A conserved mechanism for shaping visual perception by internally-generated activity in mice and humans

Anish Mitra[1,3]*, John Kochalka*, Chandan Kadur, Sean Quirin, Sam Vesuna, Tony Liu, Isaac Kauvar, Charu Ramakrishnan, Eun Young Choi, Adelaida Chibukhchyan[1,2], Lisa Yamada[1,2,4,5], Adam Fogarty[1,2,4,5], Pavithra Mukunda[1,3], Dariana Gil-Hernández[1,3], Kishandra Patron[1,2], Carolyn I. Rodríguez [1,3,6], Vivek Buch[1,4] and Paul Nuyujukian[1,2,4,5] for the Stanford Comprehensive Epilepsy Program, Shaul Druckmann, Surya Ganguli, and Karl Deisseroth[1,2,3,7]

Affiliations:

[1]Human Neural Circuitry program, Stanford University, Stanford, CA, USA.
[2]Department of Bioengineering, Stanford University, Stanford, CA, USA.
[3]Department of Psychiatry and Behavioral Sciences, Stanford University, Stanford, CA USA.
[4]Department of Neurosurgery, Stanford University, Stanford, CA, USA.
[5]Department of Electrical Engineering, Stanford University, Stanford, CA, USA.
[6]Veterans Affairs Palo Alto Health Care System, Palo Alto, CA, USA.
[7]Howard Hughes Medical Institute, Stanford University, Stanford, CA, USA.
*These authors contributed equally to this work.

* Equal contribution

[t] Corresponding author: deissero@stanford.edu

Abstract

Perception can be influenced not only by external stimuli but also by the brain's internally-generated activity. We sought to identify principles– conserved across mammalian phylogeny– that govern the internally-driven shaping of perception by designing a behavioral report of visual perception equally accessible to both mouse and human subjects alongside brainwide recording of internal representations in both species. A cortex-spanning activity screen using widefield Ca^{2+} imaging in mice revealed that spontaneous activity specifically in visual and retrosplenial cortices predicted, before any stimulus, subsequent visual detection choices, and two-photon imaging revealed single-cell fluctuations in spontaneous representational content that predicted upcoming perceptual decisions in these same regions. Upon presentation of task epochs characterized by different likelihood of target visual signals, the representational content of spontaneous activity shifted to reflect these contextual changes, indicating that spontaneous activity in this circuitry adaptively represented upcoming expected input statistics. Precisely-timed projection-specific optogenetic modulation revealed a causal role for retrosplenial cortex in organizing this spontaneous activity, conveying top-down expectation to the visual cortex to guide behavior. Providing a link to clinically-significant states of altered perception, neuromodulation with LSD correspondingly altered spontaneous activity in the same regions (visual and retrosplenial cortices), causing shifts in perceptual choice favoring report of target signals that were not in fact present. Brainwide recording in human subjects with deep intracranial electrodes during the same task revealed a similar and specific pattern of spontaneous retrosplenial activity in the human brain that predicted perceptual choice. Thus spontaneous cortical activity, in particular through top-down signaling from retrosplenial cortex, plays a crucial role in shaping perceptual outcomes in a manner that is conserved across the mammalian lineage, and that is adaptively influenced by environmental statistics as well as by clinically-significant altered states.

Introduction

Sensory inputs to the brain inform how animals perceive their environment, but perception does not simply mirror the world; the brain itself also shapes perception through internally generated processes. For example, spontaneous neural activity—that is, brain activity not directly driven by external sensory stimuli—is hypothesized to encode predictive models of the world, potentially influencing perceptual decisions upon integration with sensory-evoked responses[1-6]. However, despite extensive elucidation of stimulus-evoked mechanisms that directly drive sensory perception from primary sensory cortices of the mammalian brain[7, 8], insight remains limited into the origins and perceptual significance of spontaneous activity.

Potential adaptive or maladaptive roles of spontaneous activity continue to be debated. One model posits that spontaneous activity actively shapes sensory perception[9-11], as evidenced by findings suggesting that spontaneous activity may adapt to environmental statistics[12-15] and may modulate perceptual decisions[16-20]. But conversely, other evidence suggests that spontaneous activity is largely decoupled from perceptual processes, and is associated more with variables such as arousal and movement than with cognitive factors[21-24]. These contradictory findings could highlight species-specific differences in neural processes relating to spontaneous activity[25], as well as the variable impact of experimental conditions such as the presence or absence of structured behavioral tasks[26].

While some species-specificity in any role of spontaneous activity may be expected given the broad scope of mammalian brain evolution, we hypothesized the existence of conserved functionality for this potentially highly adaptive feature of mammalian brain activity. To identify such an ancestral function, we designed both two-alternative forced choice (2AFC) visual detection behaviors and brain-wide activity screens suitable for both mouse and human subjects, and carried out a comprehensive investigation into the relationship between spontaneous cortical activity and visual perceptual decision-making in both species. We also were able to leverage unique advantages of the mouse system by applying two-photon imaging to resolve relevant signals at the single-cell level, and by applying projection-specific optogenetics to establish causal significance of the mechanisms revealed by brainwide screening. Remarkably, findings in both species in the same structured task pointed to the same conserved circuit dynamical mechanism for spontaneous activity modulation of visual perception; moreover, providing a computationally-guided task design revealed adaptive functionality of this modulated perception, and pharmacological modulation with a clinically significant agent (LSD) revealed shifted spontaneous activity in the same conserved circuitry leading to corresponding specific perceptual errors. Together, these findings suggest an ancestral role in the mammalian lineage for spontaneous (that is, pre-stimulus) activity in integrating contextual information to generate predictive sensory detection, offering critical insights into the cognitive control processes that guide perception in mammals.

Results

Visual detection task compatible with cortex-wide activity screen

To investigate the influence of pre-stimulus cortical state on perceptual choice, we began with widefield Ca^{2+} imaging to record multiregional neural activity across dorsal cortex in mice expressing GCaMP8m under the CaMKIIα promoter, with robust expression in layer 2/3 and layer 5 principal neurons (Fig. 1a-b). Mice performed a two-alternative forced choice (2AFC) visual detection task while head-fixed (Fig. 1c-1d). In each trial, mice were presented with one of two visual stimuli for four seconds: either pure visual white noise or a drifting grating embedded in white noise. Mice reported perceptual choice by licking one of two equidistant waterspouts. Correct choices (hits: correct grating identification; correct rejects: correct noise identification) were rewarded with water, while incorrect choices (false alarms: incorrect identification of noise as grating; misses: incorrect identification of grating as noise) resulted in no reward and a subsequent black screen before the next trial.

Mice were trained over several weeks, starting with high-contrast gratings that were easy to detect and gradually progressing to low-contrast gratings. Performance stabilized at approximately 65% accuracy, with no significant bias between false alarms and misses (Extended Data Figs. 1-2).

Relationship between task difficulty and behavioral response selectivity

The simple two-alternative forced choice tasks require a behavioral response regardless of confidence in the perceptual decision, making high- and low-confidence reports indistinguishable without additional task structure. To estimate decision confidence, we introduced a one-second lick selectivity period between visual stimulus onset and reward assessment (Fig. 1d). Mice typically began licking shortly after perceiving the visual stimulus (Fig. 1c). If a mouse licked only one spout during the lick selectivity period and its perceptual report was correct, the lick selectivity was scored as +1; if only a single spout was licked and the response was incorrect, lick selectivity was scored as -1. To account for trials where both spouts were licked, we calculated lick selectivity as the proportion of licks during the one-second period that corresponded to the mouse's final choice. Trials with fewer than five licks during this period were excluded. This calculation produced values ranging from -1 to +1, with positive values for correct choices and negative values for incorrect ones (Fig. 1c). Across all trials, lick selectivity exhibited three distinct modes: +1 for correct single-spout responses, -1 for incorrect single-spout responses, and near zero for near-equal licking of both spouts (Fig. 1e). Based on this distribution, "high selectivity" trials were defined as those with $|\text{lick selectivity}| \geq 0.75$, while values below 0.75 were categorized as "low selectivity" trials.

To test whether lick selectivity index could serve as a proxy of decision confidence, we analyzed lick selectivity modulation by task difficulty (with difficulty increased by reducing contrast of the grating stimulus; Fig. 1f). We found that the magnitude of lick selectivity declined monotonically

with increasing task difficulty, consistent with prior estimates of decision confidence in 2AFC tasks[27]. Additionally, examining the distribution of high selectivity perceptual reports across task difficulties revealed that in the high contrast (less difficult) task condition, nearly all high selectivity responses were correct (Fig. 1g), whereas in the low contrast (more difficult) task condition, nearly a third of high selectivity responses corresponded to perceptual errors.

Pre-stimulus cortical state in RSC and V1 predicts perceptual choice

We next employed a brain-spanning screen to investigate whether pre-stimulus activity state in any cortical area exhibited a predictive association with perceptual choice. To do so, we computed unbiased pixel-wise statistical relationships between macroscale neural activity across the brain in the 250 ms preceding visual stimulus onset and subsequent perceptual choice in high selectivity trials (Fig. 1h-1i; $n = 5$ mice). We discovered pre-stimulus activity specifically localized to primary visual cortex (V1) and retrosplenial cortex (RSC), which both exhibited consistent associations with choice: responses reporting grating perception (arising from correct hits and false alarms) were preceded by increased activity in both RSC and V1, while responses reporting no grating perception (noise responses, arising from correct rejections and erroneous misses) were primarily preceded by V1 activity. Analysis of the post-stimulus epoch revealed that these regions were also recruited during task-evoked motor activity (Fig. 1j; $n = 5$ mice), but in this case recruitment was nonspecific along with much of the rest of cortex.

We next tested whether pre-stimulus neural activity in V1 and RSC could predict perceptual choice. Using regions of interest derived from the thresholded and symmetrized t-statistic maps (Fig. 1i), we applied a two-timeseries model with one timeseries each from bilateral RSC and bilateral V1. In high selectivity trials, RSC and V1 activity patterns prior to stimulus onset allowed prediction of perceptual choice with 64% accuracy (Figs. 1l-1n, high selectivity; $n = 5$ mice, classification performed in 250ms windows in the ITI; permutation test, $p < 0.001$). In contrast, no predictive decodability was observed in low selectivity trials (Figs. 1l-1n, low selectivity; $n = 5$ mice, classification performed in 250ms windows in the ITI; permutation test, $p > 0.05$). Together, these results indicate that pre-stimulus neural activity in RSC and V1 predicts visual perceptual decision-making, but only when the behavioral response exhibits the measured signature of decision confidence (high selectivity).

Pre-stimulus activity of single neurons in RSC and V1 predicts perceptual choice

Having found that pre-stimulus bulk activity in V1 and RSC predicted perceptual choice, we next investigated whether pre-stimulus single-neuron resolution representations in V1 and RSC could predict subsequent perceptual decisions. Using two-photon microscopy, we measured single-cell activity in cortical neurons in RSC and V1 (Fig. 2a, see Methods) on alternating days during the visual detection task, capturing activity from an average of 920 neurons (range: 630–1270) across three imaging planes per session (two in layer II/III and one in layer V).

Univariate analysis identified neurons with evoked responses selective for either noise or grating stimuli (Fig. 2b). In V1, neurons were approximately evenly split between noise- and grating-selective responses, while RSC showed a marked enrichment of grating-selective neurons, independent of spout assignment (Fig. 2c; Extended Data Figs. 3-4; $n = 6$ mice, paired t-test, $p < 0.001$). We next tested whether pre-stimulus activity in these response-selective neurons predicted perceptual choice. In high-contrast (easy) trials, no significant relationship was observed. However, in low-contrast (difficult) trials, pre-stimulus single neuron activity was significantly linked to high selectivity behavioral responses in both V1 and RSC (Extended Data Figs. 3-4).

Pre-stimulus variation in population representations shift behavioral bias

Building on our single neuron findings, we next examined whether pre-stimulus activity at the population level, defined by cellular-resolution ensembles, predicted perceptual choice. Using a coding dimension analysis[28], we computed the mean evoked-activity vector in population activity space corresponding to high selectivity hits and correct rejects. The zero-centered difference between these vectors defined the coding dimension separating grating- and noise-responses (Fig. 2d). Activity along the coding dimension during error trials revealed that false alarms aligned with hits, while misses aligned with correct rejections (Extended Data Fig. 6), indicating that the evoked coding dimension generalized across both correct and error responses.

Single-neuron weights contributing to the coding dimension exhibited a heavy-tailed distribution, with a strong skew toward grating-selective neurons, in RSC but not V1 (Extended Data Fig. 5a-5d). When examining the laminar distribution of highly weighted neurons (90th percentile or above in coding dimension variance), we observed significant enrichment in layer V of RSC but no laminar differences in V1 (Fig. 2e; $n = 6$ mice, paired t-test, $p < 0.01$). These results reveal that while V1 equally represents grating- and noise-responses across cells and layers, RSC activity favors representation of grating-responses with deeper-layer activity.

Crucially, although the coding dimension was defined based on post-stimulus activity during the lick-selectivity period, we observed substantial pre-stimulus variation in the coding dimension (Fig. 2f, Extended Data Fig. 5e). Testing whether pre-stimulus coding dimension activity predicted perceptual choice, we found significant differences in low contrast (difficult) trials for high selectivity behavioral responses, with no differences in high contrast (easy) trials (Fig. 2g-2l; $n = 6$ mice, paired t-test, $p < 0.01$ in high-selectivity trials). Pre-stimulus coding dimension activity allowed perceptual choice predictions with 68% accuracy in V1 and 62% in RSC during high selectivity trials (Figs. 2i, 2j; $n = 6$ mice, paired t-test, $p < 0.001$).

Several key control experiments supported the specificity of these findings. First, in mice that failed to learn the visual detection task but exhibited no response bias, pre-stimulus coding dimension activity did not predict perceptual choice (Extended Data Fig. 7). Second, regressing out the influence of motor behavior during the pre-stimulus period only modestly reduced decoding accuracy (Extended Data Fig. 8) and coding dimension activity offered stronger predictive power than single-dimensional total activity in V1 and RSC (Extended Data Fig. 9).

Finally, we verified that this above-chance pre-stimulus prediction of perceptual choice did not imply a relationship between pre-stimulus choice predictions and the actual visual stimulus identity (Extended Data Table 1); in other words, neural activity was not predicting the upcoming stimulus delivered by the experimental setup, but rather was successfully predicting the animal's upcoming perceptual choice.

In signal detection tasks, performance is governed by two key factors, as described by signal detection theory: sensitivity (d prime) and bias (Fig. 2m). Sensitivity measures the ability to distinguish between signal and noise, while bias reflects the threshold for deciding whether a stimulus is present. A lower threshold favors reporting a stimulus, while a higher threshold reduces such reports. To assess whether pre-stimulus coding dimension activity relates to sensitivity, bias, or both, we grouped trials into six quantiles based on coding dimension activity and calculated behavioral d prime and bias for each quantile (Figs. 2m-2n). Linear regression revealed a significant positive relationship between pre-stimulus coding dimension activity and bias (V1 slope = 0.28, $p < 0.01$; RSC slope = 0.31, $p < 0.01$), indicating that higher pre-stimulus population activity corresponds to an increased tendency to report the presence of a signal. In contrast, the slopes for behavioral d prime were indistinguishable from zero (V1 slope = -0.01, $p > 0.5$; RSC slope = -0.03, $p > 0.5$), suggesting no systematic relationship between pre-stimulus coding dimension activity and sensitivity to the signal. These findings reveal that pre-stimulus population representations modulate perceptual bias without altering sensitivity, pointing to a specific role in shaping decision-making.

Environmental statistics modulate pre-stimulus population representations

This robust relationship between pre-stimulus activity and perceptual bias raises the question of whether spontaneous representations adaptively shift to reflect the statistics of a changing environment. To investigate whether spontaneous neural representations adapt to environmental contexts, we recorded two-photon neural and behavioral data during 10-minute enriched-context sessions where mice were presented with either 75% grating or 75% noise trials (Fig. 3a). To prevent the formation of persistent behavioral biases, each enriched session was bracketed by baseline sessions with equal proportions of grating and noise stimuli.

Behavioral analysis revealed that mice began enriched sessions with no bias but gradually developed a response bias matching the enriched stimulus, becoming more likely to report seeing gratings in grating-enriched sessions and vice versa (Fig. 3b,c; $n = 6$ mice, paired t-test, $p < 0.01$). Neural coding dimension analysis further demonstrated that pre-stimulus neural representations in both V1 and RSC progressively shifted toward the enriched stimulus representation over the course of each session (Fig. 3d-3g; $n = 6$ mice, paired t-test, $p < 0.01$). These findings indicate that both behavioral bias and pre-stimulus neural activity remarkably adapt to dynamically reflect the statistical probabilities of visual stimuli within the current environmental context.

Recent response history modulates pre-stimulus population representations

Building on the link between pre-stimulus coding dimension activity and environmental context, we explored whether recent response history influences subsequent pre-stimulus coding dimensions. Signal detection theory predicts that the bias criterion will become stricter following false alarms and more lenient after misses (Fig. 2m, 3h). To test the limits of signal detection principles that could be in play, we analyzed trials immediately following high-selectivity errors (after even single outcomes).

While behavioral analysis revealed no statistically significant bias following single high-selectivity errors (Fig. 3i,j), neural data showed a significant shift in pre-stimulus coding dimension activity after these errors (Fig. 3k-n; $n = 6$ mice, paired t-test, $p < 0.05$). However, this shift was approximately 10 times smaller than the shifts observed in enriched contexts under these conditions, and did not result in detectable changes in behavioral bias following single high-selectivity errors. Nevertheless, even a single high-selectivity error was found to modulate pre-stimulus neural activity in a manner consistent with signal detection theory, consistent with a mechanism by which response history effects could achieve behavioral potency over longer timescales.

RSC exhibits a causal role in adaptive behavioral bias

These results indicate that pre-stimulus environmental context-modulated activity in RSC and V1 influences visual perceptual decisions. Prior work has suggested that RSC could contribute to detecting contexts[29] and providing top-down input to V1[30]. To test whether pre-stimulus input from RSC to V1 shapes bias in coding dimension activity and behavior, we selectively targeted activity in projection neurons from RSC to V1 using an intersectional genetic strategy (Fig. 4a,b). We injected AAVretro-Cre virus into V1 to express the Cre recombinase in neurons projecting to V1, and injected the Cre-dependent virus AAV-DIO-KCR1 into RSC (these injections together give rise to expression of the inhibitory K^+-conducting opsin KCR1 in RSC neurons projecting to V1). As a brain-region control, we injected AAV-KCR1 instead into somatosensory cortex (S1) in a separate cohort of mice.

Temporal specificity was tested with two inhibition protocols: (1) *early*, targeting RSC projection neurons before stimulus onset (-1.75 to -1.25s); and (2) *late*, targeting the same neurons closer to stimulus onset (-0.75 to -0.25s). Optogenetic KCR recruitment was achieved by delivering a 500ms pulse of 617 nm one-photon illumination through a fiber optic patch cable positioned approximately 1 mm above the cortex and affixed to the inner wall of a light-blocking cone, targeting either the retrosplenial or somatosensory cortex. Concurrently, two-photon imaging alternated between RSC and V1 during task performance (Fig. 4c–e). Visual stimuli were restricted to the eye contralateral to the inhibited hemisphere to ensure unilateral control and caused no behavioral effects after one habituation session ($n = 9$ mice, paired t-test, $p > 0.8$). Late inhibition of RSC→V1 projection neurons caused a statistically significant shift in pre-stimulus coding dimension activity toward noise representations in both RSC and V1 (Fig. 4d; mean coding dimension activity in the 250ms preceding visual stimulus onset, $n = 6$ mice, paired t-test, $p <$

0.01). This effect was absent with early inhibition of RSC and late inhibition of S1 (Fig. 4c, 4e). A transient shift toward noise representations was also observed in RSC and V1 immediately after early RSC inhibition but returned to baseline prior to stimulus onset (Fig. 4c); in contrast, not even a transient shift was seen with inhibition of S1 (Fig. 4e). These results reveal that late inhibition of RSC cells projecting to V1 disrupts the balance of coding dimension neural activity in favor of noise responses, consistent with a specific role for RSC cells in naturally promoting detection of anticipated signal responses in visual circuitry.

We next tested for modulation of perceptual behavior under these conditions. Behavioral analyses revealed no significant changes in bias or d-prime for the control conditions, including with no inhibition, with early RSC inhibition, or with late S1 inhibition (Fig. 4f-h; $n = 9$ mice for no optogenetic inhibition, $n = 3$ mice for early optogenetic inhibition, $n = 3$ mice for S1 optogenetic inhibition, t-test, $p > 0.5$ with Bonferroni correction for 3 comparisons). In contrast, late inhibition of RSC→ V1 projection neurons significantly biased perceptual decisions toward noise responses, reducing both false alarms and hit rates (Fig. 4f-h; $n = 9$ mice for no optogenetic inhibition, $n = 6$ mice for late optogenetic inhibition, t-test, $p < 0.05$ with Bonferroni correction for 3 comparisons). Control analyses confirmed there were no significant differences in evoked (not pre-stimulus) coding dimension responses between trials with and without optogenetic inhibition (Fig. 4j). Moreover, the relationships between pre-stimulus coding dimension and both behavioral bias and d-prime under unperturbed conditions (Fig. 2n) were causally recapitulated by late RSC inhibition, albeit with a shift toward noise representations (Fig. 4k; V1 bias slope = 1.15, $p < 0.01$; RSC bias slope = 0.94, $p < 0.01$, V1 d prime slope = 0.03, $p > 0.5$; RSC d prime slope = 0.01, $p > 0.5$). Finally, we conducted a high-contrast (easy task) control to test whether ambiguity of the visual stimulus could modulate the relationship between RSC inhibition and behavioral bias. We found that while late RSC inhibition consistently caused a shift in coding dimension activity toward noise representations, behavioral bias was not affected in this easy task (Extended Data Fig. 10).

Taken together, our results suggest that inhibiting RSC→V1 communication shifts network representations and specifically biases perceptual decisions, in a manner consistent with the observation that the vast majority of single response-selective neurons in RSC were tuned to the grating response (Fig. 2c). Inhibition of RSC projection neurons suppressed grating representation and perception (whether or not gratings were actually present), and these effects initiated in RSC were propagated to V1. The relationship between optogenetically shifted pre-stimulus activity and behavior notably manifested only under challenging signal-detection conditions, as no behavioral bias was observed when the identity of the visual stimulus was unambiguous (Extended Data Fig. 10).

LSD shifts behavioral bias toward grating responses

Perceptual errors, particularly false alarms in sensory detection tasks, are commonly reported under the influence of psychedelics[27] including LSD[31], an effect with highly consequential clinical and public-health outcomes. To investigate whether this clinically-relevant intervention induces

perceptual errors in our visual detection task and to explore the role of pre-stimulus brain activity in these errors, we conducted two-photon neural recordings in mice engaged in the task before, during, and after administration of either saline or LSD (0.3 mg/kg).

We first recorded 25 baseline trials prior to injection. Saline or LSD was then administered while mice remained in the experimental apparatus. After a 15-minute post-injection interval to allow drug effects to set in, we resumed neural imaging and collected ~100 trials per mouse (trial number varied slightly depending on task engagement). Finally, a 20-minute pause was introduced before initiation of the post-injection testing epoch, based on pilot experiments showing LSD-induced behavioral effects normalize within this timeframe.

We discovered that LSD induced a pronounced bias toward grating reports in perceptual decision-making (Fig. 5a-b, $n = 4$ mice, paired t-test, $p < 0.001$). This bias was absent in saline conditions and returned to baseline in the post-LSD session. Additionally, LSD was found to lengthen reaction times (Fig. 5c, $n = 4$ mice, paired t-test, $p < 0.01$), decreased overall motion (Fig. 5d, $n = 4$ mice, paired t-test, $p < 0.01$), and lowered performance accuracy (Extended Data Fig. 11). To investigate how LSD might alter pre-stimulus neural activity in the course of exerting these pronounced behavioral effects favoring perception (whether or not accurately so), we next examined the correlation structure of spontaneous neural activity in V1 or RSC.

