Anterior cingulate cortex mixes retrospective cognitive signals and ongoing movement signatures during decision-making

Lukas T. Oesch[1], Makenna C. Thomas[1], Davis Sandberg[1], João Couto[1] and Anne K. Churchland[1,*]

[1]Department of Neurobiology, David Geffen School of Medicine, University of California, Los Angeles, Los Angeles
*Correspondence: AChurchland@mednet.ucla.edu

Abstract

The anterior cingulate cortex (ACC) is believed to play a key role in reward-based learning in uncertain and dynamic environments. Recent work suggests that in these environments, ACC neurons monitor decisions and their outcomes to inform behavioral strategy. Yet, it remains unclear whether ACC neurons similarly encode behavioral history in fully deterministic perceptual decision-making tasks and, if so, how behavioral history signals interact with other signals that are represented by ACC neurons. Here, we recorded the activity of ACC neurons in freely moving mice performing a visual evidence accumulation task. Many ACC neurons had mixed selectivity: they were strongly driven by non-linear combinations of previous choices and outcomes (trial history). We observed that trial history could be decoded well from the ACC population and that the neural representations of trial history remained stable over multiple seconds. Using linear encoding models, we demonstrate that both trial history and animals' movements strongly drive neural activity in the ACC. We found no relationship between trial history encoding and behavioral biases and instead observed that the neural dynamics encoding trial history were low-dimensional, similar between sessions from the same subjects and conserved across different subjects. These findings suggest the ACC might implement a trial history monitoring process that is independent of behavioral biases and common to different subjects.

Introduction

Despite extensive interest in and a prolific literature on the anterior cingulate cortex (ACC) its functions and the information encoded by its neurons remain a topic of debate (Ebitz and Hayden, 2016; Kolling *et al.*, 2016; Shenhav, Cohen and Botvinick, 2016). One of the major proposed functions of the ACC is to predict and monitor the outcomes of ongoing behavior and to inform changes in behavioral strategy (Monosov *et al.*, 2020). Evidence supporting this idea comes from several lines of research in humans, non-human primates and rodents performing probabilistic tasks in dynamic environments. For example, ACC neurons encode predictions about expected rewards resulting from a specific behavior (Alexander and Brown, 2011; Soltani and Izquierdo, 2019). Violations of these predictions (that are sometimes called surprise signals or unsigned prediction errors) strongly drive neural activity in the ACC (Hayden *et al.*, 2011; Hyman, Holroyd and Seamans, 2017). One hypothesis is that these prediction error signals are then used to drive belief updating and the exploration of alternative behaviors (Seo and Lee, 2007; Quilodran, Rothé and Procyk,

2008; Karlsson, Tervo and Karpova, 2012; Bissonette, Powell and Roesch, 2013; Tervo *et al.*, 2014; Sarafyazd and Jazayeri, 2019; Chen *et al.*, 2024; Cole *et al.*, 2024). In support of this hypothesis, inhibition of ACC neurons disrupts rule-shifting without affecting ongoing behavior (Akam *et al.*, 2021; Tervo *et al.*, 2021; Chen *et al.*, 2024; Cole *et al.*, 2024). This view of ACC function would predict that the post-outcome activity of ACC neurons should generally be low and would only be elevated after an unexpected outcome(Quilodran, Rothé and Procyk, 2008). Surprisingly, ACC neurons not only strongly encode previous choices and outcomes after block transitions when the subjects' behavior needs to be updated but also within blocks when subjects exploit their task knowledge (Sul *et al.*, 2010; Li *et al.*, 2019; Spellman *et al.*, 2021). These findings suggest that the ACC might not only drive belief updating but might also continuously monitor trial history. However, many studies investigating the functions of ACC have employed tasks where decision-rules switch in a block-wise manner. This can lead to strong correlations between previous choices and outcomes, and the previous and upcoming reporting actions reporting movements (Sul *et al.*, 2010). It is thus unclear if the recorded signals in ACC neurons reflect trial history monitoring or postural or movement variables (Muysers *et al.*, 2024). Additionally, in environments with dynamic and uncued block changes, subjects must continuously track their trial history. However, it is unknown whether trial history representations in the ACC might also arise spontaneously in static environments with stable decision-rules.

To overcome these limitations, we trained mice in a freely moving perceptual evidence accumulation task with a fixed rule and a randomized trial structure (Raposo *et al.*, 2012; Odoemene *et al.*, 2018) while recording neural activity from ACC neurons and carefully tracking animals' posture and movements. In keeping with the hypothesis that ACC monitors trial history we found that ACC neurons strongly encoded choice- and outcome history and their different interactions. Using decoding analyses we show that these representations are stable over several seconds. Taking advantage of recently developed methods to separate the influence of often highly collinear task variables and movements on neural activity (Musall *et al.*, 2019) we show that trial history signals cannot simply be explained by animals' body position or movements. Finally, we demonstrate that trial history but not movement encoding is low-dimensional and relatively similar between different subjects.

Results

Perceptual decisions in expert performers are modulated by trial history.

We trained mice on a self-initiated freely moving visual decision-making task that required them to integrate pulsatile sensory evidence over time (Raposo *et al.*, 2012; Odoemene *et al.*, 2018). After initiating a trial by poking into the center port, mice were presented with a train of visual flashes (Figure 1A, B). On trials where the number of flashes exceeded the category boundary of 12 events per second, mice were rewarded for poking into the right port and conversely, they were rewarded for choosing the left port after a low-rate stimulus trial. Importantly, whether a trial was high- or low-rate was randomly assigned. This ensured that there was no trial structure that animals could exploit to guide their decisions.

Animals learned this task (stimulus discrimination and center fixation) over the course of 10.72 ± 1.25 weeks and achieved high performance levels (Figure 1C).

A series of recent studies have suggested that perceptual decisions might be influenced by previous choices and/or outcomes even in expert performers (Braun, Urai and Donner, 2018; Odoemene *et al.*, 2018; Jiang *et al.*, 2019; Lak *et al.*, 2020; Mochol, Kiani and Moreno-Bote, 2021; Gupta *et al.*, 2024) and that previous correct and incorrect choices might differentially impact subsequent behavior (Hermoso-Mendizabal *et al.*, 2020). To verify that mice in our study were similarly influenced by trial history, we separated trials based on whether the previous choice was to the left or right and on whether the mice were rewarded for their previous choice. We then fitted a psychometric function to trials with these different trial history contexts and found that the curves were indeed different. Figure 1D shows the psychometric curves for an example mouse whose choices became right biased after correct right choice but are largely unaffected by previous incorrect right choices and previous left choices regardless of their outcome (Figure 1D, Supplementary Figure 1A - D). To compare how strongly animals' decisions were influenced by sensory evidence or specific trial history contexts on a session-by-session basis we modeled the

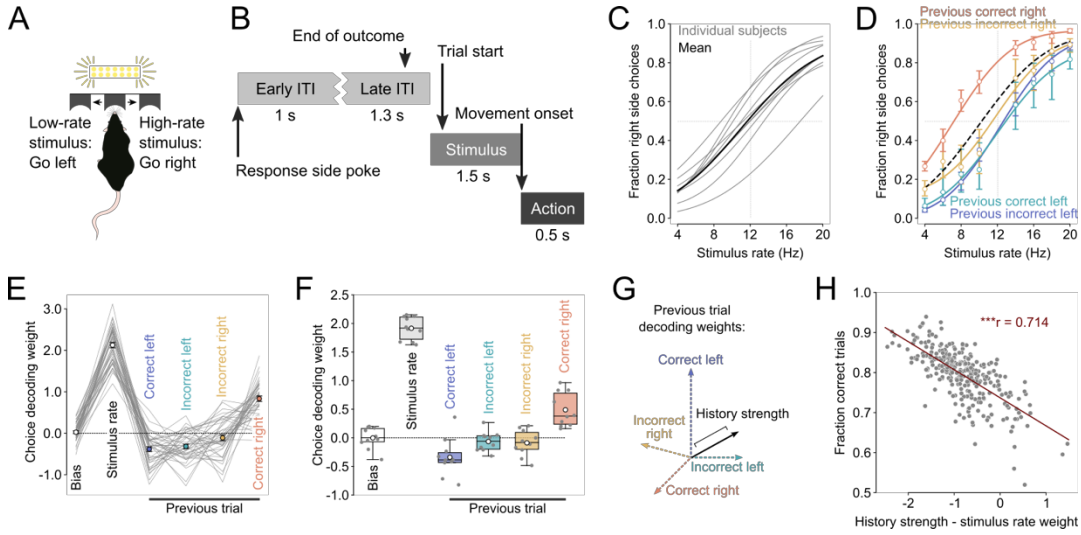

Figure 1. Perceptual decisions in expert performers are influenced by trial history. *(A) Illustration of perceptual task setup. Animals initiate trials by poking into a central port. This triggers a 1 second train of visual flashes. Animals are required to wait inside the center port for one second before reporting their decisions in either the left or the right port. Animals are rewarded for choosing the left port after low-rate stimuli and for choosing the right port after high-rate stimuli. Incorrect choices result in a 16 kHz tone and 2 s timeout. (B) Trials are defined by four distinct phases. Early ITI: the second after the presentation of the previous trial's outcome, Late ITI: 1 s before until 0.3 s after the end of the outcome, Stimulus: 0.5 s before to 1 s after stimulus onset, Action: 0.2 s before to 0.3 s after reporting movement onset. (C) Psychometric functions for individual subjects (grey, n = 10 animals, 609 sessions) and averaged across subjects (black). (D) Psychometric functions for trials that were preceded by specific combinations of previous choice and outcome (trial history) for an example subject. Dashed black line represents the overall fit. (E) Weights of a linear decoder trained to predict choice based on the stimulus rate and trial history from the subject shown in (D). Lines represent individual sessions, dots show mean ± sem weight across sessions. (F) Boxplot showing the median ± interquartile range of the behavioral decoder weights. White dots represent the mean and dark scatter points are individual subjects. (G) Behavioral influence of trial history quantified as the length of the vector in the space spanned by the four distinct combinations of previous choice and outcome. (H) The difference between history strength and stimulus rate weight is negatively correlated with task performance. ***P < 0.001.*

animals' choices as a function of the current stimulus rate plus an individual regressor for each trial history context. As expected, we found that for our expert mice the model weights for stimulus rate were significantly greater than a shuffled control and also greater than the weights for trial history, indicating the decisions were strongly influenced by the sensory stimulus that governed the rewarded side (Figure 1E, F). In addition, all the trial history weights were significantly different from the shuffled control, indicating that the animals' decisions were significantly influenced by trial history. Particularly, previous correct left and previous correct right choices strongly and consistently biased the mice's choices (negative signs indicate favoring left choices, Figure 1F). However, despite previous correct left and correct right choices consistently favoring left and right choices, respectively, the session-by-session variance of all the trial history weights was similar (Supplementary Figure 1E) and we found no evidence of clusters of sessions with similar weight patterns (Supplementary Figure 1F). Importantly, when we looked not one- but two-trials back, we found that the choices, outcomes and their interaction had a lower impact on the mice's decisions than the ones from the last trial and they were very close to zero (Supplementary Figure 1G).

We summarized the overall influence of the most recent trial history on the animals' choices (history strength) by taking the length of the vector of all trial history weights (Figure 1G). Given that the trial structure of our task was random and that only the stimulus rate determines the rewarded side, basing decisions more on trial history than stimulus rate should lead to decreased task performance. Indeed, we found that the difference between history strength and stimulus weight was negatively correlated with task performance (Figure 1H). These findings confirm that decisions in expert mice are influenced by multiple trial history contexts even when this leads to suboptimal performance (Hwang *et al.*, 2017).

Trial history can be decoded from ACC population activity

We then asked whether ACC neurons might encode trial history information in expert performers. To this end, we recorded neural activity from CaMKII-expressing neurons in the ACC using miniaturized fluorescence microscopes (Cai *et al.*, 2016, Figure 2A and B). The small, lightweight size of this scope meant that we were able to measure neural activity in freely moving mice. Further, this approach allowed us to measure many neurons simultaneously (322 ± 42 neurons per subject), allowing for the use of powerful analyses for measuring population level activity.

