Sensory History Drives Adaptive Neural Geometry in LP/Pulvinar-Prefrontal Cortex Circuits

Authors: Yi Ning Leow[1,2], Alexandria Barlowe[1], Cindy Luo[2], Yuma Osako[1,2], Mehrdad Jazayeri[2,3], Mriganka Sur[1,2]

[1] The Picower Institute for Learning and Memory (PILM)
[2] Massachusetts Institute of Technology (MIT), Department of Brain & Cognitive Sciences (BCS)
[3] McGovern Institute of Brain Research (MIBR)

Summary

Prior expectations guide attention and support perceptual filtering for efficient processing during decision-making. Here we show that during a visual discrimination task, mice adaptively use prior stimulus history to guide ongoing choices by estimating differences in evidence between consecutive trials (| Δ Dir |). The thalamic lateral posterior (LP)/pulvinar nucleus provides robust inputs to the Anterior Cingulate Cortex (ACC), which has been implicated in selective attention and predictive processing, but the function of the LP-ACC projection is unknown. We found that optogenetic manipulations of LP-ACC axons disrupted animals' ability to effectively estimate and use information across stimulus history, leading to | Δ Dir |-dependent ipsilateral biases. Two-photon calcium imaging of LP-ACC axons revealed an engagement-dependent low-dimensional organization of stimuli along a curved manifold. This representation was scaled by | Δ Dir | in a manner that emphasized greater deviations from prior evidence. Thus, our work identifies the LP-ACC pathway as essential for selecting and evaluating stimuli relative to prior evidence to guide decisions.

Introduction

Our perceptual landscape is constantly shaped by prior expectations and ongoing goals. Such prior expectations could support predictive processing, biasing cognitive mechanisms towards immediately relevant features. For example, target selection is a key driver of attention[1,2,3,4,5,6], and prior expectations exert powerful influence over the priming of attention. However, we still lack comprehensive understanding of neural circuit mechanisms that underlie such computations. The pulvinar, a higher order visual thalamic nucleus, is a critical hub that supports selective visual processing via its interactions with cortical targets. Mutual interactions between the pulvinar and the frontal-parietal attention network have been implicated in regulating sustained and selective attentional allocation during visual decision-making[3,4,5,7,8,9,10,11,12,13,14]. While previous work has provided compelling evidence for the pulvinar's contributions to visuospatial attention and filtering of ongoing distractors[11,12,14,15], how the pulvinar supports feature-based processing (but see [16]), and how its activity is adaptively influenced by prior expectations or stimulus history, is unknown.

The lateral posterior nucleus (LP) is the rodent homolog of the primate pulvinar. It shares most of the key cortical partners as the primate pulvinar[17,18], indicating that an evolutionary roadmap for the pulvinar's hub-like functions already exists in rodent species. LP inputs to the primary visual cortex (V1) have been shown not only to sharpen visual tuning and improve visual discrimination[19,20] but also encode saccadic[21] and locomotor visual mismatches[22,23]. The latter supports the pulvinar's role in predictive processing[24] particularly in the visuomotor domain, while more recent work has proposed LP-V1 circuit mechanisms as amplifying cortical responses to unexpected visual inputs representing low-level visual features[25]. However, it remains unclear how predictive processes supported by the LP/pulvinar might influence attentive processes to bias perceptual decisions. Given that LP outputs to different cortical areas could carry distinct visual and contextual information[23], a target-specific functional dissection of information carried by LP outputs will enable more precise dissection of LP contributions. Indeed, the subregions of pulvinar/LP known to couple with the frontal decision-making network appear to have less organized receptive fields and encoding of low-level visual features[17], but instead carry more abstract information such as perceptual confidence[26]. Thus, direct examination of pulvinar/LP interactions with the prefrontal areas will provide more direct evidence for how the pulvinar could support target selection using prior expectations to influence decision making. We decided to home in on LP projections to the anterior cingulate cortex (ACC), which has been described as a critical node for action selection,

predictive processing and anticipating cues that provide information and resolve uncertainty[27,28,29,30,31,32,33,34,35,36]. We have previously shown that reciprocal LP-ACC circuits in rodents integrate inputs from homologous cortical and subcortical regions to that of primate pulvinar-frontal circuits[18]. Furthermore, relative to LP inputs to other visual areas, LP-ACC projections preferentially integrate inputs from prefrontal and association areas[18] including those implicated in attentional regulation. As such, attentional biases integrated via LP to ACC can powerfully influence visual decision making.

Here, we devised a task to study the contribution of present and prior evidence in shaping responses of LP/pulvinar projections to ACC. We show that mice modulate their target selection strategy based on the similarity across consecutively presented sensory evidence — staying with a choice when the current stimulus is similar to the previous, and switching otherwise – but in a manner influenced by task parameters and performance. LP projections to ACC play an important role in filtering perceptual representations based on trial history. Computations that support this strategy based on prior expectations become disrupted with unilateral optogenetic activation of LP-ACC projections, biasing choices towards the activated side particularly when there are larger differences across consecutive stimuli. Using two-photon axonal calcium imaging, we further show that LP-ACC axons can estimate the differences across consecutive stimuli by representing the relative evidence between current and previous trials in a structured low-dimension manifold. This information is disrupted by optogenetic stimulation of the LP-ACC projection. Our work identifies the LP-ACC pathway as powerfully modulating ongoing sensory inputs through predictive expectations from prior experience, implementing a target selection strategy which guides goal-directed behavior on a trial-to-trial basis. We thus reveal a circuit mechanism through the pulvinar which prioritizes attention towards stimuli based on prior expectations.

Results

Mice make choices based on estimated differences across consecutive stimuli

We trained head-fixed mice on a two-alternative forced-choice visual motion discrimination task where they learned to discriminate between two global dot movement directions (0°, temporally or 180°, nasally) in random dot kinematograms (RDK), and reported their choice by licking either the left or right lick spout (Figure 1a). The degree of stimulus uncertainty was varied by

presenting RDKs with target directions presented at graded levels of coherence. In the full task, RDKs were presented at 100%, 64%, 32%, or 16% coherence with equal probabilities. The task was self-paced; mice had to withhold licking and run on a treadmill for 1s before a new trial could be initiated. Stimuli were presented for 1.5 seconds, followed by a 1.5-second response window. We counted two consecutive licks as a full choice. Correct choices during the stimulus period were not reinforced – two consecutive correct licks had to occur during the response period to be rewarded. However, if mice made a wrong choice at any time during the trial, it counted as an error. Reinforcement was delivered 100 ms after valid choices during the 1.5-second response window. Correct choices were rewarded with water delivery (8 ul) and wrong choices were negatively reinforced with a flashing screen and an auditory white noise for 3 seconds, which also served as an extended time-out. Mice were trained on the 100% coherence stimulus until they achieved >75% accuracy over 3 consecutive days before lower coherences were introduced. Trained mice showed coherence-dependent psychometric performance (Figure 1b) and perceptual sensitivity (Figure S1a). However, due to the delayed response window mice were trained on, we did not find reaction time (time of first lick) to scale with coherence (Figure S1b).

Although choice behavior was largely dependent on the current stimulus evidence, mice nevertheless showed serial choice biases, with a propensity to repeat a previously rewarded choice (Figure 1c). Importantly, we found that the probability of repeating a choice was highly dependent on the absolute difference between signed coherences across trials, which we denote as | Δ Dir | (Figure 1d). When the current and previous stimuli were more similar (low | Δ Dir |, Figure 1e), mice were more likely to repeat their previously rewarded choice. Conversely, when consecutive stimuli were most different (high | Δ Dir |, Figure 1f), they were least likely to repeat. The opposite was observed if the previous trial resulted in an error (Figure S1e-f). Effective use of this strategy exemplifies a potential adaptive influence of trial history information, adding an important nuance to a classical win-stay-lose-switch strategy (Figure 1d). Notably, the probability of repeating a choice showed a steeper dependence on | Δ Dir | than expected from simulated null models of mean psychometric performance that had no history dependence, or models with only simple serial choice biases that did not depend on the previous stimulus strength (Figure S1k-l). This indicated that the relationship between probability of repeating a previous choice and | Δ Dir | could not be fully explained by performance at different levels of stimulus evidence, independent of trial-by-trial stimulus differences.

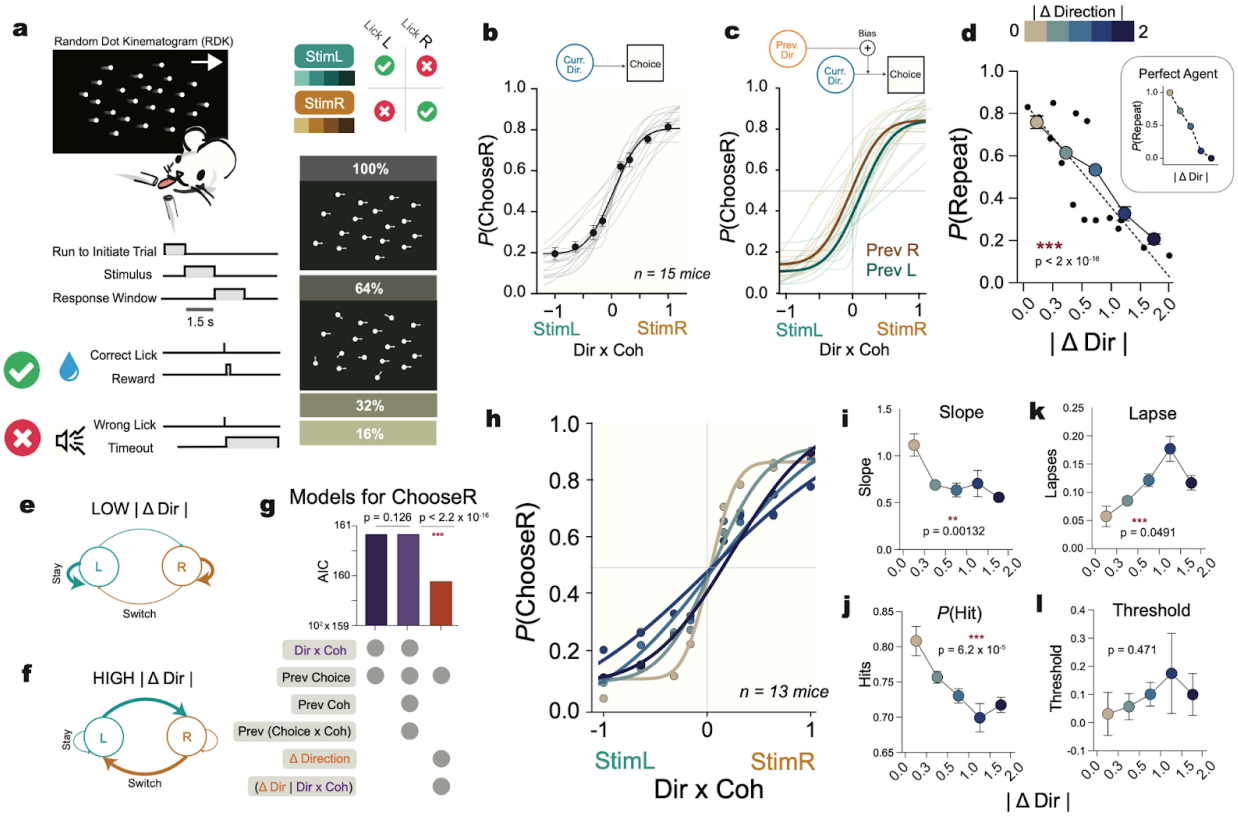

Figure 1: Mice make choices based on estimated differences across consecutive stimuli. a) Two-alternative choice visual discrimination task structure. Mice run to initiate a trial and are presented with an RDK stimulus for 1.5 seconds. Mice report their choices by licking either the left or right spout during the 1.5 second response window. 100ms upon choice, correct choices are rewarded with water delivery while wrong choices are followed by a loud white noise and flashing screen. **b)** Population level psychometric fit for P(ChooseR) for n = 15 mice. **c)** Previous choice-dependent psychometric function shifts. **d)** Rate of repetition follows change in direction (| Δ Dir |) from the previously rewarded choice ($p < 2 \times 10^{-16}$, $\chi 2 = 2505.9$ in Type II Wald Chi-Squared Test for | Δ Dir | in P(Repeat) Binomial LME with (Coherence | Mouse) as random effect, with n = 14183 trials, 13 mice) Inset: Such behavior approximates a perfect agent which achieves 100% accuracy at all coherence levels. **e & f)** Model of | Δ Dir | history-dependent behavior where mice are more likely to stay with a previous choice when there is low | Δ Dir | (e) and more likely to switch when there is a greater difference (| Δ Dir |) between consecutive stimuli (f). **g)** A binomial model (orange) incorporating Δ Dir provides superior fit over models with only previous choice, or only previous choice and coherence, as indicated by the lower Akaike's Information Criterion (AIC). **h)** | Δ Dir |-dependent psychometric function fits for probability of choosing right. **i)** | Δ Dir | modulated psychometric slopes (p = 0.00132, Repeated measures ANOVA, $F_{(4,48)} = 7.699$, $\eta p2 = 0.391$). **j)** | Δ Dir | reduced overall hit rate (| Δ Dir | coefficient = -0.218 ± 0.046, $p = 2.08 \times 10^{-6}$ Type II Wald Chi-Squared Test, $\chi 2 = 22.52$, Binomial LME for P(Hit), with (Coherence | Mouse) as random effect, n = 14183 trials, 13 mice). **k)** | Δ Dir | influenced lapse rates (p = 0.0491, Repeated measures ANOVA, $F_{(4,48)} = 2.579$, $\eta p2 = 0.177$). **l)** | Δ Dir | did not

lead to significant change in threshold for choosing right ($p = 0.471$, Repeated measures ANOVA, $F_{(4,48)} = 0.753$, $\eta p2 = 0.059$). Error bars show SEM across mice ($n = 13$).

The | Δ Dir | modulated psychometric performance (Figure 1g,h), such that increasing | Δ Dir | was associated with reduced slopes (Figure 1i) and increased lapse rates (Figure 1k). Altogether, this reduced mean hit rates (Figure 1j) with minimal impact on the threshold (Figure 1l). After errors, the effect of | Δ Dir | on hit and lapse rates was opposite of that after hits, with increasing | Δ Dir | associated with reduced hit rates and increased lapses (Figure S1e-j). This suggests that choice modulation by | Δ Dir | was adaptive and dependent on the outcome of the previous trial. Trial-by-trial shifts in choice behavior were also better fitted by models incorporating | Δ Dir | and its influence on the psychometric slope, relative to models that simply accounted for the previous choice, whether or not modulated by previous coherence (Figure 1g, S1c-d). Thus, mice performed best when consecutive stimuli were similar and when repeating their previously rewarded choices, and performed worse when they were required to switch. In fact, increasing | Δ Dir | also increased the probability of mice omitting the trial (Figure S1m). Mice also had generally greater response latency with greater | Δ Dir | (Figure S1n). Together, these results implied that mice retained a memory of the previous stimulus estimate, and the comparison with the current stimulus facilitated decisions to switch or stay, based on the outcome of the previous choice. Such short-term memory could be expected to decay with time, and indeed | Δ Dir | had a diminishing influence on the probability of repeating a choice with increasing inter-trial intervals (Figure S1o-q). Interestingly, with diminished influence of | Δ Dir | with greater inter-trial intervals was also associated with reduced mean hit rate (Figure S1q), potentially indicating that using | Δ Dir | across trials could be an adaptive strategy. Taken together, our findings demonstrate that mice use relative evidence across consecutive trials (| Δ Dir |) to modulate classical win-stay-lose-switch strategies.

Activating LP-ACC projections drives increasing ipsilateral biases with increasing | Δ Dir |

ACC has been implicated in regulating adaptive switching from choices or task rules particularly under uncertainty[31,32,33,34,35,36,37] to balance explore-exploit trade-offs[38,39,40,41,42]. What sources of evidence might ACC integrate to support such computations? Based on evidence that the pulvinar/LP activity scales with decisional confidence[26] as well as its purported role in predictive processing[24] and integrating information across multiple visual and associative cortical areas[17,18], we hypothesized that the pulvinar/LP supports ACC control of adaptive switching by tuning perception according to prior evidence and comparing relative evidence across stimuli.

To examine the functional role of LP-ACC input in visual decision-making, we activated LP-ACC axons using optogenetics. The light-activated opsin ChR2 was expressed unilaterally in the left LP, contralateral to visual stimulus presentation, and LP-ACC projections were activated by illuminating the ACC with blue light (Figure 2a). Activating LP-ACC axons during the stimulus evaluation period impaired performance (Figure 2b). This was associated with a significant reduction in psychometric slopes (Figure 2c) and a small but non-significant increase in the threshold for choosing right (Figure 2d). Accordingly, this reduced the mean hit rate (Figure 2e), without altering the overall proportion of right (contralateral) choices (Figure 2f) or trials omitted (Figure S2). These changes could not be explained by non-specific laser-dependent effects such as heating or distraction from the task since illuminating the axons with 593nm light, outside of the ChR2 excitation spectrum, had no significant influence on behavioral performance (Figure S3). The reduction in psychometric slopes instead of shift in choice thresholds indicated that activating LP-ACC axons likely led to a perceptual, rather than a motor choice bias. Thus, we next examined how LP-ACC activation could influence perceptual choices based on trial history.