LSD shifts V1 pre-stimulus activity toward grating representations

We observed no significant changes in the overall correlation structure of single neuron activity in either cortical area (Fig. 5e,f; Extended Data Fig. 11; $n = 4$ mice, paired t-test, $p > 0.5$). However, analysis of coding dimension activity revealed that LSD significantly shifted the pre-stimulus distribution of coding dimension values in V1 toward the grating representation, without altering the amplitude of evoked responses (Fig. 5g-h; mean coding dimension activity, $n = 4$ mice, paired t-test, $p < 0.01$). In contrast, LSD reduced the amplitude of evoked responses in RSC but did not significantly affect RSC pre-stimulus coding dimension activity (Fig. 5i-j; $n = 4$ mice, paired t-test, $p > 0.5$). Thus, LSD may act by selectively shifting V1 pre-stimulus activity toward perceptual representations, acting in primary sensory cortex while leaving RSC pre-stimulus activity unaffected.

To further examine the link between pre-stimulus coding dimension activity and the observed behavioral bias, we trained a linear classifier on stimulus-evoked coding dimension activity and used this classifier to predict perceptual decisions based on pre-stimulus activity. Under LSD, the V1 decoder performed significantly above chance (72% accuracy, Fig. 5k) and exhibited a strong bias toward grating responses (Fig. 5l; $n = 4$ mice, paired t-test, $p < 0.001$). In contrast, RSC decoder performance remained at chance levels and showed no significant bias (Fig. 5m, $n = 4$ mice, paired t-test, $p > 0.5$). These results indicate that the grating bias induced by LSD arises from shifts in V1 pre-stimulus activity, while pre-stimulus activity in RSC becomes decoupled from perceptual choice.

In the absence of pharmacological intervention, we had found that following high-selectivity misses, activity along the pre-stimulus coding dimension shifts toward grating representations, and toward noise representations after high-selectivity false alarms (Fig. 3e-3n). To test this dynamic under LSD, we examined post-error shifts in pre-stimulus coding dimension activity. Under saline administration, pre-stimulus activity showed significant shifts after high-selectivity errors, consistent with response history effects and signal detection theory principles. However, under LSD administration, the shifts were abolished in both V1 and RSC (Fig. 5n,o). These results show that LSD disrupts the natural relationship between erroneous response history and subsequent adaptive modification of pre-stimulus neural representations, reminiscent of the fixed quality of misperceptions in certain disease-related and drug-altered states which are resistant to adaptive modification by adverse feedback.

Neural activity in human retrosplenial/posteromedial cortex predicts perceptual choice

In prior work (as noted above), past contradictory findings regarding impact of spontaneous activity[25] have been variously attributed to species/anatomical differences or to the overall design and structure of the behavioral task[26]. We have addressed this challenge by developing a task designed for applicability across mammalian phylogeny from mouse to human, and for compatibility with brainwide activity screening for unbiased circuit localization. Here, to investigate whether pre-stimulus activity influences perceptual choice, we extended the same two-alternative forced choice (2AFC) visual detection task from the mouse experiments above (Figs. 1-5) to three human subjects implanted with deep intracranial electrodes (Extended Data Fig. 12); all subjects were undergoing monitoring for seizure localization, with brain-spanning electrode coverage determined by clinical need and unbiased by this study.

Mirroring our murine areal screen (Fig. 1), we computed the electrode-wise statistical association between neural activity in the 500ms preceding visual stimulus onset and subsequent perceptual choice (Fig. 6a–6c); to align with the murine high-selectivity focus, we restricted human analysis to trials with response times within 1 second of stimulus onset (comprising 91% of all trials). Across subjects, the electrode locations where pre-stimulus activity was most predictive of perceptual choice are shown in darker red (Fig. 6a); although we were limited to fewer trials in human beings due to clinical constraints, pre-stimulus activity, as in mouse, was found to exhibit distinct patterns anticipating upcoming perceptual choice, revealed by choice-averaged neural activity plots (Fig. 6b). Supporting the cross-species consistency of the task, we observed a distribution of correct (hit or correct reject) and incorrect (miss or false alarm) responses to low contrast trials that was strikingly similar not only across human subjects, but also from mouse to human (Fig. 6c, 1g).

To quantitatively identify specific brain areas linked to perceptual choice in this unbiased brainwide activity screen, we pooled electrodes into regions defined by the Human Connectome Project[32] and trained a linear classifier on pre-stimulus neural activity averaged within each region (Fig. 6d). Across all recording locations, the classifier revealed that RSC/posterior cingulate

exhibited the strongest statistical relationship with perceptual choice, followed by the mid-cingulate, dorsolateral prefrontal cortex, and anterior cingulate (classification performed in 250ms windows in the ITI; permutation test; $p < 0.01$ after Bonferonni correction for 11 comparisons); no other cortical regions exhibited significant predictive relationships with perceptual choice. Notably, the human posteromedial cortex (PMC) as defined by the Human Connectome Project, encompassing posterior cingulate and retrosplenial cortices, is considered the functional homologue of rodent RSC and has been implicated in top-down modulation of perceptual expectations[1, 18, 33].

The human cortex is organized into large-scale functional networks, each spanning multiple cortical areas[34]. To determine whether pre-stimulus activity within specific networks is predictive of perceptual choice, we averaged electrode activity into these predefined network parcels and trained a linear classifier on pre-stimulus activity to predict perceptual decisions (Fig. 6e-f). Analysis was restricted to networks with sufficient electrode coverage, defined as including at least three electrodes per network. We identified two cortical networks with statistically significant predictive relationships. The strongest association was observed in the default mode network, achieving 59.7% classification accuracy ($n = 3$ subjects classification performed in 250ms windows in the ITI; permutation test, $p < 0.01$ after Bonferonni correction for 6 comparisons), followed by the frontoparietal control network at 56.1% accuracy ($n = 3$ subjects classification performed in 250ms windows in the ITI; permutation test, $p < 0.01$ after Bonferonni correction for 6 comparisons); no other networks exhibited statistically significant links to perceptual choice.

Finally, we tested whether the relationship between pre-stimulus neural activity, behavioral d prime, and bias found in mice (Fig. 2n) extended also to humans. As before, we grouped trials into quantiles after orienting activity in each electrode such that increased pre-stimulus activity was associated grating responses. We then averaged over all electrode time series and calculated behavioral d prime and bias for each quantile (Figs. 6g-6h). Linear regression revealed a significant positive relationship between pre-stimulus coding dimension activity and bias in the posterior cingulate/RSC and DMN (posterior cingulate/RSC bias slope = 0.30, $p < 0.01$; DMN bias slope = 0.22, $p < 0.01$) but no significant relationship was found for behavioral d prime (posterior cingulate/RSC d prime slope = -0.02, $p > 0.8$; DMN bias slope = -0.0001, $p > 0.8$). These findings reveal a conserved relationship between perceptual bias and pre-stimulus neural activity in posterior cingulate/RSC (and its associated default mode network), across mice and humans.

Discussion

Here we have bridged mouse and human systems using brainwide activity screening and a cross-species task design to identify conserved principles by which spontaneous brain activity influences perceptual decisions. Using this framework, we were able to illuminate, at a detailed and even cellular-resolution level, the adaptive nature of spontaneous cortical activity. An initial widefield activity screen in mice revealed that pre-stimulus activity specifically in retrosplenial cortex (RSC) and visual cortex (V1) predicted perceptual choice. Higher-resolution two-photon imaging of

single cells, guided by the earlier widefield screen, demonstrated that before-stimulus spontaneous activity in RSC and V1 neural populations predicted subsequent perceptual bias. Optogenetic modulation of RSC neurons projecting to V1 was found to specifically shift neural coding-dimension activity and behavioral responses, confirming a top-down role for RSC in modulating perceptual decision-making. Finally, a screen of deep brain-spanning recordings in human subjects during the same perceptual task revealed that pre-stimulus activity– most strikingly in the posteromedial/retrosplenial cortex (the human homolog of rodent RSC)– predicted perceptual choice. Together, these findings provide cross-species evidence– from mouse to human– for a conserved mechanism of state-dependent perceptual modulation, consistent with signal detection theory and highlighting an ancestral achievement of evolution in developing the structure and function of the mammalian brain.

Favoring the hypothesis that spontaneous neural activity may be useful in part to shape perceptual decision-making under challenging conditions, we showed that pre-stimulus activity in both RSC and V1 predicted perceptual decisions but only when the sensory stimulus was ambiguous. Two-photon imaging revealed that pre-stimulus single-neuron activity and population-level coding dimensions in RSC and V1 not only were tuned to represent these perceptual decisions, but also dynamically shifted with environmental statistics in a manner consistent with signal detection theory. Specifically, under enriched contexts with increased probabilities of grating or noise stimuli, RSC and V1 coding dimensions shifted to reflect these environmental biases, aligning both neural representations and perceptual decisions with contextual expectations. This flexibility underscores the role of spontaneous activity in implementing predictive models, whereby prior experience shapes neural states to optimize sensory interpretation. Thus, spontaneous activity in the mammalian brain can operate as a flexible framework for integrating prior expectations with incoming sensory input, particularly under conditions of uncertainty[35-38].

Our experiments revealed not only an adaptive process by which the content of spontaneous activity can tune to represent environmental statistics, but also a maladaptive mechanism altering this process (specifically, pharmacological modulation via psychedelic drug action). While spontaneous activity in V1 became biased toward grating representations under LSD, we found that activity in RSC became entirely unlinked from perceptual choice. LSD was found to abolish the typical top-down influence of RSC on spontaneous activity and perceptual bias, disrupting the integration of prior experience into decision-making, and allowing perceptual decisions regarding ambiguous signals (and even of signals not present at all). These findings reveal that psychedelics may alter perception in part by shifting the representational dynamics of spontaneous activity that normally implements top-down perceptual control[39].

Our parallel experiments in humans provided evidence for the conserved and even translational relevance of these findings. Pre-stimulus activity in the human posteromedial/retrosplenial cortex– the human homolog of rodent RSC—predicted perceptual decisions during the visual detection task. The specificity of this effect to PMC/RSC and its associated default mode network suggested a conserved mechanism by which top-down signals from posterior cortical regions can influence

sensory processing and decision-making, with implications for understanding psychiatric conditions such as schizophrenia, PTSD, and depression, wherein disruptions in predictive modeling are well-documented[40-44]. By linking spontaneous neural dynamics to context-based decision-making, these findings provide a framework for investigating how altered cortical states contribute to perceptual and cognitive symptoms in these disorders.

Critical mechanistic experiments (optogenetic modulation of RSC projections to V1, establishing behavioral importance specifically for late pre-stimulus activity in RSC) revealed both a causal role and important timing constraints for RSC in shaping perceptual choices through top-down modulation. These results are consistent with prior evidence implicating RSC in integrating contextual information with transmission to sensory regions[29, 30]; here we find that RSC influences perceptual decisions even in the absence of overt task cues, supporting a role as a key node in the default mode network for generating predictive models of the external world[45] before any stimulus is presented, choice required, or action initiated. In sum, this work reveals an ancestral mammalian mechanism by which spontaneous cortical activity, particularly in RSC, integrates environmental context to shape perceptual decisions. This dynamic interplay between spontaneous and evoked activity highlights the importance of internal neural states in cognitive processes, and opens up avenues for considering spontaneous cortical dynamics in the understanding of neuropsychiatric conditions.

Acknowledgments:

We are grateful to all the study participants and their families, without whom this work would not have been possible, and to all the nursing, IT, and clinical research support staff. The Human Neural Circuitry (HNC) Program was founded with the support of the President and Provost of Stanford University, as well as the Dean of the Stanford Medical School and the leadership of both Stanford Hospital and Lucille Packard Children's Hospital. In addition, this work was supported by the Gatsby Foundation (K.D.). AM was supported by the NIMH (5K08MH13188802) and a Career Award for Medically Trained Scientists from the Burroughs-Wellcome Fund. IK was supported as a Merck Awardee of the Life Sciences Research Foundation and a Wu Tsai Stanford Neurosciences Institute Interdisciplinary Scholar. TXL was supported by a Knight-Hennessy Scholars fellowship and NSF GRFP. DGH was supported by the Stanford Medicine REACH Initiative. Anatomical data were provided (in part) by the Human Connectome Project, WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54MH091657) funded by the 16 NIH Institutes and Centers that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University.

Non-author group contributors:

Non-author group membership followed the recommendations of the International Committee of Medical Journal Editors (ICMJE) and Council of Science Editors (CSE).

Stanford Comprehensive Epilepsy Program: A full list of non-author contributors is pending for inclusion in the supplemental material.
https://med.stanford.edu/neurology/divisions/epilepsy/ourteam.html#adult

Author contributions:

AM, JK, and KD designed mouse behavioral experiments.

AM and JK conducted and analyzed mouse behavioral experiments.

AM, JK, and KD designed mouse imaging experiments.

AM and JK conducted and analyzed the mouse behavioral experiments.

The human-subject IRB protocol was designed by KD and CIR with clinical research coordinator support from PM, AC, KP, and DGH.

VB, CIR, AM, and KD carried out planning and coordination on the human subject side.

VB implanted intracranial electrodes.

AM and JK analyzed human electrophysiology data.

JK and EY generated surface visualizations for human electrophysiology findings

AM, SV, and KD designed human behavioral experiments with input from TXL, SV, AC, CIR, and PM.

IK, AC, and TXL collected human behavioral data during neural recording.

CIR, AM and KD carried out psychiatric interviews during intracranial recording.

PN designed and constructed the PHI-compatible and secure human subject intracranial recording pipeline with support from LY.

LY, AF, and PN recorded and processed intracranial human recordings.

AM, JK, SG, and SD designed the computational analyses.

KD conducted all medical and psychiatric screening and consented all human subjects, directed the clinical protocol, and supervised all aspects of the project.

AM, JK, and KD wrote the paper, and all authors commented on the manuscript.

Competing interests:

In the last three years, CIR has been a consultant for Biohaven Inc., Osmind, and Biogen; received research grant support from Biohaven Inc.; received royalties from American Psychiatric Association Publishing; and received a stipend from APA Publishing for her role as Deputy Editor at The American Journal of Psychiatry and a stipend for her role as Deputy Editor of Neuropsychopharmacology. KD is a founder and scientific advisor for Maplight Therapeutics and Stellaromics, and a scientific advisor to RedTree LLC and Modulight. Other authors declare that they have no competing interests.

Data and materials availability: All processed data not subject to clinical privacy restrictions will be made available upon publication of this manuscript, along with associated code.

References

1. de Lange, F.P., Heilbron, M. & Kok, P. How Do Expectations Shape Perception? *Trends Cogn Sci* **22**, 764-779 (2018).
2. Raichle, M.E. & Gusnard, D.A. Intrinsic brain activity sets the stage for expression of motivated behavior. *J Comp Neurol* **493**, 167-176 (2005).
3. Arieli, A., Shoham, D., Hildesheim, R. & Grinvald, A. Coherent spatiotemporal patterns of ongoing activity revealed by real-time optical imaging coupled with single-unit recording in the cat visual cortex. *J Neurophysiol* **73**, 2072-2093 (1995).
4. Arieli, A., Sterkin, A., Grinvald, A. & Aertsen, A. Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. *Science* **273**, 1868-1871 (1996).

5. Kenet, T., Bibitchkov, D., Tsodyks, M., Grinvald, A. & Arieli, A. Spontaneously emerging cortical representations of visual attributes. *Nature* **425**, 954-956 (2003).
6. Petersen, C.C., Hahn, T.T., Mehta, M., Grinvald, A. & Sakmann, B. Interaction of sensory responses with spontaneous depolarization in layer 2/3 barrel cortex. *Proc Natl Acad Sci U S A* **100**, 13638-13643 (2003).
7. Siegle, J.H. et al. Survey of spiking in the mouse visual system reveals functional hierarchy. *Nature* **592**, 86-92 (2021).
8. Fox, P.T. et al. Mapping human visual cortex with positron emission tomography. *Nature* **323**, 806-809 (1986).
9. Ringach, D.L. Spontaneous and driven cortical activity: implications for computation. *Curr Opin Neurobiol* **19**, 439-444 (2009).
10. 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).
11. Luczak, A., Bartho, P. & Harris, K.D. Spontaneous events outline the realm of possible sensory responses in neocortical populations. *Neuron* **62**, 413-425 (2009).
12. Berkes, P., Orban, G., Lengyel, M. & Fiser, J. Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. *Science* **331**, 83-87 (2011).
13. Fiser, J., Chiu, C. & Weliky, M. Small modulation of ongoing cortical dynamics by sensory input during natural vision. *Nature* **431**, 573-578 (2004).
14. Fiser, A. et al. Experience-dependent spatial expectations in mouse visual cortex. *Nat Neurosci* **19**, 1658-1664 (2016).
15. Han, F., Caporale, N. & Dan, Y. Reverberation of recent visual experience in spontaneous cortical waves. *Neuron* **60**, 321-327 (2008).
16. Davis, Z.W., Muller, L., Martinez-Trujillo, J., Sejnowski, T. & Reynolds, J.H. Spontaneous travelling cortical waves gate perception in behaving primates. *Nature* **587**, 432-436 (2020).
17. Hesselmann, G., Kell, C.A., Eger, E. & Kleinschmidt, A. Spontaneous local variations in ongoing neural activity bias perceptual decisions. *Proc Natl Acad Sci U S A* **105**, 10984-10989 (2008).
18. Sadaghiani, S., Hesselmann, G. & Kleinschmidt, A. Distributed and antagonistic contributions of ongoing activity fluctuations to auditory stimulus detection. *J Neurosci* **29**, 13410-13417 (2009).
19. Pajani, A., Kok, P., Kouider, S. & de Lange, F.P. Spontaneous Activity Patterns in Primary Visual Cortex Predispose to Visual Hallucinations. *J Neurosci* **35**, 12947-12953 (2015).
20. Ress, D., Backus, B.T. & Heeger, D.J. Activity in primary visual cortex predicts performance in a visual detection task. *Nat Neurosci* **3**, 940-945 (2000).
21. Reimer, J. et al. Pupil fluctuations track rapid changes in adrenergic and cholinergic activity in cortex. *Nat Commun* **7**, 13289 (2016).
22. Fu, Y. et al. A cortical circuit for gain control by behavioral state. *Cell* **156**, 1139-1152 (2014).

23. Niell, C.M. & Stryker, M.P. Modulation of visual responses by behavioral state in mouse visual cortex. *Neuron* **65**, 472-479 (2010).
24. Vinck, M., Batista-Brito, R., Knoblich, U. & Cardin, J.A. Arousal and locomotion make distinct contributions to cortical activity patterns and visual encoding. *Neuron* **86**, 740-754 (2015).
25. Talluri, B.C. et al. Activity in primate visual cortex is minimally driven by spontaneous movements. *Nat Neurosci* **26**, 1953-1959 (2023).
26. Yin, C. et al. Spontaneous movements and their impact on neural activity fluctuate with latent engagement states. *bioRxiv* (2024).
27. Schmack, K., Bosc, M., Ott, T., Sturgill, J.F. & Kepecs, A. Striatal dopamine mediates hallucination-like perception in mice. *Science* **372** (2021).
28. Li, N., Daie, K., Svoboda, K. & Druckmann, S. Robust neuronal dynamics in premotor cortex during motor planning. *Nature* **532**, 459-464 (2016).
29. Franco, L.M. & Goard, M.J. A distributed circuit for associating environmental context with motor choice in retrosplenial cortex. *Sci Adv* **7** (2021).
30. Makino, H. & Komiyama, T. Learning enhances the relative impact of top-down processing in the visual cortex. *Nat Neurosci* **18**, 1116-1122 (2015).
31. Liechti, M.E. Modern Clinical Research on LSD. *Neuropsychopharmacology* **42**, 2114-2127 (2017).
32. Glasser, M.F. et al. A multi-modal parcellation of human cerebral cortex. *Nature* **536**, 171-178 (2016).
33. Vann, S.D., Aggleton, J.P. & Maguire, E.A. What does the retrosplenial cortex do? *Nat Rev Neurosci* **10**, 792-802 (2009).
34. Yeo, B.T. et al. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *J Neurophysiol* **106**, 1125-1165 (2011).
35. Powers, A.R., III, Kelley, M. & Corlett, P.R. Hallucinations as top-down effects on perception. *Biol Psychiatry Cogn Neurosci Neuroimaging* **1**, 393-400 (2016).
36. Knill, D.C. & Pouget, A. The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends Neurosci* **27**, 712-719 (2004).
37. Vilares, I. & Kording, K. Bayesian models: the structure of the world, uncertainty, behavior, and the brain. *Ann N Y Acad Sci* **1224**, 22-39 (2011).
38. Bach, D.R. & Dolan, R.J. Knowing how much you don't know: a neural organization of uncertainty estimates. *Nat Rev Neurosci* **13**, 572-586 (2012).
39. Singleton, S.P. et al. Receptor-informed network control theory links LSD and psilocybin to a flattening of the brain's control energy landscape. *Nat Commun* **13**, 5812 (2022).
40. Miller, D.R. et al. Default Mode Network Subsystems are Differentially Disrupted in Posttraumatic Stress Disorder. *Biol Psychiatry Cogn Neurosci Neuroimaging* **2**, 363-371 (2017).
41. Barch, D.M. & Ceaser, A. Cognition in schizophrenia: core psychological and neural mechanisms. *Trends Cogn Sci* **16**, 27-34 (2012).
42. Hamilton, J.P., Chen, M.C. & Gotlib, I.H. Neural systems approaches to understanding major depressive disorder: an intrinsic functional organization perspective. *Neurobiol Dis* **52**, 4-11 (2013).

43. Northoff, G. The brain's spontaneous activity and its psychopathological symptoms - "Spatiotemporal binding and integration". *Progress in neuro-psychopharmacology & biological psychiatry* **80**, 81-90 (2018).
44. Sylvester, C.M. et al. Functional network dysfunction in anxiety and anxiety disorders. *Trends Neurosci* **35**, 527-535 (2012).
45. Raichle, M.E. et al. A default mode of brain function. *Proc Natl Acad Sci U S A* **98**, 676-682 (2001).

