexts across 5 days for the experiments selected in B. Each day is shown in a different color. Blue, day 1; red, day 2; yellow, day 3; purple, day 4; green, day 5.

B. Heatmap showing the lowest number of responses in any of the 4 possible trial outcomes (CL, IR, IL or CR) for the 4,701 neurons recorded in all 5-day experiments across 6 mice. Only experiments with at least 2 responses in all 5 days were considered for further analyses (shown in green), unless otherwise indicated.

C. Recoding sites (425 x 425 μm) for the included (top, green), and excluded (bottom, red) experiments.

Fig. S2. Related to Fig. 1.

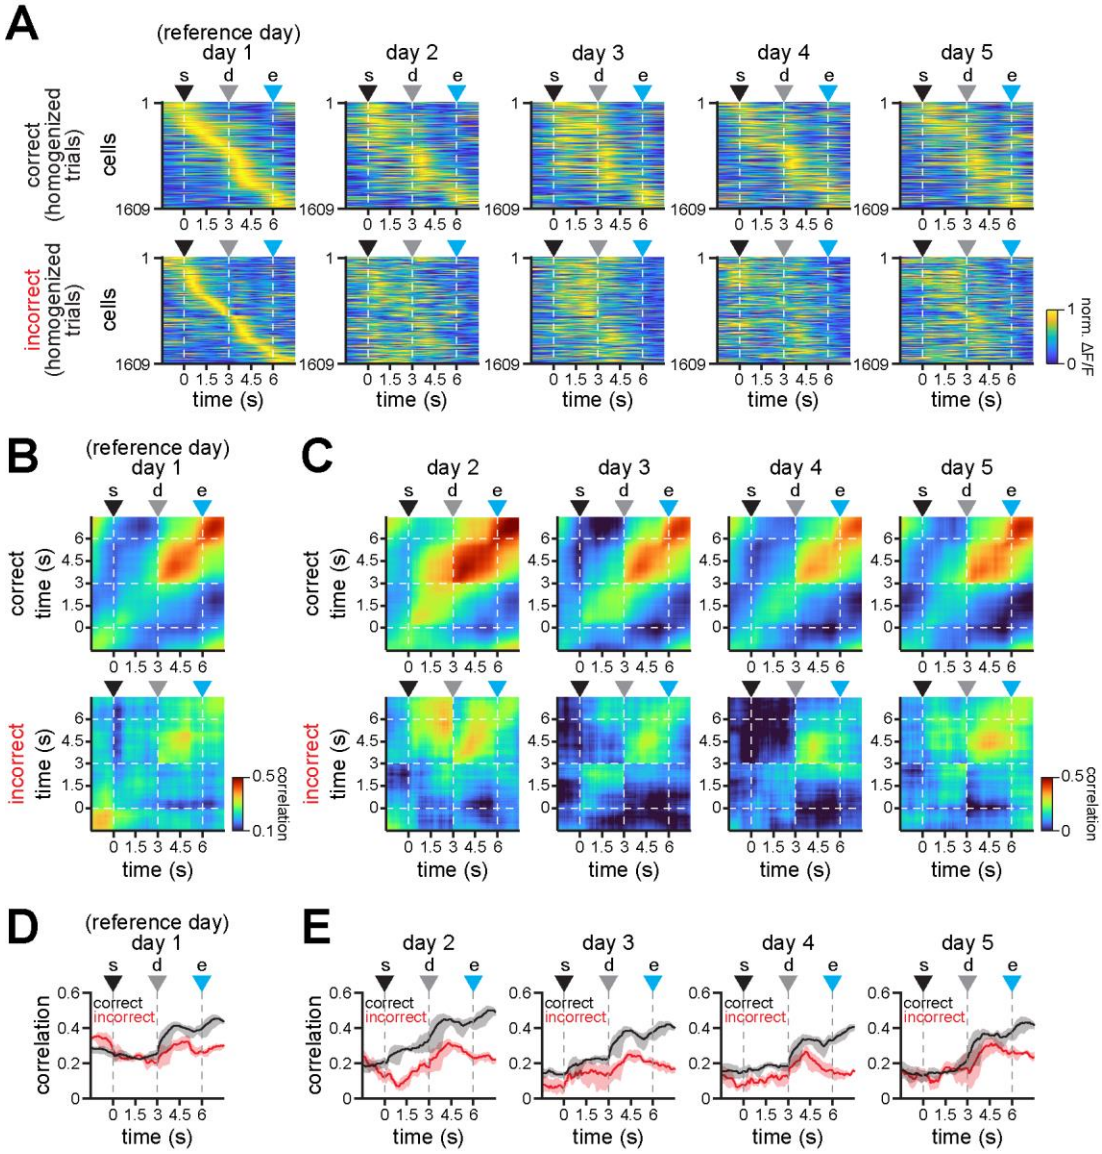

Fig. S2. Stability of population activity in equally sampled correct and incorrect trials. Related to Fig. 1.