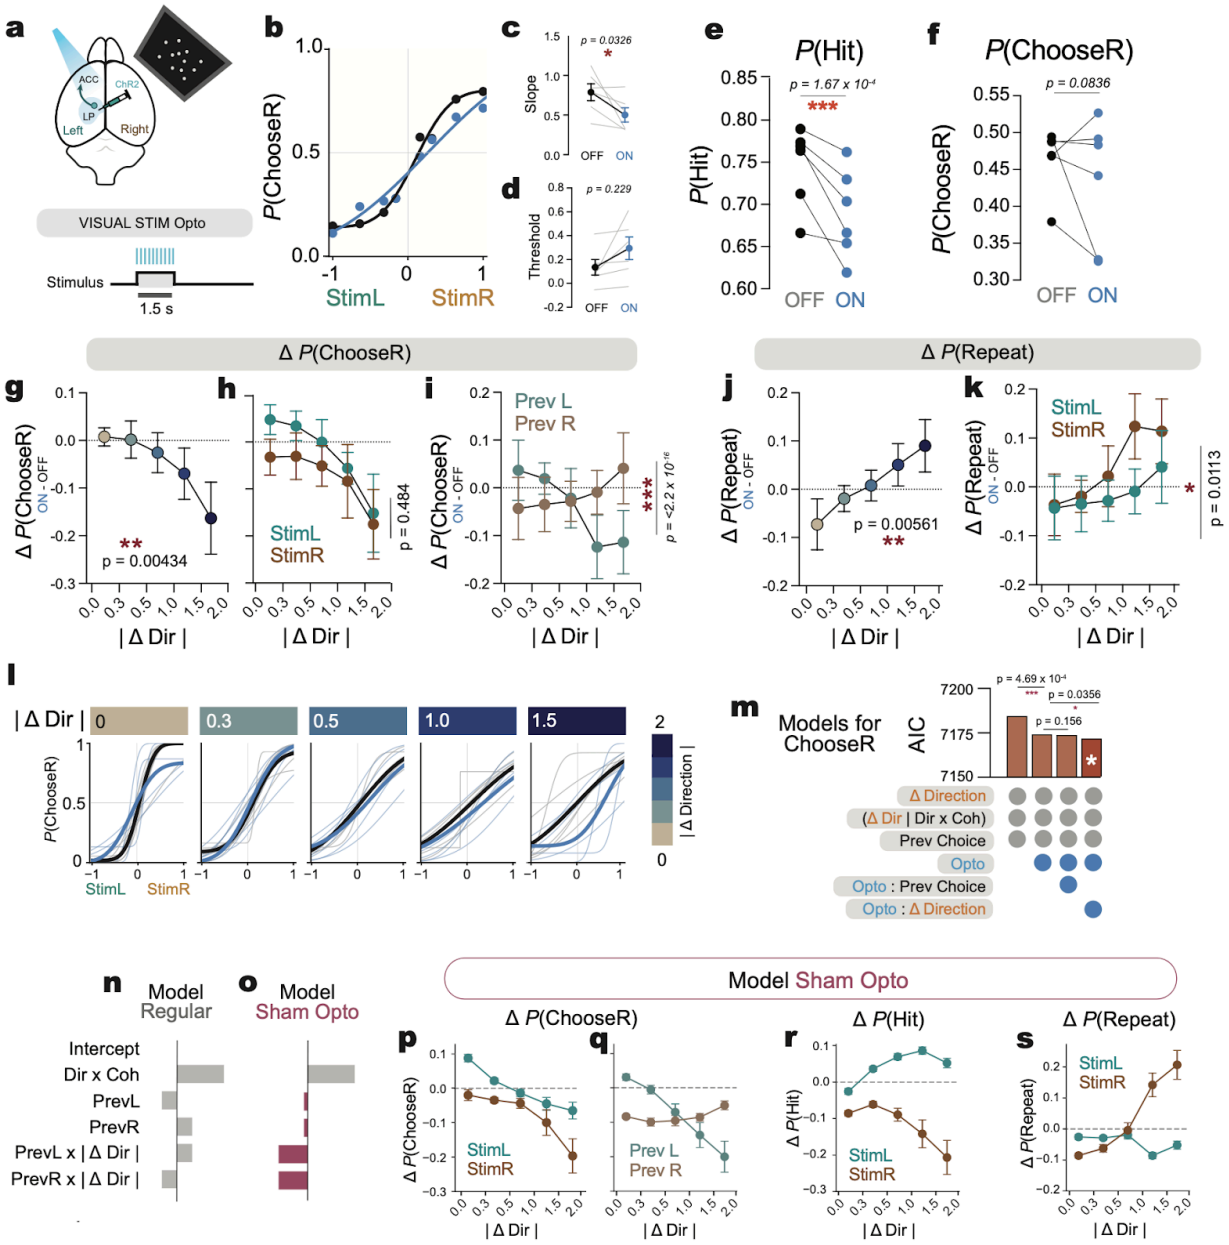

Figure 2: Activating LP-ACC drives increasing ipsilateral biases with increasing | Δ Dir |. a) Experimental schematic for optogenetic experiments. AAV carrying ChR2 was injected into the left medial LP, contralateral to the visual presentation. LP-ACC axons were activated by illuminating the axons over the ACC. (Bottom) Optogenetic stimulation for 1.5s during the presentation of the visual stimulus for 25-30% of trials. **b)** Population (n = 6 mice) psychometric function fits for laser off (black) and on (cyan) trials. n = 7835 total trials. **c)** Activating LP-ACC axons reduced psychometric slopes (p = 0.0326, one-tailed paired t-test. Mean ON-OFF: -0.264 ± 0.112, t = 2.355, n = 6). **d)** Threshold for choosing right was not significantly altered by optogenetic activation (p = 0.229, one-tailed paired t-test. Mean ON-OFF: 0.0930 ± 0.116, t = 0.803, n = 6 mice). **e)** Reduced psychometric slopes led to reduced hit rates (p = 1.67 x 10^{-4}, LRT for trial factor Opto in P(Hit) Binomial LME, $\chi 2$ = 14.174, n = 5744 trials, 6 mice). **f)** No change in P(ChooseR) (p = 0.0836, LRT for trial factor

Opto in P(ChooseR) Binomial LME, $\chi2$ = 2.993, n = 5744 trials, 6 mice). g-i) Change in P(ChooseR) evoked by optogenetic stimulation **g)** as a function of | Δ Dir | overall (Opto : | Δ Dir |, $\chi2$ = 8.137), and **h)** based on each stimulus direction (Opto : | Δ Dir | : StimDir, $\chi2$ = 1.453) and **i)** previous rewarded choice (Opto : | Δ Dir | : PrevChoice, $\chi2$ = 76.22). For g-i, Type II Wald Chi-Squared Test were done on terms in a P(ChooseR) Binomial LME, with n = 5744 trials, n = 6 mice, with (Dir x Coh | Mouse) as random effect. **j-k)** Change in P(Repeat) evoked by optogenetic stimulation j) as a function of | Δ Dir | overall (Opto : | Δ Dir |, p = 0.00561, $\chi2$ = 7.670), and based on k) each stimulus direction (Opto : | Δ Dir | : StimDir, p = 0.0113, $\chi2$ = 8.963). For j & k, Type II Wald Chi-Squared Test were done on terms in a P(Repeat) Binomial LME, with n = 5744 trials, n = 6 mice, with (Coherence | Mouse) as random effect. **l)** Population (n = 6 mice) psychometric function fits for laser off (black) and on (cyan) trials, based on the binned | Δ Dir |. Thin lines show psychometric function fits for individual mice. **m)** Akaike Information Criterion (AIC) to indicate model fit for nested binomial LME models for P(ChooseR). Inclusion of trial factors with optogenetic stimulation and its interaction with Δ Dir provides superior model fit (Likelihood Ratio Test, n = 5744 trials, n = 6 mice) **n)** Schematic of model coefficients as fitted by empirical data from light off trials, including current trial evidence (Dir x Coh) and history factors (PrevL/R, PrevL/R x | Δ Dir |) to predict P(ChooseR), which serves as our baseline model for each mouse. **o)** In sham opto models, coefficients were reweighted, altering the directionality and magnitude of history-dependent factors (magenta). This sham opto model recapitulates empirical behavioral observations. Specifically, the weight of all previous trial parameters were made negative, thus decreasing probability of choosing right. Magnitude of weights for PrevL/R were reduced but weights for PrevL/R x | Δ Dir | were increased. Notably, there was otherwise no change in the sensitivity to current trial evidence. **p)** Δ P(ChooseR) as generated by sham opto model (model - light off), plotted for each stimulus direction. **q)** Δ P(ChooseR) in sham opto model, plotted based on previously rewarded choice. **r)** Δ P(Hit) in sham model, plotted based on stimulus direction. **s)** Δ P(Repeat) in sham model, plotted based on stimulus direction. For p-s), error bars represent SEM across mice and the models based on their behavior.

Interestingly, we found that with increasing | Δ Dir | across trial history, activation of LP-ACC axons induced greater reductions in right (contralateral) choices, P(ChooseR) (Figure 2g). This was true for both left and right-associated stimuli (Figure 2h). As such, LP-ACC activation promoted ipsilateral biases, with greater biases particularly when animals were required to switch away from a previous choice (higher | Δ Dir |). However, rather than simply promoting ipsilateral choice repetition, optogenetic activation affected trials after both left and right choices, albeit exerting opposite effects relative to | Δ Dir | (Figure 2i). After left rewarded choices, activating ipsilateral LP-ACC axons reduced P(ChooseR) to a greater extent with greater | Δ Dir |, suggesting that there was an impairment in accumulating the evidence to switch to a contralateral choice. Conversely, following right rewarded choices, activating LP-ACC axons instead led to greater reductions in P(ChooseR) under low | Δ Dir |. Taken together, activating LP-ACC axons drove increasing ipsilateral choices in a | Δ Dir |-dependent manner. Consequently, the hit rate for left stimulus trials increased with greater | Δ Dir |, but decreased

for right stimulus trials (Figure 2j,k). In further support for a | Δ Dir |-dependent change in choice behavior, binomial models of P(ChooseR) were best fitted (with lower AIC) when incorporating the interaction between optogenetic stimulation and the change in direction (Δ Dir) (Figure 2m).

The evidence thus far suggested that LP-ACC activation specifically biased choices to scale with | Δ Dir |. Next, we considered alternative explanations, such as LP-ACC activation exerting a non-specific influence on perceptual sensitivity or simply biasing repetition of ipsilateral choices. To simulate these possibilities, we first fitted a basic binomial model based on trials without optogenetic stimulation for each animal. The base model described the P(ChooseR) based on the current stimulus evidence (Dir x Coh), previous choice (prevL, prevR), and the | Δ Dir |-dependent modulation of choice history (Figure 2n, Figure S4a). We then simulated the effects of optogenetic stimulation by constructing sham models with altered coefficients from the base model to predict P(ChooseR). For example, reduced perceptual sensitivity could be simulated by specifically reducing the weight on current stimulus evidence (Dir x Coh). Notably, we were unable to recapitulate the trial-by-trial impact of optogenetically stimulating LP-ACC axons in a model with reduced weight of current evidence (Figure S4e). This implied that activating LP-ACC axons did not lead to a reduced hit rate due to a trivial reduction in perceptual sensitivity. Rather, it was likely that LP-ACC input influenced choice behavior based on trial history. Nevertheless, even models with increased influence of trial history failed to account for the | Δ Dir |-dependent optogenetic effects on P(ChooseR) (Figure S4d). It was only with a model with altered interaction between previous choices and | Δ Dir | that we were able to recapitulate the effects of optogenetic activation (Figure 2n-s, Figure S4b). In this sham optogenetic model, history-dependent variables were set to all have negative weights, and thus promoted left (ipsilateral) choices. Additionally, the magnitude of | Δ Dir | influence was increased, while the | Δ Dir |-independent influence of previous choice was diminished. The sham model recapitulated optogenetic effects of greater reduction in P(ChooseR) with greater | Δ Dir | regardless of the stimulus direction (Figure 2p, h), opposite | Δ Dir |-dependent changes in P(ChooseR) based on previous choice (Figure 2q, i), stimulus-dependent changes in hit-rate (Figure 2r, k) and P(Repeat) (Figure 2s, Figure S4b-vii). Taken together, our simulations suggest that the most likely explanation for the effects caused by LP-ACC activation is an altered interaction between previous choices and the difference between current and previous trial directions, leading to greater ipsilateral biases with increasing | Δ Dir |.

If LP-ACC activity served as a critical pathway for providing both current and previous trial information, the optogenetic stimulation could disrupt this information to influence subsequent post-stimulation trials. However, we found that behavioral performance recovered after optogenetic trials whether the stimulation trial led to a hit or an error (Figure S5). In other words, optogenetic activation of LP-ACC axons disrupted behavior during opto-stimulated trials but not their stimulus representations, which could be retained by other brain areas. To understand if LP-ACC axons played a critical role in retaining the representation of the previous trial information upon reinforcement, we optogenetically stimulated LP-ACC axons during reinforcement and looked at how this might influence choices in the next trial (Figure S6). We found no evidence for any influence of this reinforcement-epoch stimulation on subsequent choices, nor relative influence of previous trial information. This suggests that LP-ACC axonal activity during this reinforcement period, representing the trial outcome and stimulus, is not critical for its influence in supporting decisions in the next trial. Together, these indicate that LP-ACC neurons likely provide a control signal to bias behavior by accessing both current and previous trial information from other brain areas in the decision-making network. They thus operate more specifically downstream of integrating this information to support decisions by comparing evidence between current and previous stimuli.

LP-ACC axons represent visual stimuli with task-gated low-dimensional geometry

The influence of LP-ACC axons on | Δ Dir |-modulated choice behavior motivated us to examine what information was encoded by LP inputs to the ACC during stimulus evaluation. We labeled LP neurons with the Ca2+ sensor GCaMP7f and imaged activity of LP axons through a cranial window over the ACC (Figure 3a). After individual boutons were defined as regions of interest, boutons were clustered based on their pairwise correlations with other boutons in the same field of view to obtain the best approximation of independent axonal units from boutons that could originate from the same neuron (Methods, Figure S7). A weighted average of correlated boutons was taken for each cluster to derive the activity of single axonal units – henceforth referred to as axons.

We first sought to resolve whether LP-ACC axons better represented stimulus or choice. As a population, bulk averaged LP-ACC axon activity did not have a significant preference for left or right-associated stimuli (Figure 3b), but had a slight preference for ipsilateral (left) choices

(Figure 3c). In spite of the lack of stimulus direction preference in bulk LP-ACC activity, cross-validated Support Vector Machine (SVM) classifiers trained to decode stimulus direction performed significantly better than decoders trained to discriminate between choices taken by the mouse (Figure 3d) during the stimulus presentation period. Furthermore, a greater proportion of axons had activity that scaled monotonically with direction (Figure S8a,b), than for choice (Figure S8c,d). Both left and right stimuli or choices were comparably represented (Figure S8a-d). Of choice-selective axons, left/ipsilateral choice-preferring axons appeared to represent choice with little modulation by strength of evidence (Figure S8d). This stands in contrast to cortical areas with reciprocal connectivity with LP, such as the Posterior Parietal Cortex (PPC) and ACC which have dominantly choice encoding with contralateral biases[43,44,45,46]. Superior stimulus direction coding over choice not only indicated that LP-ACC axonal activity was unlikely to be fully explained by directional licking movement activity, it further indicated that LP-ACC axons likely operated upstream in the perceptual decision-making process, between stimulus representation and committing to a choice.

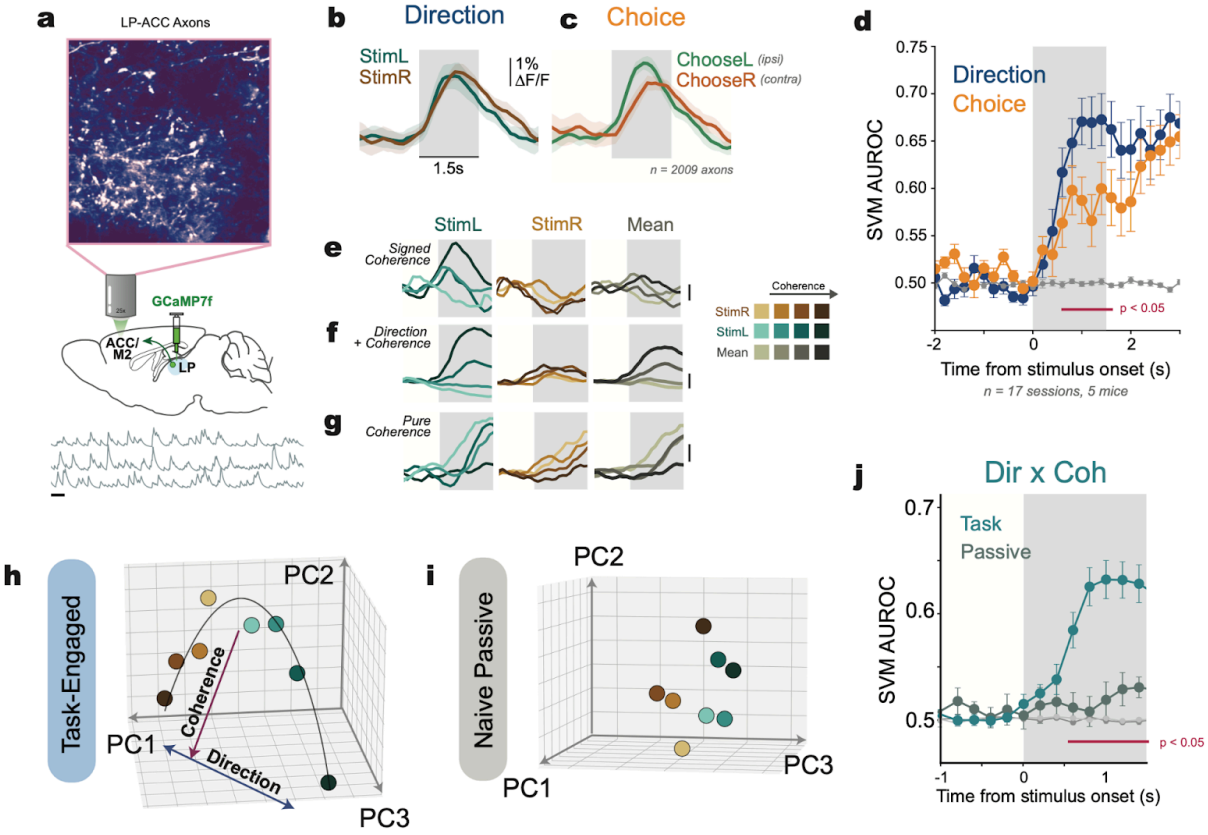

Figure 3: LP-ACC axonal populations generate low-dimensional representations of the decision variable in a manner gated by task-engagement. a) Experimental schematic of imaging LP-ACC axons: an AAV carrying GCaMP7f was injected in the left medial LP. Cranial

windows were implanted over the ACC in order to image the left LP-ACC axons. (Right) Example field of view of axons from a single session and weighted average trace of a cluster of highly correlated axonal boutons. **b)** LP-ACC population average for each stimulus direction. **c)** LP-ACC population average for each choice side. **d)** Discriminability (Area under the receiver operating characteristic curve, AUROC) of population-level decoding using cross-validated SVM classifiers trained to decode direction (blue) and choice (orange) show greater discriminability for trial stimulus direction over eventual choice taken. Red bar indicates significant difference ($p < 0.05$, two-tailed paired t-test) in decoding performance between stimulus direction and choice at that time-point. **e)** Example single axon with activity that scales with signed coherence. **f)** Example single axon with preferential activity or left stimuli but also scales monotonically with stimulus coherence. **g)** Example single axon representing pure coherence encoding without direction discriminability. **h)** LP-ACC population activity for unique presented stimuli projected onto the first 3 principal components when mice were engaged in the task. The representation organizes into a curved manifold that simultaneously represents stimulus direction and coherence. **i)** LP-ACC population does not organize into such a curved manifold when presented with the same unique stimuli under naive passive viewing conditions. **j)** Cross-validated SVM decoder performance (AUROC) trained to discriminate each of the 8 presented stimuli (2 directions x 4 coherences) perform much better in the task-engaged condition compared to in naive mice passively viewing identical stimuli. Red bar indicates significant difference ($p < 0.05$) in decoding performance at that time-point, using two-tailed t-test.