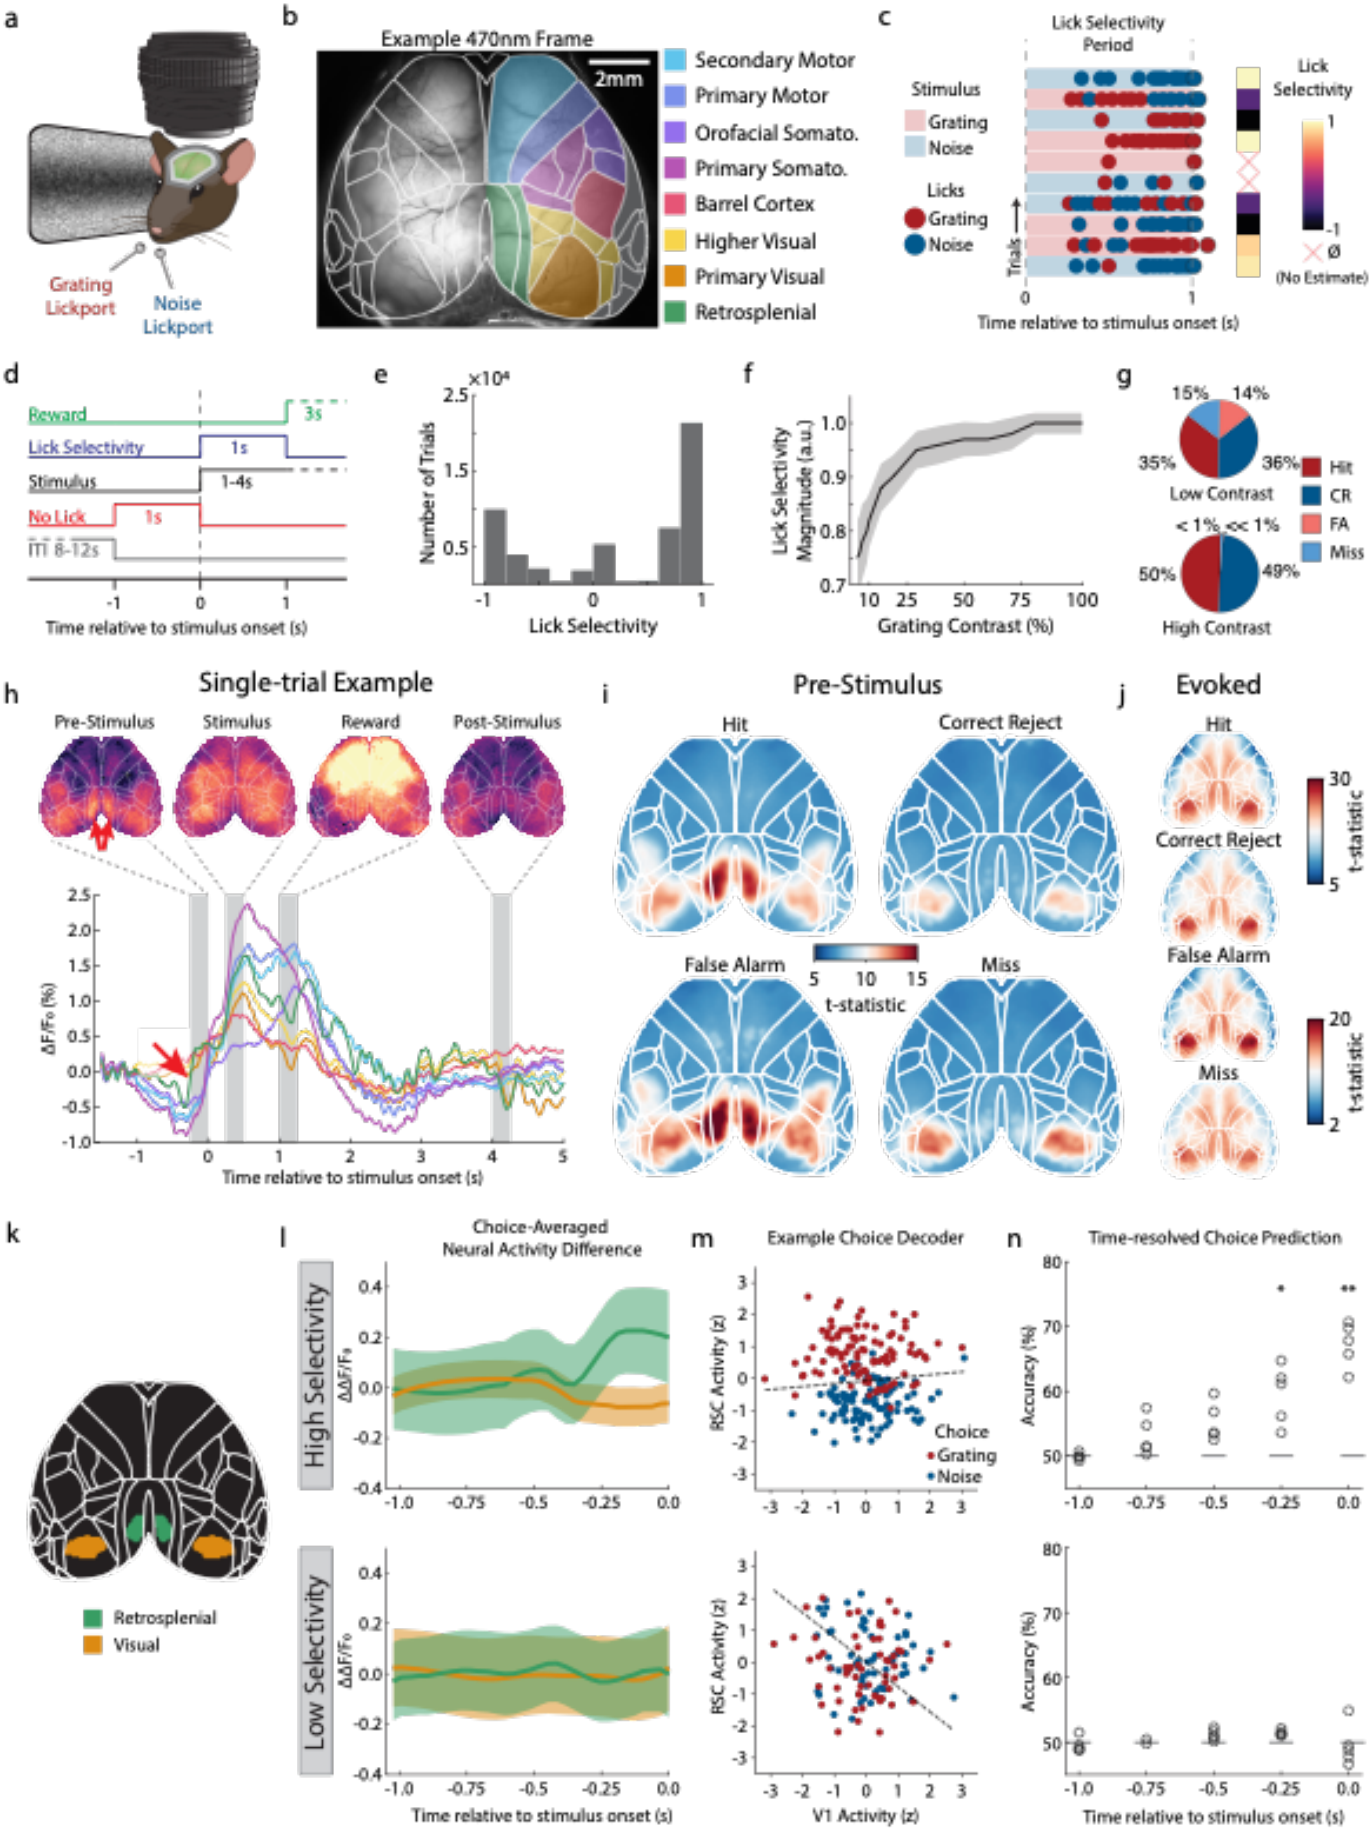
a
Grating Lickport
Noise Lickport
b
Example 470nm Frame
2mm
Secondary Motor
Primary Motor
Orofacial Somato.
Primary Somato.
Barrel Cortex
Higher Visual
Primary Visual
Retrosplenial
c
Lick Selectivity Period
Stimulus
Grating
Noise
Licks
Grating
Noise
Trials
Lick Selectivity
1
-1
Ø
(No Estimate)
Time relative to stimulus onset (s)
d
Reward
3s
Lick Selectivity
1s
Stimulus
1-4s
No Lick
1s
ITI 8-12s
Time relative to stimulus onset (s)
e
×10⁴
Number of Trials
Lick Selectivity
f
Lick Selectivity Magnitude (a.u.)
Grating Contrast (%)
g
15%
14%
35%
36%
Low Contrast
< 1%
<< 1%
50%
49%
High Contrast
Hit
CR
FA
Miss
h
Single-trial Example
Pre-Stimulus
Stimulus
Reward
Post-Stimulus
ΔF/F₀ (%)
Time relative to stimulus onset (s)
i
Pre-Stimulus
Hit
Correct Reject
False Alarm
Miss
5
10
15
t-statistic
j
Evoked
Hit
Correct Reject
False Alarm
Miss
30
5
t-statistic
20
2
t-statistic
k
Retrosplenial
Visual
l
Choice-Averaged Neural Activity Difference
High Selectivity
Low Selectivity
ΔΔF/F₀
Time relative to stimulus onset (s)
m
Example Choice Decoder
RSC Activity (z)
Choice
Grating
Noise
V1 Activity (z)
n
Time-resolved Choice Prediction
Accuracy (%)
Time relative to stimulus onset (s)

Figure 1: Wide-field cortical activity predicts perceptual choice prior to stimulus presentation in a two-alternative forced choice (2AFC) visual detection task. (a) Survey of cortical activity during the 2AFC visual detection task. Head-fixed mice were trained to lick one of two spouts for a water reward to indicate the presence or absence of a faint drifting grating embedded in visual white noise. (b) Example widefield imaging frame with anatomical atlas. (c) Computation of lick selectivity. Lick patterns during the "lick selectivity" period defined a lick selectivity index. If all licks in this period were on the correct spout for the trial, the lick selectivity index was +1. Conversely, if all licks were on the incorrect spout, the lick selectivity index was -1. Equal licking of both spouts resulted in a lick selectivity index of 0. Fewer than five licks resulted in no estimate. (d) Visual detection task structure. Mice were required to refrain from licking the spout for at least one second prior to visual stimulus onset ("no lick"). Mice could respond immediately upon visual stimulus onset, but no water reward was released for one second ("lick selectivity"). The first lick after the onset of the reward period ("reward") dictated the outcome. (e) Histogram of lick selectivity indices across all trials and all mice (n = 14 mice). The histogram illustrated three modes: high selectivity correct trials (+1), high selectivity incorrect trials (-1), and low selectivity trials (0). (f) Lick selectivity magnitude as a proxy for response confidence. Response confidence was expected to increase with decreasing task difficulty. Accordingly, the magnitude of lick selectivity (absolute value taken to normalize for task performance accuracy) increased with decreasing task difficulty (increase in grating contrast, e.g., visibility of the grating; n = 14 mice, paired t-test for grating contrast 10 vs. 100, $p < 0.01$). (g) Distribution of behavioral responses. Mice made high selectivity hits, correct rejects (CR), misses, and false alarms (FA). The pie charts illustrated the proportion of high selectivity (absolute value > 0.8) responses of each kind across all trials in high contrast (easy) and low contrast (difficult) sessions (n = 14 mice). (h) Example wide-field calcium imaging data during a single trial. Note the arrow indicating robust pre-stimulus (time < 0) activity. (i) T-statistic maps of pre-stimulus brain activity (in the 250 ms preceding visual stimulus) revealed statistically significant activity in the visual and retrosplenial cortex in high selectivity trials (n = 17,021 trials, pixel-wise t-test of baseline vs. pre-stimulus activity per condition). (j) T-statistic maps of stimulus-evoked brain activity (n = 17,021 trials, pixel-wise t-test of pre-stimulus vs. evoked activity per condition). (k) Brain atlas depicting retrosplenial and visual regions of interest identified by symmetrizing and applying TFCE (see Methods) to the results in (i). (l) Differences in pre-stimulus brain activity in visual and retrosplenial cortices across grating and noise response trials in high- and low-selectivity trials (n = 5 mice). (m) Examples of linear choice decoding on pre-stimulus retrosplenial and visual activity in high- and low-selectivity trials. (n) Training a linear choice decoder on retrosplenial and visual activity produced statistically significant decoding of choice in high selectivity trials, but not in low selectivity trials (n = 5 mice, classification over 250 ms epochs in the inter-trial interval, permutation test, $p < 0.05$ for $-0.5 < t < -0.25$ s, $p < 0.01$ for $-0.25 < t < 0$ s; all other time points were not significant). Confidence intervals are presented as mean ± s.e.m.

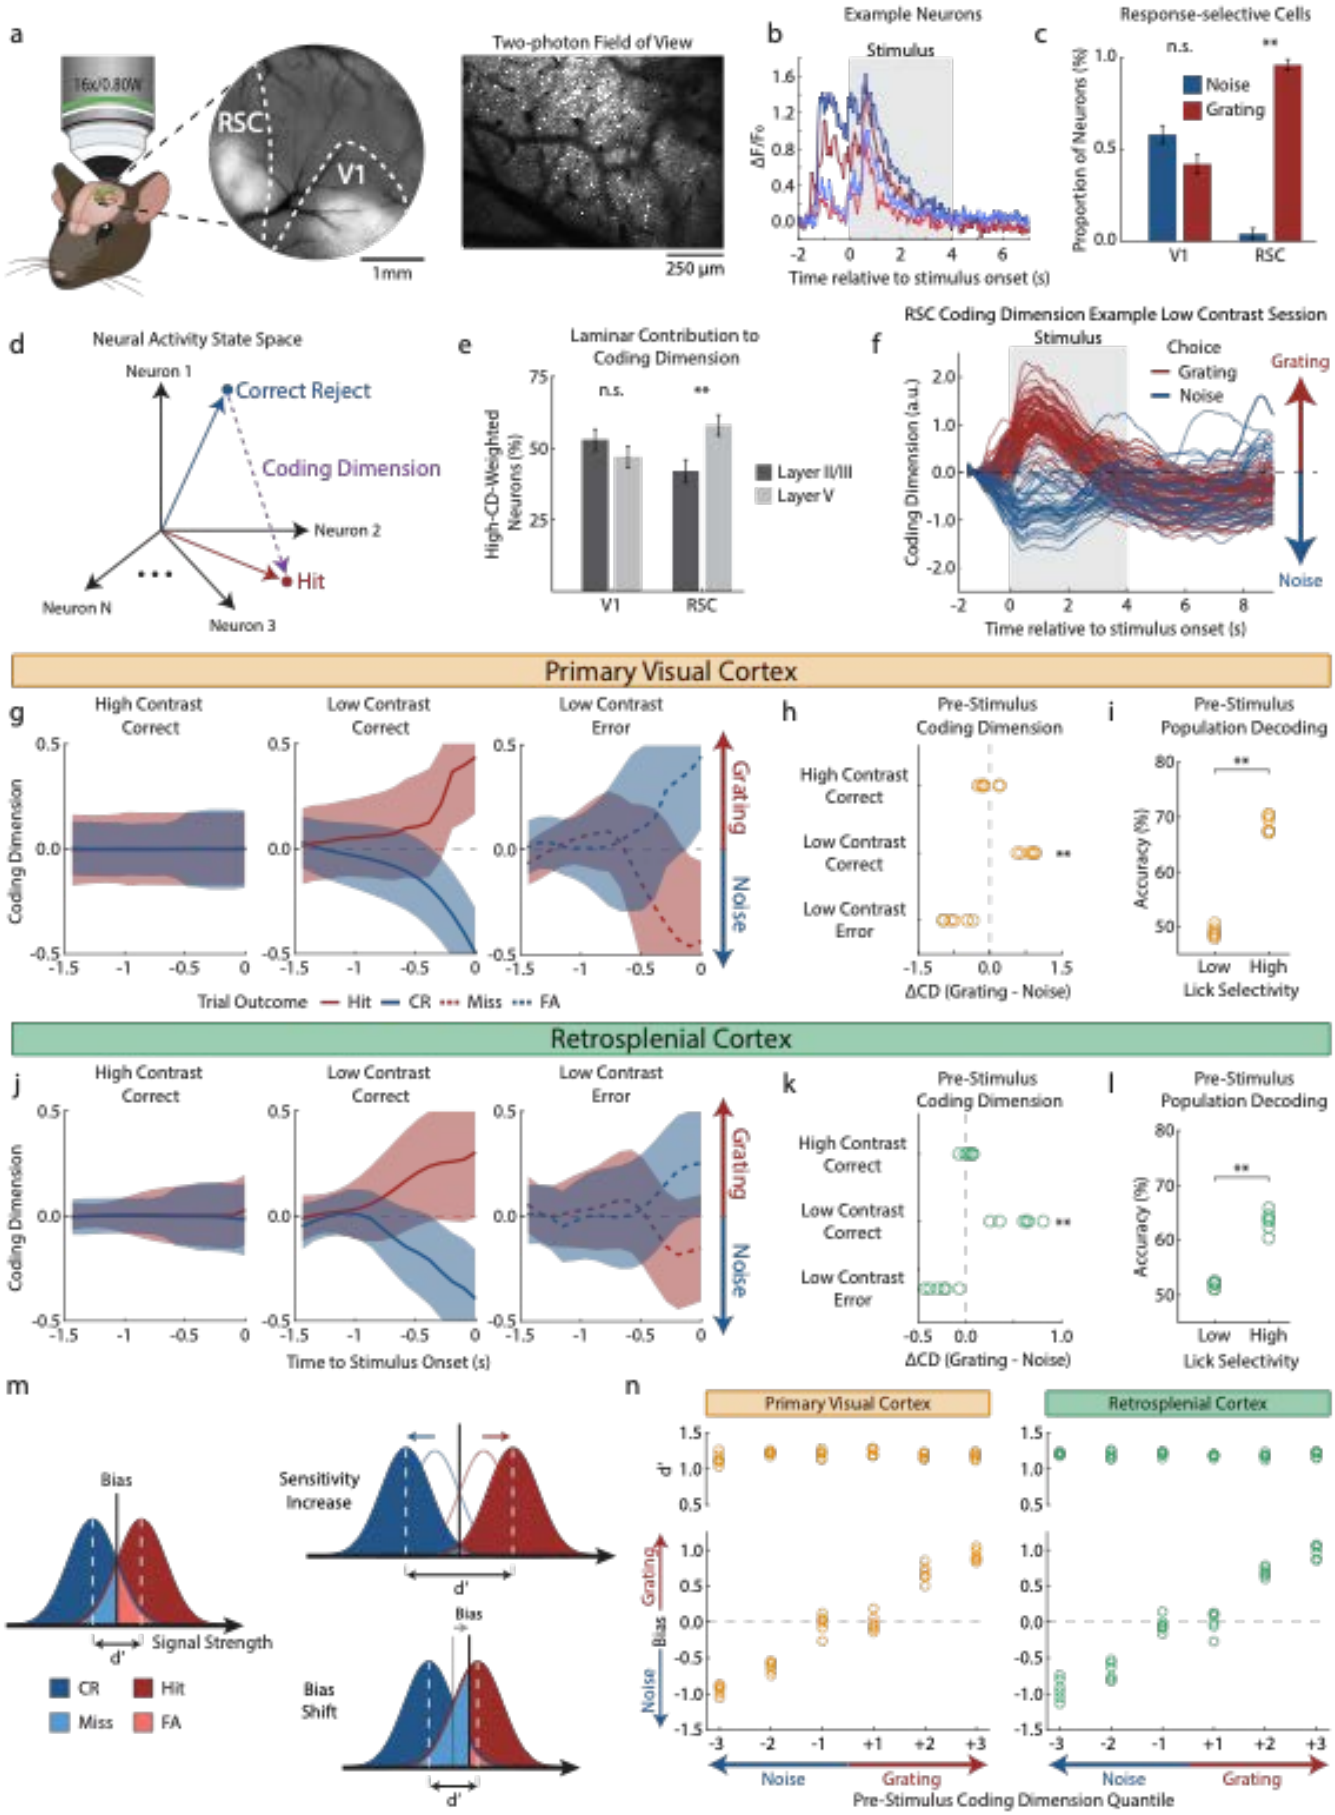
a
16x/0.80W
RSC
V1
1mm
Two-photon Field of View
250 µm
b
Example Neurons
Stimulus
ΔF/F0
Time relative to stimulus onset (s)
c
Response-selective Cells
n.s.
**
Noise
Grating
Proportion of Neurons (%)
V1
RSC
d
Neural Activity State Space
Neuron 1
Correct Reject
Coding Dimension
Neuron 2
Hit
Neuron N
Neuron 3
e
Laminar Contribution to Coding Dimension
n.s.
**
High-CD-Weighted Neurons (%)
Layer II/III
Layer V
V1
RSC
f
RSC Coding Dimension Example Low Contrast Session
Stimulus
Choice
Grating
Noise
Coding Dimension (a.u.)
Time relative to stimulus onset (s)
Primary Visual Cortex
g
High Contrast Correct
Low Contrast Correct
Low Contrast Error
Coding Dimension
Trial Outcome — Hit — CR ··· Miss ··· FA
h
Pre-Stimulus Coding Dimension
High Contrast Correct
Low Contrast Correct
Low Contrast Error
ΔCD (Grating - Noise)
i
Pre-Stimulus Population Decoding
Accuracy (%)
Low
High
Lick Selectivity
Retrosplenial Cortex
j
High Contrast Correct
Low Contrast Correct
Low Contrast Error
Coding Dimension
Time to Stimulus Onset (s)
k
Pre-Stimulus Coding Dimension
High Contrast Correct
Low Contrast Correct
Low Contrast Error
ΔCD (Grating - Noise)
l
Pre-Stimulus Population Decoding
Accuracy (%)
Low
High
Lick Selectivity
m
Bias
Signal Strength
d'
CR
Hit
Miss
FA
Sensitivity Increase
Bias Shift
n
Primary Visual Cortex
Retrosplenial Cortex
d'
Grating
Bias
Noise
Pre-Stimulus Coding Dimension Quantile

Figure 2: Two-photon imaging shows that choice-specific representations tended to be active prior to high-selectivity perceptual choices. (a) A cranial window was implanted over the posterior cortex, allowing two-photon imaging of CaMKIIa-expressing excitatory neurons in the retrosplenial cortex (RSC) and primary visual cortex (V1), obtained in separate sessions. (b) Sample single-trial responses from four neurons. (c) Mass univariate analysis revealed grating- and noise-responsive cells, showing an equivalent distribution in V1 but a bias toward grating-responsive cells in RSC (n = 6 mice, paired t-test, $p < 0.001$). (d) Schematic illustrating the computation of the coding dimension. (e) Laminar analysis revealed that neurons highly weighted for perceptual choice were biased to layer V in RSC, with no statistically significant difference between layers II/III and V in visual cortex (n = 6 mice, paired t-test, $p < 0.01$). (f) Example of single-trial coding dimension traces over a full session. Red indicates grating responses; blue indicates noise responses. (g) V1 pre-stimulus coding dimension in high contrast (easy) and low contrast (hard) sessions. Whereas there was no mean difference in pre-stimulus activity in high contrast sessions, the coding dimension exhibited pre-stimulus drift in the direction of the chosen response prior to visual stimulus onset in low contrast sessions (n = 6 mice). (h) V1 coding dimension exhibited a statistically significant shift in pre-stimulus activity in low contrast responses (n = 6 mice, paired t-test, $p < 0.01$ in high-selectivity trials). (i) V1 population decoding illustrated statistically significant predictability of choice from pre-stimulus activity in high-selectivity, but not low-selectivity, trials. (j-l) The same analyses as (g-i) conducted in RSC (n = 6 mice, paired t-test, $p < 0.001$ in high-selectivity trials). (m) Signal detection theory schematic illustrating the independent effects of bias and d' on task performance. (n) Illustration of behavioral bias and d' in perceptual choice as a function of pre-stimulus coding dimension. Signal sensitivity (behavioral d') was not significantly modulated as a function of pre-stimulus coding dimension. However, perceptual bias was strongly affected by pre-stimulus coding dimension in both V1 and RSC (n = 6 mice). Confidence intervals are presented as mean ± s.e.m.