The activity of many neurons was modulated by specific trial history contexts (Figure 2C). For example, the activity of some neurons was driven almost exclusively by previous correct right or correct left choices during the inter-trial interval (Figure 2D, N1, N4) or by both incorrect left and right choices but at different times in the trial (Figure 2D, N7), indicating clear non-linear mixed selectivity (Rigotti *et al.*, 2013; Kaufman *et al.*, 2022). To visualize the temporal evolution of trial history encoding in the ACC neural population, we sorted the neurons by the timing of their peak activity during trials where mice had made a correct left choice on the previous trial. The distribution of peak times "tiled" the entire

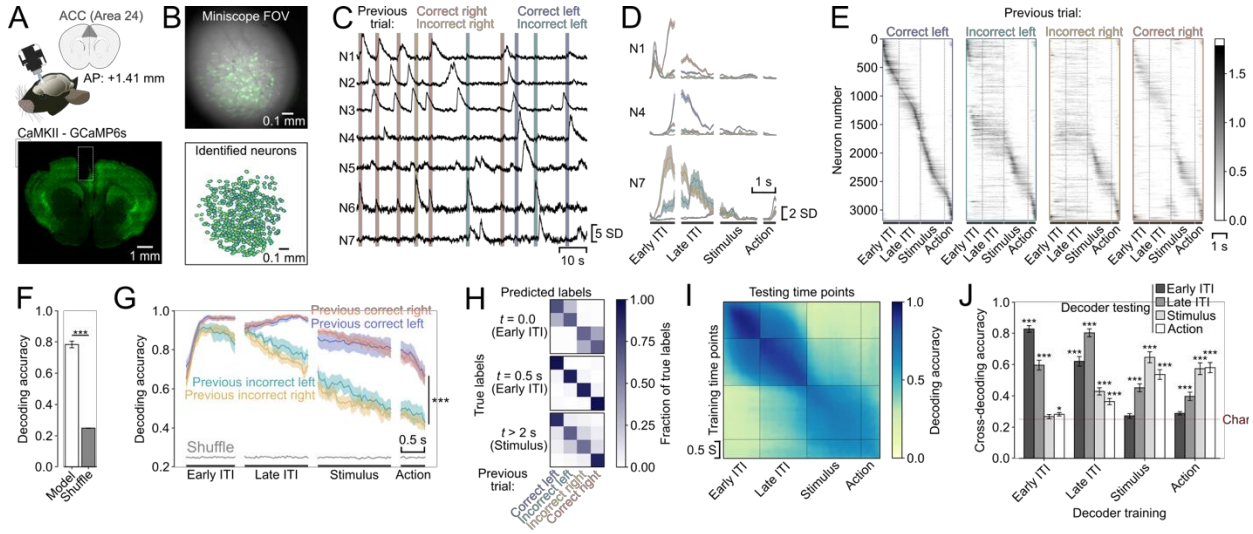

Figure 2. ACC neuron populations uniquely represent multiple trial history contexts. *(A) Transgenic mice expressing GCaMP6s under the control of the CaMKII promoter were implanted with a gradient refractive index lens above their ACC; neural activity was recorded using a miniaturized fluorescence microscope (top). Histological verification of GCaMP6 expression and lens position (bottom). (B) (Top) Field of view from an example recording. A single frame is overlaid with the variance projection of the entire recording (green). We identified individual neurons using an automatic segmentation algorithm* (Zhou *et al.*, 2018) *(bottom). (C) Example traces from seven simultaneously recorded neurons. Horizontal lines represent the first second of the ITI after respective choices and outcomes. (D) Peri-event time histograms of three neurons from (C) showing distinct responses for different trial history contexts. Lines and shading represent mean ± sem. (E) Neural responses averaged over the different trial history contexts and sorted based on their peak activity after correct left choices. Vertical lines indicate splits between the trial phases (see Figure 1B). Colormap is bounded by 0 and the 99th activity percentile. Neurons were taken from one example session for every imaged subject. (F) Mean ± sem accuracy of a linear decoder trained to predict trial history based on ACC population activity. Bars represent the average across the entire trial time (n = 9 subjects). ***P < 0.001, linear mixed-effects model with subject and session as random effects. (G) Mean ± sem decoder accuracy over trial time for different trial history contexts. ***P < 0.001, linear mixed-effects model with subject and session as random effects. (H) Decoder confusion matrices for three different time points with reference to the last choice. t > 2 s is the start of the stimulus presentation. (I) Decoding stability: trial history decoders were trained for every time point and tested on every other time point. The average over all subjects is shown. (J) Mean ± sem cross-decoding accuracy divided by trial phase. *P < 0.05, ***P < 0.001, linear mixed-effects model with subject and session as random effects and post-hoc comparison between model and shuffled accuracies.*

duration of the trial (Figure 2E, left). When we plotted the average neural activity for the other

trial history contexts using the same sorting as for previous correct left choices, we found no tiling for previous incorrect trials and rather weak structure for previous correct right trials during the early and late ITI. In contrast, during the stimulus presentation and response action phases of the trials the tiling was preserved between different trial history contexts (Figure 2E, right). This suggests that trial history information modulates neural activity heterogeneously and may dominate activity during the early and late ITI, while trial history independent factors may drive the activity of ACC neurons more strongly during stimulus presentation and reporting action.

To quantitatively assess the strength of neural representations of trial history we trained linear decoders on ACC population activity to predict the different trial history contexts (4 classes). The decoding accuracy over the entire trial time was high (78 %, chance level is 25%) and significantly higher than the shuffled control (Figure 2F). When comparing the

decoding accuracy of trial history decoders to binary decoders trained to decode previous choice and previous outcome only (with a chance level at 50%) we found no difference between them. This suggests that no information about previous choice or outcome is lost when optimizing decoders to predict their combinations. (Supplementary Figure 2). Next, we asked whether all trial history contexts could be decoded equally well at all moments in the trial. To answer this question, we separately computed the decoding accuracy for each trial history context across the entire trial duration (Figure 2G). We found that the temporal dynamics depended on the outcome of the previous trial: the decoding accuracy of previously rewarded left and right trials was similar and remained high throughout the ITI and slowly decayed thereafter, but the decoding accuracy for previous incorrect left and right choices started decaying during the late ITI and was significantly lower than the accuracy for either of the previous correct choices. We then asked whether the neural representation of certain trial history contexts was more similar than others and whether the similarity could change across trial time. The more similar a pair of trial history contexts is the higher the probability should be that the decoders misclassify them. We therefore examined the errors made by the decoders made by constructing confusion matrices over trial time. We found that the types of misclassifications changed with time. At the very beginning of the ITI, right after the mice had finished reporting their decision, the decoders frequently confounded correct with incorrect choices for either of the two choices (Figure 2H, *top*). 500 ms later into the early ITI, all the different trial history contexts were highly separable (Figure 2H, *middle*) and finally at the beginning of the stimulus presentation decoders were less accurate at distinguishing between previous left and right choice for incorrect trials, whereas previous correct left and right trials were still highly separable (Figure 2H, *bottom*). These findings suggest that neural representations evolve from initially tracking previous choices and then later to tracking combinations of previous choices and outcomes during the ITI and the stimulus and action periods.

We next asked whether similar neurons encode information about trial history over extended periods of time (Bonaiuto, De Berker and Bestmann, 2016; Akrami *et al.*, 2018) or whether these representations evolve along stable dynamics as has been found for choice information (Najafi *et al.*, 2020). To test this, we trained decoders on all time points and tested their decoding accuracy on all the other timepoints. We found that decoders were very stable within the early and late phases of the ITI and within the stimulus and action phase, with a marked split between ITI and new trial timepoints (Figure 2I). When we quantified the decoder stability we found that trial history decoders trained on the early ITI also performed very well on the late ITI time points but were at chance level or only slightly above for the later trial phases. Decoders trained on the late ITI predicted trial history better than chance across the entire trial. Finally, decoders trained on the stimulus and action phase performed above chance on all other phases but the early ITI. Taken together, these findings point to a stable trial history representation during the ITI that undergoes a major reorganization when the mice transition from the ITI to initiating a new trial.

Finally, we assessed whether trial history could be decoded better in sessions where subjects more strongly relied on trial history to guide their choices. We found no correlation between behavioral trial history strength and trial history decoding accuracy

(Supplementary Figure 4A). There was a weak positive relationship between behavioral weights for trial history and decoding accuracy only for previous incorrect right trials (Supplementary Figure 4B). Our findings thus suggest that trial history information is similarly encoded in the ACC population activity regardless of the animals' behavioral trial history biases.

Neural representations of trial history cannot solely be explained by body posture or movements.

Recent work has demonstrated that much of the neural activity across the dorsal cortex of mice can be explained by spontaneous movements (Musall *et al.*, 2019, 2023; Stringer *et al.*, 2019; Salkoff *et al.*, 2020). Some movements are highly stereotyped and aligned to task events, task variables or animal's past actions (Yin *et al.*, 2024). We therefore examined whether neural signatures of trial history were simply a reflection of the animals' movements or whether they represented a uniquely cognitive signal. To address this question we fitted linear encoding models which relate the activity of individual neurons to task variables, and instructed- and uninstructed movement regressors (Musall *et al.*, 2019). The task variable regressors included choice, outcome, trial history and visual stimulus events; instructed movements included the pokes into the three different ports and uninstructed movement regressors included lower dimensional representations of the full behavioral video (Video SVD) and the motion energy of the video (Motion energy SVD), and neurons' tuning to head-orientation angle and body position (Figure 3A, *left and center*, Supplementary Figure 3A). To avoid overfitting the models, we regularized the regressor weights using ridge regression and ran ten-fold cross validation. On held-out trials, the models could predict 32.8 ± 1.9% of the neural variance across the neurons in all included sessions (Figure 3A, right, compare red and gray traces for each neuron).

Having built a model that successfully predicted much of the single-trial neural activity, we set out to determine which variables contributed to this prediction. We first asked how much neural variance could be explained by each model regressor alone (maximal explained variance). To address this question, we fitted encoding models with all regressors shuffled except the regressor of interest plus a time-varying intercept. We then also fitted an intercept only model and subtracted the explained variance of the intercept-only model from the single-regressor model explained variance. This ensured that the explained variance was driven by the specific variable in question rather than some temporal fluctuation that happened during all trials. We found that uninstructed movements, particularly the video and video motion-energy regressors explained relatively high amounts of variance on their own (Figure 3B). This observation replicates previous reports that movements explain considerable neural activity (Musall *et al.*, 2019; Stringer *et al.*, 2019; Salkoff *et al.*, 2020) and demonstrates that movements also have a major impact on neural activity in freely moving animals (Tremblay *et al.*, 2023).