Our ability to decode stimulus direction from LP-ACC axonal populations in spite of comparable levels of bulk activity between stimulus directions reflected the information harbored by the heterogeneous tuning in individual LP-ACC axons, which together contributed information for resolving stimulus direction as a population. Indeed, individual LP-ACC axons could represent stimulus direction and coherence selectivity by scaling monotonically with signed coherence in their preferred direction (Figure 3e), with directional preference that is further scaled by unsigned coherence (Figure 3f), or with selectivity only for coherence (Figure 3g).

Given the heterogeneity in the direction selectivities across single LP-ACC axons, we asked if the axonal population could have structured low-dimensional representations of the unique stimuli that facilitate discriminability between stimulus directions. We performed dimensionality reduction of the LP-ACC population activity using Principal Component Analysis (PCA) and projected trial-averaged (held-out test trials) population activity along the first 3 principal components (accounting for 34.7% total explained variance, Figure S9c). We found that unique stimuli were represented in a structured low-dimensional geometry, where signed coherence was continuously embedded on a curved line (Figure 3h). LP-ACC axonal populations evolved to represent each stimulus as a point along a curved manifold in the 3-dimensional latent space across the stimulus evaluation period (Figure 3h, Figure S9a), as first described in primate

cortical areas involved in decision-making [47]. This curved manifold jointly encoded the dominant stimulus direction and stimulus difficulty during task-engaged contexts. Notably, this low-dimensional geometry in LP-ACC activity was not evident in naive mice passively viewing identical RDK stimuli (Figure 3i). Rather, the axons incorporated learned and/or attention-dependent influences to organize the representation of behaviorally-relevant visual stimuli in a structured low-dimensional geometry during active task engagement. In turn, the structured population representation facilitated stimulus decoding, which was superior in task-engaged contexts compared to naive passive sessions (Figure S9a-f). This low-dimensional geometry was not simply due to directional movement activity, because it was also observed in omitted trials with no licking activity (Figure S9g-h). Thus, LP inputs to ACC are gated by active task engagement or learned stimulus associations, reflecting their dynamic involvement in perceptual decision-making, rather than simply relaying bottom-up activity driven by low-level visual features.

LP-ACC axons multiplex stimulus, choice, uncertainty and outcome variables across trial history

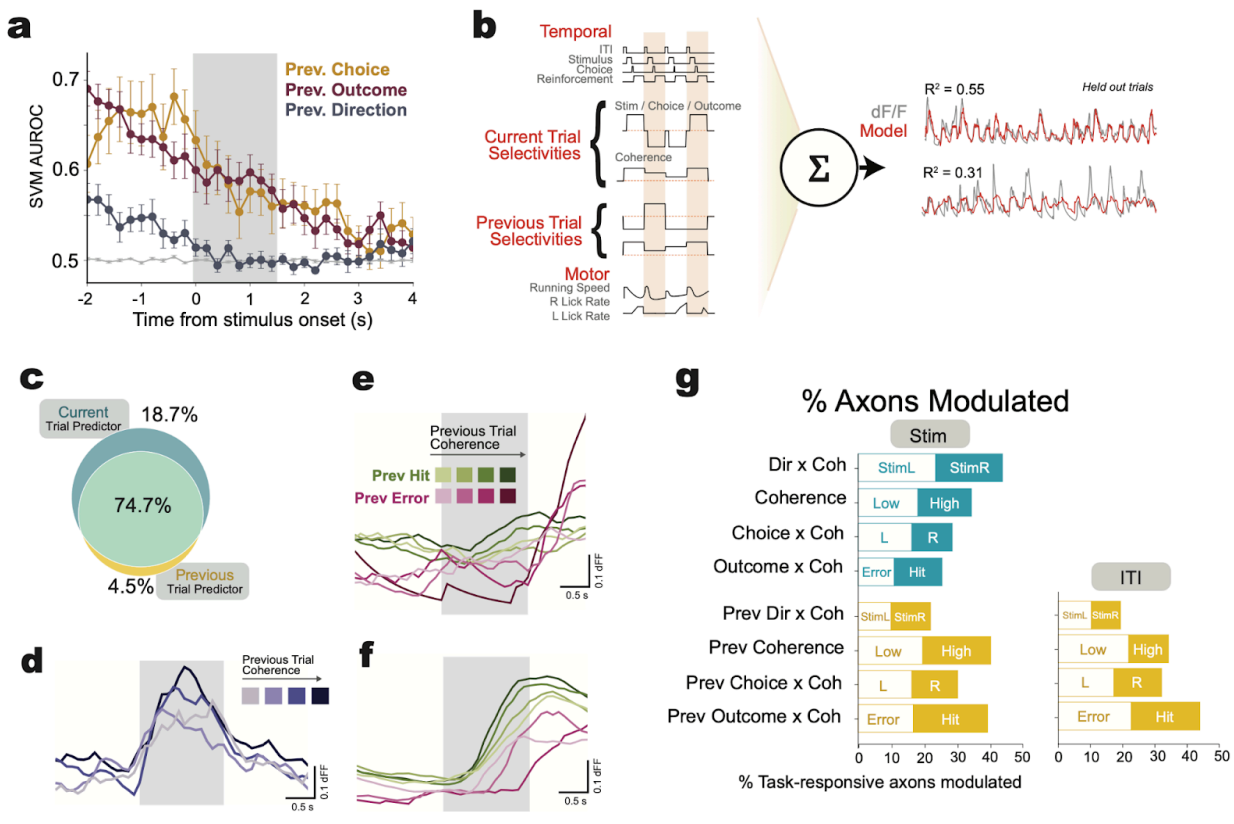

Figure 4: Single LP-ACC axons multiplex uncertainty across multiple task dimensions and across trial history. a) Information such as previous choice, outcome and direction from previous trials could be decoded from LP-ACC populations during ongoing trials. Grey line shows the SVM AUROC for dataset with shuffled previous direction. Error bars reflect SEM across 17 imaging sessions. **b)** Task factors in the design matrix included temporal factors, current and previous trial selectivities, and ongoing motor activity. The encoding model involved a regression fit on all the predictors, weighing them to best reconstruct axonal activity with a linear weighted summation of these predictors. On the right are example $\Delta F/F$ axonal traces (grey) overlaid with the model reconstruction (orange) of their activity. **c)** Most (74.7%) axons were simultaneously modulated by a current and a previous trial feature. **d)** Example single axon showing modulation by previous trial coherence, plotted for the same current stimulus coherence (32%) and only after hit trials. **e & f)** Two representative single axons showing modulation by previous outcome, scaled byprevious evidence strength. **g)** Proportion of task-responsive axons with significant modulation by each predictor, and the specific task variable they have a preference for the stimulus epoch (left) and during the ITI (right). Blue bars indicate current trial variables, yellow bars indicate previous trial variables.

The behavioral consequences of activating LP-ACC axons motivated us to next ask if LP input could encode information beyond the current stimulus. Indeed, we were able to decode multiple other task variables from LP-ACC axonal activity, including trial history (Figure 4a) and anticipated outcome even prior to reinforcement (Figure S7e,f, S10) using cross-validated linear SVMs. Representation of anticipated outcome, while correlated with stimulus difficulty,was independent of a subpopulation of LP-ACC axons that purely represented levels of coherence regardless of outcome (Figure S7g,h). Anticipated outcome appeared to be encoded by a subpopulation of LP-ACC axons with a response profile that exhibited signatures of confidence encoding, or uncertainty over choice, as predicted by statistical decision theory [48,49] (Figure S10b-g). To account for the multiple trial factors that LP-ACC axons may integrate and to determine their relative contributions at the single axon level, we fitted encoding models to reconstruct individual LP-ACC axon activity using weighted combinations of task variables, including motor activity from locomotion and licking. Each model was designed to examine the strength of LP-ACC selectivity for stimulus direction, trial difficulty, choice, and outcome during the current and previous trial (Figure 4b). We identified 1003 (53.2% of 1885 axons) task-responsive axons with a mean cross-validated R2 of 0.0904 ± 0.0022 (Figure S11a). More than half (53.3%) of the task-responsive axons were modulated by at least 4 different task predictors (Figure S11b). This highlights the highly mixed selectivity found in single LP-ACC axons, which is consistent with the mixed selectivity that has been reported in prefrontal areas like the ACC[50,51,52,53]. In fact, the majority of task-responsive axons (74.7%) simultaneously encoded variables for current and previous trial information during the stimulus evaluation

epoch (Figure 4c). Previous trial coherence (Figure 4d,g) and outcome (Figure 4e,f,g) were the dominant trial history factors encoded by LP-ACC axons. The large proportion of the population simultaneously representing current and previous trial factors implied that ongoing perceptual representations undergo considerable modulation by previous trial uncertainty and possibly reinforcement-dependent expectations. The history-dependent information was not necessarily due to a slow decay of previous trial activity. For example, in some axons modulated by the previous outcome, the binary encoding of previous trial outcome was apparent during the inter-trial interval, but activity became additionally modulated by previous trial coherence upon stimulus presentation (Figure 4e,f). This suggests that trial history information in LP-ACC axons was likely inherited from elsewhere and integrated dynamically during decision-making, which could explain our finding that optogenetically activating LP-ACC axons disrupted only the stimulated trial but not their representations in the following trial (Figure S5). Furthermore, compared to during the ITI, there was a greater proportion of axons encoding previous direction and coherence during the stimulus period. These suggested that some previous trial information was not conveyed with persistent activity that decayed passively, but instead actively emerged as they became relevant during stimulus evaluation. Indeed, optogenetic stimulation of LP-ACC axons during the reinforcement period did not influence the subsequent integration of previous trial information (Figure S6). Thus, single LP-ACC axons integrated dynamic representations of uncertainty across multiple task domains including stimulus direction, choice, and outcome anticipation not just for the ongoing trial but also across trial history.

LP-ACC axonal populations represent | Δ Dir | relative to lateralized history biases

How might the varying extents of history modulation alter the low-dimensional representation of ongoing stimulus input? To address this question, we performed targeted dimensionality reduction of LP-ACC population activity (Figure 5a), allowing us to examine population trajectories along coding dimensions that best accounted for variance (or selectivity) across stimulus direction (DirectionCD, Figure 5b) and coherence (CoherenceCD, Figure 5c). We then projected trial-averaged LP-ACC population activity onto the DirectionCD and CoherenceCD for each unique stimulus, based on the magnitude of change in stimulus direction from the previous trial (| Δ Dir |). Even just in the DirectionCD, LP-ACC projections for the same stimulus were significantly modulated by the relative change from the previous trial direction (Figure 5d,e). As

such, ongoing low dimensional representations of the current stimulus by LP-ACC projections incorporate information from trial history, and can thereby exert history-dependent biases.

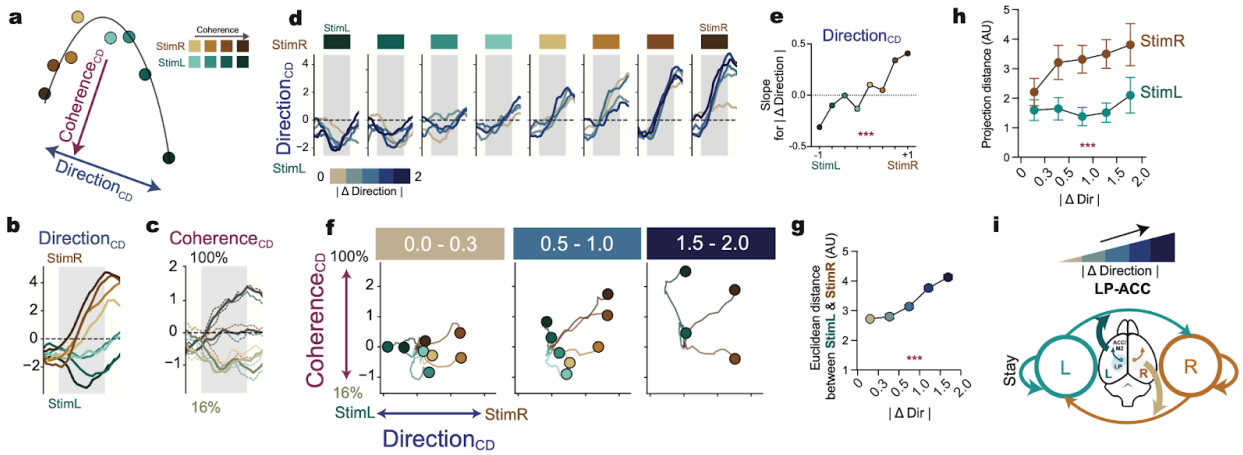

Figure 5: | Δ Dir | modulates population trajectories across the low dimensional manifolds. a) Targeted dimensionality reduction of population activity onto the 2 coding dimensions that best separated stimulus direction (DirectionCD) and stimulus coherence (CoherenceCD). **b)** Projections of individual stimuli along DirectionCD. **c)** Projections of individual stimuli along CoherenceCD. **d)** Projections of the same stimuli along DirectionCD varied depending on the | Δ Direction | from the previous stimulus. Shown only for consecutive hit trials. **e)** Modulation as quantified by slopes for | Δ Dir | for each stimulus. Slope derived from least squares fit for | Δ Dir | to mean projection. One-way repeated measures ANOVA on effect of stimulus direction on slopes ($p = 1.65 \times 10^{-43}$, $F_{(7,525)} = 37.9$, ηp2 = 0.336, over 75 iterations). **f)** In the joint Direction and Coherence subspaces, LP-ACC population trajectories traversed greater distances with increasing | Δ Dir |. Each line represents the trial-averaged population trajectory for each stimulus, starting from the stimulus onset. Each circle at the end represents the final trajectory position at 1.5s after stimulus onset. **g)** Euclidean distance between projections for left and right stimuli as a function of | Δ Dir | (One-way repeated measures ANOVA for effect of | Δ Dir |, $p = 5.379 \times 10^{-45}$, $F_{(4,300)} = 77.33$, ηp2 = 0.508, over 75 iterations) **h)** Mean projection distance across the joint Direction and Coherence coding dimensions for each stimulus direction (averaged across all coherences) based on the relative | Δ Dir |. Error bars show SEM across 100 iterations. Two-way repeated measures ANOVA for effect of Direction ($p = 4.1 \times 10^{-51}$, $F_{(1,75)} = 1478.5$, ηp2 = 0.952), | Δ Dir | ($p = 2.3 \times 10^{-43}$, $F_{(4,300)} = 73.54$, ηp2 = 0.495) and Direction : | Δ Dir | ($p = 5.2 \times 10^{-30}$, $F_{(4,300)} = 45.8$, ηp2 = 0.379) i) Schematic showing lateralized history-dependent effects of LP-ACC input that accumulates evidence in the form of comparisons between current and previous trial evidence (| Δ Dir |) to make switch-stay decisions under uncertainty.

Jointly represented in DirectionCD and CoherenceCD, we found that the curved manifold structure was most prominent with increasing | Δ Dir | (Figure 5f,g). LP-ACC population projection trajectories were more constrained when the stimulus direction was similar to the

503 previous (low | Δ Dir |) and when evidence supported repeating the previous choice (Figure 5f,g)
504 but expanded with increasing | Δ Dir |. Indeed, we found that the euclidean distance between
505 left and right stimuli along these two axes increased with | Δ Dir | (Figure 5g). This suggests that
506 the low-dimensional representation of stimuli by LP-ACC axons may preferentially represent
507 information relative to | Δ Dir |, such that stimuli are better resolved when there is a greater
508 difference from prior expectation. Furthermore, we note that the distance traveled along
509 DirectionCD and CoherenceCD were different between left and right stimuli (Figure 5h), with
510 right-associated stimuli (StimR, contralateral) associated with greater projection distances
511 (Figure 5h). Notably, this asymmetry was mirrored in the optogenetic stimulation of LP-ACC
512 axons, where contralateral-associated stimulus trials were selectively vulnerable to disruption
513 (Figure 2g-k). Thus, trial history biases the geometric representation of incoming stimuli by the
514 LP-ACC population, to further support accumulating evidence for switching to a contralateral
515 choice.

516 Discussion

517 In this study, we identified LP projections to ACC as a critical input facilitating biased target
518 selection by demonstrating its role in supporting the comparison of evidence across consecutive
519 trials. As a population, LP-ACC axonal activity represents sensory information relative to
520 previous trial evidence, with population trajectories in a low-dimensional manifold that is most
521 pronounced with greater differences across trials that require choice switching, particularly to
522 contralateral choices. Thus, disruption of LP-ACC with optogenetic activation biases actions
523 towards ipsilateral choices, particularly for trials requiring switching away from previous choices.
524
525 Our work provides important evidence that the pulvinar is critical for computing predictive
526 differences between stimuli[24,54,55]. Models of pulvinar-cortical circuits demonstrated their capacity
527 to support computing differences between competing stimuli[56] and to estimate decisional
528 confidence for a presented stimulus[26]. LP input to the visual cortices has also been shown to
529 encode mismatches between optic flow and self-generated locomotion[22,23] or eye movements[21],
530 as well as amplify cortical responses to unexpected stimuli[25]. To the best of our knowledge, we
531 provide the first evidence that LP/pulvinar projections to a prefrontal cortical region reflect the
532 degree to which the current stimulus is different from history-dependent expectations, rather
533 than between concurrently occurring visual or visuomotor stimuli. Our work focused on the
534 LP-ACC circuit as it more strongly engaged brain-wide networks involved in perceptual

decision-making[18]. Importantly, our experiments demonstrated a causal influence of these predictive differences in biasing ongoing perceptual decisions.