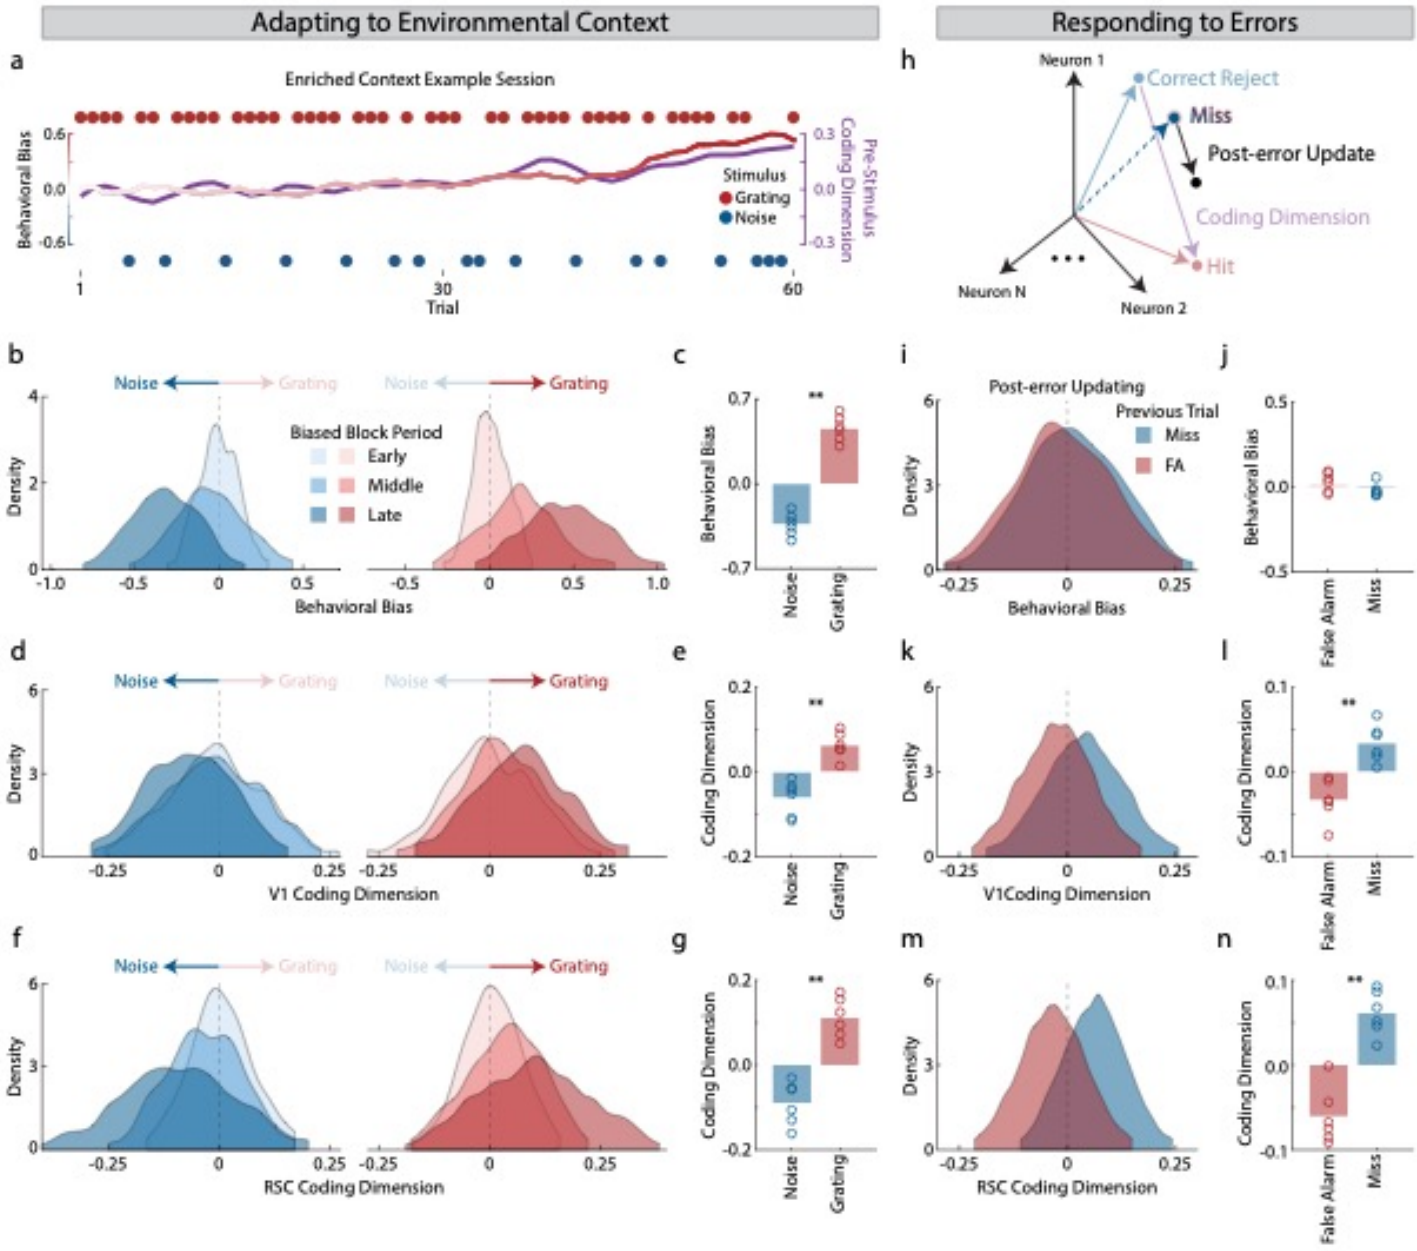
Adapting to Environmental Context
Responding to Errors
a
Enriched Context Example Session
Behavioral Bias
Pre-Stimulus Coding Dimension
Stimulus
Grating
Noise
Trial
h
Neuron 1
Correct Reject
Miss
Post-error Update
Coding Dimension
Hit
Neuron N
Neuron 2
b
Noise
Grating
Biased Block Period
Early
Middle
Late
Density
Behavioral Bias
c
Behavioral Bias
i
Post-error Updating
Previous Trial
Miss
FA
j
d
V1 Coding Dimension
e
k
V1Coding Dimension
l
f
RSC Coding Dimension
g
m
n
False Alarm

Figure 3: Context shifts pre-stimulus neural representations. (a) Example of a biased context in which gratings were present in 75% of trials. Note the shift in behavioral bias and pre-stimulus coding dimension toward the grating representation over time. (b) Kernel density estimate of the distribution of behavioral bias in noise and grating contexts, binned into early, middle, and late thirds of trials. (c) Behavioral bias summary in the late epoch of the context ($n = 6$ mice, paired t-test, $p < 0.01$). (d) Kernel density distribution of the coding dimension in V1 in noise and grating contexts, binned as before. (e) V1 coding dimension summary in the late epoch of the context ($n = 6$ mice, paired t-test, $p < 0.01$). (f-g) Same as (d-e), but in RSC ($n = 6$ mice, paired t-test, $p < 0.01$). (h) Schematic illustration of a shift in coding dimension activity after a perceptual error. (i) Kernel density estimates of distributions of behavioral biases following high-selectivity errors. (j) Behavioral bias summary in trials following high-selectivity errors ($n = 6$ mice, paired t-test, $p > 0.8$). (k) Kernel density estimates of distributions of the V1 coding dimension following high-selectivity errors. (l) V1 coding dimension summary following high-selectivity errors ($n = 6$ mice, paired t-test, $p < 0.01$). (m-n) Same as (k-l), but in RSC ($n = 6$ mice, paired t-test, $p < 0.01$).

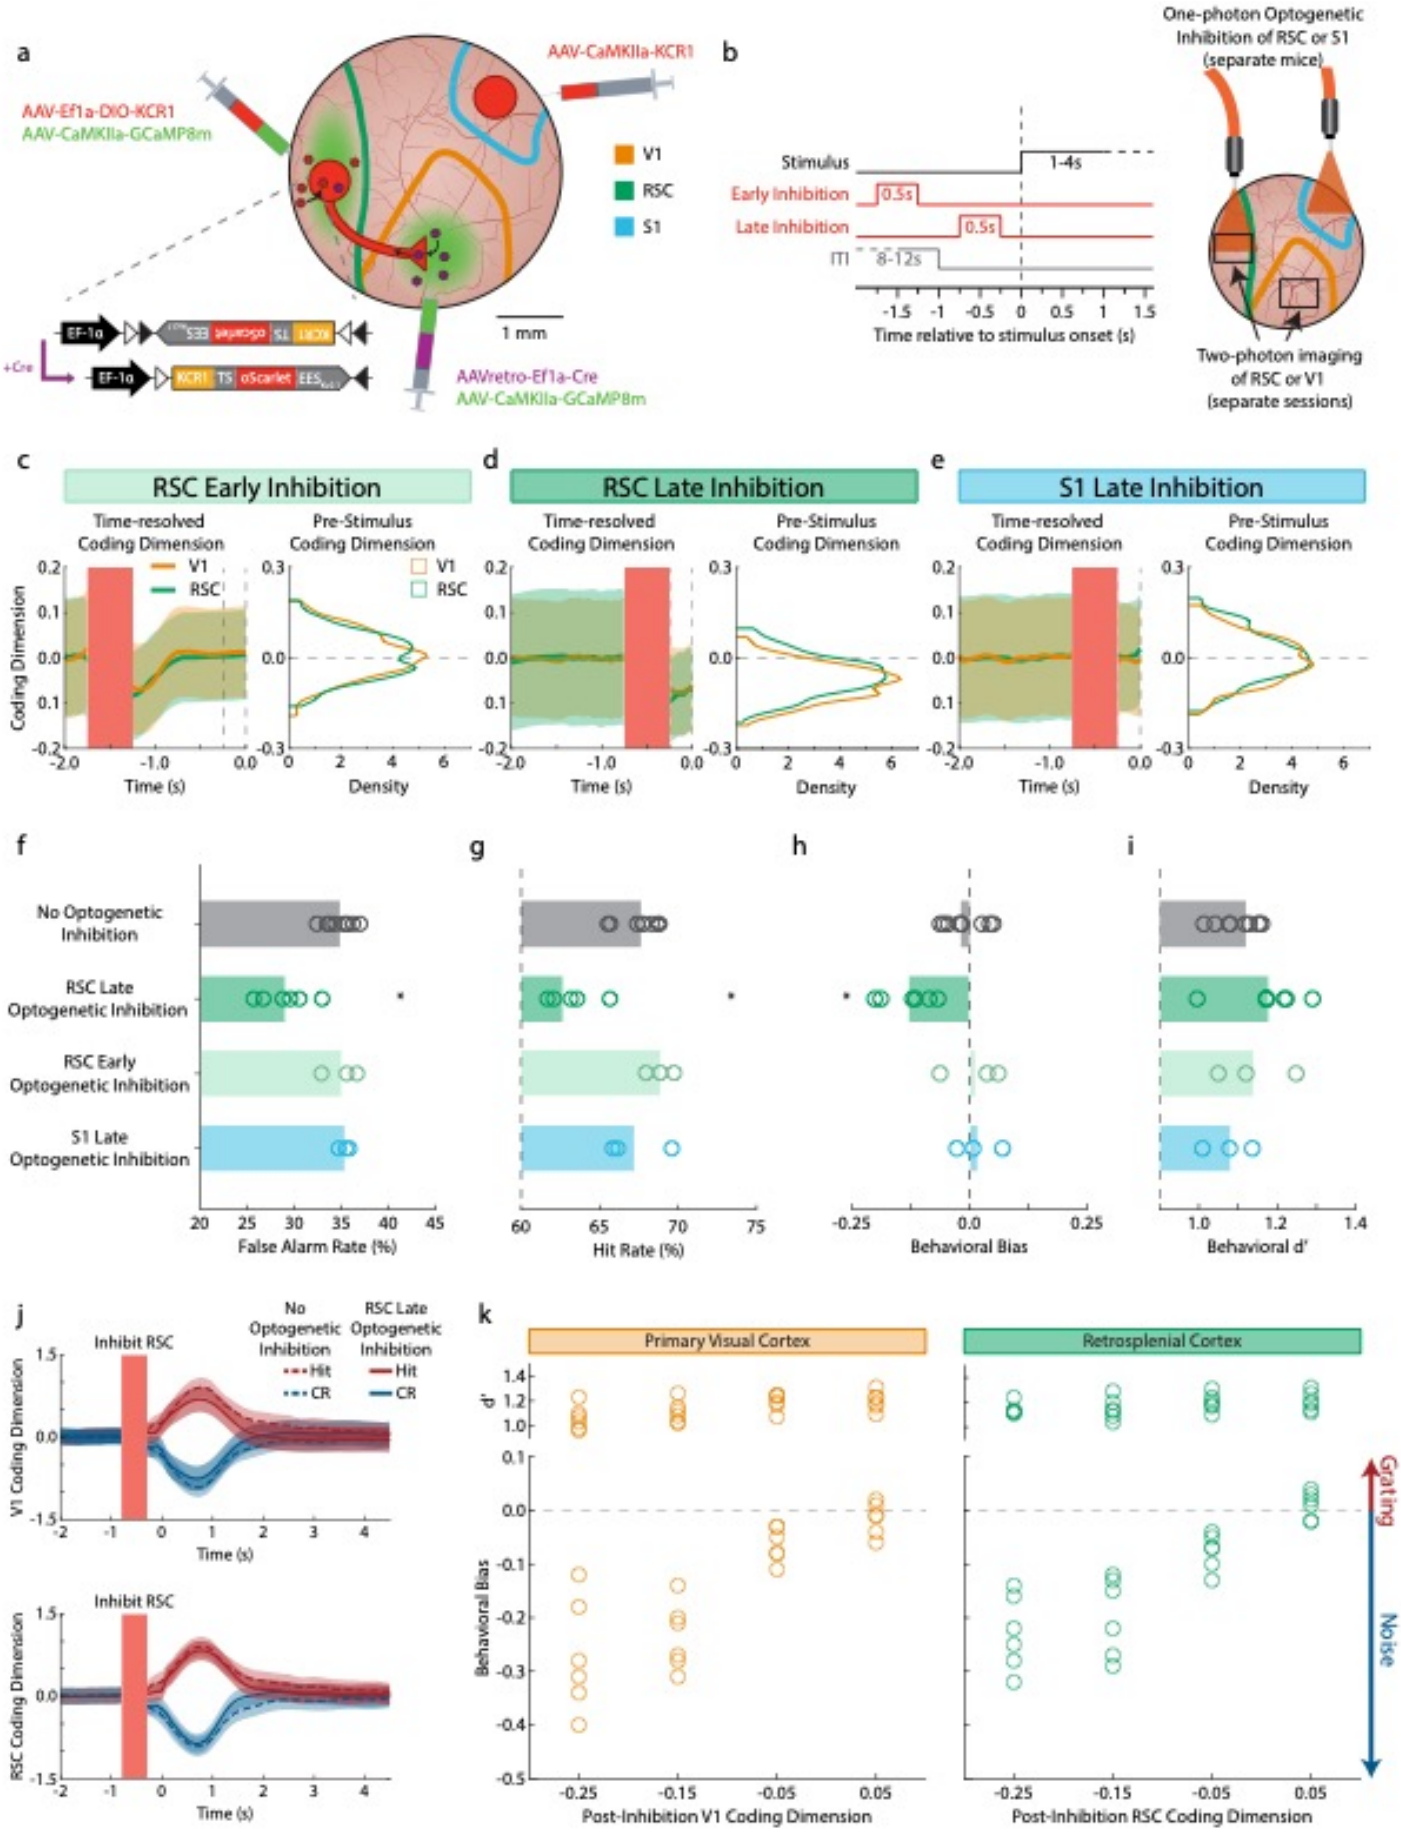
a
AAV-CaMKIIa-KCR1
AAV-Ef1a-DIO-KCR1
AAV-CaMKIIa-GCaMP8m
V1
RSC
S1
1 mm
+Cre
EF-1α
KCR1
TS
oScarlet
EES
AAVretro-Ef1a-Cre
AAV-CaMKIIa-GCaMP8m
b
Stimulus
1-4s
Early Inhibition
0.5s
Late Inhibition
0.5s
ITI
8-12s
Time relative to stimulus onset (s)
One-photon Optogenetic Inhibition of RSC or S1 (separate mice)
Two-photon imaging of RSC or V1 (separate sessions)
c
RSC Early Inhibition
Time-resolved Coding Dimension
Pre-Stimulus Coding Dimension
V1
RSC
Coding Dimension
Time (s)
Density
d
RSC Late Inhibition
Time-resolved Coding Dimension
Pre-Stimulus Coding Dimension
Time (s)
Density
e
S1 Late Inhibition
Time-resolved Coding Dimension
Pre-Stimulus Coding Dimension
Time (s)
Density
f
No Optogenetic Inhibition
RSC Late Optogenetic Inhibition
RSC Early Optogenetic Inhibition
S1 Late Optogenetic Inhibition
False Alarm Rate (%)
g
Hit Rate (%)
h
Behavioral Bias
i
Behavioral d'
j
Inhibit RSC
No Optogenetic Inhibition
RSC Late Optogenetic Inhibition
Hit
CR
V1 Coding Dimension
Time (s)
Inhibit RSC
RSC Coding Dimension
Time (s)
k
Primary Visual Cortex
Retrosplenial Cortex
d'
Behavioral Bias
Grating
Noise
Post-Inhibition V1 Coding Dimension
Post-Inhibition RSC Coding Dimension

Figure 4: Optogenetic inhibition of RSC input to V1 shifts behavioral bias. (a) Intersectional optogenetics. To selectively express the inhibitory opsin KCR1 in RSC:V1 projection neurons, we applied an intersectional strategy, injecting AAVretro-Cre in V1 and Cre-dependent KCR1 in RSC. In a separate cohort of mice used for areal control experiments, KCR1 was expressed in the primary somatosensory cortex (S1). (b) Simultaneous optogenetic inhibition and two-photon imaging during the task. A 500-ms pulse of 617 nm light was applied at one of two time periods: late inhibition denoted the period from 750 ms to 250 ms preceding visual stimulus onset, whereas early inhibition corresponded to the period from 1750 ms to 1250 ms prior to visual stimulus onset. Inhibitory optical stimulation was delivered by positioning a fiber optic patch cable directly above either RSC or S1 (in separate mice), with concurrent two-photon imaging in either RSC or V1 (in separate sessions). (c-e) Neural response to optogenetic inhibition. Early RSC inhibition pushed the coding dimension of neural activity toward the noise representation in V1 and RSC, but activity along the coding dimension recovered to zero prior to visual stimulus onset (c: time period in dotted lines, n = 6 mice). Similarly, late RSC inhibition shifted the coding dimension toward the noise representation; however, coding dimension activity did not recover before visual stimulus onset (d, n = 3 mice). S1 late inhibition did not cause any shift in the coding dimension (e, n = 3 mice). (f-i) Optogenetic inhibition caused a behavioral bias toward noise responses. Both the false alarm rate (f) and hit rate (g) declined during RSC late inhibition but not in the control conditions, causing a shift in behavioral bias toward the noise response (h). There was no statistically significant shift in behavioral d' (i). Statistics were computed with t-tests over mice for (c-e) with Bonferroni correction for three comparisons. (j) Full traces for coding dimension response revealed no differences in evoked activity compared to trials without optogenetic inhibition (see Supplemental Figure 4, n = 6, paired t-tests per time bin). (k) The direct relationship between pre-stimulus coding dimension and behavioral bias, established in Figure 2n, was recapitulated by optogenetic inhibition (n = 6 mice). Coding dimension did not modulate behavioral d'. Confidence intervals are presented as mean ± s.e.m.

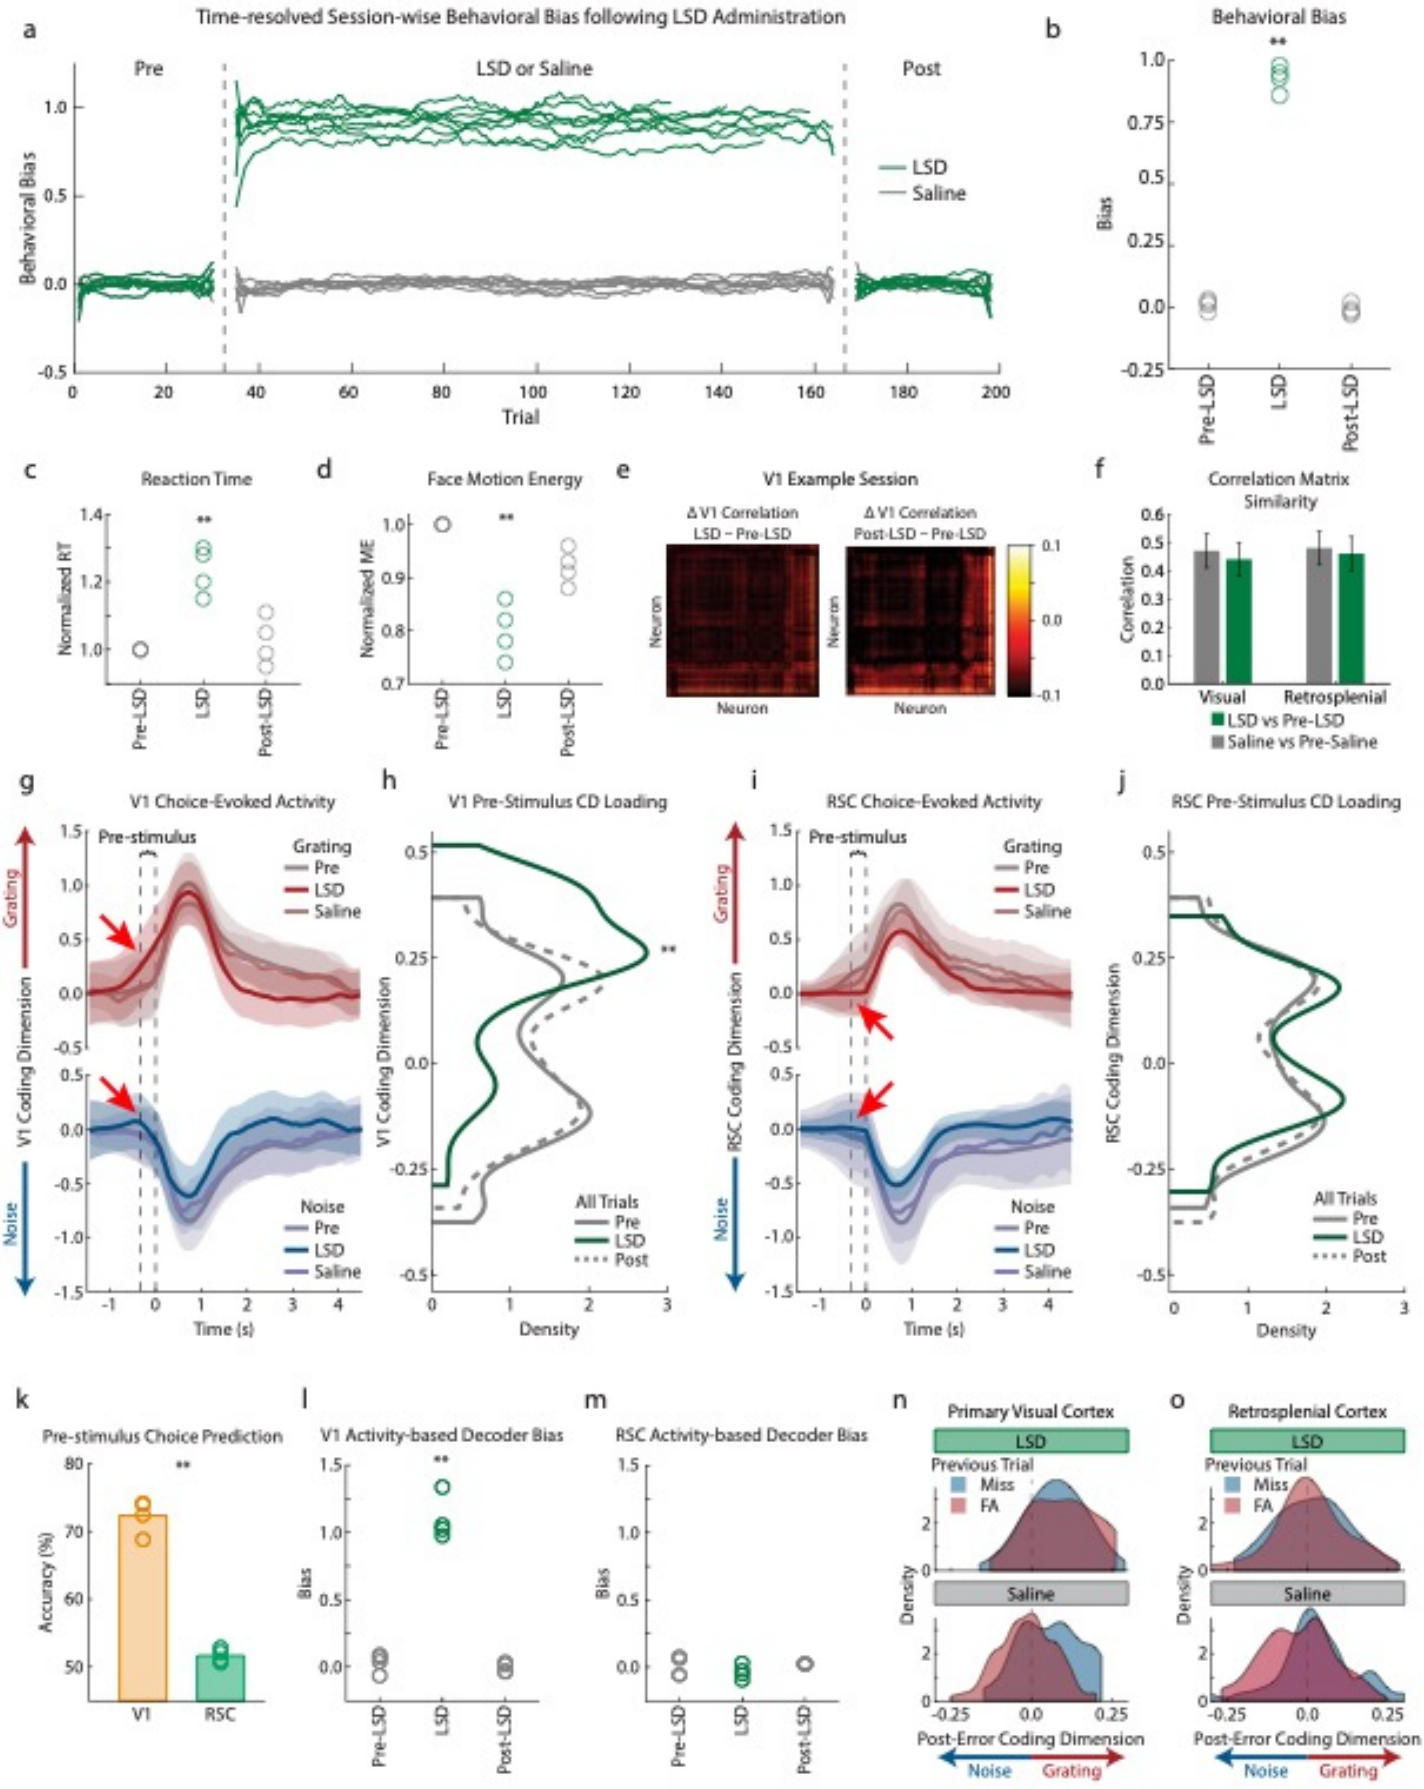
a
Time-resolved Session-wise Behavioral Bias following LSD Administration
Pre
LSD or Saline
Post
LSD
Saline
Behavioral Bias
Trial
b
Behavioral Bias
Bias
Pre-LSD
LSD
Post-LSD
c
Reaction Time
Normalized RT
d
Face Motion Energy
Normalized ME
e
V1 Example Session
Δ V1 Correlation LSD − Pre-LSD
Δ V1 Correlation Post-LSD − Pre-LSD
Neuron
f
Correlation Matrix Similarity
Correlation
Visual
Retrosplenial
LSD vs Pre-LSD
Saline vs Pre-Saline
g
V1 Choice-Evoked Activity
Pre-stimulus
Grating
Pre
LSD
Saline
Noise
V1 Coding Dimension
Time (s)
h
V1 Pre-Stimulus CD Loading
All Trials
Post
Density
i
RSC Choice-Evoked Activity
RSC Coding Dimension
j
RSC Pre-Stimulus CD Loading
k
Pre-stimulus Choice Prediction
Accuracy (%)
V1
RSC
l
V1 Activity-based Decoder Bias
m
RSC Activity-based Decoder Bias
n
Primary Visual Cortex
LSD
Previous Trial
Miss
FA
Saline
Density
Post-Error Coding Dimension
Noise
Grating
o
Retrosplenial Cortex