Because task variables can be highly correlated with each other (or with movements) the maximal explained variance is vulnerable to over-estimating the true explained variance (Supplementary Figure 3B). To account for this, we computed the unique explained

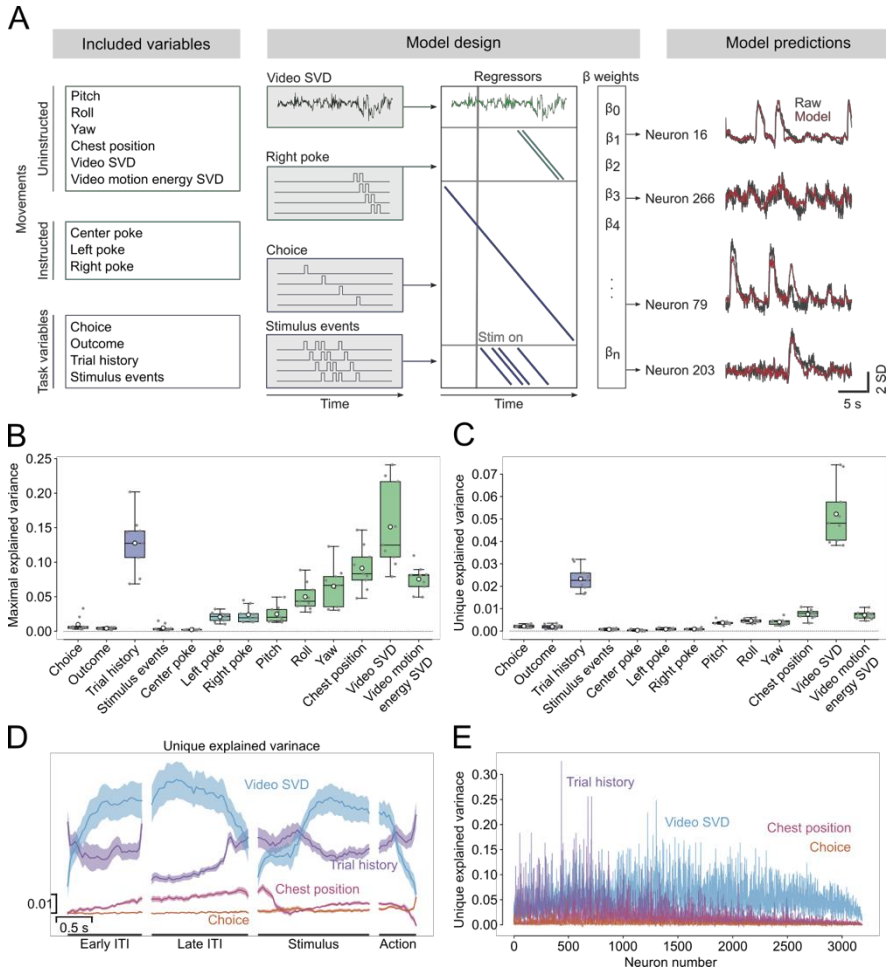

Figure 3. Trial history explains considerable unique variance on the single-cell level. *(A) Linear encoding model design. Neural activity was modeled as a function of task variables, instructed- and uninstructed movements (left and center). Right panel shows measured (grey) and reconstructed (red) activity of four example neurons. (B and C) Maximal amount of neural variance (B) or variance uniquely (C) explained by different regressors. Boxes represent median of subject means ± interquartile range, white circle depicts mean and grey scatters individual subjects. (D) Mean ± sem unique explained variance over trial time for different regressors. (E) Unique explained variance of neurons from one session per subject for different regressors. Neurons are sorted by full model explained variance.*

variance of the different regressors: we fitted the full model with only the regressor of interest shuffled and then subtracted its explained variance from the full model variance. The

outcome was unexpected: we found that trial history accounted for considerable unique neural variance (Figure 3C). This observation is in stark contrast to the impact of other task variables in other areas: typically task variables account for minimal neural variance and are dwarfed by movement signals (Musall *et al.*, 2019).

When we reconstructed the maximal explained variance for trial history across the duration of the trial, we found that it remained high and stable across the early and late ITI and then gradually decreased, similar to the gradual decrease in trial history decoding accuracy from ACC population activity. We then examined the time course of the unique variance for trial history and observed that it did not decrease over time and instead fluctuated and was anti-correlated to the unique variance of the video SVD (Figure 3D). This might suggest that these two regressors are partially correlated and compete for variance (Supplementary Figure 3B). The mice's upcoming choice, on the other hand, explained minimal unique variance at the end of the action phase while the animals' body position did so most strongly during the late ITI (Figure 3D).

We assessed whether the maximal or the unique explained variance for trial history of the ACC neuron populations were related to the behavioral trial history biases. Similar to the decoding results, we did not observe any correlation between behavioral trial history biases and maximal or unique neural variance explained by trial history (Supplementary Figure 4C and D), providing reassurance that neural trial history encoding signals are not simply driven by behavioral biases.

Previous work demonstrated that neurons in prefrontal and parietal cortices of primates and rats and mice have mixed selectivity to different task variables (Mante *et al.*, 2013; Rigotti *et al.*, 2013; Kaufman *et al.*, 2022; Diehl and Redish, 2023; Kira *et al.*, 2023; Posani *et al.*, 2024). We thus asked whether multiple regressors might explain unique variance within the same neurons. Indeed, most neurons showed mixed selectivity for multiple regressors (Figure 3E). Interestingly, video SVD accounted for similar amounts of unique variance across all neurons whereas trial history tended to explain large amounts of unique variance in a subset of neurons, particularly those neurons whose activity was captured well by the full encoding model.

Taken together these data demonstrate that ACC neurons, in contrast to neurons in many regions of the dorsal cortex, strongly encode trial history information that cannot simply be explained by movements.

Neural dynamics representing trial history are low-dimensional and similar between subjects.

Having established a method to estimate neural activity driven by trial history rather than a set of other possibly collinear variables, such as chest position or head-orientation, we then sought to understand whether there is a limited set of neural dynamics representing trial history shared across many neurons in the ACC. To address this question we first applied principal component analysis (PCA) to the weights of trial history, the video SVD components or the neuron's tuning to the chest position obtained from the full encoding models. We observed that within the top two dimensions, the four trial history contexts were well-separated (Figure 4A) while the first and third dimensions of chest point encoding showed well-defined place tuning (Figure 4B). We then compared the number of dimensions required to capture at least 90% of the variance of the weights for trial history, chest position tuning and video SVD (Figure 4C). We found that relatively few PCA dimensions were necessary to explain the trial history encoding (12.0 ± 0.6) and chest position (17.2 ± 1.7). By contrast, many more dimensions were needed to explain the same amount of variance in the video SVD (63.0 ± 6.4, Figure 4D). These findings indicate that trial history information is encoded in a small set of temporal dynamics and that there are far fewer chest tuning motifs in the ACC population than video SVD motifs.

Next, we hypothesized that if the neural dynamics for trial history represented a reliable monitoring process they should be similar across different sessions from the same subject and even between subjects. To test this hypothesis we used procrustes analysis (Horrocks, Rodrigues and Saleem, 2024) analysis. Procrustes analysis compares two geometric shapes with each other by finding a set of matrix transformations that minimize the disparity between the two shapes (Figure 4E). When we matched the trial history dynamics

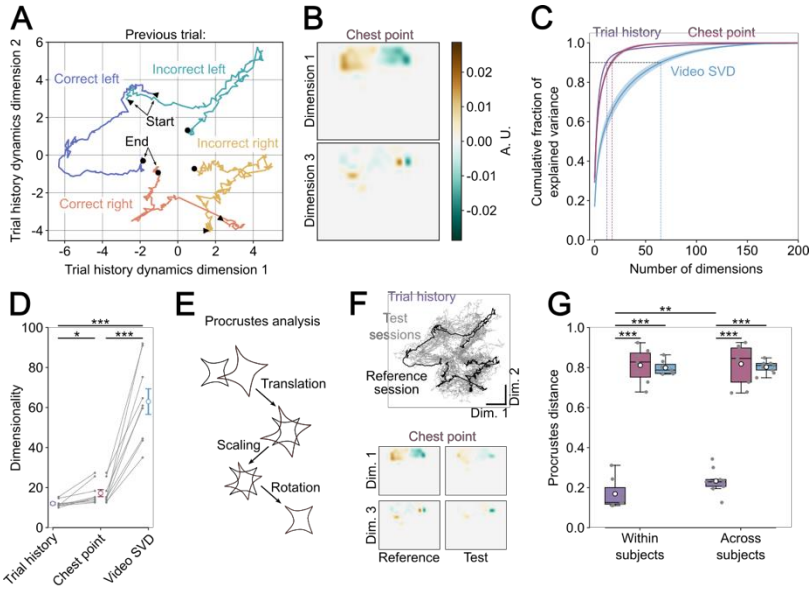

Figure 4. Neural dynamics representing trial history are low-dimensional and similar between subjects. *(A) Trial history encoding weights projected into the two first principal component dimensions and separated by trial history. (B) Two dimensions of the chest point encoding weights. (C) Mean ± sem fraction of cumulative explained encoding weight variance for different variables as a function of PCA dimensions. (D) Minimum number of PCA dimensions required to explain 90% of the encoding weight variance. Grey lines represent individual subject, with circle and error bar depicts mean ± sem over subjects. *P < 0.05, ***P < 0.001, linear mixed-effects model with subject as random effect (E) Schematic representation of matrix transformations for Procrustes analysis. (F) Trial history dynamics of all other sessions (grey) matched to the reference session (black) on A (top). Two dimensions of chest point encoding weights of an example session (right) matched to the reference session in B (left). (G) Dissimilarity (procrustes distance) of different regressors for different sessions within subjects and across subjects. **P < 0.01, ***P < 0.001, linear mixed-effects model with random effect of subject.*

across different sessions we observed very similar shapes (Figure 4F, *top*), whereas not all dimensions of chest point encoding matched well between two example sessions (Figure 4F, *bottom*) We then quantified dissimilarity (procrustes distance) of trial history, chest position and video SVD encoding between sessions from the same subjects and between different subjects (Figure 4G). We found that the within- and across subject distance of trial history encoding was significantly lower than the distance of chest and video SVD tuning. Although there was a significant difference between within- and across-subject distance for trial history the difference was relatively small. We did not find a difference between within- and across-subject encoding distance of chest and video SVD.

These results provide evidence for the hypothesis that trial history dynamics discovered by the linear encoding model might represent a monitoring process shared by different subjects.

Discussion

We set out to test whether ACC neurons also encode behavioral history in fully deterministic perceptual decision-making tasks. To do this, we trained freely moving mice to make abstract decisions about visual stimuli and measured neural activity of excitatory ACC neurons using a miniaturized, head-mounted microscope. This approach allowed us to record neural dynamics during unrestrained and self-initiated perceptual decision-making. We found that trial history significantly influenced mice's decision-making and that subjects weighed trial history differently across sessions. We observed that the activity of

individual ACC neurons was strongly modulated by trial history. Using linear decoders we showed that the neural representations of distinct trial history contexts are well separated across the entire trial time and that these representations remained stable over several seconds (Figure 2I and J). We showed that trial history explained large fractions of neural variance on the individual neuron level and that much of this variance could not be accounted for by other correlated task variables or movements (Figure 3B-E). Finally, found that trial history encoding is low-dimensional and that neural dynamics representing trial history over trial time are relatively similar across different subjects, unlike representations of posture and movement (Figure 4F, G).

Neural signals in the anterior cingulate cortex are proposed to represent a large variety of cognitive signals including surprise, effortful decisions, and contextual-, task rule- or strategy representations (Seo and Lee, 2007; Sul *et al.*, 2010; Hayden *et al.*, 2011; Holroyd and Yeung, 2012; Tervo *et al.*, 2021, 2014; Sarafyazd and Jazayeri, 2019; Stolyarova *et al.*, 2019; Hart *et al.*, 2020; Akam *et al.*, 2021; Reinert *et al.*, 2021). We find that ACC neurons in expert mice performing a fully deterministic perceptual task encode trial history more strongly than other task variables. This is remarkable for two reasons. First, the trial structure of our perceptual task was random and it lacked a block structure. This ensured that the only informative feature for the animals was the sensory evidence and any reliance on trial history hampered optimal performance. Nevertheless, individual ACC neurons and the ACC population as a whole faithfully tracked trial history. Second, a series of recent studies have shown that neural activity in many cortical regions of mice can be better explained by movement signals rather than true internal or cognitive variables (Musall *et al.*, 2019; Stringer *et al.*, 2019; Salkoff *et al.*, 2020). Indeed, internal states seem to be highly intertwined with the stereotypy of movement patterns (Yin *et al.*, 2024) and the two might be hard to separate (Hasnain *et al.*, 2024). In line with these observations, we find prominent representations of postural- and movement signals in many ACC neurons. However, we also observe strong trial history encoding that was not readily explainable by movements (in contrast to encoding of upcoming choice, for example). This encoding might thus represent a truly cognitive signal.