Role of LP in history-dependent decisional biases

Recent work has shown that serial biases in rodents can be best explained by history-dependent updating of initial states in drift-diffusion models of perceptual decisions[57]. Spatial priors have been reported to be widely distributed across the rodent brain[58]. Furthermore, reinforcement can also bias decisions in the next trial by tuning perceptual discriminability[59,60], or in a confidence-dependent manner, exerting greater influence under greater uncertainty[61,62]. With the pervasive influence of sensory and choice history, there has been interest in further dissecting neural mechanisms mediating history-dependence[57,63,64]. History-dependent biases can influence decisions at perceptual and motor levels, which may also be attributable to different neural circuits. A cortical area that has been implicated in history-dependent biases is the PPC[45,65,66,67,68,69]. Inactivation of the PPC, specifically during the pre-stimulus interval or intertrial interval (ITI), is often associated with reduced history bias[66,67]. Unilateral ChR2 activation of PPC during visual stimuli causes ipsilateral biases, but rats also become more likely to repeat successful contralateral choices[70]. The PPC shares reciprocal connectivity with LP[17,18]. However, unlike with the critical importance of the ITI for history-dependent influences of the PPC, we found that activating LP-ACC axons during the ITI did not lead to appreciable changes in performance accuracy nor history-dependent biases (Figure S6) in spite of the previous trial information being robustly represented in LP-ACC axons during the ITI (Figure 4a). Thus, neural encoding or decodability of a task variable does not necessarily imply a causal contribution to its behavioral impact. In fact, recent large-scale neural recordings have pointed to brain-wide representations of priors and evidence accumulation[58,71]. Yet, as argued for the pervasive brain-wide representations of movement across different brain areas[72], each brain area can contribute differently to the encoding of a variable, be critical at different stages of decision-making, and have distinct functional consequences. We find that LP-ACC axons were specifically important for evaluating ongoing stimuli relative to expectations from trial history, serving as a control signal that likely retrieves information about both current and previous trial evidence from other brain regions on the fly to bias choices. As LP serves as one of many parallel pathways carrying trial history information, this likely also explains why optogenetic stimulation of LP-ACC axons had a disruptive effect limited to only the stimulated trial (Figure 2, Figure S5).

History-dependent geometric representation of decision variables

We showed that LP-ACC axons represent the current decision variable in a structured geometry that is dynamically modulated by trial history. Primate parietal and prefrontal cortical regions engaged in visual decision-making have been reported to represent decision variables along a low-dimensional curved manifold[47]. We demonstrate that such a representational geometry is neither limited to the primate decision-making network, nor to cortical circuits and processing. Activity of thalamic structures, such as the pulvinar, with extensive recurrent connectivity with cortical decision-making networks, can likewise organize into low-dimensional representations that facilitate decision-making, or even contribute to cortical curved manifolds, further reflecting their intimate engagement with cortical networks for perceptual decision-making.

Our work demonstrates that this low-dimensional geometry is neither a result of bottom-up driven visual activity nor due to executed motor activity, as the geometry was absent under passive viewing by naive mice but retained in omitted trials without executed movement in trained mice. We were not able to resolve if the curved manifold representation was a result of learning or attention. While the geometric representation also held in omitted trials, it is challenging to determine if these omissions resulted from the lack of attention, with no intent to attempt a trial (as in sated mice), or sub-threshold evidence accumulation in attentive mice that nevertheless led to missed trials. The geometric properties could instead be a learned representation of the decision variable[73,74]. . Future work examining LP-ACC population activity in trained mice passively viewing the stimuli while not engaged in the task will be necessary to resolve if the low-dimensional geometry is a consequence of learning or attention during task engagement[52,75-79].

It remains an open question whether the geometric representations of cortical and thalamic structures constitute a shared network representation, or if each area could undergo specialized region-specific geometric alterations to support different aspects of perception and choice [47,75,76]. Nevertheless, organizing the decision variable along a curved manifold could be a metabolically efficient manner of affording greater dynamic range to its representation47. The curvature effectively allows greater separation between more ambiguous low-coherence stimuli relative to a linear representation. We found that the low-dimensional geometry of the decision variable in LP-ACC axons was sensitive to modulation by trial history. While we expected that the geometric representation would reflect sensory discrimination, we found that the curved geometry was most prominent when there was a greater | Δ Dir | between consecutive trials

(Figure 5). This raises the intriguing possibility that representation of the decision variable by LP-ACC axons prioritizes differences from prior expectations, such that the population trajectory under greater | Δ Dir | represents an accumulation of evidence relative to prior history, providing drive to switch from a previous choice. Indeed, we found that optogenetic stimulation of LP-ACC axons had a greater impact in trials with | Δ Dir | (Figure 1), where stimulation could introduce noise[80,81,82] or disrupt recurrent dynamics that support organized representation[56,83].

Impact of LP input within ACC

Finally, while we focused on the information sent from the pulvinar/LP to the ACC, it will be important for future work to clarify how this information is utilized by the ACC and its microcircuits. The ACC is particularly important for controlling the accumulation of sensory evidence to reach action thresholds in a manner that also requires the engagement of inhibitory microcircuits[84]. LP inputs to ACC are primarily concentrated in layer 1[18], where it could engage ACC inhibitory mechanisms in cortical microcircuits to influence perceptual decision-making[85,86,87,88]. Further dissection of how ACC circuits incorporate pulvinar/LP input will provide even deeper insight into pulvinar-frontal interactions.

Adaptive allocation of attentional resources based on prior expectations is critical for metabolically efficient perceptual processing. Our work positions pulvinar/LP inputs to the ACC as an important node for selecting and evaluating stimuli relative to prior evidence, in order to modulate behavior on a trial-by-trial basis. This has broad implications for considering thalamic contributions to perceptual filtering and attentional biases.

References

1. van Moorselaar, D., Lampers, E., Cordesius, E., & Slagter, H. A. (2020). Neural mechanisms underlying expectation-dependent inhibition of distracting information. eLife, 9. doi:10.7554/eLife.61048
2. Bartsch, M. V., Merkel, C., Schoenfeld, M. A., & Hopf, J.-M. (2021). Attention expedites target selection by prioritizing the neural processing of distractor features. Communications Biology, 4(1), 814. doi:10.1038/s42003-021-02305-9
3. Fiebelkorn, I. C., & Kastner, S. (2020). Functional specialization in the attention network. Annual Review of Psychology, 71(1), 221–249. doi:10.1146/annurev-psych-010418-103429
4. Kastner, S., & Pinsk, M. A. (2004). Visual attention as a multilevel selection process. Cognitive, Affective & Behavioral Neuroscience, 4(4), 483–500. doi:10.3758/cabn.4.4.483
5. Fiebelkorn, I. C., Pinsk, M. A., & Kastner, S. (2018). A dynamic interplay within the frontoparietal network underlies rhythmic spatial attention. Neuron, 99(4), 842-853.e8. doi:10.1016/j.neuron.2018.07.038
6. Chelazzi, L., Marini, F., Pascucci, D., & Turatto, M. (2019). Getting rid of visual distractors: the why, when, how, and where. Current Opinion in Psychology, 29, 135–147. doi:10.1016/j.copsyc.2019.02.004
7. Rafal, R. D., & Posner, M. I. (1987). Deficits in human visual spatial attention following thalamic lesions. Proceedings of the National Academy of Sciences of the United States of America, 84(20), 7349–7353. doi:10.1073/pnas.84.20.7349
8. Snow, J. C., Allen, H. A., Rafal, R. D., & Humphreys, G. W. (2009). Impaired attentional selection following lesions to human pulvinar: evidence for homology between human and monkey. Proceedings of the National Academy of Sciences of the United States of America, 106(10), 4054–4059. doi:10.1073/pnas.0810086106
9. Kastner, S., & Pinsk, M. A. (2004). Visual attention as a multilevel selection process. Cognitive, Affective & Behavioral Neuroscience, 4(4), 483–500. doi:10.3758/cabn.4.4.483
10. Kastner, S., Fiebelkorn, I. C., & Eradath, M. K. (2020). Dynamic pulvino-cortical interactions in the primate attention network. Current Opinion in Neurobiology, 65, 10–19. doi:10.1016/j.conb.2020.08.002
11. Saalmann, Y. B., Pinsk, M. A., Wang, L., Li, X., & Kastner, S. (2012). The pulvinar regulates information transmission between cortical areas based on attention demands. Science (New York, N.Y.), 337(6095), 753–756. doi:10.1126/science.1223082
12. Fiebelkorn, I. C., Pinsk, M. A., & Kastner, S. (2019). The mediodorsal pulvinar coordinates the macaque fronto-parietal network during rhythmic spatial attention. Nature Communications, 10(1), 215. doi:10.1038/s41467-018-08151-4
13. Fiebelkorn, I. C., & Kastner, S. (2019). The puzzling pulvinar. Neuron, 101(2), 201–203. doi:10.1016/j.neuron.2018.12.032
14. Fischer, J., & Whitney, D. (2012). Attention gates visual coding in the human pulvinar. Nature Communications, 3(1), 1051. doi:10.1038/ncomms2054

15. Zhou, H., Schafer, R. J., & Desimone, R. (2016). Pulvinar-cortex interactions in vision and attention. Neuron, 89(1), 209–220. doi:10.1016/j.neuron.2015.11.034
16. Schneider, K. A. (2011). Subcortical mechanisms of feature-based attention. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience, 31(23), 8643–8653. doi:10.1523/JNEUROSCI.6274-10.2011
17. Bennett, C., Gale, S. D., Garrett, M. E., Newton, M. L., Callaway, E. M., Murphy, G. J., & Olsen, S. R. (2019). Higher-order thalamic circuits channel parallel streams of visual information in mice. Neuron, 102(2), 477-492.e5. doi:10.1016/j.neuron.2019.02.010
18. Leow, Y. N., Zhou, B., Sullivan, H. A., Barlowe, A. R., Wickersham, I. R., & Sur, M. (2022). Brain-wide mapping of inputs to the mouse lateral posterior (LP/Pulvinar) thalamus-anterior cingulate cortex network. The Journal of Comparative Neurology, 530(11), 1992–2013. doi:10.1002/cne.25317
19. Hu, F., Kamigaki, T., Zhang, Z., Zhang, S., Dan, U., & Dan, Y. (2019). Prefrontal corticotectal neurons enhance visual processing through the superior colliculus and pulvinar thalamus. Neuron, 104(6), 1141-1152.e4. doi:10.1016/j.neuron.2019.09.019
20. Fang, Q., Chou, X.-L., Peng, B., Zhong, W., Zhang, L. I., & Tao, H. W. (2020). A differential circuit via retino-colliculo-pulvinar pathway enhances feature selectivity in visual cortex through surround suppression. Neuron, 105(2), 355-369.e6. doi:10.1016/j.neuron.2019.10.027
21. Miura, S. K., & Scanziani, M. (2022). Distinguishing externally from saccade-induced motion in visual cortex. Nature, 610(7930), 135–142. doi:10.1038/s41586-022-05196-w
22. Roth, M. M., Dahmen, J. C., Muir, D. R., Imhof, F., Martini, F. J., & Hofer, S. B. (2016). Thalamic nuclei convey diverse contextual information to layer 1 of visual cortex. Nature Neuroscience, 19(2), 299–307. doi:10.1038/nn.4197
23. Blot, A., Roth, M. M., Gasler, I., Javadzadeh, M., Imhof, F., & Hofer, S. B. (2021). Visual intracortical and transthalamic pathways carry distinct information to cortical areas. Neuron, 109(12), 1996-2008.e6. doi:10.1016/j.neuron.2021.04.017
24. Kanai, R., Komura, Y., Shipp, S., & Friston, K. (2015). Cerebral hierarchies: predictive processing, precision and the pulvinar. Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences, 370(1668), 20140169. doi:10.1098/rstb.2014.0169
25. Furutachi, S., Franklin, A. D., Aldea, A. M., Mrsic-Flogel, T. D., & Hofer, S. B. (2024). Cooperative thalamocortical circuit mechanism for sensory prediction errors. Nature, 633(8029), 398–406. doi:10.1038/s41586-024-07851-w
26. Komura, Y., Nikkuni, A., Hirashima, N., Uetake, T., & Miyamoto, A. (2013). Responses of pulvinar neurons reflect a subject's confidence in visual categorization. Nature Neuroscience, 16(6), 749–755. doi:10.1038/nn.3393
27. White, M. G., Panicker, M., Mu, C., Carter, A. M., Roberts, B. M., Dharmasri, P. A., & Mathur, B. N. (2018). Anterior cingulate cortex input to the claustrum is required for top-down action control. Cell Reports, 22(1), 84–95. doi:10.1016/j.celrep.2017.12.023
28. Jezzini, A., Bromberg-Martin, E. S., Trambaiolli, L. R., Haber, S. N., & Monosov, I. E. (2021). A prefrontal network integrates preferences for advance information about uncertain rewards and punishments. Neuron, 109(14), 2339-2352.e5. doi:10.1016/j.neuron.2021.05.013

29. Bromberg-Martin, E. S., & Monosov, I. E. (2020). Neural circuitry of information seeking. Current Opinion in Behavioral Sciences, 35, 62–70. doi:10.1016/j.cobeha.2020.07.006
30. Bichot, N. P., & Schall, J. D. (1999). Effects of similarity and history on neural mechanisms of visual selection. Nature Neuroscience, 2(6), 549–554. doi:10.1038/9205
31. Sarafyazd, M., & Jazayeri, M. (2019). Hierarchical reasoning by neural circuits in the frontal cortex. Science (New York, N.Y.), 364(6441), eaav8911. doi:10.1126/science.aav8911
32. Hyman, J. M., Holroyd, C. B., & Seamans, J. K. (2017). A novel neural prediction error found in anterior cingulate cortex ensembles. Neuron, 95(2), 447-456.e3. doi:10.1016/j.neuron.2017.06.021
33. Stolyarova, A., Rakhshan, M., Hart, E. E., O'Dell, T. J., Peters, M. A. K., Lau, H., … Izquierdo, A. (2019). Contributions of anterior cingulate cortex and basolateral amygdala to decision confidence and learning under uncertainty. Nature Communications, 10(1), 4704. doi:10.1038/s41467-019-12725-1
34. Kéri, S., Decety, J., Roland, P. E., & Gulyás, B. (2004). Feature uncertainty activates anterior cingulate cortex. Human Brain Mapping, 21(1), 26–33. doi:10.1002/hbm.10150
35. Cole, N., Harvey, M., Myers-Joseph, D., Gilra, A., & Khan, A. G. (2024). Prediction-error signals in anterior cingulate cortex drive task-switching. Nature Communications, 15(1), 7088. doi:10.1038/s41467-024-51368-9
36. Hayden, B. Y., Heilbronner, S. R., Pearson, J. M., & Platt, M. L. (2011). Surprise signals in anterior cingulate cortex: neuronal encoding of unsigned reward prediction errors driving adjustment in behavior. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience, 31(11), 4178–4187. doi:10.1523/JNEUROSCI.4652-10.2011
37. Economides, M., Guitart-Masip, M., Kurth-Nelson, Z., & Dolan, R. J. (2014). Anterior cingulate cortex instigates adaptive switches in choice by integrating immediate and delayed components of value in ventromedial prefrontal cortex. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience, 34(9), 3340–3349. doi:10.1523/jneurosci.4313-13.2014
38. Karlsson, M. P., Tervo, D. G. R., & Karpova, A. Y. (2012). Network resets in medial prefrontal cortex mark the onset of behavioral uncertainty. Science (New York, N.Y.), 338(6103), 135–139. doi:10.1126/science.1226518
39. Tervo, D. G. R., Kuleshova, E., Manakov, M., Proskurin, M., Karlsson, M., Lustig, A., … Karpova, A. Y. (2021). The anterior cingulate cortex directs exploration of alternative strategies. Neuron, 109(11), 1876-1887.e6. doi:10.1016/j.neuron.2021.03.028
40. Kolling, N., Wittmann, M., & Rushworth, M. F. S. (2014). Multiple neural mechanisms of decision making and their competition under changing risk pressure. Neuron, 81(5), 1190–1202. doi:10.1016/j.neuron.2014.01.033
41. Kolling, N., Wittmann, M. K., Behrens, T. E. J., Boorman, E. D., Mars, R. B., & Rushworth, M. F. S. (2016). Value, search, persistence and model updating in anterior cingulate cortex. Nature Neuroscience, 19(10), 1280–1285. doi:10.1038/nn.4382
42. Blanchard, T. C., & Gershman, S. J. (2018). Pure correlates of exploration and exploitation in the human brain. Cognitive, Affective & Behavioral Neuroscience, 18(1), 117–126. doi:10.3758/s13415-017-0556-2

43. Goard, M. J., Pho, G. N., Woodson, J., & Sur, M. (2016). Distinct roles of visual, parietal, and frontal motor cortices in memory-guided sensorimotor decisions. eLife, 5. doi:10.7554/elife.13764
44. Pho, G. N., Goard, M. J., Woodson, J., Crawford, B., & Sur, M. (2018). Task-dependent representations of stimulus and choice in mouse parietal cortex. Nature Communications, 9(1). doi:10.1038/s41467-018-05012-y
45. Scott, B. B., Constantinople, C. M., Akrami, A., Hanks, T. D., Brody, C. D., & Tank, D. W. (2017). Fronto-parietal cortical circuits encode accumulated evidence with a diversity of timescales. Neuron, 95(2), 385-398.e5. doi:10.1016/j.neuron.2017.06.013
46. Runyan, C. A., Piasini, E., Panzeri, S., & Harvey, C. D. (2017). Distinct timescales of population coding across cortex. Nature, 548(7665), 92–96. doi:10.1038/nature23020
47. Okazawa, G., Hatch, C. E., Mancoo, A., Machens, C. K., & Kiani, R. (2021). Representational geometry of perceptual decisions in the monkey parietal cortex. Cell, 184(14), 3748-3761.e18. doi:10.1016/j.cell.2021.05.022
48. Hangya, B., Sanders, J. I., & Kepecs, A. (2016). A mathematical framework for statistical decision confidence. Neural Computation, 28(9), 1840–1858. doi:10.1162/NECO_a_00864
49. Sanders, J. I., Hangya, B., & Kepecs, A. (2016). Signatures of a statistical computation in the human sense of confidence. Neuron, 90(3), 499–506. doi:10.1016/j.neuron.2016.03.025
50. Rigotti, M., Barak, O., Warden, M. R., Wang, X.-J., Daw, N. D., Miller, E. K., & Fusi, S. (2013). The importance of mixed selectivity in complex cognitive tasks. Nature, 497(7451), 585–590. doi:10.1038/nature12160
51. Mante, V., Sussillo, D., Shenoy, K. V., & Newsome, W. T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. Nature, 503(7474), 78–84. doi:10.1038/nature12742
52. Kaufman, M. T., Benna, M. K., Rigotti, M., Stefanini, F., Fusi, S., & Churchland, A. K. (2022). The implications of categorical and category-free mixed selectivity on representational geometries. Current Opinion in Neurobiology, 77(102644), 102644. doi:10.1016/j.conb.2022.102644
53. Christensen, A. J., Ott, T., & Kepecs, A. (2022). Cognition and the single neuron: How cell types construct the dynamic computations of frontal cortex. Current Opinion in Neurobiology, 77(102630), 102630. doi:10.1016/j.conb.2022.102630
54. O'Reilly, R. C., Russin, J. L., Zolfaghar, M., & Rohrlich, J. (2021). Deep predictive learning in neocortex and pulvinar. Journal of Cognitive Neuroscience, 33(6), 1158–1196. doi:10.1162/jocn_a_01708
55. Saalmann, Y. B., & Kastner, S. (2015). The cognitive thalamus. Frontiers in Systems Neuroscience, 9, 39. doi:10.3389/fnsys.2015.00039
56. Jaramillo, J., Mejias, J. F., & Wang, X.-J. (2019). Engagement of pulvino-cortical feedforward and feedback pathways in cognitive computations. Neuron, 101(2), 321-336.e9. doi:10.1016/j.neuron.2018.11.023
57. Gupta, D., DePasquale, B., Kopec, C. D., & Brody, C. D. (2024). Trial-history biases in evidence accumulation can give rise to apparent lapses in decision-making. Nature Communications, 15(1), 662. doi:10.1038/s41467-024-44880-5