Figure 5: LSD shifts behavioral bias and V1 coding dimension representation in the grating direction. (a) Behavioral bias in every mouse, every session plotted before, during, and after the influence of LSD or saline. Sliding window behavioral bias traces were shown for LSD (green) and saline (gray) sessions, illustrating a shift toward grating responses under LSD. (b) The shift in behavioral bias toward grating responses was statistically significant under LSD ($n = 4$ mice, paired t-test, $p < 0.001$). (c-d) Reaction time and motion energy were significantly altered under LSD but normalized in the post-LSD session ($n = 4$ mice, paired t-test, $p < 0.01$). (e-f) Analysis of correlation structure. The change in correlation structure in spontaneous neural activity was examined as a function of time after LSD or saline administration. Sample delta correlation matrices are shown in (e). Panel (f) illustrates comparable changes in correlation structure across saline and LSD injections ($n = 4$ mice, paired t-test, $p > 0.5$). (g) V1 coding dimension for grating (red hues) and noise (blue hues) responses in pre-injection, LSD, and saline conditions. Evoked activity was highly similar across conditions. (h) V1 pre-stimulus coding dimension shifted significantly toward the grating representation under LSD but not under saline ($n = 4$ mice, paired t-test, $p < 0.01$). (i) RSC coding dimension for grating (red hues) and noise (blue hues) responses in pre-injection, LSD, and saline conditions. (j) There were no statistically significant shifts in pre-stimulus RSC activity under LSD compared to saline. (k) V1 activity, but not RSC, exhibited pre-stimulus predictivity over perceptual choice under LSD. (l) A linear choice decoder trained on pre-LSD coding dimension activity in V1 illustrated a statistically significant bias toward grating responses under LSD ($n = 4$ mice, paired t-test, $p > 0.5$). (m) In addition to chance-level performance, a choice decoder trained on RSC activity exhibited no bias under LSD, diverging from the mouse's behavioral bias ($n = 4$ mice, paired t-test, $p > 0.5$). (n) The distribution of coding dimension values in the pre-stimulus period following a high-selectivity error did not exhibit any differences across error type, indicating a loss of error-corrective activity ($n = 4$ mice, Wilcoxon signed-rank test, $p < 0.001$ saline, $p > 0.25$ LSD). (o) As in V1, there was a loss of post-error correction in the RSC coding dimension following high-selectivity errors ($n = 4$ mice, Wilcoxon signed-rank test, $p < 0.001$ saline, $p > 0.25$ LSD). Confidence intervals are presented as mean ± s.e.m.

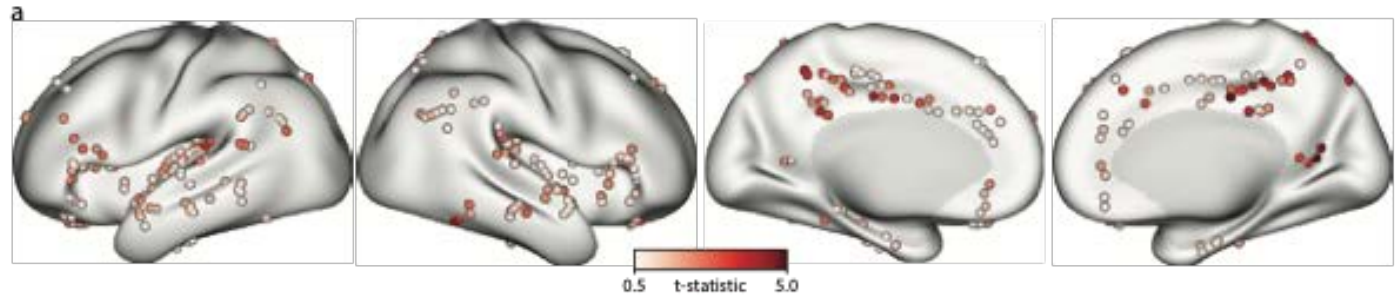
a
0.5
t-statistic
5.0

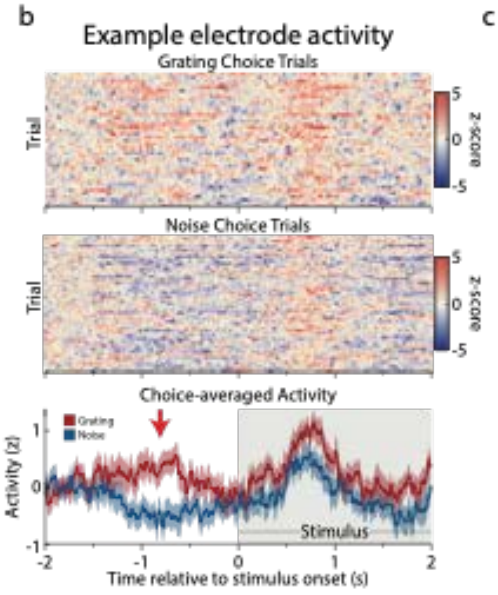
b
Example electrode activity
Grating Choice Trials
Trial
z-score
Noise Choice Trials
Choice-averaged Activity
Grating
Noise
Activity (z)
Stimulus
Time relative to stimulus onset (s)

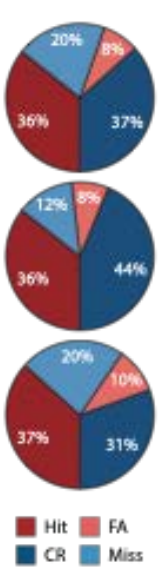
c
20%
8%
36%
37%
12%
8%
36%
44%
20%
10%
37%
31%
Hit
FA
CR
Miss

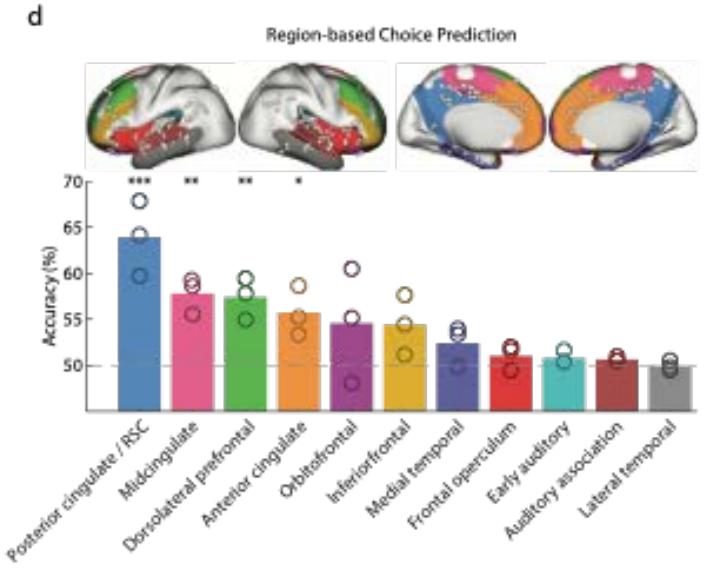
d
Region-based Choice Prediction
Accuracy (%)
Posterior cingulate / RSC
Midcingulate
Dorsolateral prefrontal
Anterior cingulate
Orbitofrontal
Inferiorfrontal
Medial temporal
Frontal operculum
Early auditory
Auditory association
Lateral temporal

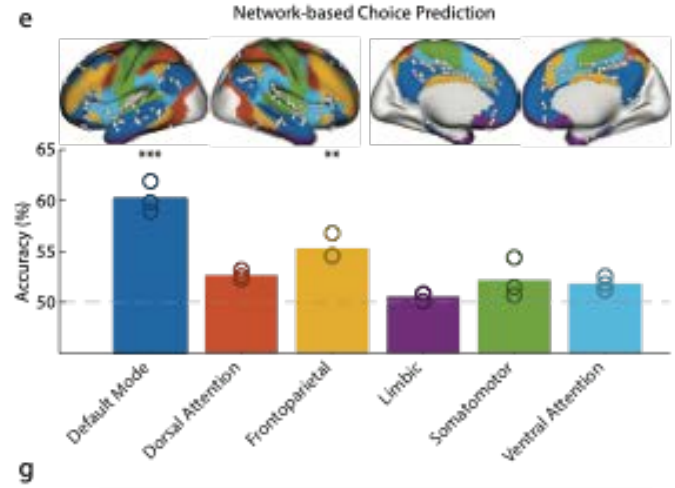
e
Network-based Choice Prediction
Accuracy (%)
Default Mode
Dorsal Attention
Frontoparietal
Limbic
Somatomotor
Ventral Attention

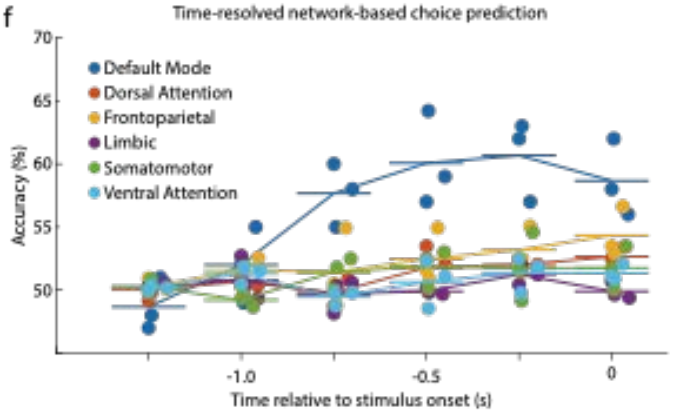
f
Time-resolved network-based choice prediction
Default Mode
Dorsal Attention
Frontoparietal
Limbic
Somatomotor
Ventral Attention
Accuracy (%)
Time relative to stimulus onset (s)

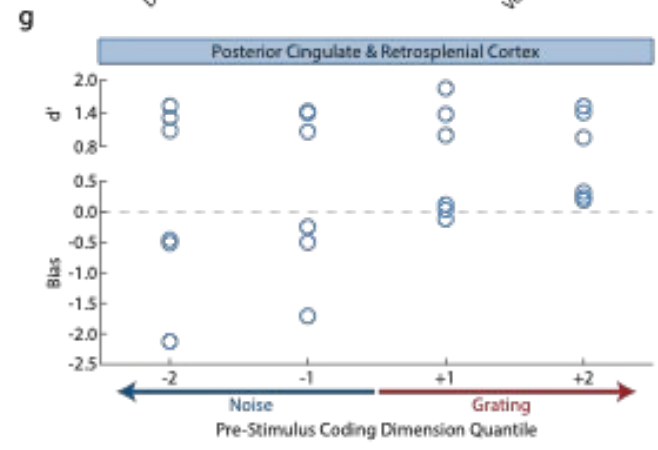
g
Posterior Cingulate & Retrosplenial Cortex
d'
Bias
Noise
Grating
Pre-Stimulus Coding Dimension Quantile

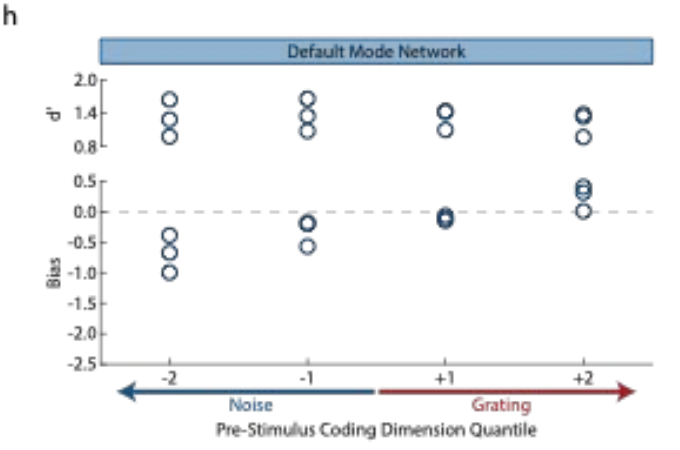
h
Default Mode Network
d'
Bias
Noise
Grating
Pre-Stimulus Coding Dimension Quantile

Figure 6: Pre-stimulus activity in human posterior medial cortex predicts perceptual report. (a) T-statistic map illustrating differences in pre-visual-stimulus activity across grating and noise responses across all electrodes in all subjects. Higher values indicate electrodes with greater differences in pre-stimulus activity across perceptual response types. (b) Single-electrode example data. The top raster illustrates single-trial activity in a single electrode from the posterior cingulate cortex in all trials with grating responses. The middle raster illustrates the same in all trials with noise responses. The bottom time series plots the mean over the grating and noise trials, illustrating separation in pre-stimulus spontaneous activity. (c) Distribution of behavioral responses in three human subjects. (d) Areal choice decoding. Each electrode was assigned to a Human Connectome Project cortical parcel, which we hierarchically grouped into cortical areas[32]. For areas with coverage in all three subjects, we trained a linear classifier to predict perceptual choice based on mean pre-stimulus activity in each subject. The posterior cingulate cortex exhibited the strongest pre-stimulus predictivity over choice (classification was performed over 250-ms epochs in the 250 ms preceding stimulus onset and compared to classification performance at randomly selected epochs during the inter-trial interval using a permutation test; * designates $p < 0.01$ after Bonferroni correction for 11 comparisons). (e-f) Network-based choice decoding. To test the role of brain networks, we trained a linear classifier on mean network activity in six networks[34] with electrode mapping coverage. The default mode network (blue) and frontoparietal control network exhibited statistically significant predictivity over choice (n = 3 subjects; classification was performed over 250-ms pre-stimulus epochs and compared to randomly selected epochs in the inter-trial interval using a permutation test; * designates $p < 0.01$ after Bonferroni correction for six comparisons). (g-h) Behavioral bias and d' in perceptual choice as a function of pre-stimulus activity quantiles, with each electrode sign-oriented such that increased activity predicted a grating response. Perceptual bias, but not behavioral d', was modulated by pre-stimulus neural activity in both the posterior cingulate/RSC and the default mode network (n = 3 subjects).

a

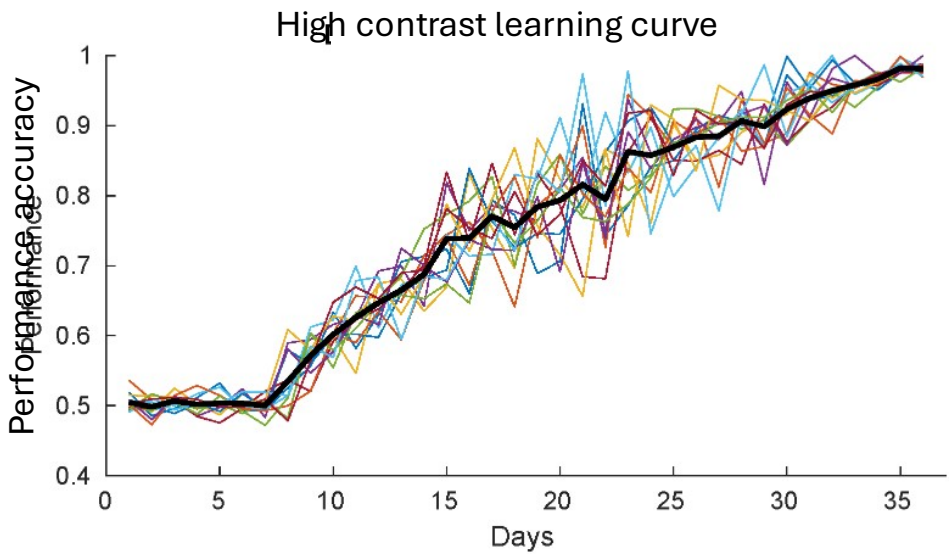
High contrast learning curve
Performance accuracy
1
0.9
0.8
0.7
0.6
0.5
0.4
0
5
10
15
20
25
30
35
Days

b

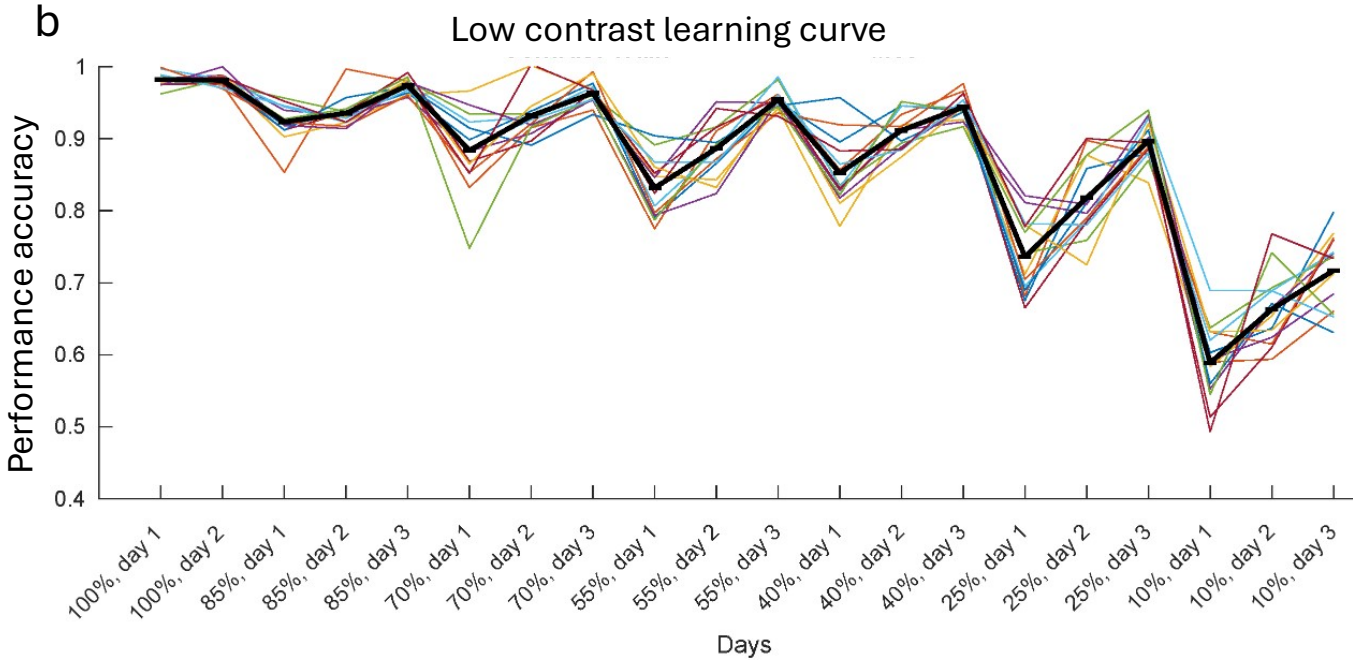
Low contrast learning curve
Performance accuracy
1
0.9
0.8
0.7
0.6
0.5
0.4
100%, day 1
100%, day 2
85%, day 1
85%, day 2
85%, day 3
70%, day 1
70%, day 2
70%, day 3
55%, day 1
55%, day 2
55%, day 3
40%, day 1
40%, day 2
40%, day 3
25%, day 1
25%, day 2
25%, day 3
10%, day 1
10%, day 2
10%, day 3
Days

c

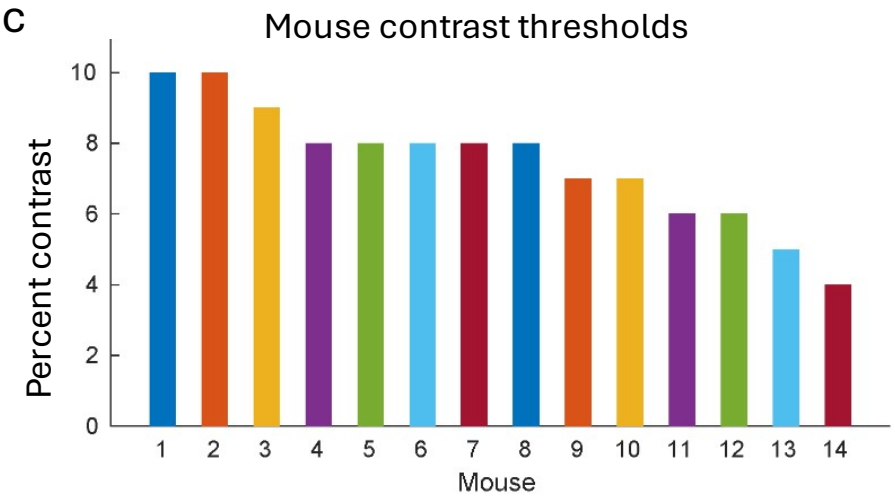
Mouse contrast thresholds
Percent contrast
10
8
6
4
2
0
1
2
3
4
5
6
7
8
9
10
11
12
13
14
Mouse

Extended Data Figure 1: Mice learn the two-alternative forced choice task in two training phases. (a) During the pre-training stage, mice are trained to lick both spouts without bias before advancing to discriminating between white noise and high-contrast (100% contrast) drifting gratings. Mice achieve criterion performance, defined as 90% accuracy for three consecutive days, after approximately 35 days of training (*n = 14* mice). (b) Mice that reach criterion for high-contrast stimuli are then trained to perform the task with progressively lower contrast gratings. Contrast levels are reduced step-wise every three days until performance stabilizes at 65% accuracy over ~20 days (*n = 14* mice). (c) The contrast detection threshold for each mouse, defined as the mean contrast at which mice reach stable performance, is shown. Stimuli within ±2% contrast of the mean threshold are used in subsequent sessions. Notably, mice accurately detect lower contrast levels when only very low-contrast gratings are presented within a session, compared to sessions with mixed contrast levels.