We found that many neurons that encoded trial history responded to specific combinations of previous choice and outcome suggesting that ACC neurons display non-linear mixed-selectivity with respect to trial history similar to neurons in M2 (Siniscalchi, Wang and Kwan, 2019). Non-linear mixed-selectivity support flexibles readout of multiple parameters (Rigotti *et al.*, 2013; Kaufman *et al.*, 2022) and signatures of mixed-selectivity have been found in many frontal (Mante *et al.*, 2013; Diehl and Redish, 2023; Posani *et al.*, 2024) and parietal (Raposo, Kaufman and Churchland, 2014; Morcos and Harvey, 2016; Kira *et al.*, 2023) regions. Whether these neural representations are categorical (Hirokawa *et al.*, 2019; Hocker *et al.*, 2021) or not (Raposo, Kaufman and Churchland, 2014; Posani *et al.*, 2024) it is still a matter of debate.

Here we show that trial history dynamics in ACC neurons are low-dimensional. This is in stark contrast to the high-dimensional neural encoding of movements as captured by the video SVD. Importantly, we also found that, unlike for postural- and movement encoding, trial history encoding was very similar between sessions from the same subjects and

between subjects, pointing to neural dynamics for trial history that are shared across different subjects.

Our findings further identified that trial history representations remained remarkably stable over seconds of trial time within the ITI or the new trial. This finding is consistent with reports suggesting that ACC neurons have long intrinsic- and task-related timescales that may enable them to retain relevant task variable information into and beyond subsequent trials (Murray *et al.*, 2014; Maisson *et al.*, 2021; Song *et al.*, 2024). In line with this idea, choices and outcomes extending one or multiple trials back can be decoded from ACC and prelimbic populations (Sul *et al.*, 2010; Spellman *et al.*, 2021). Whether trial history would also be decodable over multiple trials in our task without any defined trial structure remains to be explored.

Previous work demonstrates that despite readily using sensory evidence to guide choices, expert perceptual decision makers across species are influenced by previous choices and outcomes (Busse *et al.*, 2011; Choe *et al.*, 2017; Akrami *et al.*, 2018; Braun, Urai and Donner, 2018; Odoemene *et al.*, 2018; Hermoso-Mendizabal *et al.*, 2020; Lak *et al.*, 2020; Mendonça *et al.*, 2020; Mochol, Kiani and Moreno-Bote, 2021; Gupta *et al.*, 2024). Our findings that expert subjects tend to repeat previously rewarded choices on average are in line with these studies. However, we observed a large session-to-session variability in how subjects weighted different trial history contexts and we frequently found sessions where subjects equally strongly weighted unrewarded and rewarded trials, challenging the view that trial history biases are mainly driven by positive reinforcement (Lak *et al.*, 2020).

We took advantage of the between- and within-subject variability in trial history biases to assess whether stronger behavioral biases were associated with stronger neural encoding of trial history. We did not observe any relationship between behavioral usage and neural encoding of trial history. This is in agreement with recent studies in non-human primates and mice that found that ACC and prelimbic activity did not predict choice biases (Li *et al.*, 2019) and disruption of ACC activity had no impact on ongoing decision-making (Akam *et al.*, 2021; Spellman *et al.*, 2021; Chen *et al.*, 2024; Cole *et al.*, 2024) in tasks with dynamic block switching while it may decrease task directed effort (Vázquez *et al.*, 2024; González *et al.*, 2025). This raises two important questions: First, it remains unclear where trial history information might ultimately be transformed into behavioral biases. Recent studies have proposed that for example the posterior parietal cortex may be involved in forming trial history biases (Hwang *et al.*, 2017; Akrami *et al.*, 2018). Second, it remains unclear what the function of the trial history signals in ACC may be. Given that the neural signatures of trial history remain stable across a range of different levels of task performance and bias it seems unlikely that these signals directly reflect an adaptive process, such as the trade-off between exploration and exploitation or a tunable learning rate parameter. Rather they might reflect an ongoing monitoring process that continuously unfolds as animals experience their environment. Thanks to their multi-second stability, trial history representations in the ACC retain information about past experience that might serve to form new associations between previous and current experiences in down-stream areas.

There are several limitations to our study. First, we do not know what aspects of the trial history representations might be exclusive to the ACC. Strong trial history signals have

also been observed in other prefrontal regions, such as the prelimbic (Spellman *et al.*, 2021) or orbitofrontal cortex (Hocker *et al.*, 2021), the M2 (Jiang *et al.*, 2019) or the PPC (Morcos and Harvey, 2016; Hwang *et al.*, 2017; Akrami *et al.*, 2018). It remains to be elucidated whether trial history signals in these regions are similarly independent of movement signals or of other tracked task variables and whether they display similar temporal stability and between subject similarity. Second, while our correlational evidence suggests that the ACC might not directly be involved in the performance of our perceptual task we have not directly tested the effects of ACC disruption on task performance. Third, mice in our task actively initiated trials and reported their choices by freely moving within an arena adding strong volitional and spatial components to the task and making the behavior costly. This is in contrast to most head-restrained tasks where trials are automatically triggered and reporting actions commonly consist of simply licking a spout. On the other hand, head-restraint may induce a significant amount of discomfort on the animals. This especially important for the ACC that is also an important brain region controlling the affective component of pain (Acuña *et al.*, 2023). It remains an open question how these differences would impact the neural representations of trial history and movements in ACC neurons.

In summary, our study provides evidence that freely moving mice performing a perceptual task maintain neural representations of trial history in their ACC over time spans of seconds that are not simply reflecting movements, but rather a cognitive monitoring process that is similar between different subjects.

Methods

Animal subjects

We used transgenic mice on a C57/B6 background that expressed the genetically encoded calcium indicator GCaMP6s in calmodulin-kinase II expressing cells under the control of a tetracycline response element (GCaMP6s was constitutively expressed in the absence of doxycycline). We obtained these subjects by crossing TRE-GCaMP6s mice (B6;DBA-Tg(tetO-GCaMP6s)2Niell/J, Jax stock no. 024742) (Wekselblatt *et al.*, 2016) with CaMKII-tTA mice (B6.Cg-Tg(Camk2a-tTA)1Mmay/DboJ, Jax stock no. 007004) (Mayford *et al.*, 1996). For our study we used both male (6 for behavior, 6 for imaging) and female (4 for behavior, 3 for imaging) mice that were 8-25 weeks old when handling and water restriction started. Mice were housed in a room with controlled humidity (30 - 70%) and temperature (71 - 76 °F / 21.67 - 24.44°C). Mice were kept on a reverse dark-light cycle, with lights on at 10 am and lights off at 10 pm. Animals were provided with cotton nestlets and received an acrylic shelter and running wheel as additional enrichment. They were fed regular laboratory chow (NIH-31 Modified Open Formula Mouse/Rat Diet 7013) *ad libitum* and given additional treats (Mini Yogurt Drops, Bio-Serv) at the beginning of the day and after performing the task. Animals were water restricted during behavioral training and testing (see below). All behavioral and surgical procedures adhered to the guidelines established by the National Institutes of Health and were approved by the Institutional

Animal Care and Use Committee of the University of California, Los Angeles, David Geffen School of Medicine. No statistical methods were used to pre-determine sample sizes. Sample sizes are similar to previous publications.

Habituation and water restriction

A minimum of one week before the start of the behavioral training we began habituating the mice to being handled by the experimenters and let them walk and climb on the experimenters' hands. During the handling animals were also presented with a treat (Mini yoghurt drop). At the same time we restricted the mice's access to water to 1 ml per subject and day and the water was given right after the mouse handling. To ensure that the subjects maintained healthy hydration levels we measured their body weight daily during the handling or before running the behavioral task. If their weight fell beneath 80% of their body weight before water restriction they were supplemented with 0.4 ml of water after the behavioral session (Guo *et al.*, 2014).

Perceptual evidence accumulation task

We trained animals on an audio-visual evidence accumulation task (Raposo 2012, Odoemene 2017). In this task mice were required to integrate pulsatile sensory evidence and compare the number of experienced stimuli to an implicit category boundary (12 Hz). Trials were randomly assigned to be high- or low-rate and we did not impose any limitations as to how many times in a row certain trial types could repeat. Mice self-initiated trials by poking into a centrally located port (and thereby breaking an infrared laser beam). After a very brief random delay this triggered the playback of a train of visual (or auditory stimuli). The train was composed of a set of 15 ms long stimulus events that occurred pseudo-randomly distributed within a duration of 1 second. Stimulus events were required to be separated by at least 25 ms. We generated stimulus trains with 4, 6, 8, 10, 14, 16, 18, and 20 stimulus events, omitting 12 Hz stimulus trains, which would lie on the category boundary. After the stimulus train had finished playing we repeated the same stimulus train again. This made that stimulus information was still available when the mice were allowed to make their decisions after center fixation. Mice had to remain poking in the central port for at least 1 second plus a random delay drawn from an exponential distribution, after which an auditory go-cue (7kHz) signaled to them that they could report their choice. The random delay was added to make sure that the animals could not simply time their response and to minimize movement preparation signals in the neural activity before the go-cue. If animals left the central port before the go-cue the trial was judged as an early withdrawal and the mice were presented with a white noise stimulus and 2 second timeout. Animals then reported their decision by poking into one of two ports either located on the left or the right side. Pokes in the left port after a <12 Hz stimulus or in the left port after a >12 Hz stimulus were rewarded with a 4 - 6 µl drop of water, while left pokes after high-rate trials and right pokes after low-rate trials were punished with a tone (15 KHz) and 2 second timeout. If animals failed to report a choice within 30 s the trial ended and was deemed a no-choice trial.

All the behavioral testing took place in sound attenuated booths that were illuminated only with infrared light (for video recordings). Animals performed the task inside a custom-made modular transparent red acrylic enclosure (20 x 20 cm, 24 high). Ports were installed along the back wall of the enclosure and separated by two 5 cm dividers made from clear acrylic. The task was implemented using custom written matlab functions and an arduino-based state machine (Bpod, Sanworks). All auditory and visual cues were delivered via a low-latency soundcard (Fenix, HT Omega) and sounds were played from two speakers located on either side of the central port (Harman-Kardon) while visual flashes were played from two LED panels attached above the center port. Water was delivered to the left and right ports through suspended syringes and silicone tubing and the volume was controlled via the opening time of a solenoid valve (Lee Company). Animal behavior was filmed through the transparent acrylic floor from underneath the animals (Chameleon3, Teledyne-FLIR, Fujinon 3 MP Varifocal Lens). The videos were acquired at 100 fps and the camera was controlled using labcams (https://github.com/jcouto/labcams). Camera acquisition was synchronized to behavioral events via TTL pulses at the beginning of every trial. All the code to run the behavioral task, and assembly instructions and a parts list are publicly available (https://github.com/churchlandlab/chipmunk).

Behavioral training procedure

One day before the start of the behavioral training water restricted mice were habituated to the task enclosure. They were presented with 1 ml water inside the left and right ports and in a weighing boat in front of the ports and allowed to drink and explore the enclosure for 5 - 10 minutes. At the beginning of the first full training sessions the inside of all the three ports was baited with a drop of water. For the first 3 - 5 sessions the wait time and all the delays were set to 0 s, so that the mice did not need to fixate inside the central port after triggering the stimulus. Similarly, for these sessions we removed the dividers between the ports to facilitate exploration and we set the stimulus to keep repeating until mice made their decision or for up to 30 s (which is the maximum allowed time to choose a response port). We did this to help the mice temporally bind causal task events (center poke -> stimulus, stimulus playing -> response port poke). Once animals exhibited stable center-to-side sequences we introduced the dividers and started increasing the required center fixation time. For most animals we slowly increased the wait time by 0.3 ms after each successfully completed trial (i.e. no early withdrawal). It took the animals up to 3 months (10.72 ± 1.25 weeks) to reach > 1 s wait-times.