58. Findling, C., Hubert, F., Acerbi, L., Benson, B., Benson, J., Birman, D., … International Brain Laboratory. (2023). Brain-wide representations of prior information in mouse decision-making. doi:10.1101/2023.07.04.547684
59. Law, C.-T., & Gold, J. I. (2009). Reinforcement learning can account for associative and perceptual learning on a visual-decision task. Nature Neuroscience, 12(5), 655–663. doi:10.1038/nn.2304
60. Reato, D., Steinfeld, R., Tacão-Monteiro, A., & Renart, A. (2023). Response outcome gates the effect of spontaneous cortical state fluctuations on perceptual decisions. eLife, 12. doi:10.7554/eLife.81774
61. Lak, A., Hueske, E., Hirokawa, J., Masset, P., Ott, T., Urai, A. E., … Kepecs, A. (2020). Reinforcement biases subsequent perceptual decisions when confidence is low, a widespread behavioral phenomenon. eLife, 9. doi:10.7554/eLife.49834
62. Pedrosa, V., Menichini, E., Pajot-Moric, Q., Vincent, P., Zhou, L., Teachen, L., … Akrami, A. (2023). Humans, rats and mice show species-specific adaptations to sensory statistics in categorisation behaviour. doi:10.1101/2023.01.30.526119
63. Hachen, I., Reinartz, S., Brasselet, R., Stroligo, A., & Diamond, M. E. (2021). Dynamics of history-dependent perceptual judgment. Nature Communications, 12(1), 6036. doi:10.1038/s41467-021-26104-2
64. Le, N. M., Yildirim, M., Wang, Y., Sugihara, H., Jazayeri, M., & Sur, M. (2023). Mixtures of strategies underlie rodent behavior during reversal learning. PLoS Computational Biology, 19(9), e1011430. doi:10.1371/journal.pcbi.1011430
65. Morcos, A. S., & Harvey, C. D. (2016). History-dependent variability in population dynamics during evidence accumulation in cortex. Nature Neuroscience, 19(12), 1672–1681. doi:10.1038/nn.4403
66. Hwang, E. J., Dahlen, J. E., Mukundan, M., & Komiyama, T. (2017). History-based action selection bias in posterior parietal cortex. Nature Communications, 8(1), 1242. doi:10.1038/s41467-017-01356-z
67. Akrami, A., Kopec, C. D., Diamond, M. E., & Brody, C. D. (2018). Posterior parietal cortex represents sensory history and mediates its effects on behaviour. Nature, 554(7692), 368–372. doi:10.1038/nature25510
68. Braun, A., Urai, A. E., & Donner, T. H. (2018). Adaptive history biases result from confidence-weighted accumulation of past choices. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience, 38(10), 2418–2429. doi:10.1523/jneurosci.2189-17.2017
69. Urai, A. E., & Donner, T. H. (2022). Persistent activity in human parietal cortex mediates perceptual choice repetition bias. Nature Communications, 13(1), 6015. doi:10.1038/s41467-022-33237-5
70. Licata, A. M., Kaufman, M. T., Raposo, D., Ryan, M. B., Sheppard, J. P., & Churchland, A. K. (2017). Posterior parietal cortex guides visual decisions in rats. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience, 37(19), 4954–4966. doi:10.1523/JNEUROSCI.0105-17.2017
71. Khilkevich, A., Lohse, M., Low, R., Orsolic, I., Bozic, T., Windmill, P., & Mrsic-Flogel, T. D. (2024). Brain-wide dynamics linking sensation to action during decision-making. Nature. doi:10.1038/s41586-024-07908-w

72. Wang, Z. A., Chen, S., Liu, Y., Liu, L. D., Svoboda, K., Li, N., & Druckmann, S. (2023). Not everything, not everywhere, not all at once: a study of brain-wide encoding of movement. doi:10.1101/2023.06.08.544257
73. Mease, R. A., Metz, M., & Groh, A. (2016). Cortical sensory responses are enhanced by the higher-order thalamus. Cell Reports, 14(2), 208–215. doi:10.1016/j.celrep.2015.12.026
74. Peters, A. J., Marica, A.-M., Fabre, J. M. J., Harris, K. D., & Carandini, M. (2022). Visuomotor learning promotes visually evoked activity in the medial prefrontal cortex. Cell Reports, 41(3), 111487. doi:10.1016/j.celrep.2022.111487
75. Okazawa, G., & Kiani, R. (2023). Neural mechanisms that make perceptual decisions flexible. Annual Review of Physiology, 85, 191–215. doi:10.1146/annurev-physiol-031722-024731
76. Bernardi, S., Benna, M. K., Rigotti, M., Munuera, J., Fusi, S., & Salzman, C. D. (2020). The geometry of abstraction in the hippocampus and prefrontal cortex. Cell, 183(4), 954-967.e21. doi:10.1016/j.cell.2020.09.031
77. Nieh, E. H., Schottdorf, M., Freeman, N. W., Low, R. J., Lewallen, S., Koay, S. A., … Tank, D. W. (2021). Geometry of abstract learned knowledge in the hippocampus. Nature, 595(7865), 80–84. doi:10.1038/s41586-021-03652-7
78. Kobak, D., Brendel, W., Constantinidis, C., Feierstein, C. E., Kepecs, A., Mainen, Z. F., … Machens, C. K. (2016). Demixed principal component analysis of neural population data. eLife, 5. doi:10.7554/elife.10989
79. Courellis, H. S., Minxha, J., Cardenas, A. R., Kimmel, D. L., Reed, C. M., Valiante, T. A., … Rutishauser, U. (2024). Abstract representations emerge in human hippocampal neurons during inference. Nature, 632(8026), 841–849. doi:10.1038/s41586-024-07799-x
80. Fetsch, C. R., Kiani, R., Newsome, W. T., & Shadlen, M. N. (2014). Effects of cortical microstimulation on confidence in a perceptual decision. Neuron, 83(4), 797–804. doi:10.1016/j.neuron.2014.07.011
81. Rahnev, D., Maniscalco, B., Graves, T., Huang, E., de Lange, F. P., & Lau, H. (2011). Attention induces conservative subjective biases in visual perception. Nature Neuroscience, 14(12), 1513–1515. doi:10.1038/nn.2948
82. Rahnev, D., Lau, H., & de Lange, F. P. (2011). Prior expectation modulates the interaction between sensory and prefrontal regions in the human brain. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience, 31(29), 10741–10748. doi:10.1523/JNEUROSCI.1478-11.2011
83. Sauerbrei, B. A., Guo, J.-Z., Cohen, J. D., Mischiati, M., Guo, W., Kabra, M., … Hantman, A. W. (2020). Cortical pattern generation during dexterous movement is input-driven. Nature, 577(7790), 386–391. doi:10.1038/s41586-019-1869-9
84. Kim, J.-H., Ma, D.-H., Jung, E., Choi, I., & Lee, S.-H. (2021). Gated feedforward inhibition in the frontal cortex releases goal-directed action. Nature Neuroscience, 24(10), 1452–1464. doi:10.1038/s41593-021-00910-9
85. Cruikshank, S. J., Ahmed, O. J., Stevens, T. R., Patrick, S. L., Gonzalez, A. N., Elmaleh, M., & Connors, B. W. (2012). Thalamic control of layer 1 circuits in prefrontal cortex. The

Journal of Neuroscience: The Official Journal of the Society for Neuroscience, 32(49), 17813–17823. doi:10.1523/JNEUROSCI.3231-12.2012

86. Lee, A. J., Wang, G., Jiang, X., Johnson, S. M., Hoang, E. T., Lanté, F., … Julius Zhu, J. (2015). Canonical organization of Layer 1 neuron-led cortical inhibitory and disinhibitory interneuronal circuits. Cerebral Cortex (New York, N.Y.: 1991), 25(8), 2114–2126. doi:10.1093/cercor/bhu020
87. Schuman, B., Machold, R. P., Hashikawa, Y., Fuzik, J., Fishell, G. J., & Rudy, B. (2019). Four unique interneuron populations reside in neocortical layer 1. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience, 39(1), 125–139. doi:10.1523/JNEUROSCI.1613-18.2018
88. Fan, L. Z., Kheifets, S., Böhm, U. L., Wu, H., Piatkevich, K. D., Xie, M. E., … Cohen, A. E. (2020). All-optical electrophysiology reveals the role of lateral inhibition in sensory processing in cortical layer 1. Cell, 180(3), 521-535.e18. doi:10.1016/j.cell.2020.01.001
89. Friedrich, J., Zhou, P., & Paninski, L. (2017). Fast online deconvolution of calcium imaging data. PLoS Computational Biology, 13(3), e1005423. doi:10.1371/journal.pcbi.1005423
90. Giovannucci, A., Friedrich, J., Gunn, P., Kalfon, J., Brown, B. L., Koay, S. A., … Pnevmatikakis, E. A. (2019). CaImAn an open source tool for scalable calcium imaging data analysis. eLife, 8. doi:10.7554/eLife.38173

STAR Methods

Animals and Surgery

All experiments were performed under protocols approved by the MIT Institutional Animal Care and Use Committee (IACUC) and conformed to NIH guidelines. Stereotaxic surgeries were performed under isofluorane anaesthesia (4% for induction and 1% for maintenance). Viral injections were performed using glass micropipettes and injected at a rate of 50-75nl/minute. 75nl of AAV1-syn-jGCaMP7f (Addgene:104488-AAV1) was injected in the left medial LP (Relative to Bregma: -1.8mm AP, -1.2mm ML, -2.6DV) contralateral to the visual stimulus screen. 3mm wide cranial windows (Two 3mm glass coverslips glued under a 5mm glass coverslip for stability) were im- planted centered on the midline at AP +1mm (relative to Bregma) over ACC/M2. A titanium head-plate was implanted over window to allow head-fixation during imaging. Implants were secured with dental cement (CB Metabond) mixed with black ink to render the dental cement opaque. After surgery, mice recovered on a heated water-circulation pad for at least an hour. Mice were injected subcutaneously with slow-release (72hr) Buprenorphine (1mg/kg) before surgery and Meloxicam (5mg/kg) after surgery for analgesia on the day of surgery. Meloxicam (5mg/kg) was administered once every 24 hours for 2 days after surgery or longer until mice begin nesting. After surgery, mice were pair- (if there is a compatible cage-mate) or single-housed. For mice imaged under passive viewing contexts, imaging was performed 3-6 weeks after surgery.

Mice trained on the visual discrimination task recovered post-surgery for 7-14 days before being placed on water regulation. A running wheel was also introduced into their home-cage at this time as enrichment and to habituate them to the running wheel on the rig. Mice were closely monitored to ensure they do not lose more than 10% of their body weight over the span of a week. If so, or if demonstrating signs of poor health, mice were returned to ad lib water for at least a week before a slow taper returning to water regulation with an increased daily water allowance. If mice did not receive all their daily water during a training session, they were supplemented up to their daily water (ranging 1-1.5ml depending on the starting weight of the mouse) with hydrogel.

Visual Stimulus Presentation

Random dot kinematogram (RDK) stimuli were generated and presented through PsychoPy on a 22.5cm x 14 cm screen at 1360 x 768 pixel resolution at 60 frames per second. The screen was placed contralateral to the imaged hemisphere, 8 cm away from the right eye of the mouse and oriented approximately 35 degrees relative to the mouse head-body axis. The size of each dot corresponded to 2° visual angle. Approximately 280 dots were presented in a full field during each trial. Each random dot had a lifetime of 500ms, after which another dot reappeared randomly in the field moving in the same direction. In presentations where dot motion coherence was varied, the noise dots moved at the same speed in a random (out of 8 possible directions) but with constant direction for its lifetime.

In the task, RDKs were presented at uniform speed of 46°/s. In naive passive conditions, sessions contained presentations with random dots at three different speeds (27°, 46°, or 57° / s). However, for the purpose of comparison with the task context, we only used data from trials with RDKs moving at 46°/s. In the naive passive condition, each stimulus presentation lasted 1.5s following a blank inter-stimulus period of 1.5 seconds.

In the task condition, the presentation of the visual stimulus was self-paced, with at least a 3-second inter-trial interval, a new trial began when the mouse ran at a threshold speed for 1s and had no detectable licks for at least 1s. An RDK stimulus was then presented for 1.5s. Stimuli were presented in pseudo-randomly shuffled blocks of 8 unique stimuli (2 directions and for each, 4 coherences: 16, 32, 64 and 100%), to ensure sessions contained comparable proportions of every unique stimulus, which is presented at least once every 8 trials.

Two-photon imaging and pre-processing

LP-ACC axons were imaged through a 3mm cranial window over the ACC with an Olympus objective (25x, 1.05NA) at 4x zoom. GCaMP7f was excited at 930nm wavelength with a Ti Sapphire laser. Images were acquired at 512 x 381 pixels with resonant scanning at 40Hz, with frames averaged to a working frame-rate of 10Hz with which we performed motion correction. Videos were motion-corrected with a correlation-based rigid template registration using ImageJ. Individual axonal bouton regions of interest (ROIs), were manually selected from the maximum Z-projection of the video. A 2-pixel wide annulus is selected around each ROI as an estimate of the background neuropil signal. The fluorescence signal from the ROI is taken as $F_{ROI} = F_{ROIraw} - 0.7 * F_{neuropil}$. The signals were smoothed by filtering with a causal Gaussian kernel. The session

signals were detrended by fitting a linear regression across the entire time series. For computing dF/F0 for each ROI, we fit a histogram of the fluorescence signal with a Gaussian kernel and take the fluorescence signal corresponding to the peak of the fitted kernel as the baseline F0. dF/F0 was then computed as (F - F0)/F0. Signals were denoised with the OASIS algorithm [89] as implemented in the CaImAn package [90].

Axonal Clustering

In order to cluster boutons that may potentially be from the same neuronal unit, we performed pairwise correlations between every bouton pair in a session. We then performed hierarchical clustering based on these pairwise temporal correlations. The number of clusters was determined empirically for each session by iteratively testing decreasing distance thresholds to satisfy a cluster correlation threshold where 1) The minimum correlation within a cluster was >0.2, 2) at least two times that of the 85th percentile between/inter-cluster correlation, and 3) there is no statistical difference between inter- and intracluster correlations when the same distance threshold was applied to a shuffled dataset where individual units were temporally randomly shifted by 500-1000 frames (50-100 seconds). We weighted boutons of each cluster by the signal-to-noise ratio of the ROI, taking a weighted average of fluorescence signals to represent bouton clusters (referred to as axons throughout the manuscript).

Optogenetics

Laser stimulation was delivered either via an optical fiber placed over 3mm-wide cranial window or through a 400um wide optic fiber implanted over the ACC. ChR2-expressing axons were stimulated with a 493nm laser during the visual stimulus period (1.5s) at 20 Hz (10mW/mm2, 493nm). For the post-reinforcement activation, stimulation was delivered at 10 Hz (10mW/mm2, 493nm) over 3 seconds. As within-subject controls, we performed optogenetic stimulation experiments with a 593nm laser (10mW/mm2), which fell outside of the range of ChR2 excitation spectrum. For stimulation during the post-reinforcement period, we ensured that a new trial could not be initiated until 500ms after the laser stimulation ended. Laser stimulation was triggered through WaveSurfer (Janelia). 2 or 3 trials were randomly selected in each stimulus block (8 different stimuli) to be laser On trials (~20-25% of total trials).

Statistics

Statistical analyses were performed using the Python libraries pingouin or statsmodels. Regular and binomial linear mixed-effects models were fitted in R using the lme4 package. Where there were multiple comparisons, P-values were reported after Benjamini-Hochberg's adjustment.

SVM Decoding

Decoding of stimulus features (direction, coherence etc) was performed by training cross-validated linear SVM classifiers on axonal activity in 200 ms time bins. Trials from each session were randomly shuffled, then stratified to have equal proportions of each of the unique stimuli in training and testing sets, split in a 70:30 ratio. Unless otherwise stated, training sets contained both hit and error trials for task-engaged conditions. Decoding efficacy was evaluated by area under the receiver operator characteristic (AUROC) curve using held out test sets (30% of trials). For omitted trials, only trials where no early (during stimulus) or late (after response period) licking detected were included to exclude motor confounds. We performed 100 iterations for decoding of each session, with random training and testing trial splits in each iteration. Where decoding involves more than two targets, decoders were trained using a one-versus-one strategy to maximally discriminate between all targets, while reported AUROC represents the average of all one-versus-rest AUROC scores. For each iteration, decoders were additionally trained and tested on axonal activity with shuffled target labels to serve as controls.

Encoding Models

Trials were shuffled and then split 3-fold into training and testing sets, stratified to have equal proportions of different trial stimuli (8 different types) in training and testing sets. Trials spanning 1s before and 5s after stimulus onset were then concatenated to construct the full design matrix. As predictors, we include 3 temporal factors (intertrial-interval (1s, 10 frames), stimulus (1.5s, 15 frames), choice (variable number of frames between stimulus offset and reinforcement) and reinforcement (variable number of frames from reinforcement to 5s after stimulus onset) that spanned the time bins in each epoch. The motor parameters we included were right lick rates, left lick rates, and running speed, which were included as rolling averages (500 ms window averaging). We included trial-based selectivity for stimulus direction, choice, and trial outcome,

each ranging -1 to +1 and graded by coherence. To capture the temporal dynamics of each trial predictor, predictors were expanded in time in the form of a Toeplitz matrix. A coefficient was obtained for every time bin of the trial for every predictor, together constituting the fitted response kernel for that predictor. All task selectivity factors were represented from stimulus onset, while previous trial parameters were represented from the ITI onwards. Each predictor in the design matrix was mean-centered. The dFF of each axon was normalized by its maximum value. Regression fits were regularized using ‘elasticnet‘ penalty with a 1:99 L1:L2 regularization ratio. For each axonal unit, we ran 100 iterations of GLM fits with different combinations of training and testing sets, reporting the mean R2 and model coefficients derived from all iterations. We considered a unit task-responsive if the task-dependent model, on held-out test trials, had an R2 significantly greater than 0.01 across all 100 iterations (one-tailed, one-sample t-test, $p < 0.05$). A unit was regarded to encode a specific task variable during that epoch if dropping the task variable from the model caused a significant reduction in explained variance R2 (one-tailed, one-sample t-test, $p<0.05$) relative to the full model, across 100 iterations.