a

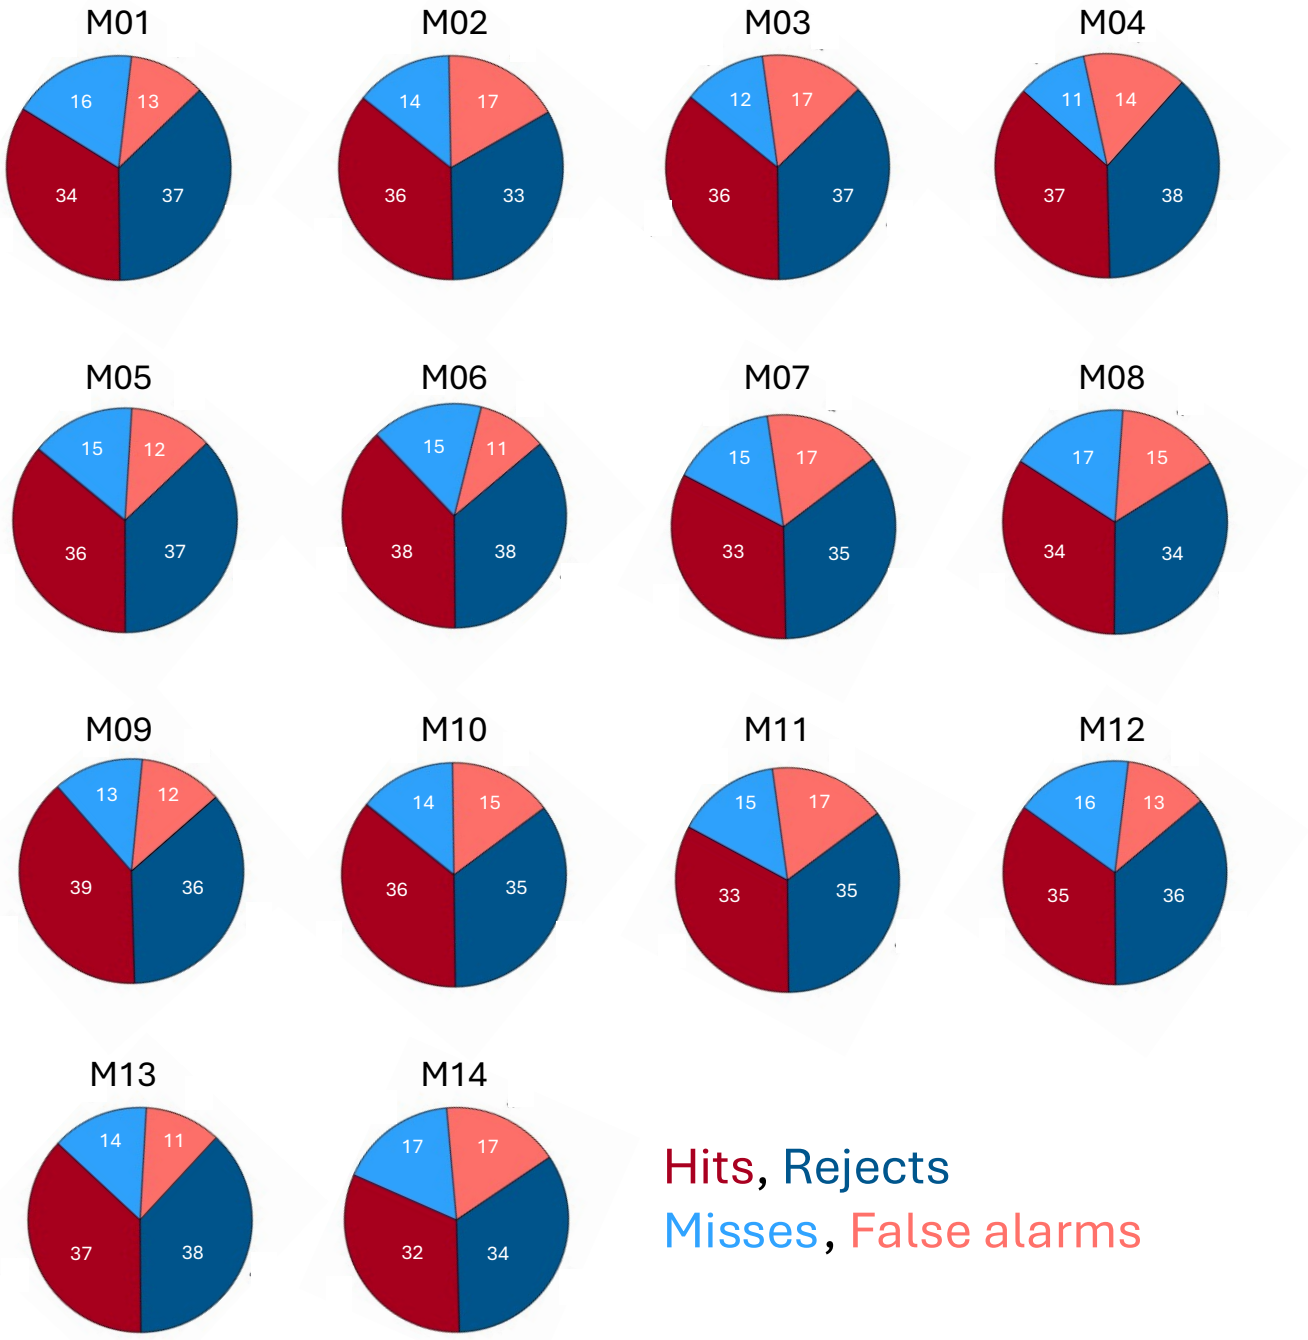

Hits, Rejects
Misses, False alarms

b

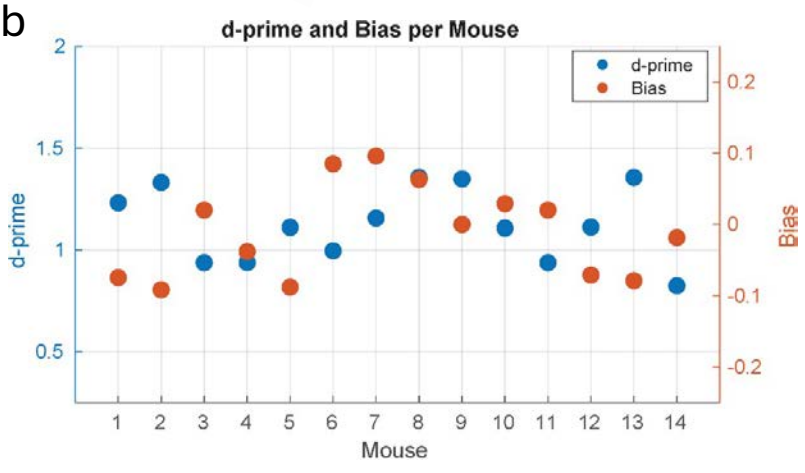

Extended data Figure 2: Distribution of behavioral responses and accuracy across all mice. (a) Percentage of high-selectivity trials that are hits, rejects, misses, and false alarms in each mouse. Although we picked grating contrast thresholds per mouse to aim for 65% accuracy in performance, mice often improved their performance over the weeks during which the experiments were conducted. (b) Behavioral d prime and bias across each mouse in the study.

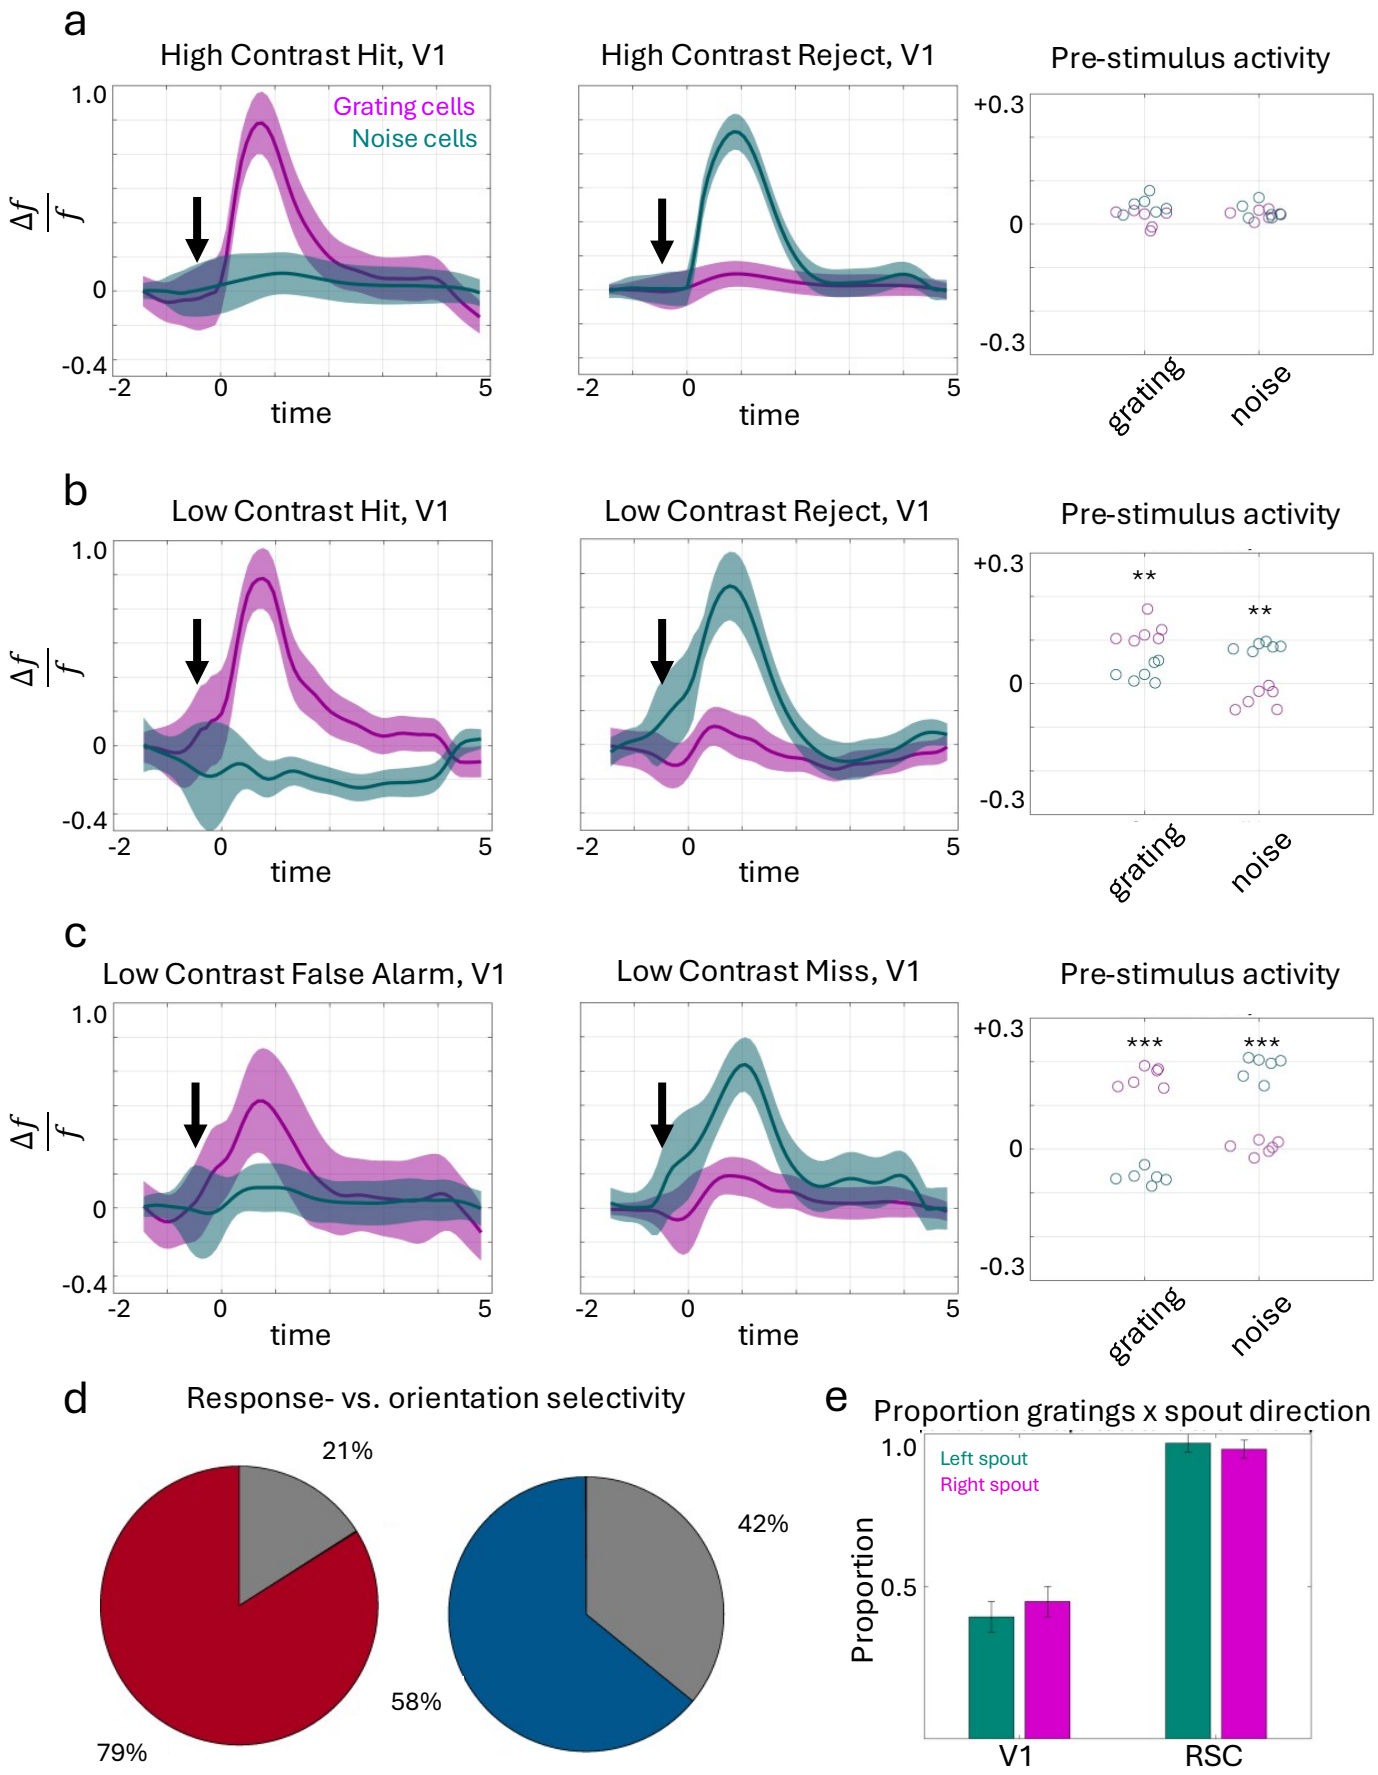
a
High Contrast Hit, V1
High Contrast Reject, V1
Pre-stimulus activity
Grating cells
Noise cells
1.0
0
-0.4
-2
0
5
time
+0.3
0
-0.3
grating
noise
b
Low Contrast Hit, V1
Low Contrast Reject, V1
Pre-stimulus activity
**
**
c
Low Contrast False Alarm, V1
Low Contrast Miss, V1
Pre-stimulus activity
***
***
d
Response- vs. orientation selectivity
21%
79%
42%
58%
e
Proportion gratings x spout direction
Left spout
Right spout
Proportion
0.5
V1
RSC

Extended data Figure 3: Single cell analysis of response selective neurons in visual cortex (V1). (a) Grating- and noise-selective neurons in high contrast hits and rejects. There is no significant difference in mean pre-stimulus activity ($n = 6$, paired t-test, $p > 0.8$). (b) Grating- and noise-selective neurons in low contrast hits and rejects. There is a significant difference in pre-stimulus activity ($n = 6$, paired t-test, $p < 0.01$). (c) Grating- and noise-selective neurons in low contrast false alarms and misses. There is a significant difference in pre-stimulus activity ($n = 6$, paired t-test, $p < 0.001$). (d) Orientation selectivity in response-selective cells. In a subset of sessions, following 2AFC task performance, lick spouts were retracted, and mice were shown drifting gratings in 8 directions to compute neuron-specific orientation specificity. In these sessions, we then examined the proportion of response-selective neurons that exhibited orientation selectivity for visual stimuli. For grating-response selective neurons, we found that on average, 79% of neurons exhibited orientation selectivity for downward drifting gratings (the same orientation as the grating stimulus in the 2AFC task). For noise-response selective neurons, we found that on average 58% of neurons (blue) exhibited orientation selectivity for one of the other 7 directions. (e) Grating-response neuron selectivity for left-grating and right-grating responses. We analyzed whether the proportion of grating-selective neurons varied depending on whether the grating stimulus was reward by the left- or right-spout. Computing the mean grating-selective neuron to noise-selective neuron proportion in two separate cohorts of 3 mice, assigned to left-grating or right-grating responses, we found no statistically significant difference ($n = 6$, paired t-test, $p > 0.5$ in V1 and RSC). Confidence intervals are presented as mean $\pm$ s.e.m.

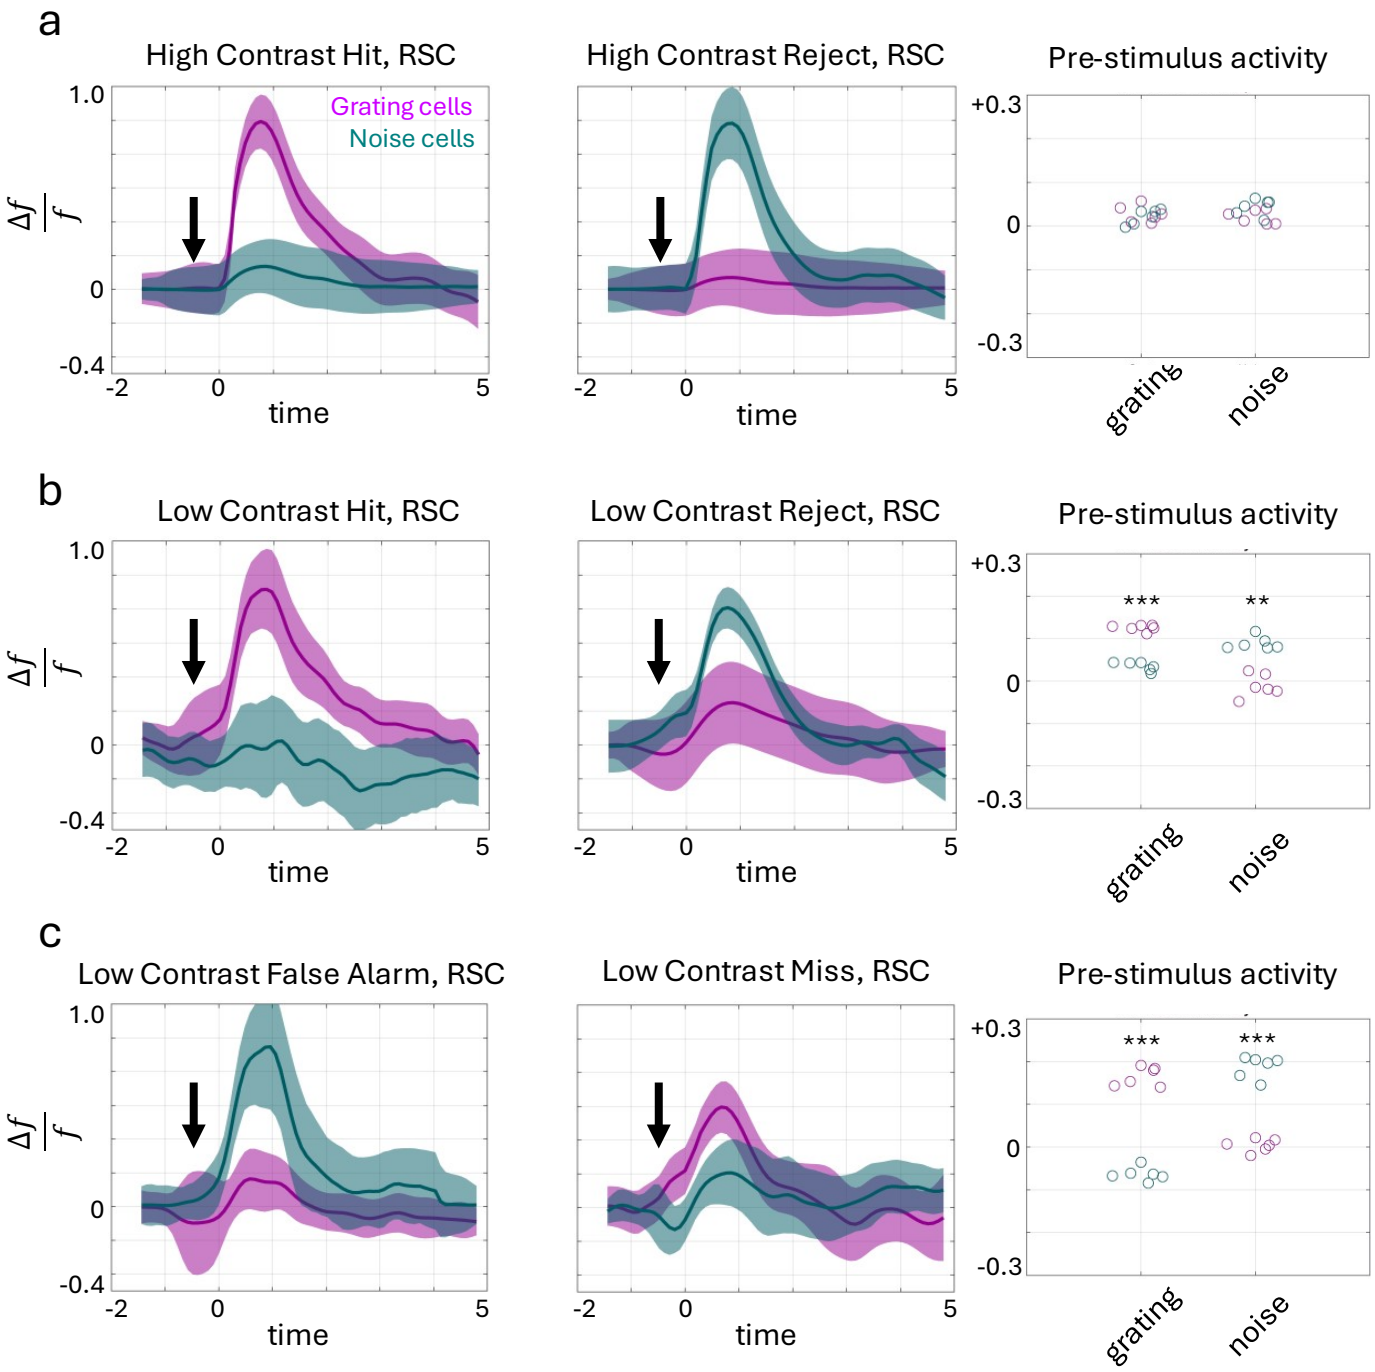
a
High Contrast Hit, RSC
High Contrast Reject, RSC
Pre-stimulus activity
Grating cells
Noise cells
1.0
0
-0.4
-2
0
5
time
+0.3
0
-0.3
grating
noise
b
Low Contrast Hit, RSC
Low Contrast Reject, RSC
Pre-stimulus activity
***
**
c
Low Contrast False Alarm, RSC
Low Contrast Miss, RSC
Pre-stimulus activity
***
***

Extended data Figure 4: Single cell analysis of response selective neurons in retrosplenial cortex (RSC). (a) Grating- and noise-selective neurons in high contrast hits and rejects. There is no significant difference in mean pre-stimulus activity (*n* = 6, paired t-test, $p > 0.8$). (b) Grating- and noise-selective neurons in low contrast hits and rejects. There is a significant difference in pre-stimulus activity (*n* = 6, paired t-test, $p < 0.001$). (c) Grating- and noise-selective neurons in low contrast false alarms and misses. There is a significant difference in pre-stimulus activity (n = 6, paired t-test, $p < 0.001$). Confidence intervals are presented as mean $\pm$ s.e.m.