Most of the subjects were exposed to rewards and punishments from the first day of training. This allowed them to learn the task rule (with 4 and 20 Hz stimuli) within 2 - 3 weeks (10 - 15 sessions) and avoided the animals developing strong side-biases as has been reported before (Odoemene *et al.*, 2018). One cohort of mice (LO028 and LO032) included in this study was allowed to revise their choices and be rewarded on the correct port after initially selecting the incorrect one. This procedure, however, was more prone to producing first a switching bias that would have to be counteracted by manipulating the trial structure and present trials rewarded on one side more frequently leading in turn to a side bias. This side bias had then to be counteracted by presenting trials for the other side more frequently before returning to a balanced number of trials on both sides and

removing any correlations in the trial structure**.** This approach therefore slowed down the animals' learning of the stimulus rate - side association. We started introducing the more difficult stimulus rates once the animals waited for > 1 s and performed at above 80% correct on the easy stimulus rates.

Behavioral analysis (Figure 1 and Supplementary Figure 1)

We included subjects for the behavioral analyses based on a two-step selection process: First, we only included subjects that learned to wait for > 1 s (including the random delay) and that experienced all the different stimulus rates at least during one session and that performed above 80% on easiest trials at least once. Second, from these subjects we included only sessions where the animals actually waited > 1 s plus delay, experienced at least two different pairs of stimulus rates (for example 4 -20 and 6 - 18 Hz), performed at least 100 trials and where the early withdrawal rate did not exceed 55%. We deliberately did not filter individual sessions for high performance to capture the session-by-session variability in biases that might affect performance levels.

We fitted individual psychometric curves for each subject using all the completed trials from all their sessions and then averaged the recovered psychometric functions (Figure 1C). Psychometric functions were modeled using cumulative Gaussian functions parametrized by four parameters, perceptual bias (α), sensitivity (β), lower (γ)- and upper lapse rate (λ):

$$p_{right\ side\ choice} = \gamma + (1 - \gamma - \lambda) \cdot F(x;\ \alpha, \beta)$$

where $F(x;\ \alpha, \beta)$ is the cumulative Gaussian function:

$$F(x;\ \alpha, \beta) = \frac{\beta}{\sqrt{2\pi}} \int_{-\infty}^{x} e^{\frac{-\beta^2 (x-\alpha)^2}{2}}$$

We also fitted separate psychometric functions to trials with different preceding trial history (Figure 1D). To do this we only included completed trials that were also directly preceded by a completed trial (as opposed to an early withdrawal). Again, to have sufficient amounts of data to estimate the psychometric parameters we pooled trials from multiple sessions. We used a custom python package to fit the four model parameters. (https://github.com/jcouto/fit_psychometric.git).

To assess the overall influence of sensory evidence and trial history on animals' decisions and to analyze session-by-session variability we constructed logistic regression models that sought to predict the mice's choice based on stimulus rate and trial history (Figure 1E and F). Our models featured a regressor for stimulus rate plus a categorical regressor for each combination of previous choice and outcome (for example previous correct left choice) and an intercept (sensory bias):

$$ln\left(\frac{p_{right\ side\ choice}}{1-\ p_{right\ side\ choice}}\right) = \beta_0 + \beta_1 \cdot x_{Stimulus\ rate} + \vec{\beta}_2 \cdot \vec{x}_{Trial\ history}$$

where $\vec{x}_{Trial\ history}$ is given by:

$$x_{Previous\ correct\ left} + x_{Previous\ incorrect\ left} + x_{Previous\ incorrect\ right} + x_{Previous\ correct\ right}$$

For the stimulus rate regressor the original stimulus rates were centered around 0 and scaled to their minimum and maximum. A 4 Hz stimulus was therefore encoded as -1 whereas a 20 Hz stimulus assumed a value of 1 with 0 representing a 12 Hz stimulus. The four trial history regressors were categorical regressors that were set to 1 whenever the trial history matched and were 0 otherwise. In these models the intercept represents the log odds of choosing the right side during a 12 Hz stimulus in the absence of any trial history. When fitting the models we balanced the number of left- and right-choice trials by subsampling from the class with more observations. We iteratively sampled observations from the majority class 20 times, making sure to include all the observations with similar frequencies. For each of these 20 rounds of sampling we applied 10-fold cross-validation. Models were fitted with an l2 penalty set to 1. We chose not to optimize this penalty to be able to more easily compare the model weights between different subjects and sessions.

To capture the overall influence of trial history on animals' choices we calculated the length of the vector of the four trial history weights (history strength, Figure 1G and H). All these analyses were performed using custom python code (https://github.com/churchlandlab/chiCa). We also modeled how choices and outcomes from the last and second last trials influenced animals' decisions (Supplementary Figure 1G). We did this by including a regressor for the choice, the outcome and their interaction for the last (t-1) and second last (t-2) trial instead of the trial history regressors. These main effects were coded as 1 when the choice at trial t-1 or t-2 was to the right side or when the choice was correct and 0 otherwise. The interaction regressor was set to 1 when the choices were to the right and correct and 0 otherwise. The stimulus regressor was modeled the same way as in the trial history model. Please note that when modeling the main effects of previous choices and previous outcomes the intercept term no longer reflects a trial history free bias (at 12 Hz stimulus rate) but rather represents the condition with all dummy variables set to 0, an incorrect left choice at t-1 and t-2 in this case. We also assessed whether sessions cluster with respect to their behavioral decoding weights (Supplementary Figure 1F). To this end we used UMAP to find a non-linear embedding to project the session's 6 regressor weights into a 2-dimesional space and plotted all the sessions color-coded by subject (Supplementary Figure 1F).

Surgical procedures

For our neural recordings, we targeted an anterior and ventral portion of anterior cingulate cortex, which was proposed to be a homolog of the primate Brodmann Area 24 (van Heukelum *et al.*, 2020). To gain optical access to ACC on the right hemisphere we implanted 4 x 1 mm gradient refractive index (GRIN) lens centered at the following target coordinates (with respect to Bregma): AP: +1.3 mm, ML: 0.5 mm, VD: -1.3 mm (-1.5 mm from the skull surface).

Before the surgery animals were subcutaneously injected with meloxicam (2 mg/kg) for analgesia and enrofloxacin (5 mg/kg) as antibiotic treatment. Mice were then first anesthetized in 3% isoflurane before being maintained on 1 - 1.5% isoflurane throughout the surgical procedures. Body temperature was monitored and maintained close to 37°C using a feedback-controlled heating pad (PhysioSuite, Kent Scientific). Animals were placed in a stereotaxic frame (Kopf, Stoelting). We cut the hair on the scalp with scissors and subsequently used a hair removal cream (Nair) to remove remaining hair. After disinfecting the scalp an incision was made and a diamond-shaped part of the skin was removed. We then carefully detached the neck muscles from the skull using forceps and scraped the skull clean with a scalpel. We next applied a layer of tissue adhesive glue (Vetbond, 3M) around the edges of the skull to seal off any exposed soft tissue. We then marked the position of the lens, drilled three very small craniotomies arranged in a triangle around the lens site, and inserted three micro screws into the skull (Fine Science Tools). We made a circular craniotomy with 1.2 mm diameter to later insert the GRIN lens. The dura underneath the craniotomy was carefully removed with sharp 30-gauge needles. We aspirated tissue using blunt 28- or 30-gauge needles using a custom vacuum suction apparatus. We were careful to aspirate a cylindrical volume rather than a conical one and to create a flat surface at the bottom of the aspirated volume. All excess blood was aspirated before lowering the lens into the craniotomy using a suction-based lens holder. After placing the lens we cemented it to the skull and anchored it to the three screws using a three-component dental cement (C&B metabond, Parkell). We then built a support structure around the lens using black dental cement (Orthojet, Lang Dental). The tip of the lens was protected using the lid of a 1 ml eppendorf tube glued to the dental cement structure using acrylic glue (Zap-A-Gap). Mice were allowed to recover for 7 days after the surgery and they were checked daily. On the first 4 days after surgery they were injected with meloxicam and enrofloxacin.

A minimum of 14 days animals were implanted with a miniscope baseplate. To this end animals were again anesthetized in 3% and maintained on 1 -1.5% isoflurane anesthesia, and their body temperature was held constant. The eppendorf lid was detached using a drop of acetone and the lens surface was cleaned with ethanol. To provide additional contrast we painted the area around the GRIN lens with black nail polish. We then positioned the miniaturized microscope with a baseplate attached above the lens surface. At this point the anesthesia was lowered to 0.5% to decorrelated cortical activity in order to identify the best imaging focal plane. The baseplate was then attached using black dental cement. After the cement had dried we removed the miniature microscope and covered the baseplate with a protective cap. We monitored the animals behavior

immediately after recovery. This procedure was non-invasive and did not require administration of analgesics or antibiotics.

Histology

At the end of the experiments animals were deeply anesthetized with an intraperitoneal injection of pentobarbital (Euthasol) and transcardially perfused with phosphate-buffered saline followed by 4% paraformaldehyde. Subsequently the brain was extracted and stored in 4% paraformaldehyde overnight for post-fixation before being stored in phosphate-buffered saline. The brains were then embedded in 4% agarose and cut into 80 μm sections. Sections were mounted and nuclei were stained using DAPI inside the mounting medium (Fluoromount-G with DAPI, Invitrogen). Photomicrographs of GCaMP6s expression and DAPI were taken using a Nikon Eclipse Ti2 microscope equipped with a Yokogawa CSU-22 Spinning Disk unit to enable confocal fluorescence microscopy. Approximate coordinates of the GRIN lens center were obtained by comparing major landmarks on the sections (position of the anterior commissure, presence of corpus callosum) with a standard mouse brain atlas (Paxinos and Keith B. J. Franklin, 2007, 4th edition).

Miniature microscope imaging and image processing

We imaged the neural activity of ACC neurons using a head-mounted miniaturized fluorescence microscope (Cai *et al.*, 2016, UCLA miniscope V4, https://github.com/Aharoni-Lab/Miniscope-v4). For every subject we determined the best focal plane by performing brief recordings at different imaging depths (by adjusting the electro tunable lens). We then kept this depth constant across recording sessions. We used excitation powers ranging from 22 to 37 mW and adjusted them to optimize imaging quality. Along with the images we also acquired data from gyroscopes on the miniaturized microscope and later reconstructed the animals head-orientation angles from these data. Images and gyroscope data were acquired at 30 frames per second. We synchronized the behavioral events and the imaging by logging the TTL pulses from the behavioral state machine indicating the start of a new trial and TTLs from the miniscope signaling the acquisition of a frame on the same computer using a Teensy (Teensy 3.2).

To segment images and identify individual neurons we used a custom pipeline data handling- and processing pipeline (https://github.com/jcouto/labdata-tools) and built a plugin to run caiman-based image segmentation routines within this framework (Zhou *et al.*, 2018, https://github.com/flatironinstitute/CaImAn). We first spatially downsampled the videos by a factor of 2 using the imresize function in Matlab (which yielded better results than comparable methods in python). We then corrected the videos for motion along the x and y axes using the normCorre implementation in caiman and subsequently ran the segmentation using constrained non-negative matrix factorization (CNMF-E). Raw traces were temporally detrended with the built-in caiman methods. We then visually inspected the identified putative neurons and excluded signals that were not likely originating from neurons. We used the following exclusion criteria: Location outside of the lens perimeter, aberrant shape and size, absence of activity above local background fluorescence. We

then aligned the neural recording to the behavior by identifying the recording frames during which new trials started. At this stage we also identified frame drops (unexpectedly long intervals between frame acquisition TTLs) and accounted for these by interpolating the neural signals.

Movement and posture analysis

From the gyroscope data acquired using the miniature microscope we reconstructed the animals' head orientation at every imaging time point, which was given by the three angles: pitch, roll and yaw.

To track the position of the animals in space we used deeplabcut (Mathis *et al.*, 2018, , https://github.com/DeepLabCut/DeepLabCut). We tracked the following 12 body parts that were visible at the underside of the mice: tail tip, tail center, rectum, genital, chest, nose tip, left foot, right foot, left hand, right hand, left ear, right ear. For our analysis we only used the chest point as it represents the most central tracked body part (to estimate position in space) and because it was rarely occluded. Models were trained with a total of 160 labeled frames from two different sessions, 80 frames each.