A. Averaged randomly sampled activity (100 iterations) to the minimum number of trials in either correct or incorrect responses for each cell. Only responses to their preferred context are shown. Responses are sorted in one half of resampled correct (top row) or incorrect (bottom row) trials, and plotted for the other half of resampled trials. Note the thicker yellow ridge and the darker blue background in the population of correct responses, indicating higher stability across days.

B-C. Heatmaps showing the similarity of population activity vectors along trial duration as defined by Pearson's correlation coefficients. These matrices represent average correlation coefficients obtained within each of the 100 iterations during trial resampling of correct and incorrect trials as in A. Calculations

are cross-validated, i.e. population vectors in one halve of resampled trials in day 1 are compared to non-overlapping population vectors in the other halve of resampled trials in day 1 (B). In addition, all resampled trials in day 1 are compared to all resampled trials in days 2 to 5 (C). Note the higher similarity of population vectors in correct trials across days, consistent with our findings in the original dataset (**Fig. 1D-E**).

D-E. Pearson's correlation coefficient along the diagonal in B and C ± 1 second perpendicular to the diagonal. Note higher stability and similarity in population activity of correct trials. Solid line, mean; shaded area, 95% CI.

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

Fig. S3. Related to Fig. 1.

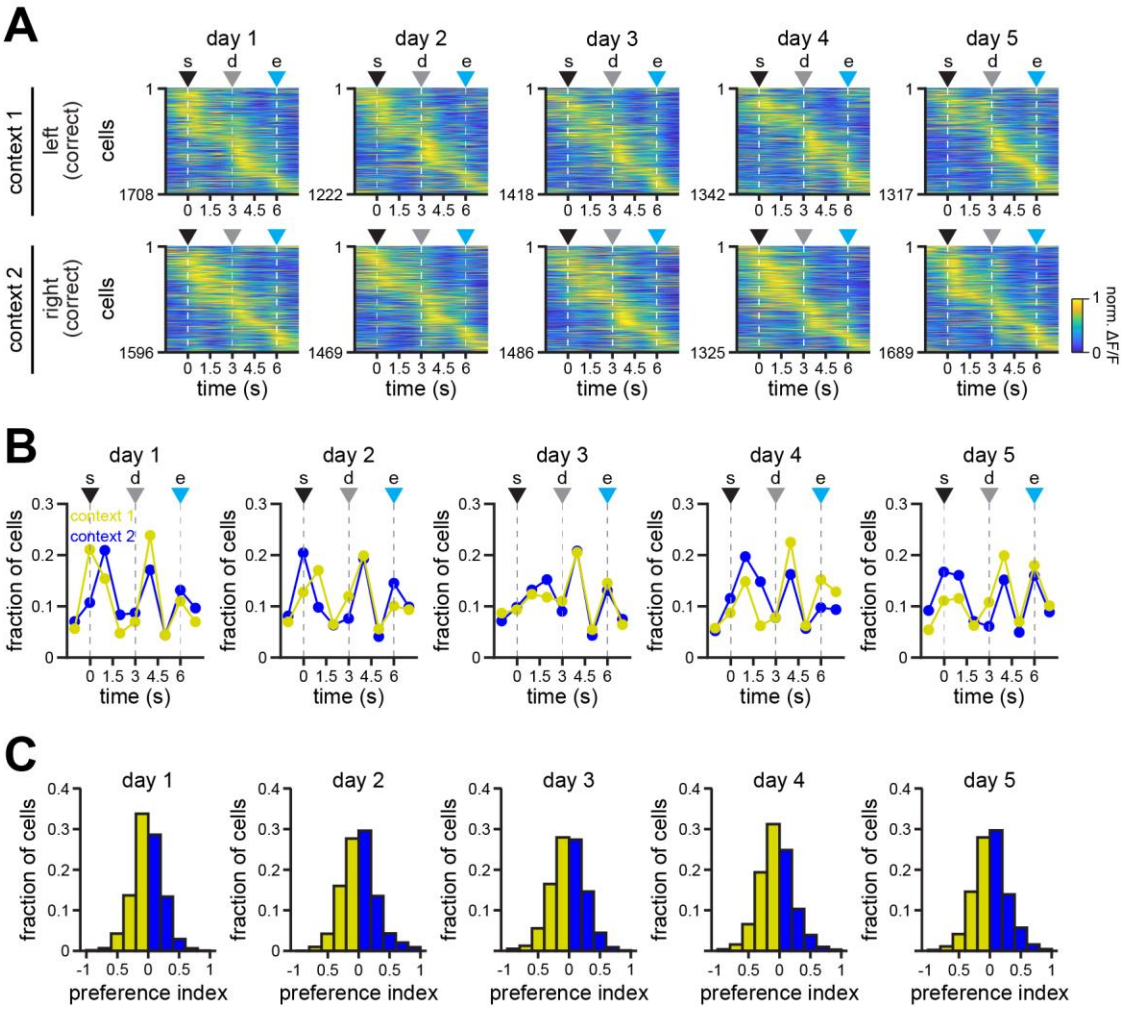

Fig. S3. Similar statistics in neural activity on independent days. Related to Fig. 1.

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 of the trial duration by each population of cells determined independently in each day.

B. Fraction of neurons with peak activity along trial duration for context 1 (yellow) and for context 2 (blue).

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

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

Fig. S4. Related to Fig. 2.