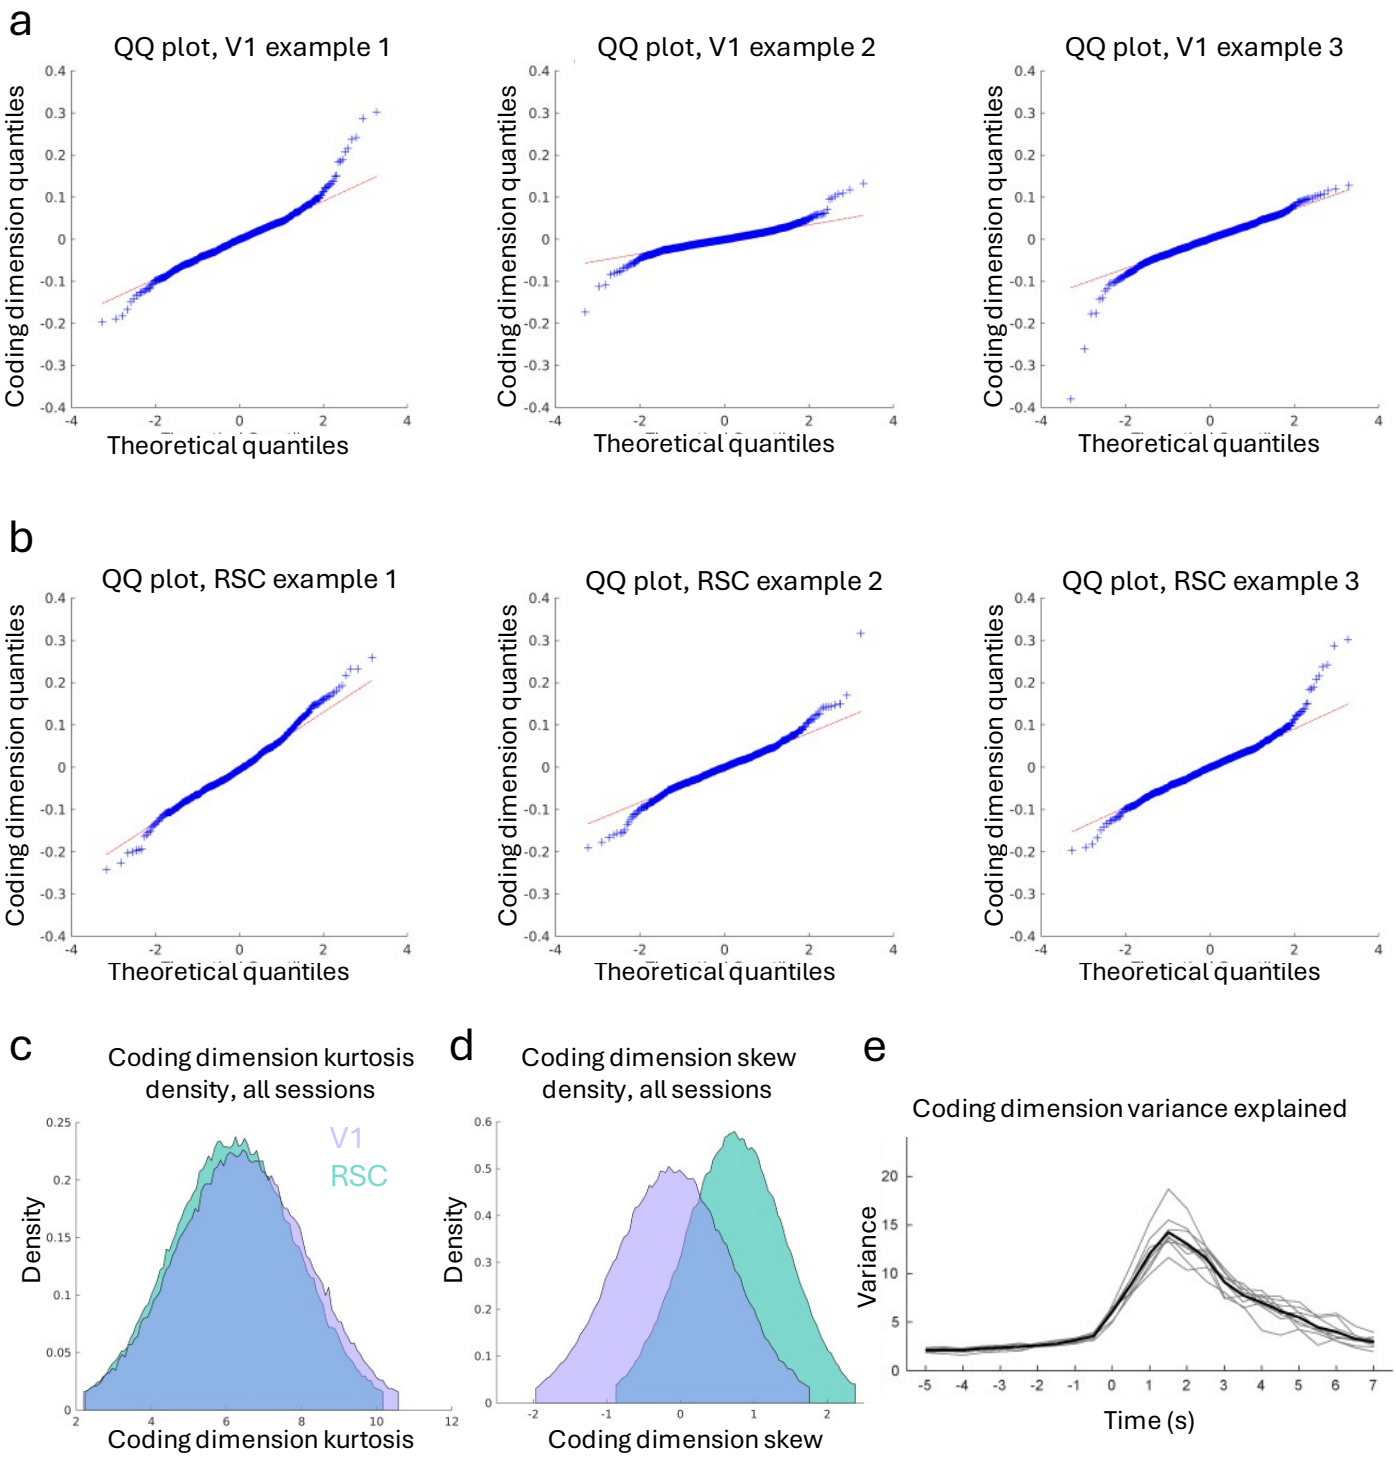
a
QQ plot, V1 example 1
QQ plot, V1 example 2
QQ plot, V1 example 3
Coding dimension quantiles
Theoretical quantiles
b
QQ plot, RSC example 1
QQ plot, RSC example 2
QQ plot, RSC example 3
Coding dimension quantiles
Theoretical quantiles
c
Coding dimension kurtosis density, all sessions
V1
RSC
Density
Coding dimension kurtosis
d
Coding dimension skew density, all sessions
Density
Coding dimension skew
e
Coding dimension variance explained
Variance
Time (s)

Extended data Figure 5: Distribution of coding dimension weights in V1 and RSC. (a) QQ plots from three representative V1 sessions illustrating computed coding dimension weights (y-axis) versus quantiles expected in a gaussian distribution (x-axis). Note heavy tails in both positive and negative weights. (b) QQ plots from three representative RSC sessions illustrating computed coding dimension weights (y-axis) versus quantiles expected in a gaussian distribution (x-axis). Note skew toward positive values. (c) Distribution of kurtosis values for CD weights in V1 and RSC. High kurtosis indicates heavy-tailed distributions. There is no difference in the heavy-tailedness between V1 and RSC ($n = 6$ mice, Wilcoxon signed-rank test, $p > 0.8$). (d) Distribution of skew values for CD weights in V1 and RSC. Whereas V1 weights do not exhibit substantial skew on average, the mean skew in RSC coding dimension weights is positive. The difference in skew values for CD weights between V1 and RSC is statistically significant ($n = 6$ mice, Wilcoxon signed-rank test, $p < 0.01$). (e) Illustration of total neural variance explained along the coding dimension over time. Black line represents the mean across mice, grey lines depict individual mouse means.

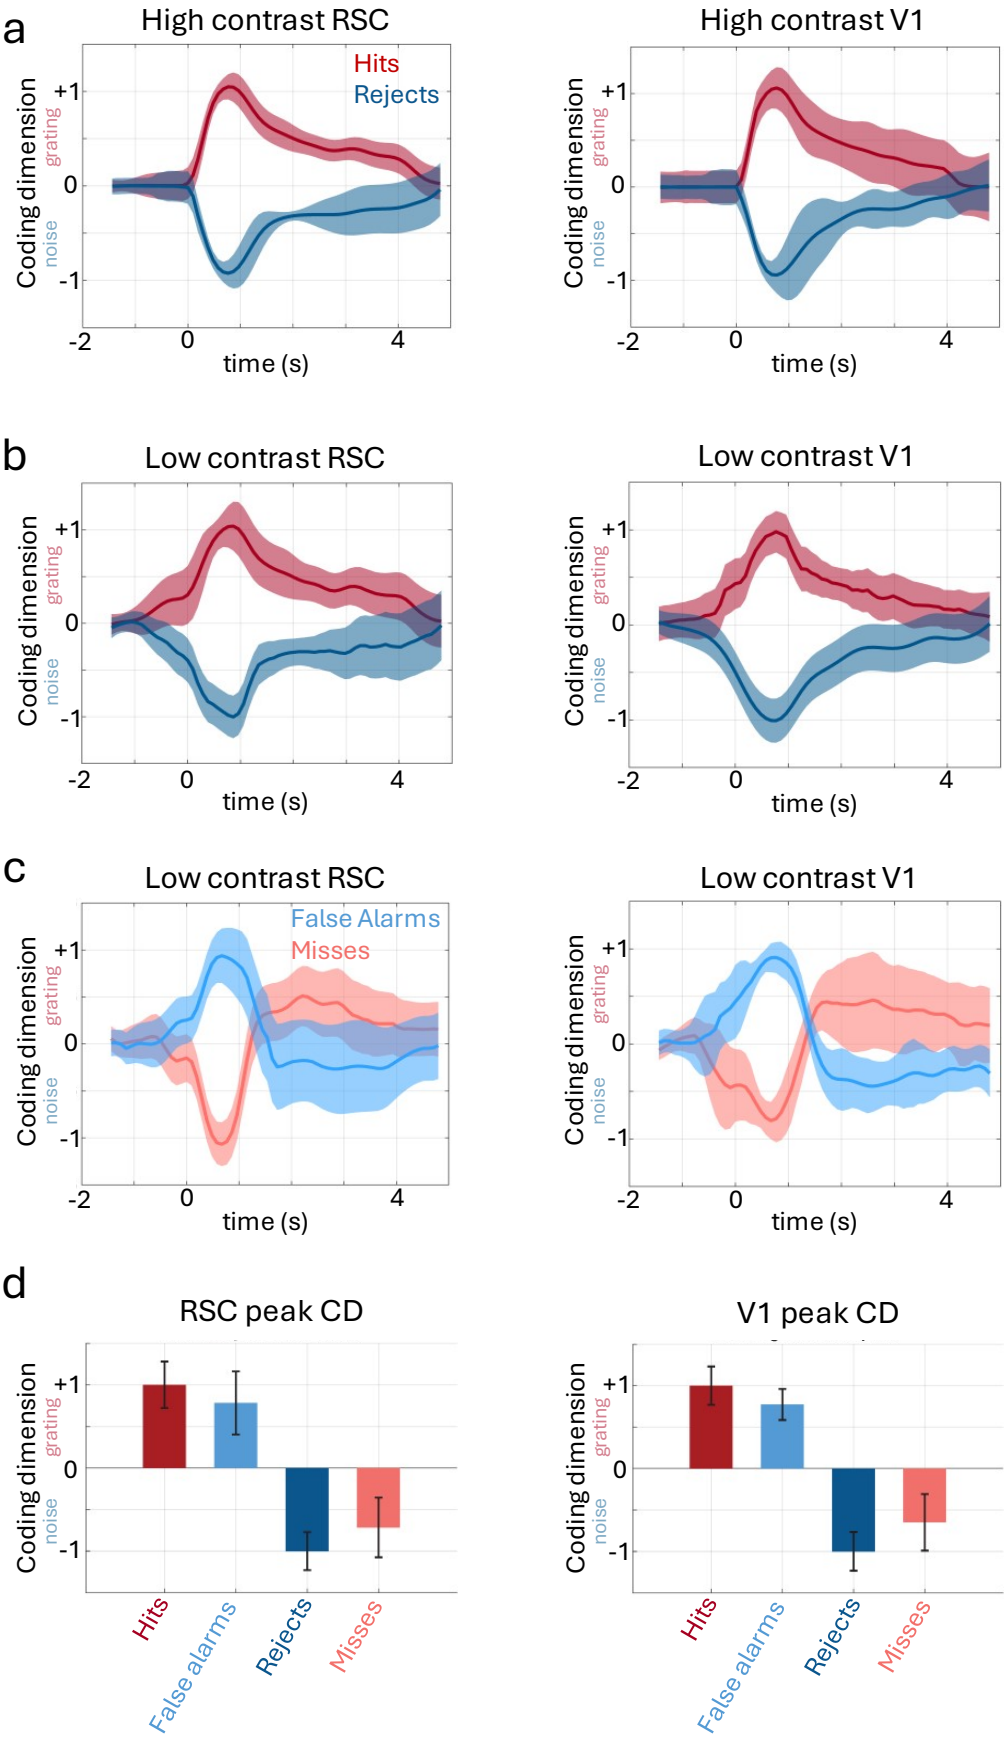
a
High contrast RSC
Hits
Rejects
Coding dimension
grating
noise
+1
0
-1
-2
0
4
time (s)
High contrast V1
b
Low contrast RSC
Low contrast V1
c
Low contrast RSC
False Alarms
Misses
Low contrast V1
d
RSC peak CD
V1 peak CD
Hits
False alarms
Rejects
Misses

Extended data Figure 6: Full evoked coding dimension time series in RSC and V1 exhibit similar shapes and peaks across consistent response types. Evoked RSC and V1 coding dimension time series are depicted during high contrast hits and rejects (a), low contrast hits and rejects (b), and low contrast false alarms and misses (c). Quantification of peak responses are shown in (d), illustrating that error responses exhibit the same sign in peak correct responses. Confidence intervals are presented as mean $\pm$ s.e.m.

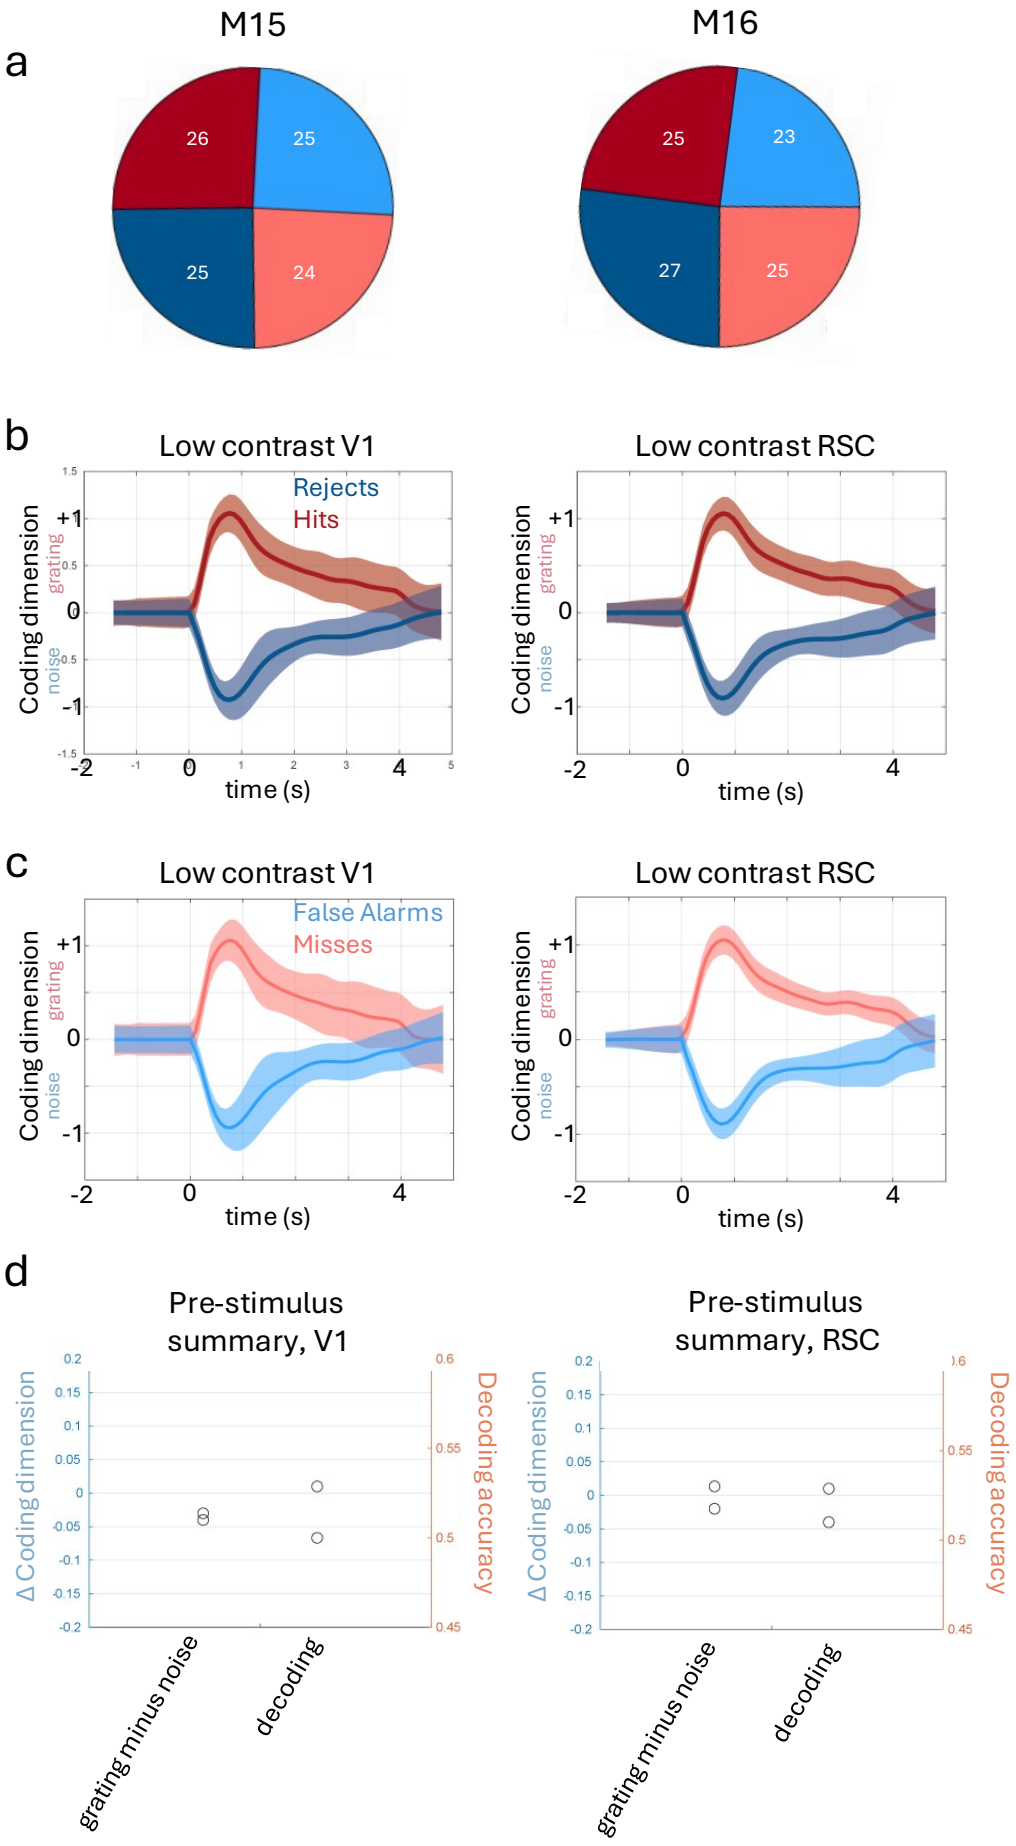
a
M15
26
25
25
24
M16
25
23
27
25
b
Low contrast V1
Rejects
Hits
Coding dimension
grating
noise
time (s)
Low contrast RSC
c
Low contrast V1
False Alarms
Misses
Low contrast RSC
d
Pre-stimulus summary, V1
Δ Coding dimension
Decoding accuracy
grating minus noise
decoding
Pre-stimulus summary, RSC

Extended data Figure 7: Control analysis of data from two mice who did not learn the task and did not exhibit a response bias. (a) Distribution of hits, rejects, false alarms, and misses in two mice. Note that there is no bias amongst response types (grating and noise). In contrast, all other mice who failed to learn the task developed a bias for one of the two lick spouts. (b) Coding dimension computed for hits and rejects for low contrast trials in V1 and RSC. (c) Coding dimension computed for false alarms and misses for low contrast trials in V1 and RSC. (d) Difference in pre-stimulus coding dimension activity between grating and noise trials, and linear decoding from pre-stimulus activity. The mean difference in coding dimension activity across response type is near zero, with no predictive decodability from pre-stimulus activity. These results indicate that the predictive decodability reported in the main text relies on successfully learning the 2AFC visual detection task. Confidence intervals are presented as mean $\pm$ s.e.m.

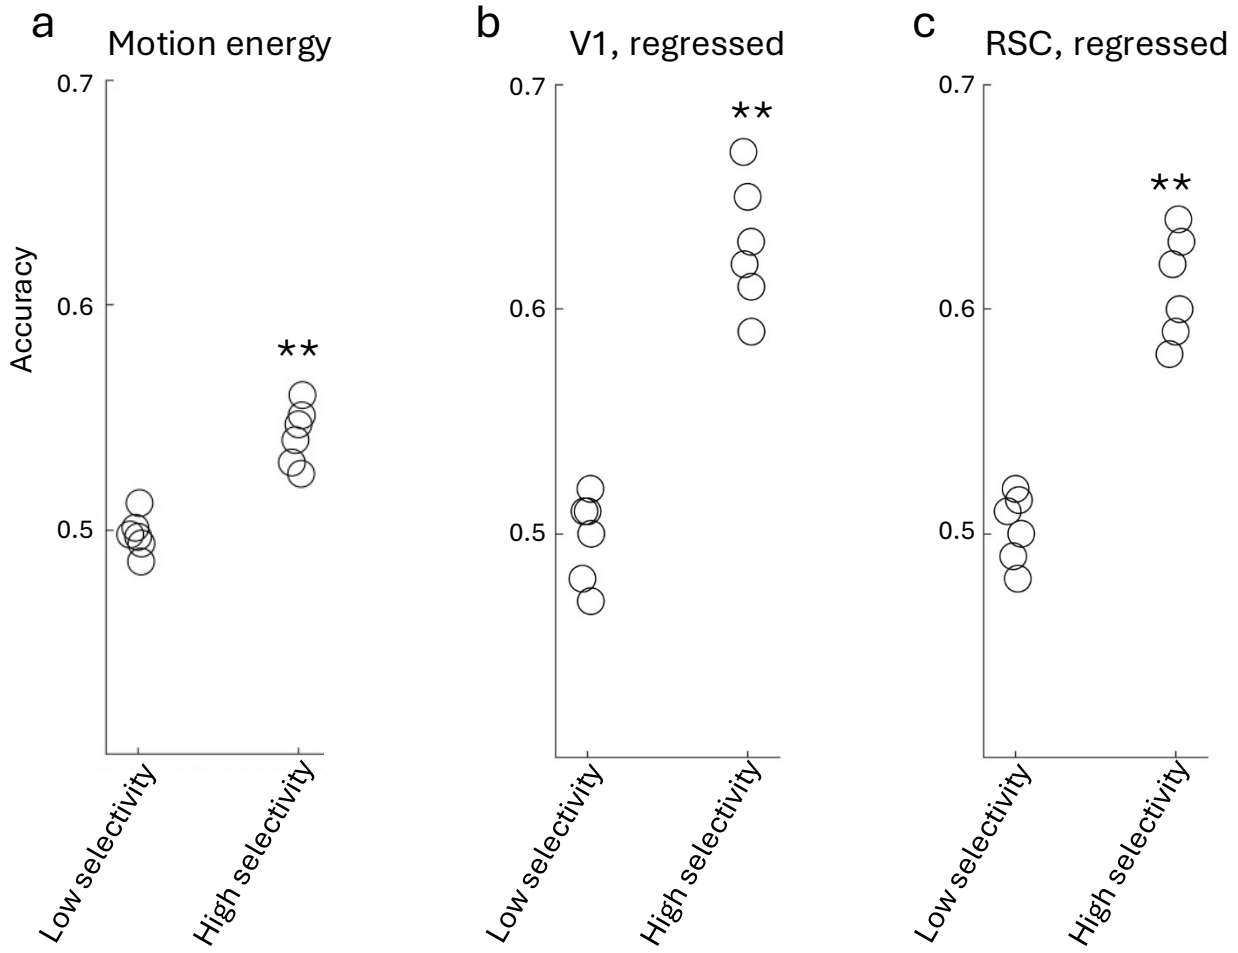

Extended data Figure 8: Motion energy accounts for only a minority portion of pre-stimulus neural choice decodability. (a) Predictive decoding of choice from motion energy in mouse behavioral videos in low selectivity and high selectivity trials. Note: due to the requirement that mice refrain from licking the water spouts in the one second preceding visual stimulus onset, the amount of motion in the pre-stimulus period is modest. High selectivity trials exhibit statistically significant decodability, with 54.5% predictive accuracy on average. Motion energy was computed pixel-wise from behavior of ongoing mouse behavior from underneath the mouse face. To avoid overfitting, an optimal model order was computed by concatenating 250ms epochs preceding visual stimulus and computing a covariance matrix. Laplacian estimate of the Bayesian information criterion was computed on the eigenspectrum leading to an optimal model order of 6. Motion energy was then reconstructed from 6 eigenvalue weighted eigenvectors, and a linear classifier was applied to compute decodability. There is a statistically significant difference in predictive accuracy between low- and high-selectivity trials (classification was performed over 250-ms epochs in the 250 ms preceding stimulus onset and compared to classification performance at randomly selected epochs during the inter-trial interval using a permutation test; $p < 0.01$). (b) Model-order reconstructed motion energy was regressed from V1 coding dimension activity and linear classification was applied to the motion regressed data. Statistically significant predictive accuracy was still present in high selectivity trials (classification was performed over 250-ms epochs in the 250 ms preceding stimulus onset and compared to classification performance at randomly selected epochs during the inter-trial interval using a permutation test; $p < 0.01$). (c) Decodability applied to motion-energy regressed RSC coding dimension activity. Again, statistically significant predictive

accuracy was present in high selectivity trials (classification was performed over 250-ms epochs in the 250 ms preceding stimulus onset and compared to classification performance at randomly selected epochs during the inter-trial interval using a permutation test; $p < 0.01$).