Dimensionality reduction

For dimensionality reduction by Principal Components Analysis (PCA), trials in each session were randomly split, with a half of trials serving as a training set, while the other half was held out for cross-validation. The split was stratified by the signed coherence of the stimulus direction to enforce equal proportions of unique stimuli in both fitting and validation sets. Each set was trial-averaged for every unique stimulus, aligned to the onset.. Every unit was mean subtracted and scaled to have unit variance. We then pooled axons from across sessions and mice into a pseudo-population (N units). PCA was performed on a response matrix composed of pseudo-population activity in a time window spanning 200-1200 ms (T, 10 time bins of 100 ms each) after stimulus onset, concatenating the time windows across each of the 8 stimuli (S) (40 x N). The averaged population activity of axons in the held-out cross-validation trials was then projected onto the low dimensional PC space.

For targeted dimensionality reduction, we performed cross-validated targeted dimensionality reduction onto stimulus direction and coherence axes, using a procedure adapted from Mante et al., 2013 [51]. We first mean-centered and scaled the Ca2+ data to have zero mean and unit variance. The traces across the entire session were then denoised by applying PCA. Ca2+ activity was projected onto the subspace defined by the minimum number of principal components required to retain 75% of the total explained variance, followed by an inverse

transformation to reconstruct the denoised data. As for PCA, we performed stratified splits of each session into random halves, with half the trials serving as a fitting set while the other half served as a validation set. For the DirectionCD, we computed an average difference vector (N x 1) by subtracting the population activity vectors for left versus right stimulus directions, averaged over the same time window as described above (200–1200 ms after stimulus onset). For the CoherenceCD, the difference vector was computed between 100% coherence trials and 16% coherence trials. The selectivity vectors for DirectionCD and CoherenceCD were then orthogonalized using QR decomposition. Following orthogonalization, we projected the population activity from the cross-validation set onto the discriminant axes. This procedure was repeated for 100 iterations, with each iteration using a different random split of the data for fitting and cross-validation.

Acknowledgments

We thank Taylor Johns for lab management and members of the Sur laboratory for many discussions and comments. We also thank Leopoldo Petreanu and former lab members for pointers on beginning the 2AFC RDK discrimination behavior. This work was supported by NIH grants R01MH126351, R01MH133066, R01NS130361, MURI Grant W911NF2110328, The Picower Institute Innovation Fund, and the Simons Foundation Autism Research Initiative through the Simons Center for the Social Brain (MS); a fellowship from A*STAR Singapore (YNL); and fellowships from the Japan Society for the Promotion of Science and Uehara Memorial Research Fund (YO).

Supplementary Figures

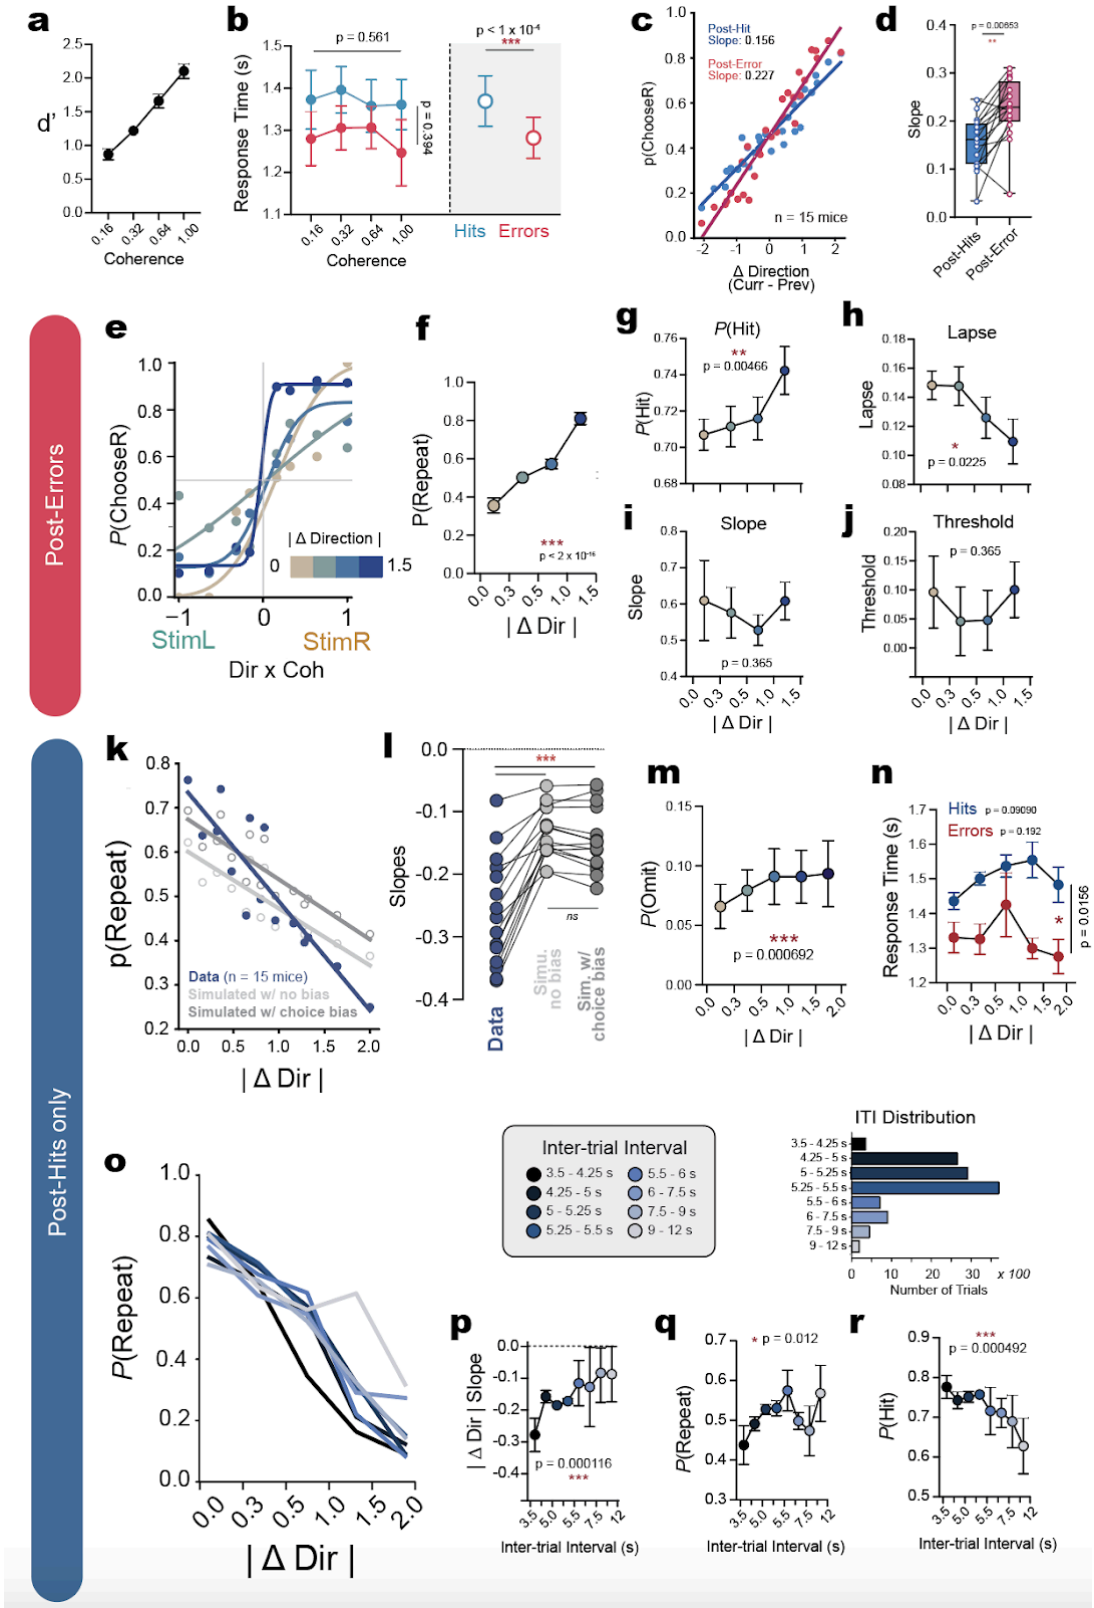

Supplementary Figure 1: Coherence- and | Δ Dir |-dependent behavioral performance during 2AFC RDK direction discrimination task. a) Perceptual sensitivity (d-prime/d') scales with coherence of the RDK stimuli. **b)** The reaction time of mice did not scale with RDK coherence for hits (blue) and errors (red), likely due to the post-stimulus response window enforced. However, errors were associated with faster reaction times. p-values are reported for the effect of coherence (p=0.561), interaction of coherence and trial outcome (p=0.394) and outcome ($p < 1 \times 10^{-4}$) as modeled in a LME. **c)** P(ChooseR) showed strong dependence on the change in direction between current and previous trials (Δ Direction). This relationship was stronger after errors than after hits. Slopes are reported for a linear least squares fit to population data from n = 15 mice. **d)** Slopes for the influence of Δ Direction on P(ChooseR), from linear least squares fit for individual mice for post-hit and post-error trials. Dependence on Δ Direction was generally higher after an error trial than after a hit trial. p-value reported for two-tailed paired t-test. **e)** | Δ Direction |-dependent influences on psychometric curves (population fit from n = 15 mice). **f)** After an erroneous choice, mice are least likely to repeat the choice when the difference between the current and previous stimulus (| Δ Dir |) is low. This relationship is the opposite of trials after hits (Figure 1d). Thus, mice adapt the | Δ Dir | strategy depending on the reinforcement feedback from the previous trial. **g)** Mice were more likely to achieve hits with greater differences in | Δ Dir |. p-value reported for binomial GLME, additionally accounting for current and previous trial coherences as covariates. **h-j)** Parameters from the psychometric fit: lapse, slope, threshold. p-values reported for one-way repeated measures ANOVA. For g-j, error bars reflect SEM across individual mice (n = 15). For e-j (post-error trials), | Δ Dir | is reported only up to 1.5 (maximum of 2) due to low numbers of post-error trials falling between 1.5 and 2 | Δ Dir |. **k)** The dependence of P(Repeat) on | Δ Dir | in the data (blue) is steeper than expected from behavioral models with no bias (light grey) or simple previous choice biases that does not regard strength of relative evidence (dark grey), suggesting the relationship between P(Repeat) and | Δ Dir | is not a simply an artefact of consecutive stimulus regularities. **l)** Comparison of slopes from empirical data fitted for each mouse (blue), matched with simulations based on the data without (light grey) and with (dark gray) previous choice bias. *** $p < 0.0001$ paired t-test. **m)** Trials with greater | Δ Dir | are more likely to be omitted. p-value reported for the effect of | Δ Dir | in a binomial LME, after accounting for influences of previous and current trial coherence. **n)** Influence of | Δ Dir | on reaction times. As in b), reaction times were faster in error trials. The influence of | Δ Dir | on response times on hit and error trials followed an inverted-U shape trend, being fastest in trials of least ambiguity in | Δ Dir | (the poles: most similar - 0, and most different - 2) and slower in the face of more ambiguous relative evidence, and was not significant in a linear model (hits: p = 0.09090, errors: 0.192). **o)** The influence of | Δ Dir | on P(Repeat) of rewarded trials, plotted for trials binned by duration of ITIs. **p)** The strength of | Δ Dir | influence (negative slope) is greatest with short ITIs, and decays with longer ITIs. p-value reported for the effect of ITI on slope describing the relationship between | Δ Dir | and P(Repeat) (as in o), repeated measures ANOVA. **q)** Effect of ITI on P(Repeat), p-value reported for variance explained by continuous values of ITIs in binomial GLME of P(Repeat) which accounts for previous and current trial difficulty. **r)** P(Hit) was greater with shorter ITIs. p-value reported for variance explained by continuous values of ITIs in binomial GLME for P(Hit) which accounts for previous and current trial difficulty. k-r include only trials after rewarded choices. Data from n = 15 mice.

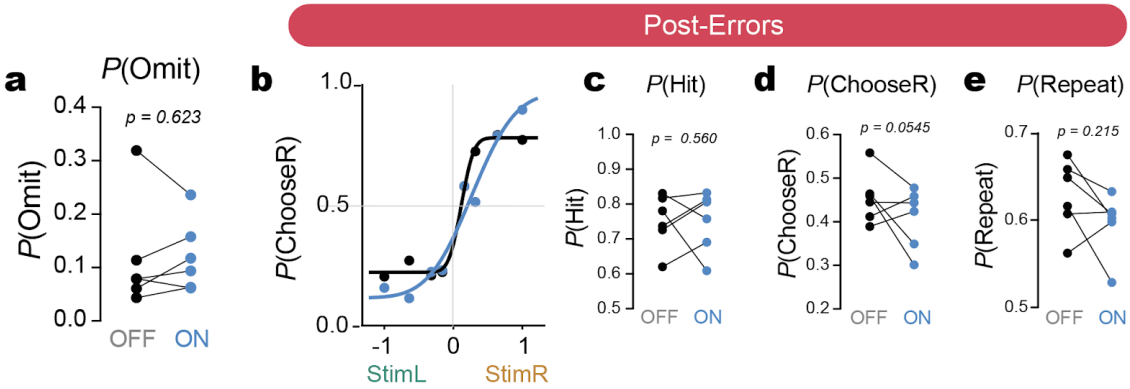

Supplementary Figure 2: Optogenetic stimulation during the stimulus period: Omissions and Post-Errors. a) Stimulation did not lead not have a significant effect on probability of omissions after accounting for coherence, considering all trials regardless of previous outcome. **b)** Psychometric fits to the average perceptual performance with (blue) and without (black) optogenetic stimulation only after error trials, for n = 6 mice. **c)** No change in P(Hit) during optogenetic stimulation trials after errors. **d)** Optogenetic stimulation during trials after a previous error led to small but not statistically significant reduction in P(ChooseR). **e)** Optogenetic stimulation during trials after a previous error did not influence P(Repeat). Note that samples for post-error trials are generally more limited. For b - e, n = 1209 trials (light off) and n = 436 (light on) trials across 66 sessions for n = 6 mice. All p-values are reported for the binary optogenetic trial factor in Binomial Linearized Mixed Effects Models of the variable in question. All models account for the previous trial difficulty/coherence as a covariate which would otherwise be disproportionately represented in post-error trials since errors are more frequent for difficult stimuli.

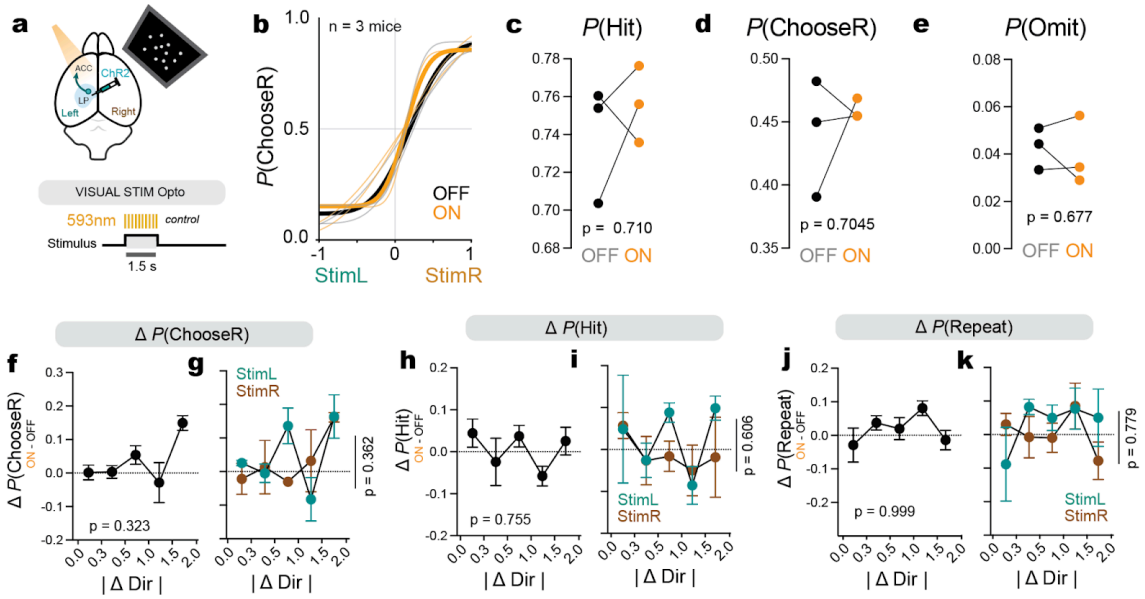

Supplementary Figure 3: Illuminating ChR2-expressing LP-ACC axons with 593nm light had no effect on behavior performance. a) Experimental schematic for control experiment with 593nm light delivered at the same frequency (20 Hz) as with blue light in the true experiment, during the stimulus presentation epoch. **b)** Psychometric fits for performance during laser off (black, n = 3081 trials) and laser on (orange, n = 1404 trials) trials for n = 3 mice. **c)** Illuminating LP-ACC axons with 593nm light did not have a significant impact on P(Hit), **d)** P(ChooseR), or **e)** P(Omit), based on two-tailed Wilcoxon signed-rank test. Each averaged value for a single mouse samples equally from all stimulus levels. f-k) 593nm light also did not influence | Δ Dir |-dependent changes in **f)** P(ChooseR) **g)** for both left- and right-associated stimuli. **h & i)** Same as f & g, but for P(Hit), **j & k)** Same as f & g, but for P(Repeat). p-values reported in f-k derived from weights of coefficients for the interaction between optogenetic stimulation and the | Δ Dir | factor, in binomial linearized mixed effects (LME) models.