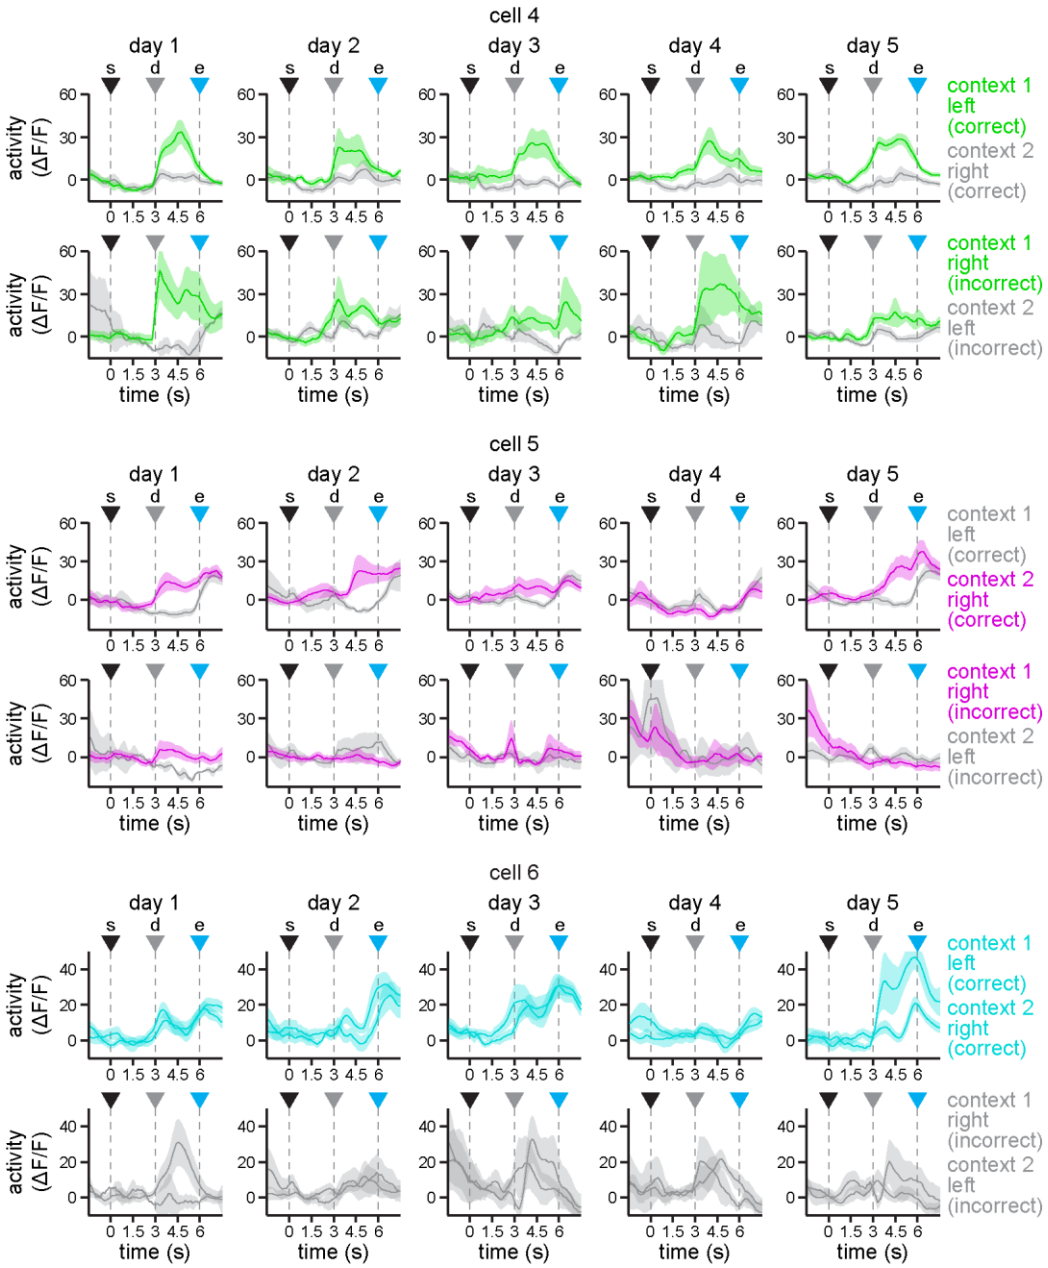

Fig. S4. Additional examples of RSC neuron responses during task performance. Related to Fig. 2. Example neurons exhibiting higher responses for context 2 (cell 4), right turns (cell 5), and correct trials (cell 6), respectively. Traces are color-coded to highlight the difference between the preferred 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. S5. Related to Fig. 3.