We also generated low-dimensional versions of the full behavioral video. To do this we used an approximate singular value decomposition (https://github.com/jcouto/wfield) on the mean-centered videos and retained the top 200 components. For every video we also generated a motion-energy video by subtracting each frame from its preceding one and then taking the absolute value of this difference. We then similarly decomposed these motion energy videos and retained the top 200 components for further analysis.

Neural decoding analyses (Figure 2 and Supplementary Figure 2)

For the neural decoding analysis we aligned the neural activity to four different phases of the inter-trial interval and the trial time. We defined the start of the ITI as the moment when the animals had reported their choice and we included neural activity for up to 1 s after (Early ITI). The next phase (Late ITI) was based on when the outcome ended. This moment was defined as the animal leaving the rewarded port after a correct trial and at the end of the punishment noise and timeout (2 s after choice report) when the animals had made an incorrect choice. For the analysis we included the 1 s preceding this moment to 0.3 s after. We defined the stimulus phase as lasting from 0.5 s before the onset of the visual stimuli to 1 s after and aligned the neural activity accordingly. Finally, we included data from 0.2 s before to 0.3 after animals left the fixation in the center port (Action).

We decoded the specific trial history context from the neural population activity at the different aligned time points. We fitted logistic regression models using the inferred spike rate signal convolved with a Gaussian kernel (1 frame standard deviation) to classify combinations of previous choice and outcome (Figure 2F, G), hence the decoders distinguished between 4 different classes using a one-versus-all approach. We balanced the number of classes in a similar way as described in behavioral analysis, subsampling from the majority classes. We then partitioned the subsampled data into a training set containing 80% of the data and a test set containing the remaining 20% while maintaining balanced labels and a test set containing the remaining 20% of the data. The activity of

individual neurons in both sets were standardized by subtracting the mean of the training set and dividing it by the standard deviation of the training set. This avoided data leakage from the training to the test set. We then first optimized the magnitude of the L2 penalty by finding the value that yielded best model predictions in the test fold within an exponential series from 10^{-10} to 10^{12}. The best parameter was then used to fit the logistic regression models. These procedures were repeated for the other 7 training-test splits. Finally, this procedure was repeated 20 times using different subsampled trials. The final model accuracies represent the averages over all sub-samplings and cross-validation folds. As a control, in parallel we fitted models where we shuffled the labels yielding chance level decoding accuracy. Decoding analyses to determine the neural decoding of previous choice or previous outcome only rather than all of their combinations were performed in the same way (Supplementary Figure 2A).

To analyze the similarity of the neural representations between the different classes we constructed confusion matrices (Figure 2H). To do this, we counted how many times decoders correctly and incorrectly classified trials as each of the different trial history contexts for every given true trial history. Values are reported as the number of predicted trials divided by the number of true trials for each class.

We also looked at how stable the neural encoding of trial history was on the population level (Figure 2I and J). We did this by predicting the trial history contexts from neural activity at a given time point using the relationship (the logistic regression model) found at another time point. We did this for every pair of included time points. We made sure to do the predictions with the originally trained models from all folds and sub-samplings and then averaged the prediction accuracies rather than averaging the model coefficients first and then performing the predictions.

Linear encoding models (Figure 3 and Supplementary Figure 3)

To assess what task variables or movements drive neural variance we constructed a linear encoding model (Musall *et al.*, 2019). Importantly, our model was built such that it could not only capture relationships that were fixed across trial time, such as neural tuning for specific locations in the task arena, but also relationships that could vary over time, such as how choice influences neural activity. This was achieved by including kernel regressors that could span parts or the entire aligned trial time (Figure 3A). These regressors were binary vectors that were set to 1 only at the time of an event. To for example capture how choice affected neural activity we added time-shifted versions of these vectors to the model until the time of the events spanned the entire trial time. Variables for which the kernel regressors spanned the entire trial time included choice, outcome and trial history (one for each combination of previous choice and outcome). Importantly, we also included a trial time variable, which consisted of regressors that were 1 for the same time point across trials. This trial time regressor effectively served as an intercept for the other time-shifted regressors. We also time-shifted regressors that tracked the stimulus events as kernels, so that we could detect neural responses to the stimulus events for up to 0.5 s after they occurred. We also included pokes into all the different ports as regressors for instructed movements and allowed them to capture neural variance from 0.5 s before to 1

s after the poke happened. Head-orientation angle tuning was assessed by dividing the full range of 360° into 60 bins (6° per bin). We then built a binary vector whose value was one whenever the animal's head-orientation angle was within that bin and zero otherwise. We did this separately for pitch, roll and yaw. In an analogous manner we constructed regressors that tracked whether a mouse's chest was within a spatial bin or not. The bins cover 1.28 x 1.28 cm. Finally, we included the top 200 video and video motion energy components as analog regressors in the design matrix. To make sure that the video and video motion energy components only contained information about the movements rather than the visual flashes that were also visible on the video we orthogonalized these regressors against the stimulus regressors using QR-decomposition before fitting the linear models. We then aligned the raw fluorescence traces for all the neurons to the specified trial events in the four different trial phases (as described in decoding analysis).

To fit the models we used ridge regression, which regularizes the model weights using an L2 penalty. We did this to handle the high degree of multi-collinearity of the different regressors (trial history and chest point, for example) and to counteract over-fitting due to the large number of included regressors. To fit the models we first split the data in 10 folds (see decoding analysis above for detailed description of the procedure) making sure that we equally sampled all trial time points. We then standardized our analog regressors (the video SVD and video motion energy SVD) in the training set and used the mean and standard deviation from the training set to standardize the regressors in the test set. We then standardized the neural activity in a similar fashion. We then first determined the optimal regularization strength for each neuron by fitting the models using an exponential scale of regularization strengths from 10^{-3} to 10^{5} and selecting the value that yielded the highest explained variance in the test set as assessed using the coefficient of determination (R^2). We then refitted the models with that value before repeating this procedure in the next fold. Besides fitting the full models we also fitted single-variable only models, where the regressors for all but one variable were shuffled, and one-removed models, where the regressors of only a single variable were shuffled. We then computed the maximal explained variance per neuron for each variable (Figure 3B) as the single-variable only variance of that variable minus the trial time-only model variance (representing the intercept). Analogously, we computed the unique variance (Figure 3C) as the difference between the full model and the one-removed model variances. To reconstruct the explained variance for specific variables over trial time (Figure 3D) we retained the predicted activity for each neuron of all the folds and then re-ordered them to reflect original trial time. We then calculated the coefficient of determination for each time point separately. Note however, that the models were fitted and regularized to explain variance across the entire trial time and so the estimates of explained variance for individual trial times tends to be lower than the full model explained variance.

Variable dimensionality and procrustes analysis (Figure 4)

We assessed the dimensionality of the neural representation of specific movements or task variables by separately applying principal component analysis to the encoding model weights (from the full model) of the different variables. We then defined the dimensionality as the minimum number of dimensions required to explain 90% of the weight variance

across the neuron population (Figure 4C and D). We measured the similarity of the population representation of task variables within and between subjects using procrustes analysis (Horrocks, Rodrigues and Saleem, 2024, Figure 4E, F and G). Procrustes analysis seeks to minimize the (least-squares) distance between two shapes by identifying the optimal matrix translation, scaling and rotation of the target shape. To match the shapes of the matrices for procrustes analysis and to ensure that we were considering meaningful low-dimensional dynamics or behavioral tuning motifs we determined the sessions with maximal dimensionality (i.e. dimensions necessary to explain 90% of the variance) and used this number of dimensions for all the other sessions. We did this for every variable separately. The dissimilarity between shapes was expressed as the procrustes distance, where identical shapes have a distance of 0, whereas a value of 1 indicates maximal dissimilarity.

Statistical analyses

All statistical comparisons were performed using linear-mixed effects models (lmerTest, R) with subject identity as a random effect. When multiple measures within a session were analyzed we included the session identity as a random effect nested within subject. We only included random intercepts in our mixed-effects models. Post-hoc group means were compared by estimating the marginal means from the mixed-model fits using emmeans (R). Pearson correlation coefficients and their significance were computed using scipy.stats.pearsonr.

Data and code availability

All code is publicly available. Wherever specific packages are used the packages are cited in the corresponding methods section. For all decoding- and encoding analyses the code can be found here (https://github.com/churchlandlab/chiCa). All the data used in this study will be made available upon publication.

Acknowledgements

We thank Letizia Ye and Marvin Vasquez for animal training and care, Marvin Vasquez for acquiring the histology photomicrographs and Anup Khanal for technical help. We thank Pinping Zhao for guidance on tissue aspiration and GRIN lens implantation procedures and Federico Sangiuliano Jimka for technical assistance with the V4 minscopes. We would also like to thank all the lab members for their feedback and discussions. This work was supported by NSF-NCS collaborative award (2219946), the Simons Global Brain to AKC and the Swiss National Science Foundation Early Postdoc Mobility fellowship (P2BEP3) to LTO.

Author contributions

LTO and AKC designed the project, LTO and JC adapted the behavioral task from a previous version, LTO, MCT and DS trained animals, LTO performed surgeries and

collected imaging data, LTO and JC wrote analysis tools, LTO analyzed and visualized the data, LTO and AKC wrote the manuscript with inputs from all other authors.

References

Acuña, M. A. *et al.* (2023) ‘Principles of nociceptive coding in the anterior cingulate cortex’, *Proceedings of the National Academy of Sciences*, 120(23), p. 2017. doi: 10.1073/pnas.2212394120.

Akam, T. *et al.* (2021) ‘The Anterior Cingulate Cortex Predicts Future States to Mediate Model-Based Action Selection’, *Neuron*. Elsevier Inc., 109(1), pp. 149-163.e7. doi: 10.1016/j.neuron.2020.10.013.

Akrami, A. *et al.* (2018) ‘Posterior parietal cortex represents sensory history and mediates its effects on behaviour’, *Nature*. Nature Publishing Group, 554(7692), pp. 368–372. doi: 10.1038/nature25510.

Alexander, W. H. and Brown, J. W. (2011) ‘Medial prefrontal cortex as an action-outcome predictor’, *Nature Neuroscience*. Nature Publishing Group, 14(10), pp. 1338–1344. doi: 10.1038/nn.2921.

Bissonette, G. B., Powell, E. M. and Roesch, M. R. (2013) ‘Neural structures underlying set-shifting: Roles of medial prefrontal cortex and anterior cingulate cortex’, *Behavioural Brain Research*. Elsevier B.V., 250, pp. 91–101. doi: 10.1016/j.bbr.2013.04.037.

Bonaiuto, J. J., De Berker, A. and Bestmann, S. (2016) ‘Response repetition biases in human perceptual decisions are explained by activity decay in competitive attractor models’, *eLife*, 5(DECEMBER2016), pp. 1–28. doi: 10.7554/eLife.20047.

Braun, A., Urai, A. E. and Donner, T. H. (2018) ‘Adaptive history biases result from confidence-weighted accumulation of past choices’, *Journal of Neuroscience*, 38(10), pp. 2418–2429. doi: 10.1523/JNEUROSCI.2189-17.2017.

Busse, L. *et al.* (2011) ‘The detection of visual contrast in the behaving mouse’, *Journal of Neuroscience*, 31(31), pp. 11351–11361. doi: 10.1523/JNEUROSCI.6689-10.2011.

Cai, D. J. *et al.* (2016) ‘A shared neural ensemble links distinct contextual memories encoded close in time’, *Nature*. Nature Publishing Group, 534(7605), pp. 115–118. doi: 10.1038/nature17955.

Chen, W. *et al.* (2024) ‘Anterior cingulate cortex provides the neural substrates for feedback-driven iteration of decision and value representation’, *Nature Communications*. Springer US, 15(1), pp. 1–15. doi: 10.1038/s41467-024-50388-9.

Choe, I. H. *et al.* (2017) ‘Mice in social conflict show rule-observance behavior enhancing long-term benefit’, *Nature Communications*. Springer US, 8(1). doi: 10.1038/s41467-017-01091-5.

Cole, N. *et al.* (2024) ‘Prediction-error signals in anterior cingulate cortex drive task-switching’, *Nature Communications*. Springer US, 15(1), p. 7088. doi: 10.1038/s41467-024-51368-9.