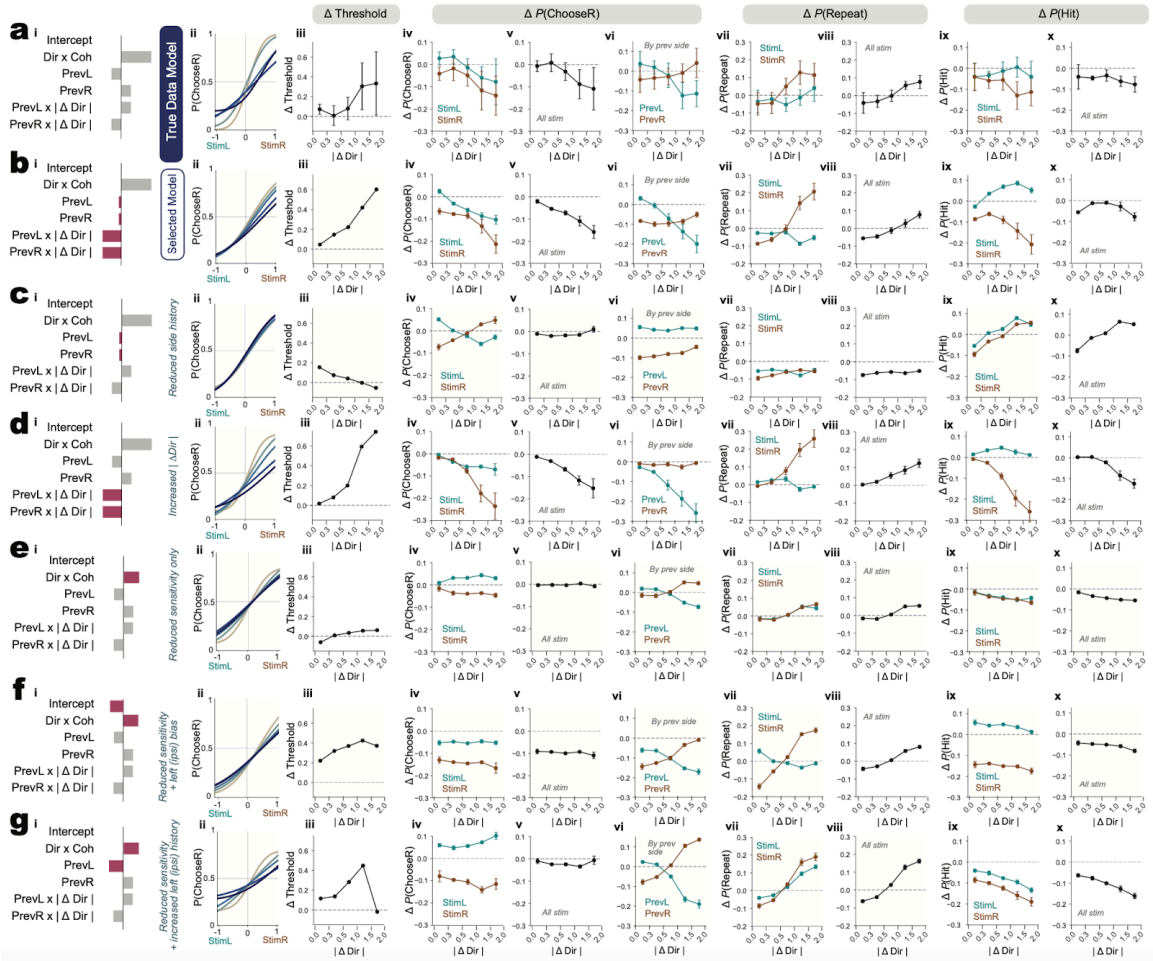

Supplementary Figure 4: Model-simulated impact of altering different task parameters. Rows **a-g** each represent a binomial model for P(ChooseR), and consequences of altered model coefficients shown across columns ii-x. The first column **(i)** indicates the weights of coefficients in these models. Grey bars indicate retention of fitted coefficients in the base model, while magenta bars indicate coefficients that have been altered to simulate optogenetic effects. The coefficients were increased or decreased, always relative to the base model (a). **ii)** Psychometric curves fitted to the simulated model ("mock opto"), plotted by | Δ Dir |. **iii)** Change (Δ) in psychometric thresholds, from psychometric function fits. **iv-vi)** Δ P(ChooseR) as a function of | Δ Dir |, **iv)** by current stimulus side, **v)** mean across all stimuli and **vi)** by previous stimulus/choice. **vii)** Δ P(Repeat) as a function of | Δ Dir |, **vii)** by current stimulus side and, **viii)** mean across all stimuli. **ix-x)** Δ P(Hit) as a function of | Δ Dir |, **ix)** by current stimulus side and, **x)** mean across all stimuli. **a)** Base Model output, after fitting to empirical data of non-opto trials across mice. All coefficients for the binomial model were retained. iii-x show the changes based on empirical optogenetic data. **b)** The final best fit model (as in Figure 2-s). **c)** Model with reduced influence of previous choice, but both supporting a left biases. **d)** Model with increased influence of how | Δ Dir | interactions with previous choice, but both scaling to promote left biases. **e)** Model with simple reduction in sensitivity to current stimuli (Dir x Coh). Model with reduced sensitivity to current stimuli as in e, but **f)** with an added left shift (intercept) or **g)** increased influence of previous left choices.

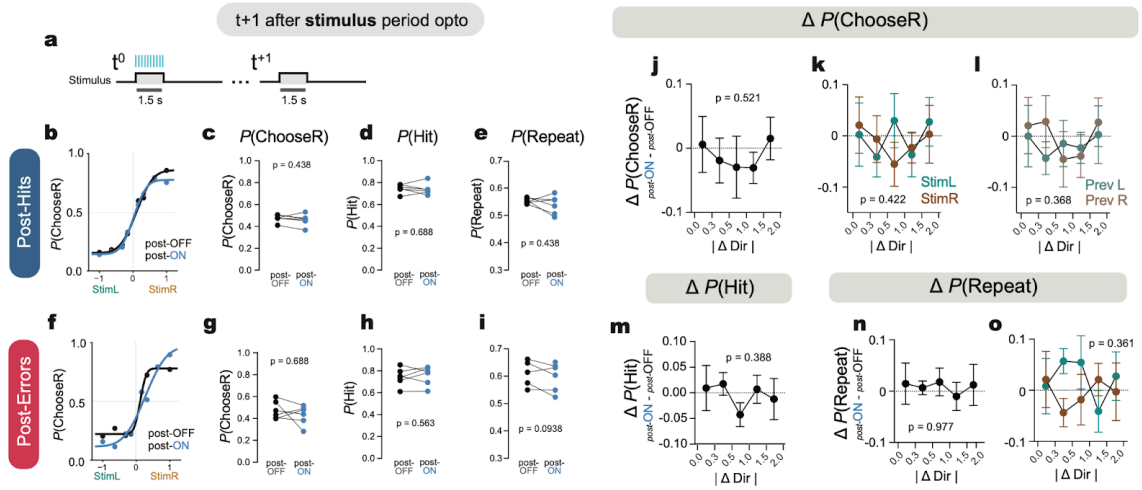

Supplementary Figure 5: Optogenetic activation during the visual stimulus evaluation epoch has no impact on the subsequent trial. a) Schematic showing optogenetic activation period during the visual stimulus period in the prior trial (t_0). **b)** Psychometric fits with (blue) and without (black) optogenetic activation in the previous trial. No change in **c)** P(ChooseR), **d)** P(Hit), **e)** P(Repeat) after hit trials with optogenetic stimulation. **f-i)** Same as b-e but after error trials. j-o) The relationship between | Δ Dir | and **j-l)** P(ChooseR), **m)** P(Hit) or, **n-o)** P(Repeat) was not disrupted in the trial after optogenetic activation. p-values indicated in c-e & g-i are for two-tailed Wilcoxon's test. p-values reported in j-o were from interaction terms for post-opto coefficients in binomial linear mixed effects model for the respective dependent variable, with unbinned continuous range of | Δ Dir |. Data reported for only consecutively attempted trials. 7315 total trials, 76 sessions, 6 mice.

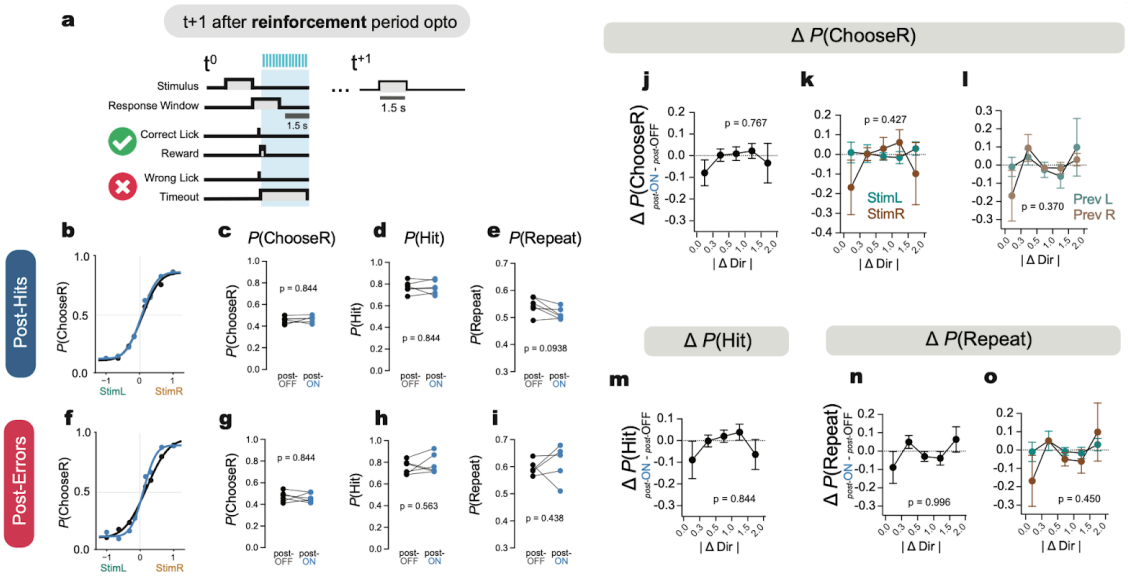

Supplementary Figure 6: Optogenetic activation during reinforcement period has no impact on the subsequent trial. a) Schematic showing optogenetic activation period during reinforcement in the previous trial (hits or errors, t_0). **b)** Psychometric fits with (blue) and without (black) optogenetic activation in the previous trial. No change in **c)** P(ChooseR), **d)** P(Hit), **e)** P(Repeat) after hit trials with optogenetic activation. **f-i)** Same as b-e but after error trials. j-o) The relationship between | Δ Dir | and **j-l)** P(ChooseR), **m)** P(Hit) or, **n-o)** P(Repeat) was not disrupted in the trial after optogenetic activation. p-values indicated in c-e & g-i are for two-tailed Wilcoxon's test. p-values reported in j-o were from interaction terms for post-opto coefficients in binomial linear mixed effects model for the respective dependent variable, with unbinned continuous range of | Δ Dir |. Data reported for only consecutively attempted trials. 9527 total trials, 88 sessions, 6 mice.

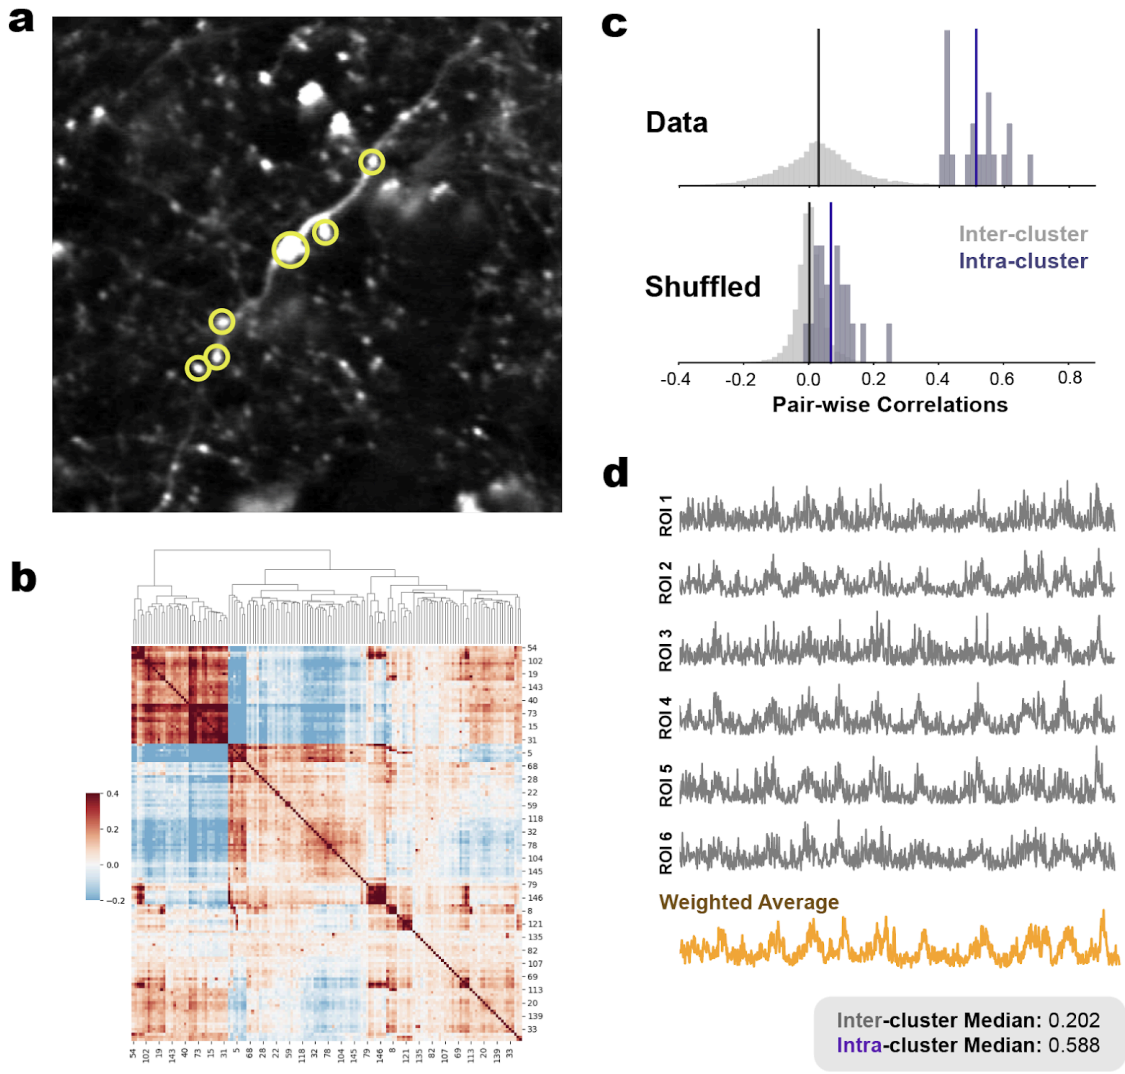

Supplementary Figure 7: Processing steps in clustering individual bouton ROIs into axonal units. a) Maximum Projection of a representative field of view. Axonal boutons found to be correlated as a cluster are highlighted. **b)** Correlation matrix for each axonal bouton in the field of view, with every other bouton. This correlation matrix then undergoes hierarchical clustering based on their similarity. **c)** An empirical threshold is determined for each session/field-of-view such that the resulting axon clusters have a threshold minimum for intra-cluster correlation, and maximal inter-cluster correlation. **d)** Raw traces (before denoising) for each axonal bouton identified in a cluster (highlighted in a). A weighted (by their signal-to-noise ratio) average of these boutons determined to belong in a cluster is taken to represent a single axonal cluster.

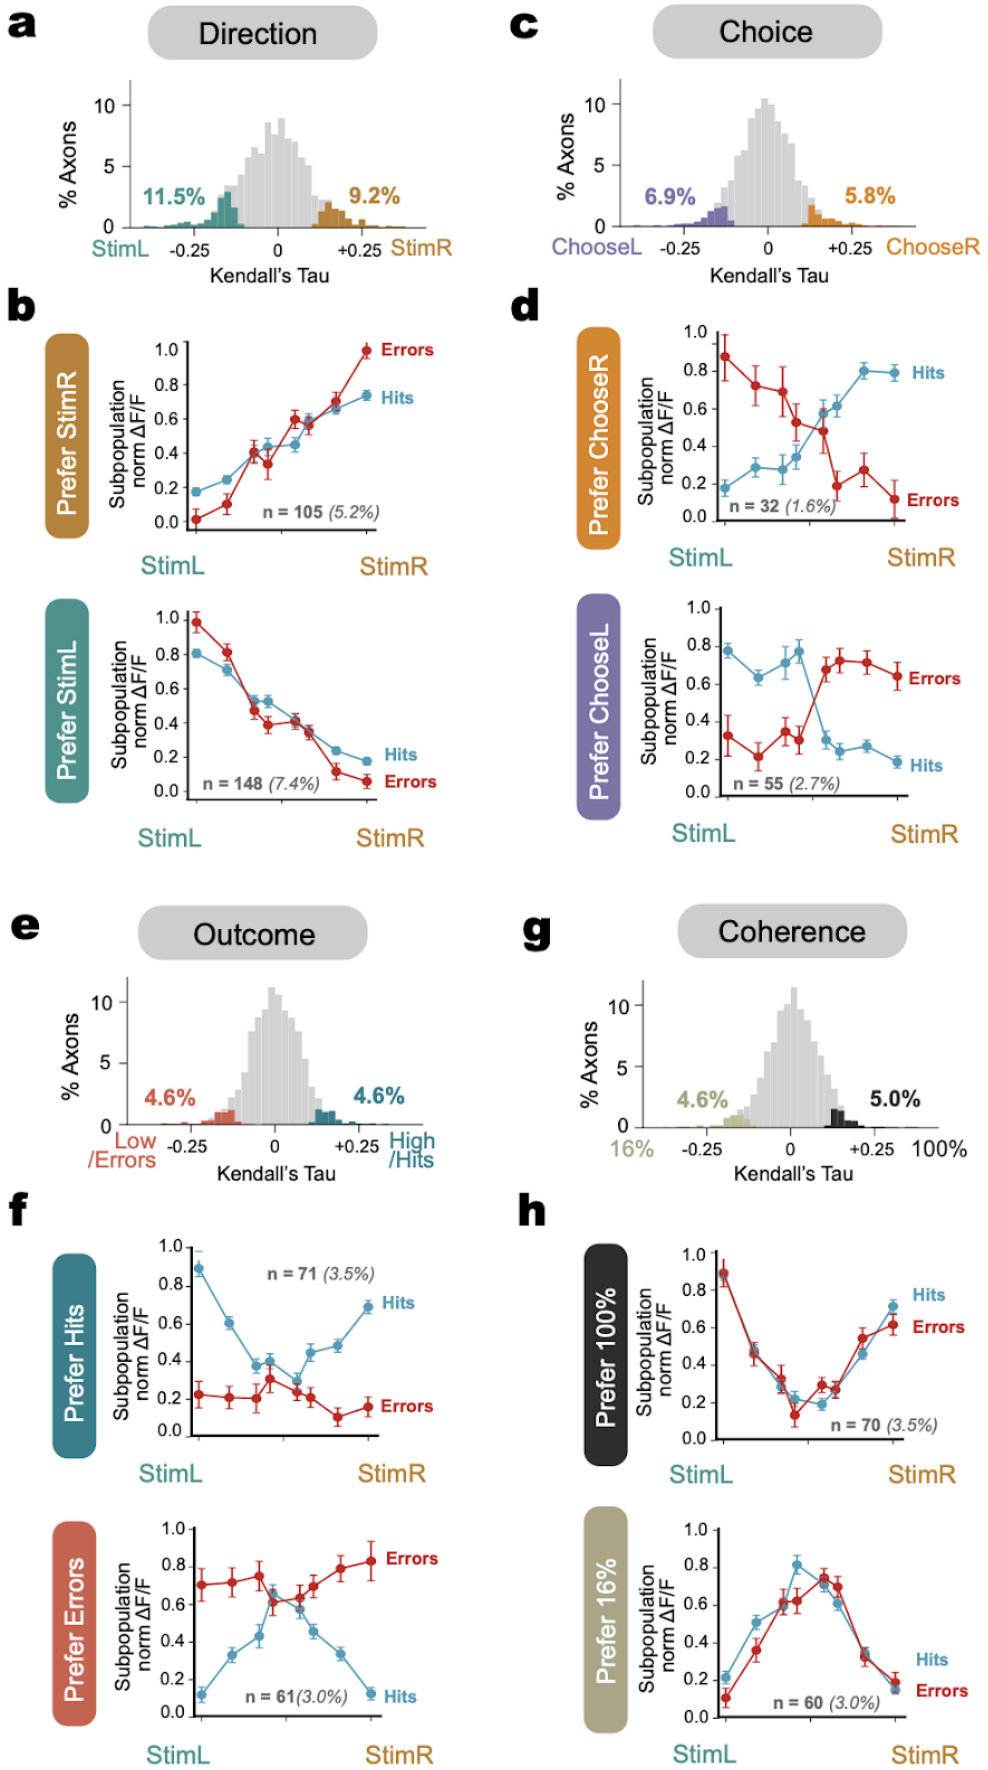
a
Direction
% Axons
11.5%
9.2%
StimL
-0.25
0
+0.25
StimR
Kendall's Tau
b
Prefer StimR
Subpopulation norm ΔF/F
Errors
Hits
n = 105 (5.2%)
StimL
StimR
Prefer StimL
Hits
Errors
n = 148 (7.4%)
c
Choice
6.9%
5.8%
ChooseL
ChooseR
d
Prefer ChooseR
n = 32 (1.6%)
Prefer ChooseL
n = 55 (2.7%)
e
Outcome
4.6%
4.6%
Low /Errors
High /Hits
g
Coherence
4.6%
5.0%
16%
100%
f
Prefer Hits
n = 71 (3.5%)
Prefer Errors
n = 61(3.0%)
h
Prefer 100%
n = 70 (3.5%)
Prefer 16%
n = 60 (3.0%)