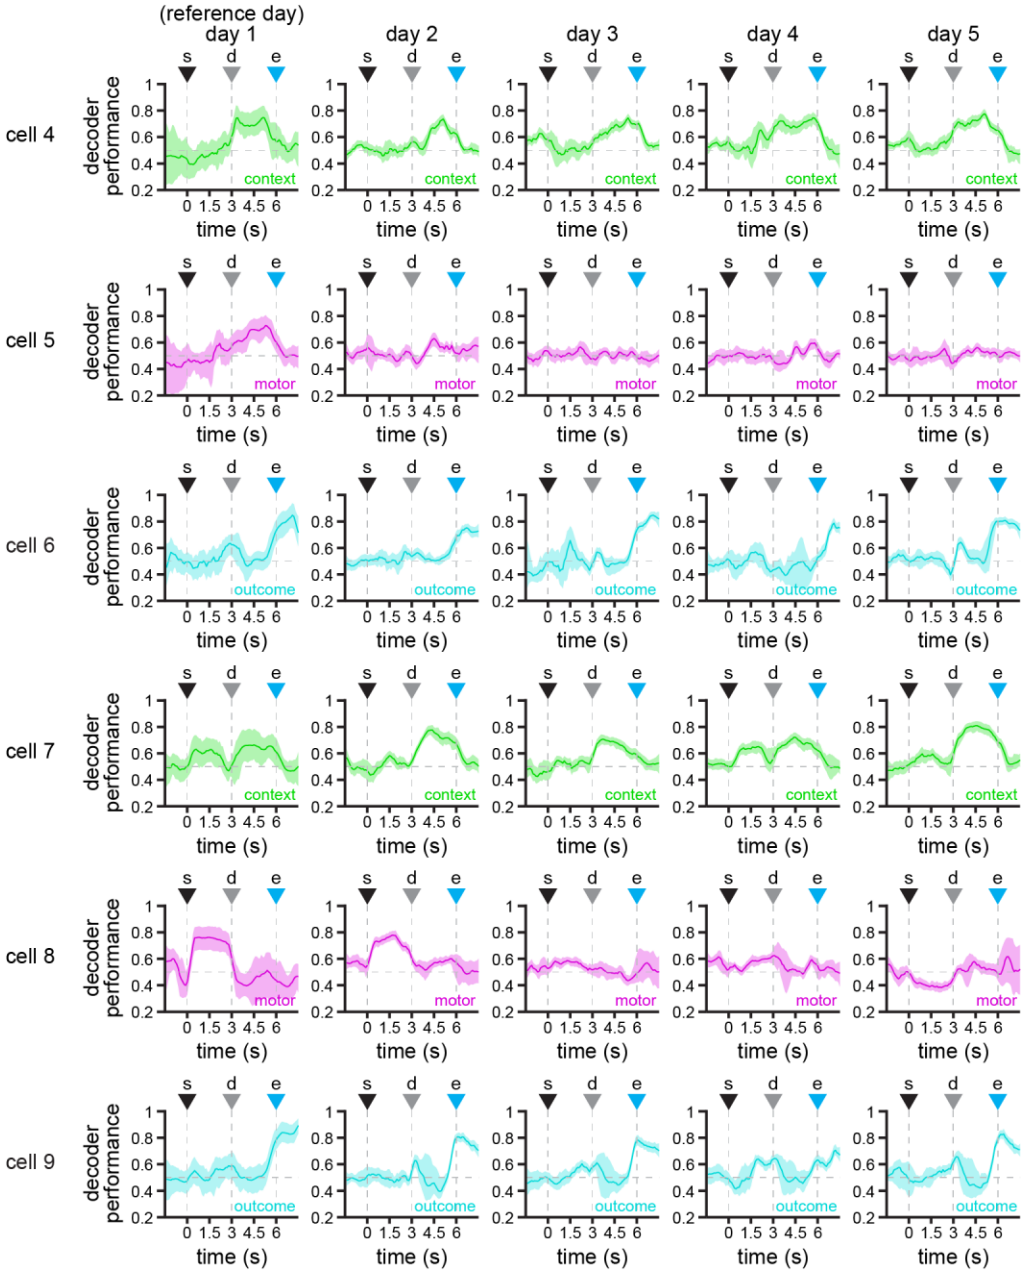

Fig. S5. Additional examples of single cell coding of task variables in RSC. Related to Fig. 3. Example cells with significant decoding of context (cells 4 and 7), motor choice (cells 5 and 8), and trial outcome (cells 6 and 9) in day 1. Solid line, mean; shaded area, 95% CI.

s = trial start; d = decision point; e = trial end.

Fig. S6. Related to Fig. 3.

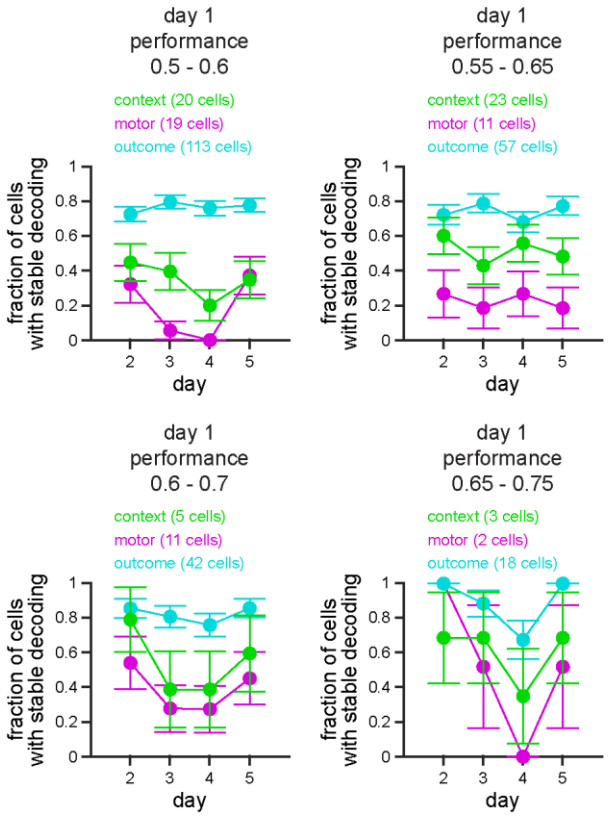

Fig. S6. The differential stability gradient is conserved in clusters of neurons defined by their decoding performance in day 1. Related to Fig. 3.

Fraction of cells that maintain significant decoding across days, clustered by their decoding performance in reference day 1 (mean ± bootstrapped s.e.m.). Their performance in day 1 is calculated as the average around ±0.5s of their peak performance. Note the gradient in decoding performance where trial outcome is more stable than environmental context, which in turn is more stable than motor choice.