Diehl, G. W. and Redish, A. D. (2023) ‘Differential processing of decision information in subregions of rodent medial prefrontal cortex’, *eLife*, 12. doi: 10.7554/eLife.82833.

Ebitz, R. B. and Hayden, B. Y. (2016) ‘Dorsal anterior cingulate : a Rorschach test for cognitive neuroscience’, *Nature Publishing Group*. Nature Publishing Group, 19(10), pp. 1278–1279. doi: 10.1038/nn.4387.

González, V. V *et al.* (2025) ‘A common stay-on-goal mechanism in anterior cingulate cortex for information and effort choices’, *eneuro*, p. ENEURO.0454-24.2025. doi: 10.1523/ENEURO.0454-24.2025.

Guo, Z. V. *et al.* (2014) ‘Procedures for behavioral experiments in head-fixed mice’, *PLoS ONE*, 9(2). doi: 10.1371/journal.pone.0088678.

Gupta, D. *et al.* (2024) ‘Trial-history biases in evidence accumulation can give rise to apparent lapses in decision-making’, *Nature Communications*, 15(1), p. 662. doi: 10.1038/s41467-024-44880-5.

Hart, E. E. *et al.* (2020) ‘Chemogenetic modulation and single-photon calcium imaging in anterior cingulate cortex reveal a mechanism for effort-based decisions’, *Journal of Neuroscience*, 40(29), pp. 5628–5643. doi: 10.1523/JNEUROSCI.2548-19.2020.

Hasnain, M. A. *et al.* (2024) ‘Separating cognitive and motor processes in the behaving mouse’, *Nature Neuroscience*. Springer US. doi: 10.1038/s41593-024-01859-1.

Hayden, B. Y. *et al.* (2011) ‘Surprise signals in anterior cingulate cortex: Neuronal encoding of unsigned reward prediction errors driving adjustment in behavior’, *Journal of Neuroscience*, 31(11), pp. 4178–4187. doi: 10.1523/JNEUROSCI.4652-10.2011.

Hermoso-Mendizabal, A. *et al.* (2020) ‘Response outcomes gate the impact of expectations on perceptual decisions’, *Nature Communications*, 11(1), pp. 1–13. doi: 10.1038/s41467-020-14824-w.

van Heukelum, S. *et al.* (2020) ‘Where is Cingulate Cortex? A Cross-Species View’, *Trends in Neurosciences*. The Authors, 43(5), pp. 285–299. doi: 10.1016/j.tins.2020.03.007.

Hirokawa, J. *et al.* (2019) ‘Frontal cortex neuron types categorically encode single decision variables’, *Nature*. Springer US, 576(7787), pp. 446–451. doi: 10.1038/s41586-019-1816-9.

Hocker, D. *et al.* (2021) ‘Subpopulations of neurons in LOFC encode previous and current rewards at time of choice’, *eLife*, 10, pp. 1–26. doi: 10.7554/eLife.70129.

Holroyd, C. B. and Yeung, N. (2012) ‘Motivation of extended behaviors by anterior cingulate cortex’, *Trends in Cognitive Sciences*. Elsevier Ltd, 16(2), pp. 122–128. doi: 10.1016/j.tics.2011.12.008.

Horrocks, E. A. B., Rodrigues, F. R. and Saleem, A. B. (2024) ‘Flexible neural population dynamics govern the speed and stability of sensory encoding in mouse visual cortex’, *Nature Communications*. Springer US, 15(1). doi: 10.1038/s41467-024-50563-y.

Hwang, E. J. *et al.* (2017) ‘History-based action selection bias in posterior parietal cortex’, *Nature Communications*. Springer US, 8(1), pp. 1–14. doi: 10.1038/s41467-017-01356-z.

Hyman, J. M., Holroyd, C. B. and Seamans, J. K. (2017) ‘A Novel Neural Prediction Error Found in Anterior Cingulate Cortex Ensembles’, *Neuron*. Elsevier Inc., 95(2), pp. 447-456.e3. doi: 10.1016/j.neuron.2017.06.021.

Jiang, W. *et al.* (2019) ‘Short-Term Influence of Recent Trial History on Perceptual Choice Changes with Stimulus Strength’, *Neuroscience*. IBRO, 409, pp. 1–15. doi: 10.1016/j.neuroscience.2019.04.010.

Karlsson, M. P., Tervo, D. G. R. and Karpova, A. Y. (2012) ‘Network resets in medial prefrontal cortex mark the onset of behavioral uncertainty’, *Science*, 338(6103), pp. 135–139. doi: 10.1126/science.1226518.

Kaufman, M. T. *et al.* (2022) ‘The implications of categorical and category-free mixed selectivity on representational geometries’, *Current Opinion in Neurobiology*. The Authors, 77, p. 102644. doi: 10.1016/j.conb.2022.102644.

Kira, S. *et al.* (2023) ‘A distributed and efficient population code of mixed selectivity neurons for flexible navigation decisions’, *Nature Communications*, 14(1). doi: 10.1038/s41467-023-37804-2.

Kolling, N. *et al.* (2016) ‘Value, search, persistence and model updating in anterior cingulate cortex’, *Nature Neuroscience*, 19(10), pp. 1280–1285. doi: 10.1038/nn.4382.

Lak, A. *et al.* (2020) ‘Reinforcement biases subsequent perceptual decisions when confidence is low: A widespread behavioral phenomenon’, *eLife*, 9, pp. 1–26. doi: 10.7554/eLife.49834.

Li, Y. S. *et al.* (2019) ‘Individual neurons in the cingulate cortex encode action monitoring, not selection, during adaptive decision-making’, *Journal of Neuroscience*, 39(34), pp. 6668–6683. doi: 10.1523/JNEUROSCI.0159-19.2019.

Maisson, D. J. N. *et al.* (2021) ‘Choice-relevant information transformation along a ventrodorsal axis in the medial prefrontal cortex’, *Nature Communications*. Springer US, 12(1), pp. 1–14. doi: 10.1038/s41467-021-25219-w.

Mante, V. *et al.* (2013) ‘Context-dependent computation by recurrent dynamics in prefrontal cortex’, *Nature*. Nature Publishing Group, 503(7474), pp. 78–84. doi: 10.1038/nature12742.

Mathis, A. *et al.* (2018) ‘DeepLabCut: markerless pose estimation of user-defined body parts with deep learning’, *Nature Neuroscience*. Springer US, 21(9), pp. 1281–1289. doi: 10.1038/s41593-018-0209-y.

Mayford, M. *et al.* (1996) ‘Control of memory formation through regulated expression of a CaMKII transgene’, *Science*, 274(5293), pp. 1678–1683. doi: 10.1126/science.274.5293.1678.

Mendonça, A. G. *et al.* (2020) ‘The impact of learning on perceptual decisions and its implication for speed-accuracy tradeoffs’, *Nature Communications*. Springer US, 11(1), pp. 1–15. doi: 10.1038/s41467-020-16196-7.

Mochol, G., Kiani, R. and Moreno-Bote, R. (2021) ‘Prefrontal cortex represents heuristics that shape choice bias and its integration into future behavior’, *Current Biology*, 31(6), pp. 1234-1244.e6. doi: 10.1016/j.cub.2021.01.068.

Monosov, I. E. *et al.* (2020) ‘Anterior Cingulate Cortex and the Control of Dynamic Behavior in Primates’, *Current Biology*. Elsevier Inc., 30(23), pp. R1442–R1454. doi: 10.1016/j.cub.2020.10.009.

Morcos, A. S. and Harvey, C. D. (2016) ‘History-dependent variability in population dynamics during evidence accumulation in cortex’, *Nature Neuroscience*, 19(12), pp. 1672–1681. doi: 10.1038/nn.4403.

Murray, J. D. *et al.* (2014) ‘A hierarchy of intrinsic timescales across primate cortex’, *Nature Neuroscience*. Nature Publishing Group, 17(12), pp. 1661–1663. doi: 10.1038/nn.3862.

Musall, S. *et al.* (2019) ‘Single-trial neural dynamics are dominated by richly varied movements’, *Nature Neuroscience*. Springer US, 22(10), pp. 1677–1686. doi: 10.1038/s41593-019-0502-4.

Musall, S. *et al.* (2023) ‘Pyramidal cell types drive functionally distinct cortical activity patterns during decision-making’, *Nature Neuroscience*, 26(3), pp. 495–505. doi: 10.1038/s41593-022-01245-9.

Muysers, H. *et al.* (2024) ‘A persistent prefrontal reference frame across time and task rules’, *Nature Communications*, 15(1), pp. 1–14. doi: 10.1038/s41467-024-46350-4.

Najafi, F. *et al.* (2020) ‘Excitatory and Inhibitory Subnetworks Are Equally Selective during Decision-Making and Emerge Simultaneously during Learning’, *Neuron*. Elsevier Inc., 105(1), pp. 165-179.e8. doi: 10.1016/j.neuron.2019.09.045.

Odoemene, O. *et al.* (2018) ‘Visual evidence accumulation guides decision-making in unrestrained mice’, *Journal of Neuroscience*, 38(47), pp. 10143–10155. doi: 10.1523/JNEUROSCI.3478-17.2018.

Paxinos, G. and Keith B. J. Franklin, M. A. (2007) *The Mouse Brain in Stereotaxic Coordinates*. Elsevier Science.

Posani, L. *et al.* (2024) ‘Rarely categorical, always high-dimensional: how the neural code changes along the cortical hierarchy’, *bioRxiv*. doi: 10.1101/2024.11.15.623878.

Quilodran, R., Rothé, M. and Procyk, E. (2008) ‘Behavioral Shifts and Action Valuation in the Anterior Cingulate Cortex’, *Neuron*, 57(2), pp. 314–325. doi: 10.1016/j.neuron.2007.11.031.

Raposo, D. *et al.* (2012) ‘Multisensory decision-making in rats and humans’, *Journal of Neuroscience*, 32(11), pp. 3726–3735. doi: 10.1523/JNEUROSCI.4998-11.2012.

Raposo, D., Kaufman, M. T. and Churchland, A. K. (2014) ‘A category-free neural population supports evolving demands during decision-making’, *Nature Neuroscience*. Nature Publishing Group, 17(12), pp. 1784–1792. doi: 10.1038/nn.3865.

Reinert, S. *et al.* (2021) ‘Mouse prefrontal cortex represents learned rules for categorization’, *Nature*. Springer US, 593(7859), pp. 411–417. doi: 10.1038/s41586-021-03452-z.

Rigotti, M. *et al.* (2013) ‘The importance of mixed selectivity in complex cognitive tasks’, *Nature*, 497(7451), pp. 585–590. doi: 10.1038/nature12160.

Salkoff, D. B. *et al.* (2020) ‘Movement and Performance Explain Widespread Cortical Activity in a Visual Detection Task’, *Cerebral Cortex*, 30(1), pp. 421–437. doi: 10.1093/cercor/bhz206.

Sarafyazd, M. and Jazayeri, M. (2019) ‘Hierarchical reasoning by neural circuits in the frontal cortex’, *Science*, 364(6441). doi: 10.1126/science.aav8911.

Seo, H. and Lee, D. (2007) ‘Temporal filtering of reward signals in the dorsal anterior cingulate cortex during a mixed-strategy game’, *Journal of Neuroscience*, 27(31), pp. 8366–8377. doi: 10.1523/JNEUROSCI.2369-07.2007.

Shenhav, A., Cohen, J. D. and Botvinick, M. M. (2016) ‘Dorsal anterior cingulate cortex and the value of control’, *Nature Neuroscience*, 19(10), pp. 1286–1291. doi: 10.1038/nn.4384.

Siniscalchi, M. J., Wang, H. and Kwan, A. C. (2019) ‘Enhanced Population Coding for Rewarded Choices in the Medial Frontal Cortex of the Mouse’, *Cerebral Cortex*, 29(10), pp. 4090–4106. doi: 10.1093/cercor/bhy292.