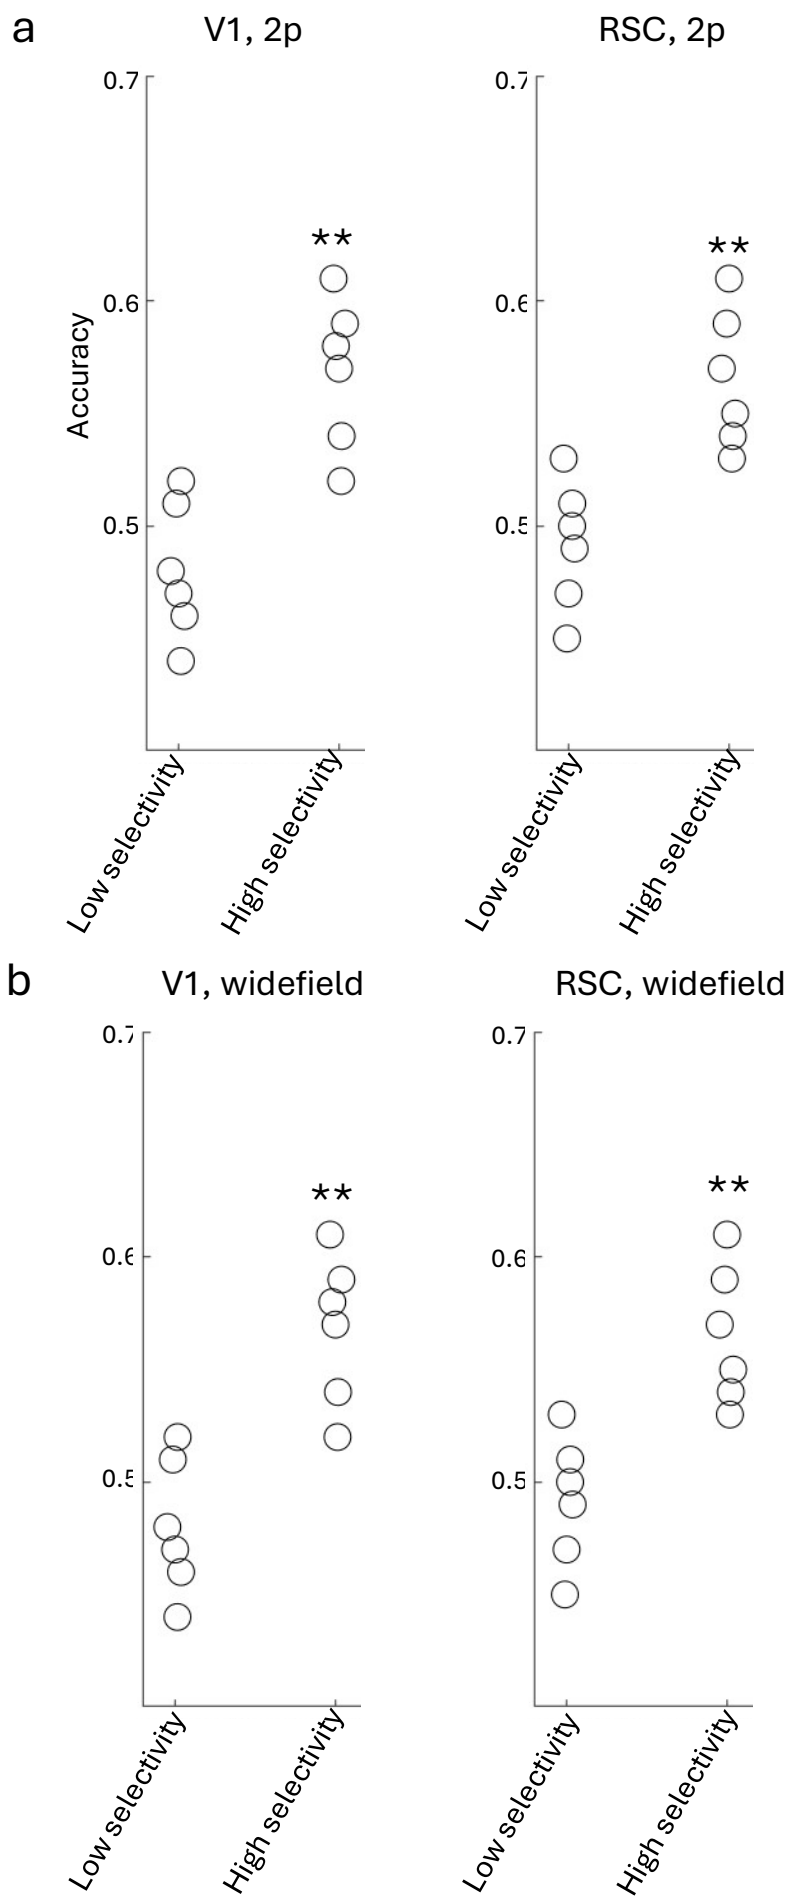
a
V1, 2p
RSC, 2p
Accuracy
0.7
0.6
0.5
**
Low selectivity
High selectivity
b
V1, widefield
RSC, widefield

Extended data Figure 9: Single area decodability. (a) Pre-stimulus decodability with mean single-area activity, over all z-scored neurons in each session, in V1 and RSC. There is statistically significant decodability in V1 and RSC in high-selectivity trials, but the magnitude of this decodability is less than that found in the coding dimension (classification was performed over 250-ms epochs in the 250 ms preceding stimulus onset and compared to classification performance at randomly selected epochs during the inter-trial interval using a permutation test; $p < 0.01$). (b) Pre-stimulus decodability using single areas in widefield data, in V1 and RSC. Although there is statistically significant decodability in each single area, the two-area model reported in the main text performs substantially better (classification was performed over 250-ms epochs in the 250 ms preceding stimulus onset and compared to classification performance at randomly selected epochs during the inter-trial interval using a permutation test; $p < 0.01$).

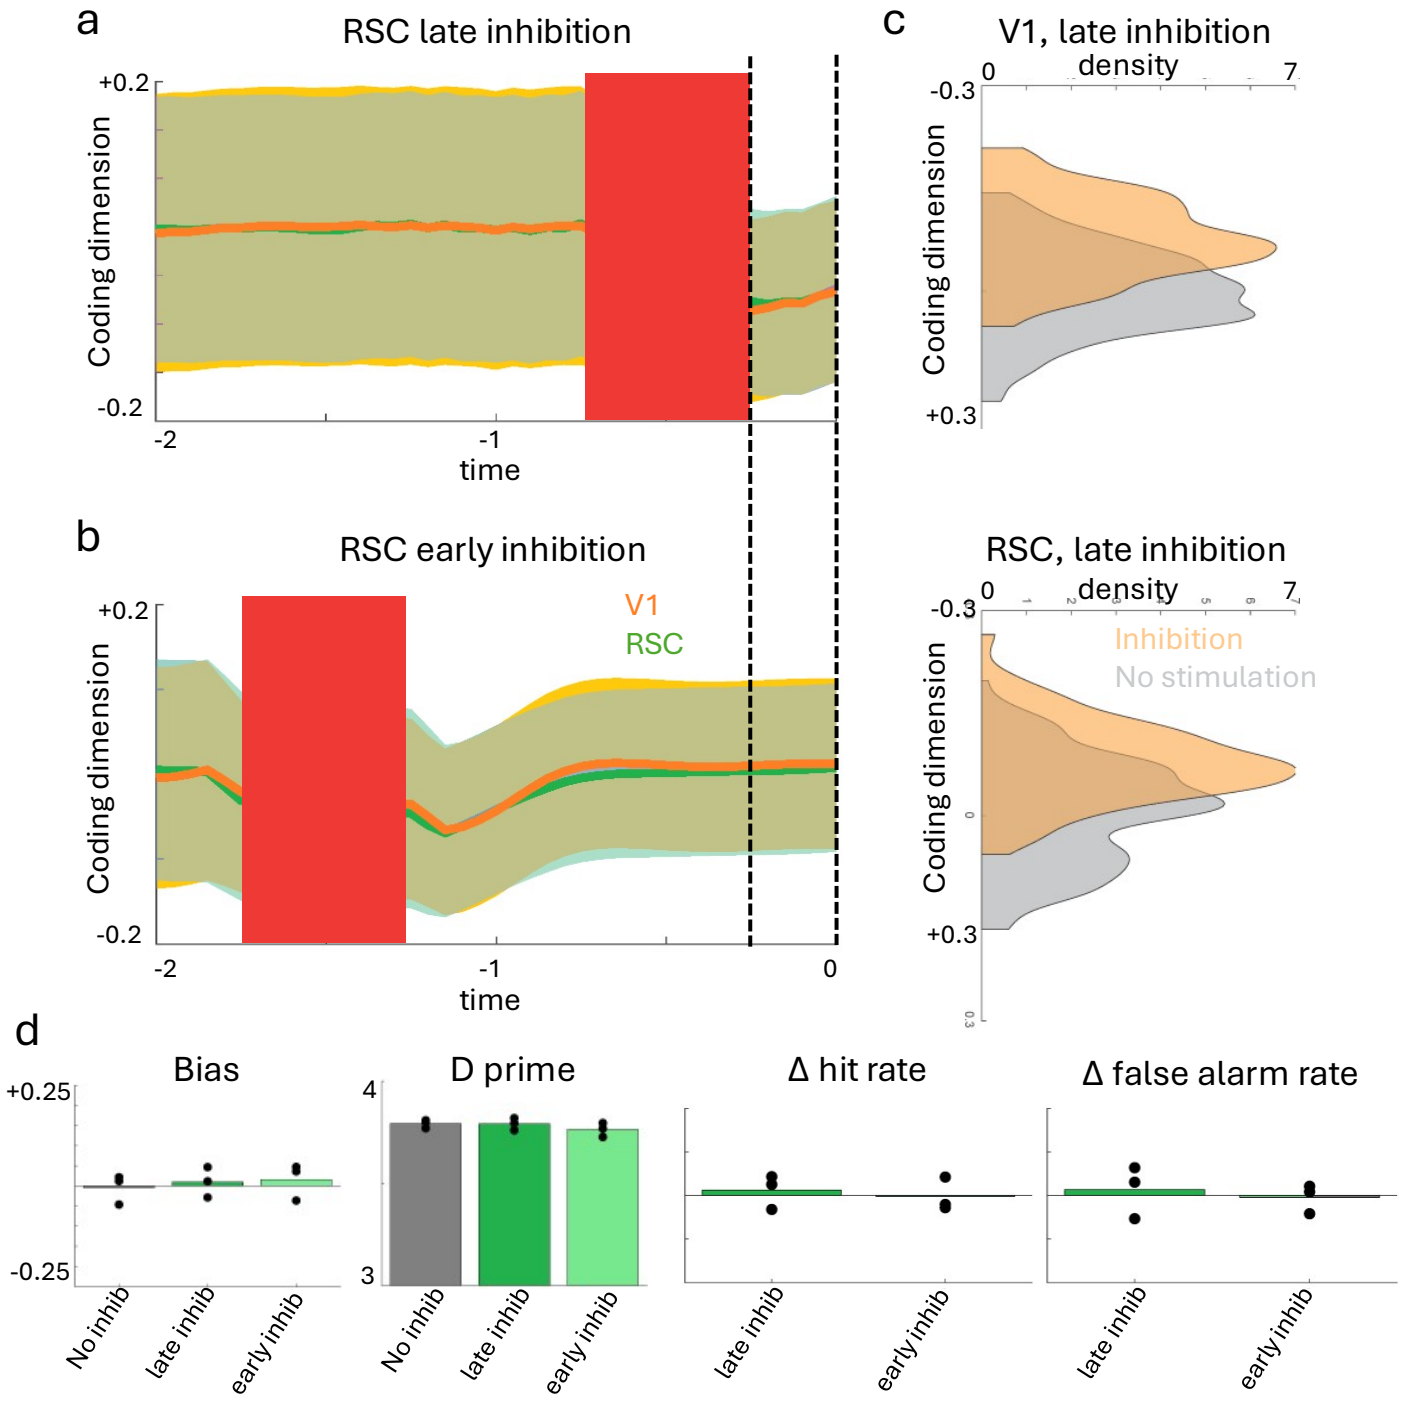
a
RSC late inhibition
+0.2
-0.2
Coding dimension
-2
-1
time
b
RSC early inhibition
V1
RSC
+0.2
-0.2
Coding dimension
-2
-1
0
time
c
V1, late inhibition
density
0
7
-0.3
+0.3
Coding dimension
RSC, late inhibition
density
0
7
-0.3
Inhibition
No stimulation
Coding dimension
+0.3
d
Bias
+0.25
-0.25
No inhib
late inhib
early inhib
D prime
4
3
No inhib
late inhib
early inhib
Δ hit rate
late inhib
early inhib
Δ false alarm rate
late inhib
early inhib

Extended data Figure 10: Optogenetic inhibition of RSC during high-contrast (easy) trials. (A-B) Pre-stimulus coding dimension activity shifts toward the noise direction with both late- and early-inhibition, as with the low-contrast trials shown in the main text. (c)The shift toward noise-representations remains statistically significant in the 250ms preceding visual stimulus in the late-inhibition condition in both V1 and RSC ($n = 3$ mice, Wilcoxon signed-rank test, $p < 0.01$). (d) Despite the shift in neural activity, the behavioral bias, behavioral d prime, change in hit rates, and change in false alarm rates are not statistically significant across no-inhibition, late inhibition, and early inhibition ($n = 3$ mice, paired t-tests, $p > 0.3$). Thus, the influence of pre-stimulus brain state on visual perceptual decisions depends on the ambiguity of the visual stimulus. Confidence intervals are presented as mean $\pm$ s.e.m.

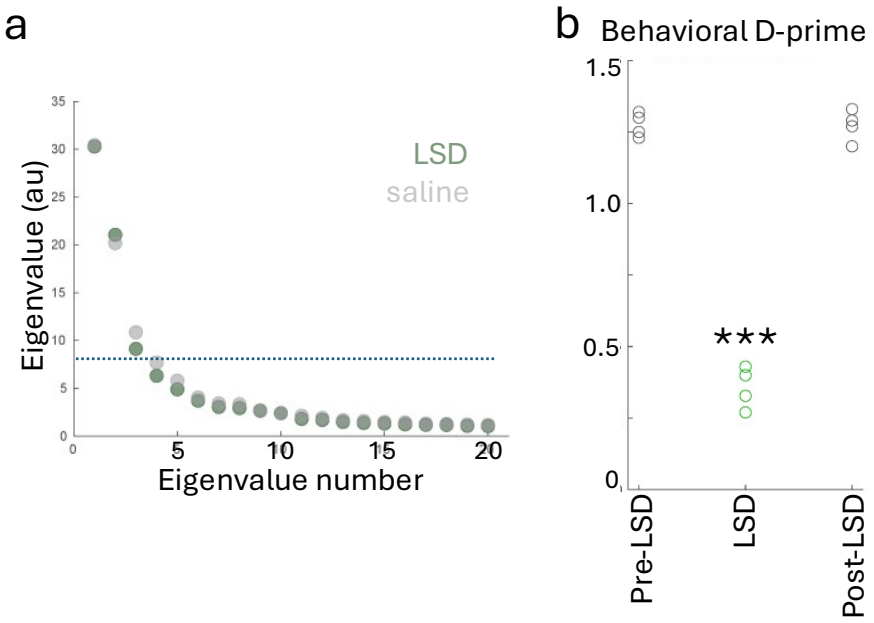

Extended data figure 11: Additional analyses of LSD neural activity and behavior. (a) Comparison of eigenspectrum of difference in inter-trial spontaneous correlation structure between experimental condition (saline or LSD) and pre-injection neural activity, example shown in V1. As we did not obtain the same neurons across sessions, we cannot directly compare correlation structure across saline and LSD conditions. Instead, we compared the dimensionality of the change in correlation structure across conditions, by computing the Laplacian estimate of the Bayesian information criterion across conditions. We found no statistical difference in both V1 and RSC (permutation re-sampling across $n = 16$ sessions, $p > 0.8$). (b) Behavioral D prime declines during LSD. Note that this decline in behavioral discrimination is dictated by bias: given that mice are substantially more prone to reporting a grating response under the influence of LSD, this definitionally reduces perceptual discrimination accuracy ($n = 4$ mice, paired t-tests, $p > 0.001$).

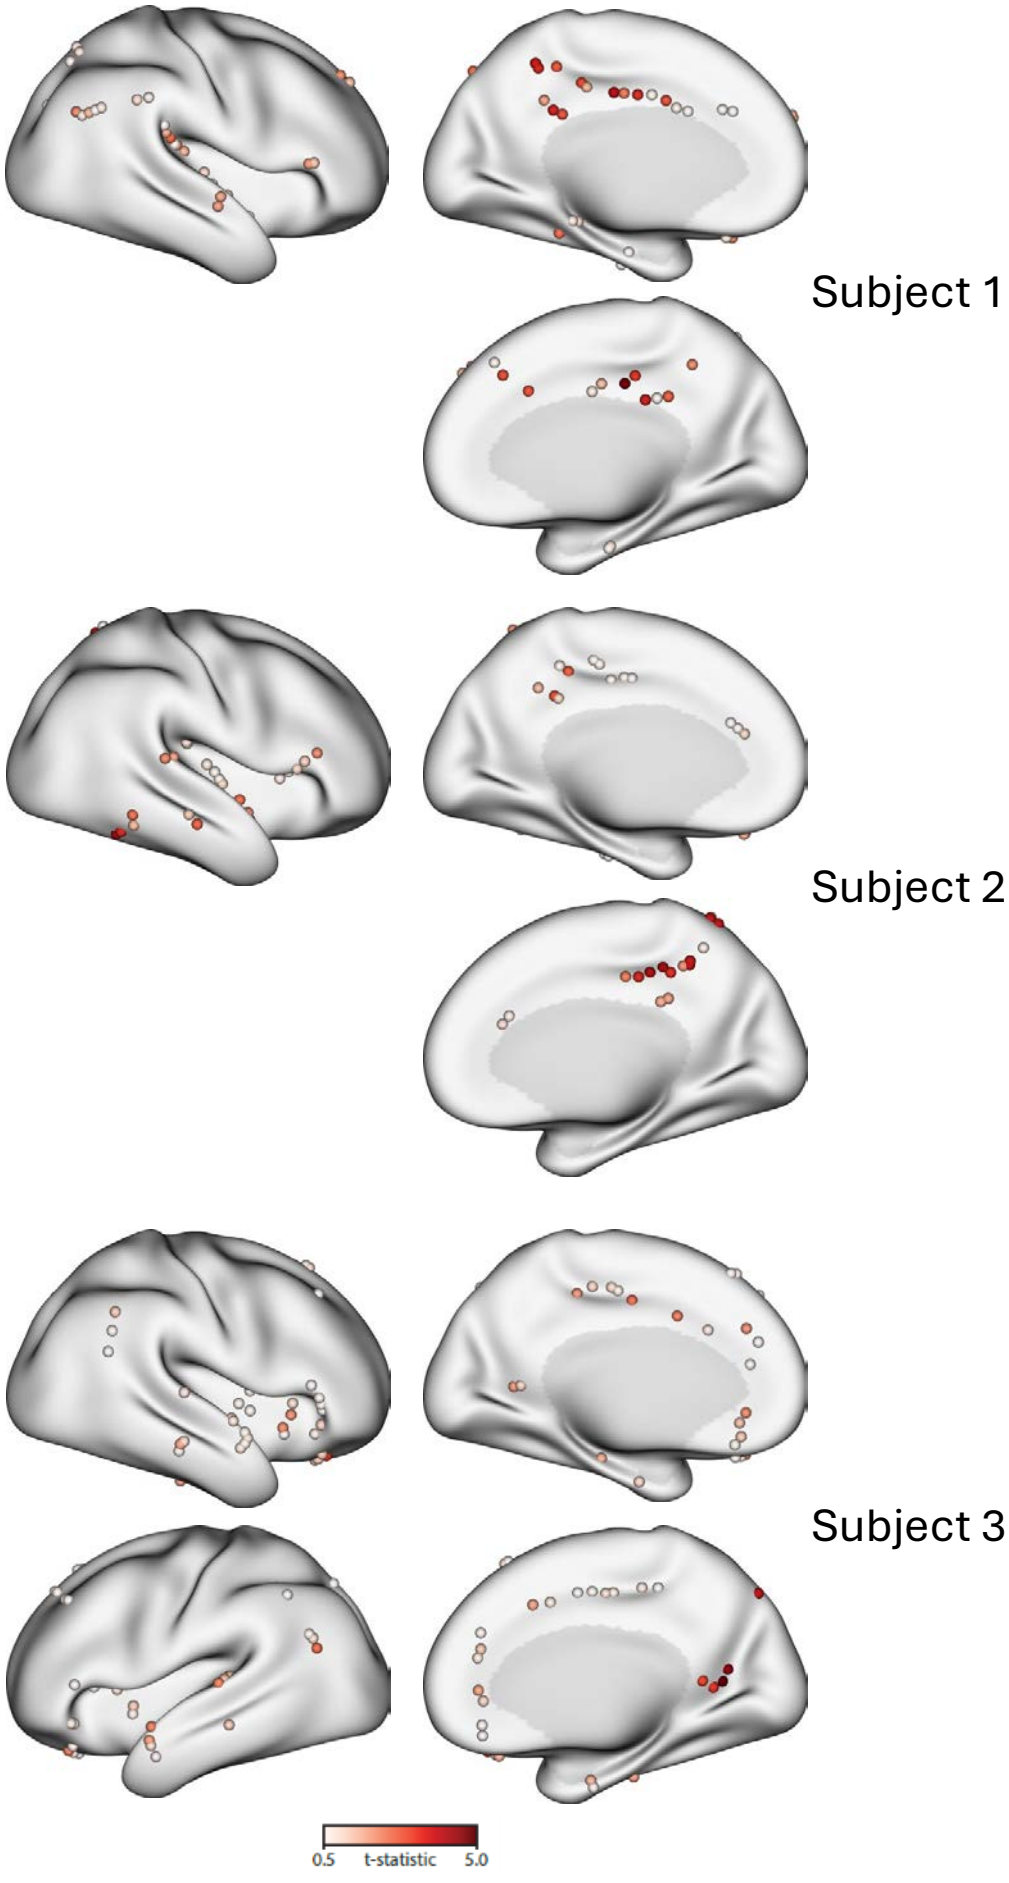
Subject 1
Subject 2
Subject 3
0.5
t-statistic
5.0

Extended data figure 12: T-statistic map illustrating differences in pre-visual-stimulus activity across grating and noise responses across all electrodes in individual subjects. Only hemispheric surfaces with electrode coverage are shown.

Prediction & Choice Category	Trials	Description
prediction = choice = stimulus	741	Prediction matches choice; mouse correct
prediction = choice ≠ stimulus	336	Prediction matches choice; mouse incorrect
prediction ≠ choice = stimulus	445	Prediction does not match choice, mouse correct
prediction ≠ choice ≠ stimulus, prediction = stimulus	94	Prediction does not match choice, mouse incorrect Prediction matches stimulus
	1616	Total trials
Mouse choice accuracy (choice = stimulus)		73.3% = (741 + 445)/1616 $p < 0.0001$
Prediction accuracy over choice (prediction = choice)		66.6% = (741 + 336)/1616 $p < 0.0001$
Prediction accuracy over stimulus (prediction = stimulus)		51.6% = (741 + 94)/1616 $p = 0.20$

Extended data Table 1: The combination of above-chance performance accuracy and above-chance pre-stimulus choice prediction does not imply that pre-stimulus activity predicts visual stimulus identity. Single mouse example illustrating that the combination of above chance accuracy in behavioral performance and above chance pre-stimulus predictability of behavioral choice does not imply above predictive accuracy over visual stimulus. In other words, the pre-stimulus state of the mouse brain does not predict the identity of the subsequent visual stimulus.