Fig S7. Related to Fig. 3.

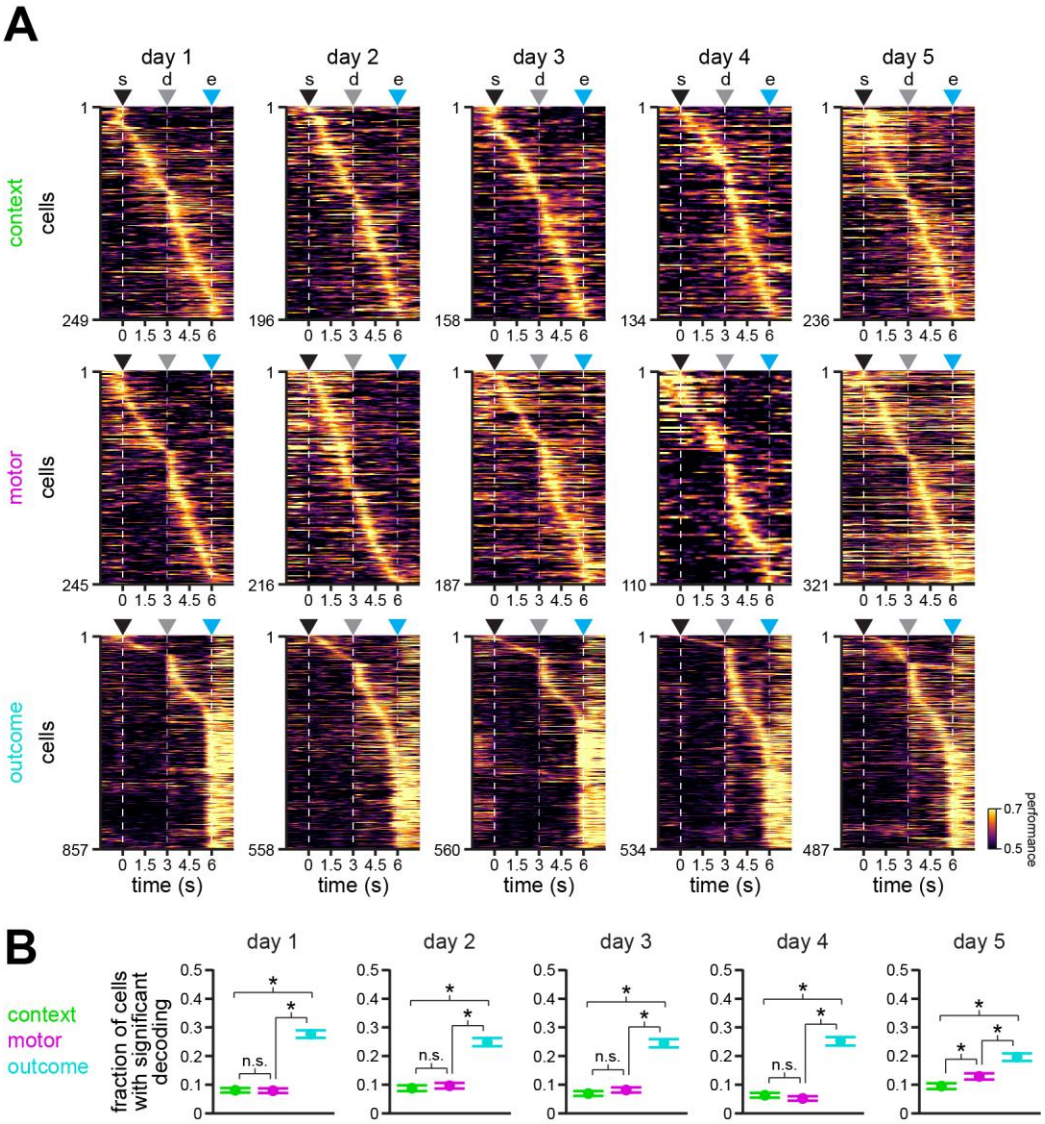

Fig. S7. Similar single-cell encoding of task variables on independent days. Related to Fig. 3.