Soltani, A. and Izquierdo, A. (2019) ‘Adaptive learning under expected and unexpected uncertainty’, *Nature Reviews Neuroscience*. Springer US, 20(10), pp. 635–644. doi: 10.1038/s41583-019-0180-y.

Song, M. *et al.* (2024) ‘Hierarchical gradients of multiple timescales in the mammalian forebrain’, *Proceedings of the National Academy of Sciences*, 121(51), p. 2017. doi: 10.1073/pnas.2415695121.

Spellman, T. *et al.* (2021) ‘Prefrontal deep projection neurons enable cognitive flexibility via persistent feedback monitoring’, *Cell*. Elsevier Inc., 184(10), pp. 2750-2766.e17. doi: 10.1016/j.cell.2021.03.047.

Stolyarova, A. *et al.* (2019) ‘Contributions of anterior cingulate cortex and basolateral amygdala to decision confidence and learning under uncertainty’, *Nature Communications*. Springer US, 10(1). doi: 10.1038/s41467-019-12725-1.

Stringer, C. *et al.* (2019) ‘Spontaneous behaviors drive multidimensional, brainwide activity’, *Science*, 364(6437). doi: 10.1126/science.aav7893.

Sul, J. H. *et al.* (2010) ‘Distinct roles of rodent orbitofrontal and medial prefrontal cortex in decision making’, *Neuron*. Elsevier Ltd, 66(3), pp. 449–460. doi: 10.1016/j.neuron.2010.03.033.

Tervo, D. G. R. *et al.* (2014) ‘Behavioral variability through stochastic choice and its gating by anterior cingulate cortex’, *Cell*. Elsevier Inc., 159(1), pp. 21–32. doi: 10.1016/j.cell.2014.08.037.

Tervo, D. G. R. *et al.* (2021) ‘The anterior cingulate cortex directs exploration of alternative strategies’, *Neuron*. Elsevier Inc., 109(11), pp. 1876-1887.e6. doi: 10.1016/j.neuron.2021.03.028.

Tremblay, S. *et al.* (2023) ‘Neural cognitive signals during spontaneous movements in the macaque’, *Nature Neuroscience*. Springer US, 26(2), pp. 295–305. doi: 10.1038/s41593-022-01220-4.

Vázquez, D. *et al.* (2024) ‘Optogenetic Inhibition of Rat Anterior Cingulate Cortex Impairs the Ability to Initiate and Stay on Task’, *Journal of Neuroscience*, 44(20), pp. 1–11. doi: 10.1523/JNEUROSCI.1850-23.2024.

Wekselblatt, J. B. *et al.* (2016) ‘Large-scale imaging of cortical dynamics during sensory perception and behavior’, *Journal of Neurophysiology*, 115(6), pp. 2852–2866. doi: 10.1152/jn.01056.2015.

Yin, C. *et al.* (2024) ‘Spontaneous movements and their impact on neural activity fluctuate with latent engagement states’, *bioRxiv*, p. 2023.06.26.546404. Available at: https://www.biorxiv.org/content/10.1101/2023.06.26.546404v2%0Ahttps://www.biorxiv.org/content/10.1101/2023.06.26.546404v2.abstract.

Zhou, P. *et al.* (2018) ‘Efficient and accurate extraction of in vivo calcium signals from microendoscopic video data’, *eLife*, 7, pp. 1–37. doi: 10.7554/eLife.28728.

Supplementary figures

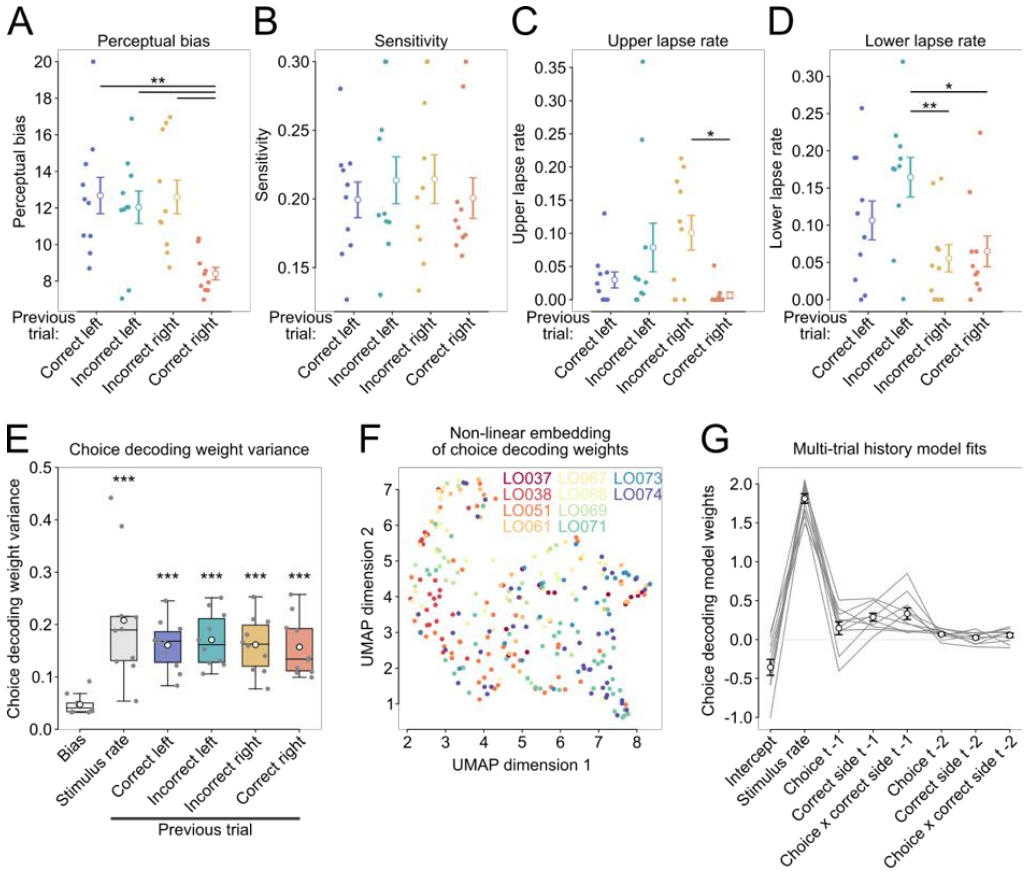

Supplementary figure 1. The behavioral influence of most recent trial history varies between sessions. *(A, B, C, D) Estimates of the psychometric parameters, bias (A), sensitivity (B), upper lapse rate (C), and lower lapse rate(D). Scatter points on the left represent individual subjects while dots with errorbars depict mean ± sem over subjects. *P < 0.05, **P < 0.01. (E) Variance of the sessions-by-session weights of the behavioral choice decoding. Boxes represent median ± interquartile range and dots depict the mean. All comparisions are made agains the bias (intercept). ***P < 0.001. (F) Non-linear embedding of the behavioral choice decoding weights. Scatter points represent sessions color-coded by subject. (G) Behavioral decoding models including regressors for most recent trial history (t-1) and trial history from two trials back (t-2). Lines are individual subjects while dots with errorbars show mean ± sem over subjects.*

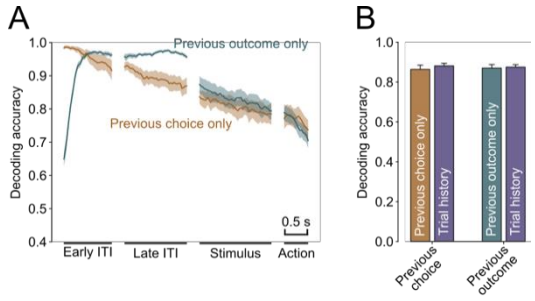

Supplementary Figure 2. Decoders trained on trial history decode previous choice and outcome equally well as corresponding binary classifiers. *(A) Mean ± sem decoding accuracy across trial time for decoders trained on previous choice- (orange) or previous outcome (teal) only. (B) Mean ±sem neural decoding accuracy for previous choice or previous outcome for their corresponding binary decoders or trial history decoders (trained to decode combinations of previous choices and outcomes).*

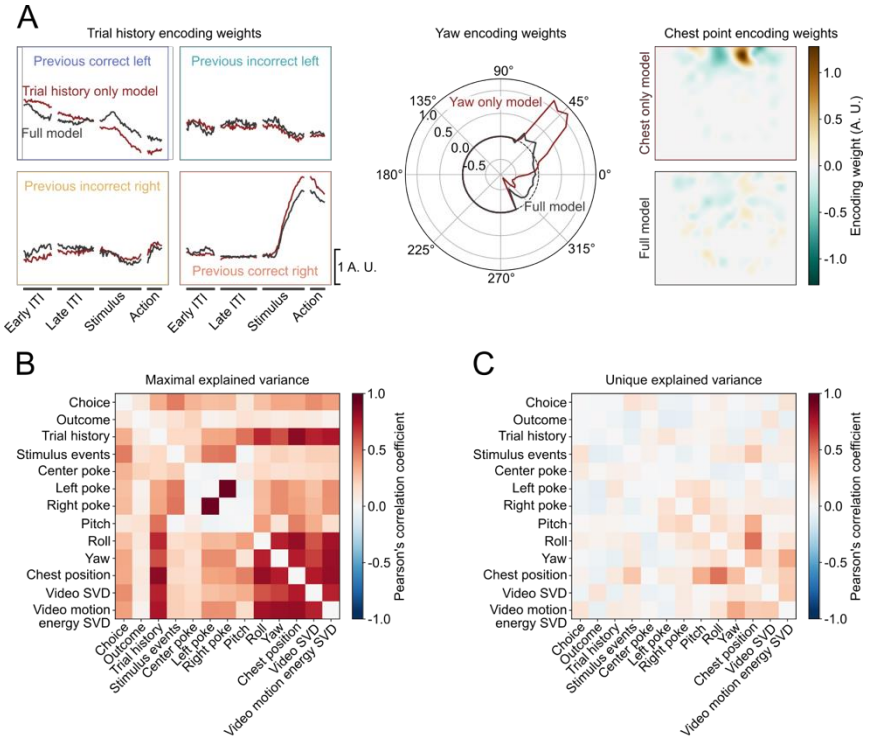

Supplementary Figure 3. The linear encoding model effectively shrinks the weight of uninformative regressors. *(A) Comparison of encoding model weights from an example neuron for model where the regressors for all except a single variable are shuffled (red) and the full models without shuffling (dark grey). Trial history weights (left), yaw tuning (center) and chest point tuning (right) are shown. (B) Correlation between maximal explained variance for different variables in all neurons from an example session. Note the high correlation between trial history and yaw and chest point. (C) Correlation of the unique explained variance between variables for the same session as in (B).*

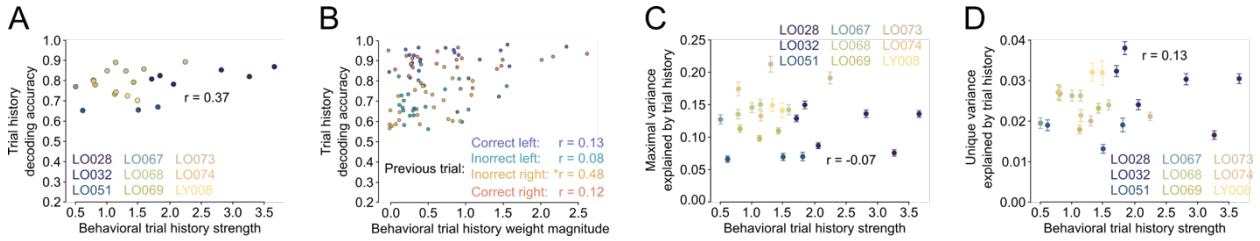

Supplementary figure 4. Neural encoding of trial history in ACC is not correlated with the strength of trial history biases. *(A) Neural population decoding of trial history is not correlated with the behavioral usage of trial history information. Each point is a single animal. (B) Decoding accuracy for specific trial history contexts is independent of the magnitude of the corresponding behavioral model weight. (C and D) Neither maximal nor unique variance of individual neurons explained by trial history are correlated with behavioral trial history biases.*