1232

Supplementary Figure 8: Subpopulation of LP-ACC axons with strong selectivity and monotonic scaling with coherence relative to task variables. a) Kendall's Tau correlation for monotonic scaling with direction x coherence. **b)** Direction-selective axons that preferred stimulus associated with right (top) and left (bottom) choice as a function of directed coherence. Errors and hits had similar responses as these axons showed greater fidelity to the true stimulus. **c)** Kendall's Tau correlation for monotonic scal- ing with spout side x coherence. **d)** Choice-selective axons that preferred right (top) and left (bottom) side evidence as a function of directed coherence. Errors and hits had a crossed profile when plotted at a function of directed coherence, as axons reflected the absolute strength of evidence when selecting a side, whether or not evidence supported the choice. **e)** Kendall's Tau correlation for monotonic scaling with spout side x coherence. **f)** Outcome-predictive axons that preferred hits (or high evidence for choice) (top), and errors (or low evidence for choice) (bottom) as a function of directed coherence. Errors and hits had mirrored v-shaped profiles when plotted at a function of directed coherence. The same stimulus was associated with different responses based on the outcome but in a way that is not dependent on the choice side, as the responses were vertically symmetric regardless of the stimulus or choice. **g)** Kendall's Tau correlation for monotonic scaling pure coherence. **h)** Outcome-predictive axons that preferred high coherence (100%, top), and low coherence (16%, bottom) as a function of directed coherence. In spite of showing a similar vertically symmetric v-shaped profile as f, these axons did not discriminate between outcomes.

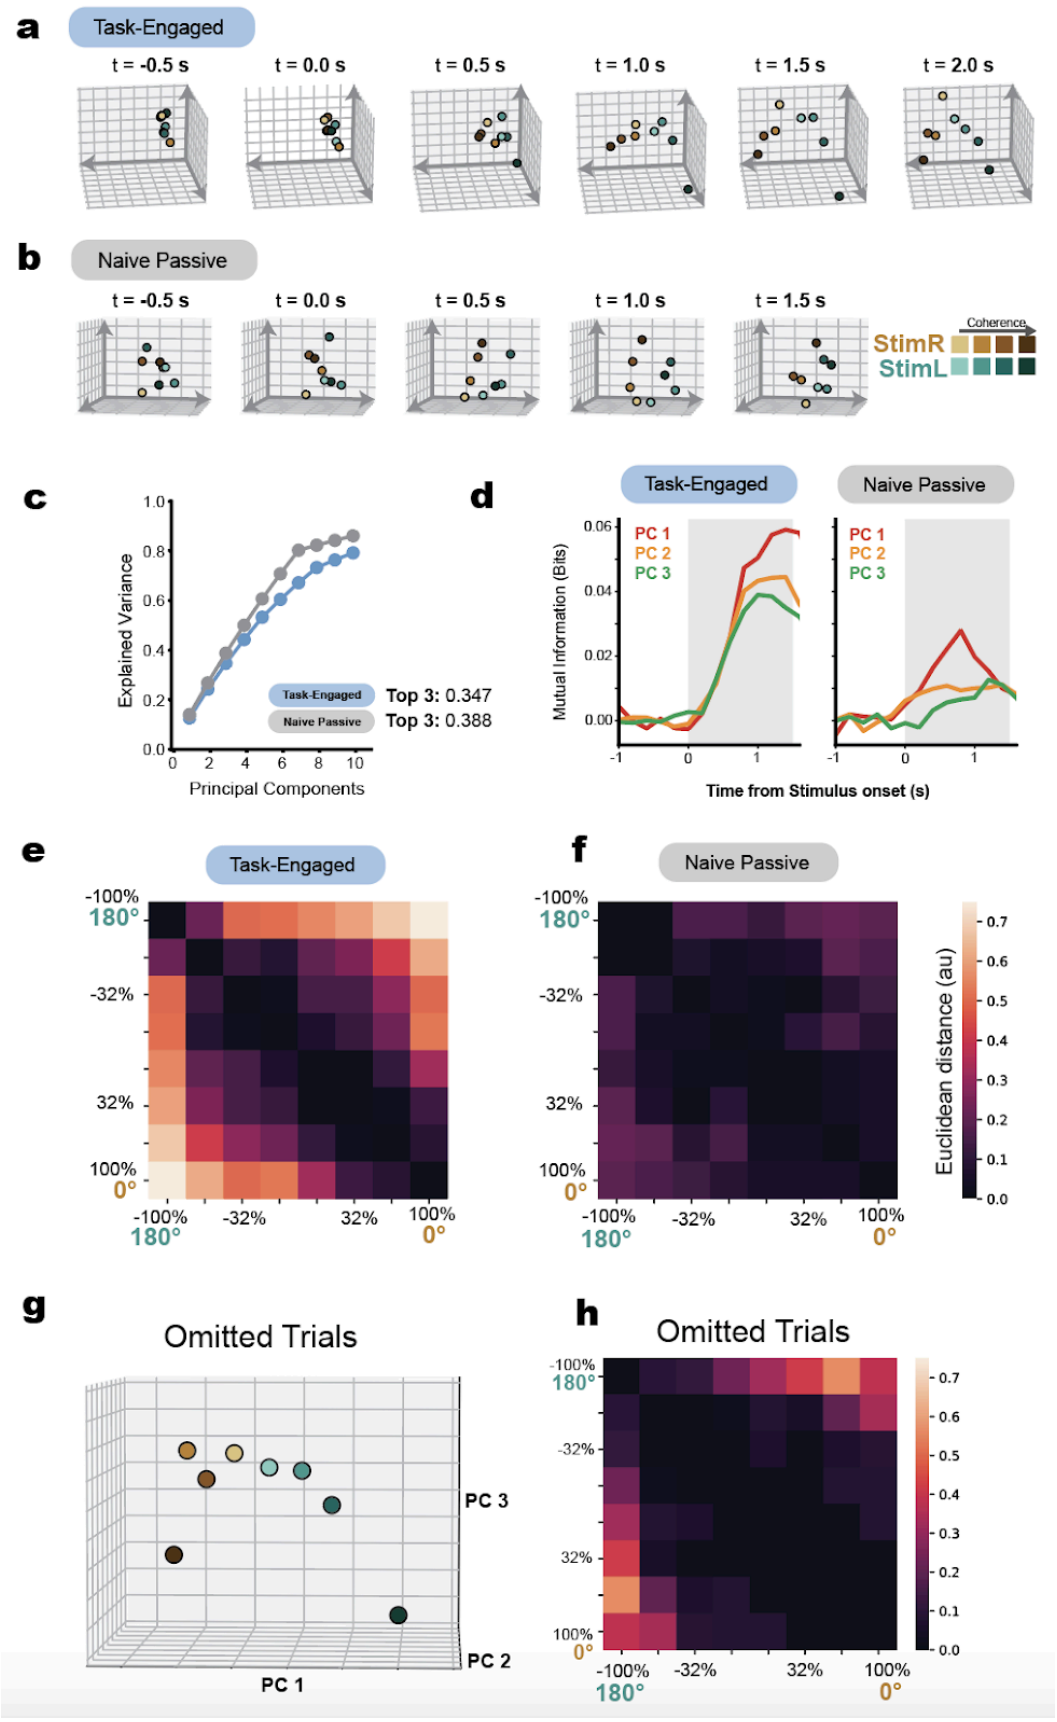
a
Task-Engaged
t = -0.5 s
t = 0.0 s
t = 0.5 s
t = 1.0 s
t = 1.5 s
t = 2.0 s
b
Naive Passive
t = -0.5 s
t = 0.0 s
t = 0.5 s
t = 1.0 s
t = 1.5 s
Coherence
StimR
StimL
c
Explained Variance
Principal Components
Task-Engaged
Naive Passive
Top 3: 0.347
Top 3: 0.388
d
Task-Engaged
Naive Passive
Mutual Information (Bits)
PC 1
PC 2
PC 3
Time from Stimulus onset (s)
e
Task-Engaged
-100%
180°
-32%
32%
100%
0°
f
Naive Passive
Euclidean distance (au)
g
Omitted Trials
PC 1
PC 2
PC 3
h
Omitted Trials

Supplementary Figure 9: Task engagement reorganizes population representation of visual stimuli in a low-dimensional curived manifold. Projecting activity onto the first 3 principal components (PCs). **a)** During task engagement, population activity evolves across the stimulus presentation to form a low dimensional representation of each unique stimulus along a curved manifold. Along this curved manifold, every stimulus is rep- resented with reference to the decision variable. **b)** Under passive conditions, the population activity shows brief direction discriminability, with much less structure. Both low dimensional representations in a and b were derived from PCA on a trial- averaged activity across a pseudo-population, pooling all sessions. We focused on the window 200ms - 1200ms upon stimulus onset for fitting in PCA. **c)** In the task context, the curved manifold allows joint representation of stimulus direction (the decision variable) and the stimulus difficulty on an orthogonal axis. **d)** The mean mu- tual information between directed coherence (-100% 180° to 100% 0°) and first 3 principal components is greater during task engagement. **e)** Pair-wise Euclidean dis- tances (in 3D space formed by first 3 PCs) between each stimulus during task en- gagement. **f)** Pair-wise Euclidean distances between each stimulus in naive mice passively viewing stimuli. **g)** Example of low-dimensional representation along first 3 PCs for attempted trials at 1s from stimulus onset. **h)** Low-dimensional representa- tion along first 3 PCs for omitted trials retain curved manifold structure but differ in curvature and rotation in state space. **i)** Pair-wise Euclidean distances between each stimulus for omitted trials.

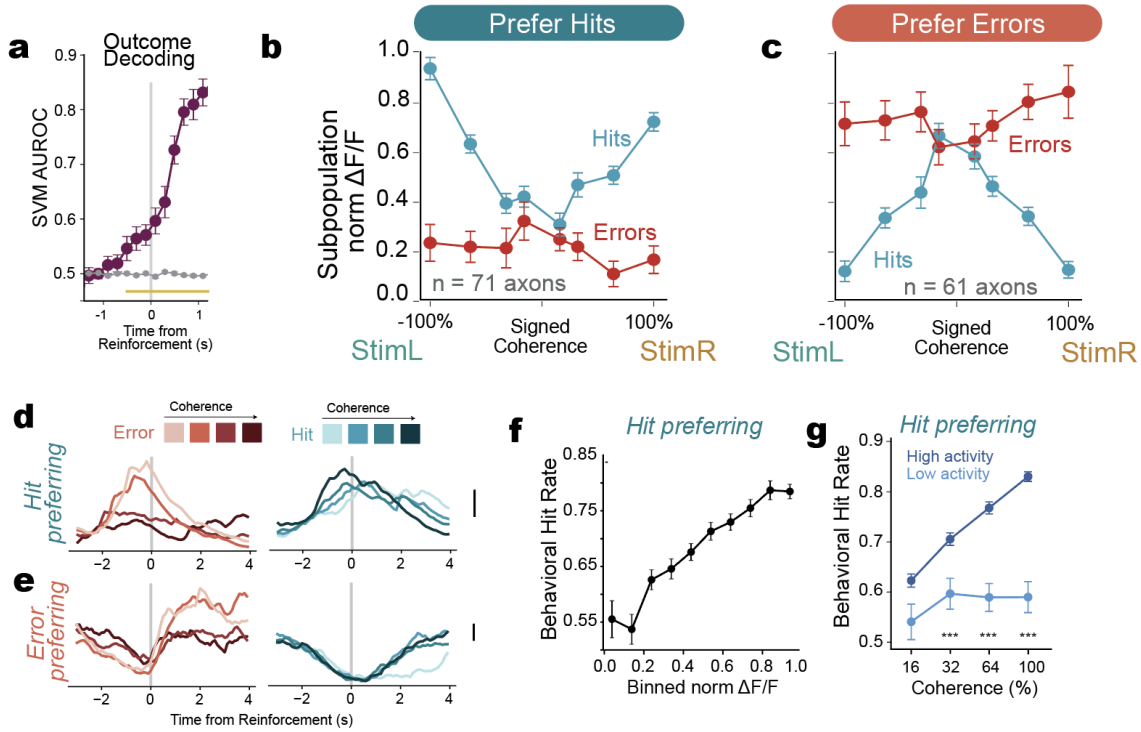

Supplementary 10: LP-ACC axons anticipate outcomes, scaled by uncertainty. a) SVM decoding of the trial outcome (purple) was significantly greater than for shuffled data (grey), starting 600 ms prior to time of reinforcement (paired t-test for each 200ms bin, $p < 0.05$). Error bars reflect SEM across 17 imaging sessions. **b)** Hit- (n = 71 axons) and **c)** Error-preferring (n = 61 axons) axonal subpopulations have mean activity during the stimulus period (0.5s - 1.5s) that reflected anticipated outcome in a manner scaled by coherence, regardless of the chosen side or presented stimulus direction. Error bars reflect SEM across axons within the subpopulation. **d)** Mean dF/F traces aligned to the time of reinforcement for example of hit-preferring and **e)** error-preferring single axons. Scale bar represents 10% of maximum dF/F. **f)** Binned activity of hit preferring axonal population was predictive of the mean hit rate across the trials. **g)** Activity of the hit-preferring axons also predicted trial hit rates even for the same strength of evidence. Asterisks indicate significant difference (paired t-test, $p < 0.001$) in associated hit rate between high (> 70% percentile dF/F) and low (< 30% percentile dF/F) activity trials within each axon. For f and g: Error bars reflect SEM across 71 axons.

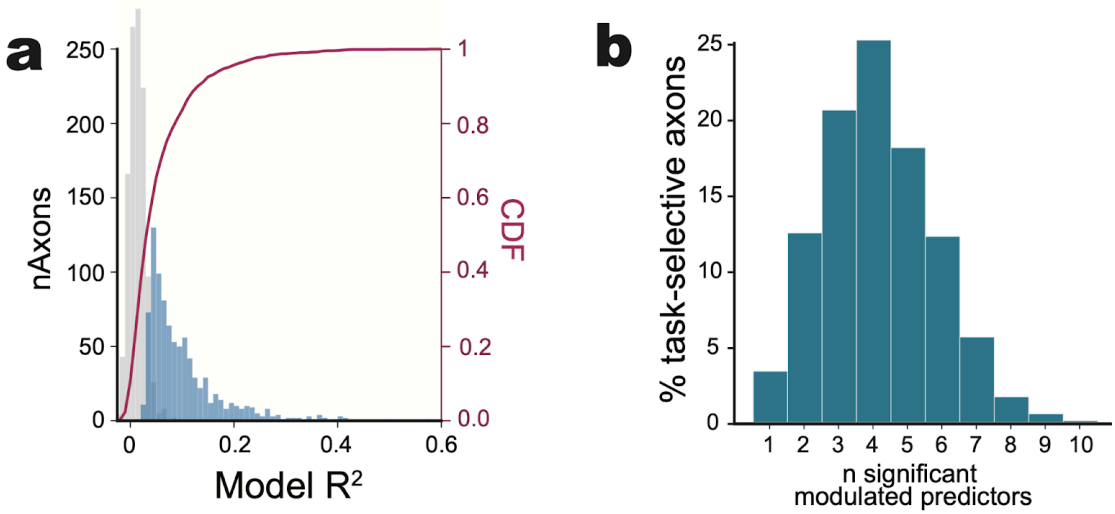

Supplementary 11: Encoding models of LP-ACC axonal activity reveal high mixed selectivity. a) Distribution of model variance explained for individual axon models (R2). Histogram shows count of axons. In grey bars are axons that were not deemed task-responsive due to poor model fits, while blue bars show task responsive proportion of the population. Magenta line shows cumulative distribution. **b)** Distribution of task-responsive axons with number of predictors encoded.