A. Average performance of a support vector machine decoder on the activity of cells exhibiting significant decoding of environmental context (top row), motor choice (middle row), or trial outcome (bottom row). Models were computed and tested using activity within each day independently. Note similar information decoding for each task variable in each of the 5 days of experimentation. These subpopulations of cells with significant decoding were obtained from independently sampled populations showing reliable responses, and at least 2 trials in each of the 4 possible trial outcomes (correct left or right, incorrect left or right) in each day (day 1, n = 3,304 cells; day 2, n = 2,691 cells; day 3, n = 2,904 cells; n = 4, 2,667 cells; day 5, n = 3,006 cells). s = trial start; d = decision point; e = trial end.

B. Fraction of cells exhibiting significant decoding for environmental context, motor choice or trial outcome (mean ± bootstrapped 95% CI). Asterisks denote significantly higher fractions of cells. Significance was defined as lower 5% CI higher than higher 95% CI among task variables.

Fig. S8. Related to Fig. 4.

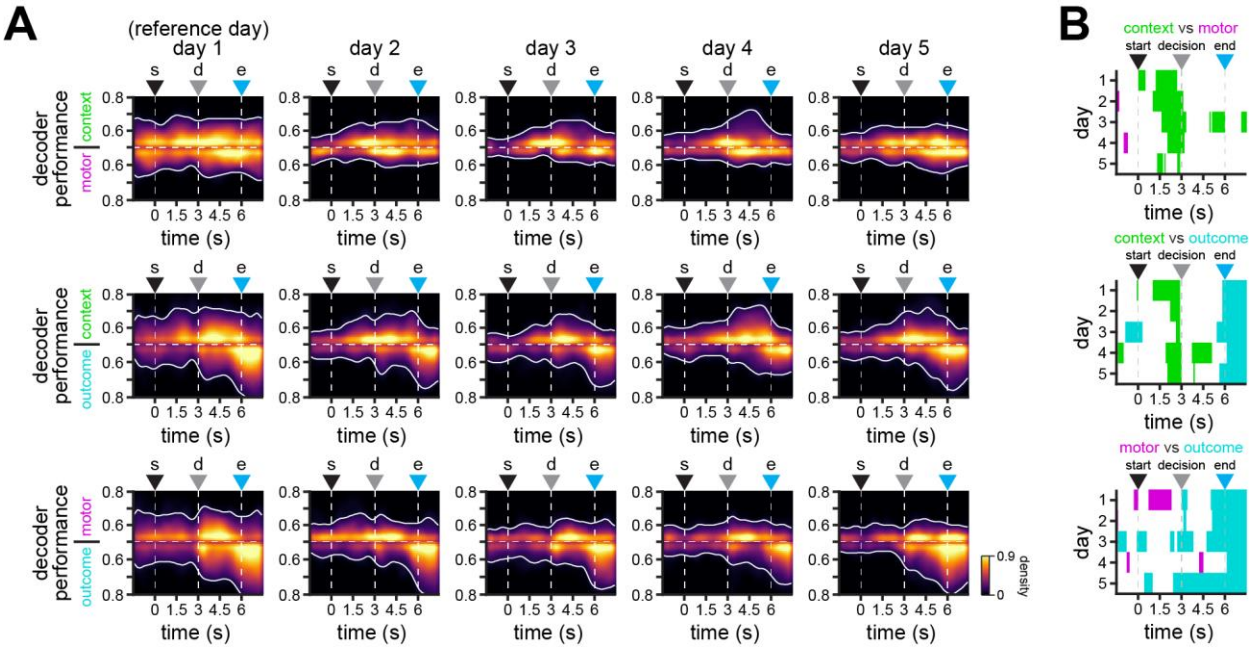

Fig. S8. Differential stability in multiplexing cells as a function of trial duration. Related to Fig. 4.

A. Heatmaps showing the density of the area under the curve of decoding performance for the 3 subpopulations shown in Fig. 4 A-B. 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 trial outcome are more stable than motor choice. White lines indicate >10% density.

B. Heatmaps showing significant differences between the encoding of the task variables indicated in each plot for the subpopulations of cells shown in Fig. 4 A-B (Wilcoxon rank-sum test). Everything shown in color (not white) is significantly different. Colors indicate the task variable decoded with better performance (green = context; magenta = motor; cyan = reward).

Fig. S9. Related to Fig. 5.

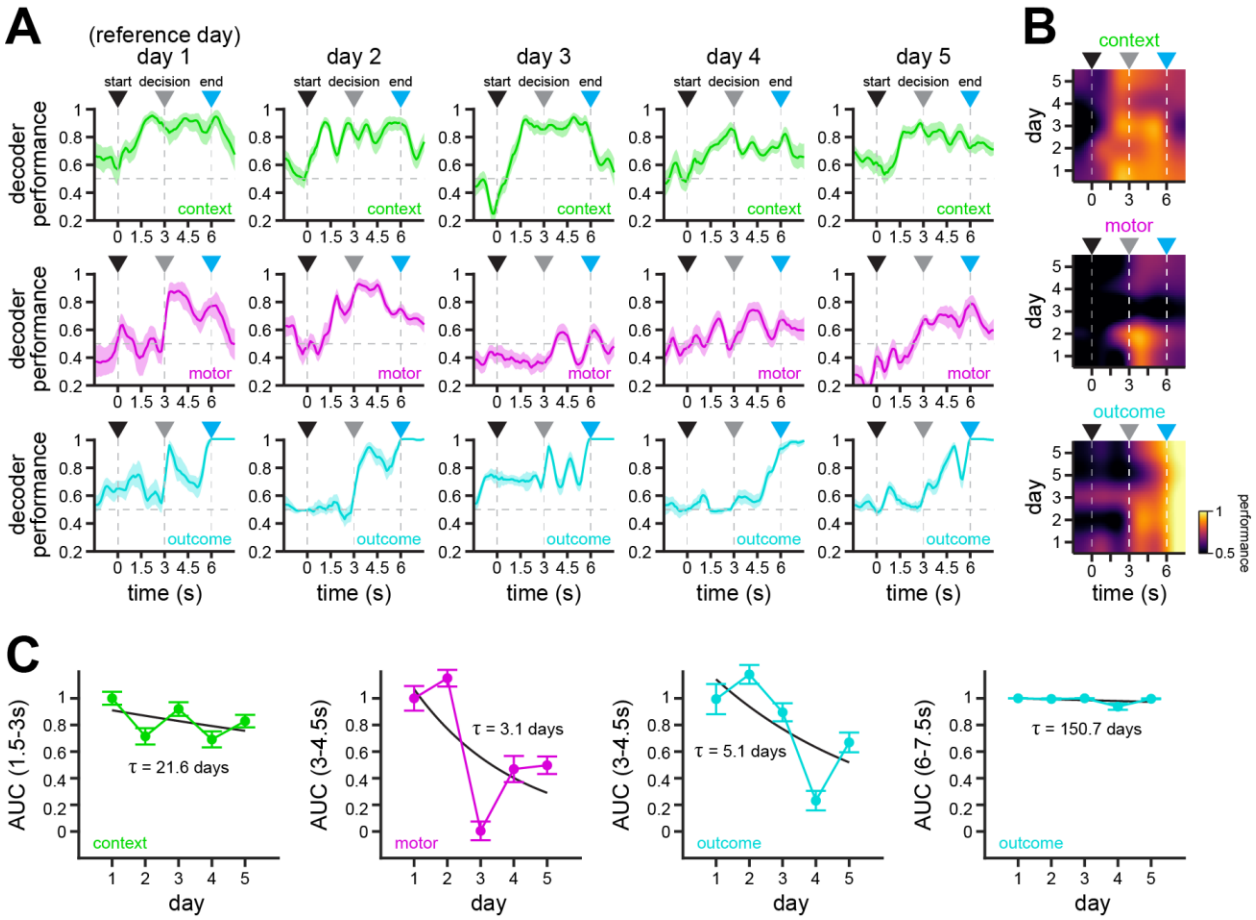

Fig. S9. Differential stability in population encoding of task variables in the RSC (1,609 cells). Related to Fig. 5.

A. Encoding of context (top row), motor choice (middle row), and trial outcome (bottom row) by all 1,609 cells shown in **Fig. 1** across 5 consecutive days of experimentation. Solid line, mean; shaded area, 95% CI. Note the gradient in population encoding, where outcome is more stable than context, which in turn is more stable than motor choice.

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

C. Area under the curve (AUC) for the indicated 1.5 s time windows relevant for each task variable (mean ± s.d.). An exponential decay function was fit to estimate decay in population encoding across days. Note the fast decay in motor choice encoding, and the slower decay in context and in trial outcome encoding around the decision point. Also note the very stable encoding of trial outcome after trial end.