Posterior parietal cortex represents sensory stimulus history and is necessary for its effects on behavior.

Athena Akrami[1,2,3], Charles D. Kopec[1,2], Mathew E. Diamond[4], Carlos D. Brody[1,2,3*]

1. Princeton Neuroscience Institute, Princeton University, Princeton, NJ 08544, USA.
2. Department of Molecular Biology, Princeton University, Princeton, NJ 08544, USA
3. Howard Hughes Medical Institute.
4. Tactile Perception and Learning Laboratory, International School for Advanced Studies (SISSA), 34136 Trieste, Italy

*. Corresponding author (brody@princeton.edu)

Many models of cognition and of neural computations posit the use and estimation of prior stimulus statistics[1–4]: it has long been known that working memory and perception are strongly impacted by previous sensory experience, even when that sensory history is irrelevant for the current task at hand. Nevertheless, the neural mechanisms and brain regions necessary for computing and using such priors are unknown. Here we report that the posterior parietal cortex (PPC) is a critical locus for the representation and use of prior stimulus information. We trained rats in an auditory Parametric Working Memory (PWM) task, and found that rats displayed substantial and readily quantifiable behavioral effects of sensory stimulus history, similar to those observed in humans[5,6] and monkeys[7]. Earlier proposals that PPC supports working memory[8,9] predict that optogenetic silencing of the PPC would lead to a behavioral impairment in our working memory task. Contrary to this prediction, silencing PPC produced a significant performance improvement. Quantitative analyses of behavior revealed that this improvement was due to the selective reduction of the effects of prior sensory stimuli. Electrophysiological recordings showed

that PPC neurons carried far more information about sensory stimuli of previous trials than about stimuli of the current trial. Furthermore, the more information about previous trial sensory history in the neural firing rates of a given rat's PPC, the greater the behavioral effect of sensory history in that rat, suggesting a tight link between behavior and PPC representations of stimulus history. Our results indicate that the PPC is a central component in the processing of sensory stimulus history, and open a window for neurobiological investigation of long-standing questions regarding how perception and working memory are affected by prior sensory information.

Finding long-term regularities in the environment, and exploiting them, is a critical brain function in a complex yet structured world. But little is known about the neural mechanisms involved in estimating the regularities, or involved in their impact on memory. The history of sensory stimuli is known to affect working memory (WM)[10,11], as it does in many other tasks involving sensory percepts[12,13]. One salient example, discovered over a century ago[14] and then repeatedly observed in human cognition[5,14,15] is known as "contraction bias," which has been conceptualized as an effect in which the representation of a stimulus held in working memory shifts towards the center of the distribution of stimuli observed in the past (the "prior distribution"). Despite the ubiquity of this phenomenon, and much psychophysical and theoretical research into the use and effects of prior stimulus distributions[2,3], the neural mechanisms of contraction bias have not been identified.

Based on previous work using somatosensory stimuli[6], and inspired by Parametric Working Memory (PWM) tasks in primates[7], we developed a computerized protocol to train rats, in high-throughput facilities, to perform a novel auditory PWM task (behavioral shaping code available

at http://brodylab.org/auditory-pwm-task-code). PWM tasks involve the sequential comparison of two graded (i.e., analog) stimuli that are separated by a delay period of a few seconds, here two auditory pink noise stimuli, 's_a' and 's_b'; rats were rewarded for correctly reporting which of the two was the louder stimulus (**Fig. 1a**). Following[16], the set of [s_a, s_b] pairs used across trials in a session was chosen so that neither s_a nor s_b, taken alone,

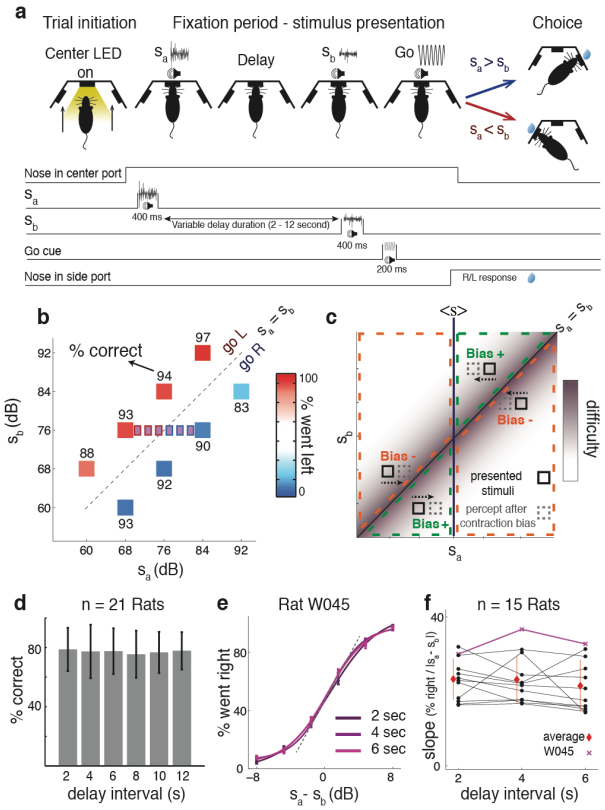

Figure 1. Rat performance and contraction bias in an auditory Parametric Working Memory Task. **a**, Schematic of a task trial. Trained rats inserted their nose in a central port until an auditory 'Go' signal indicated release. During this "nose fixation" period, a 400-ms auditory noise stimulus 's_a' was presented, followed by a multi-second variable-duration silent delay, then by a second auditory noise stimulus 's_b'. If the overall loudness of s_a was greater than that of s_b, the subjects were rewarded for poking into the right port after the 'Go' signal; if s_a<s_b, they were rewarded for poking left. **b**, Stimulus set and performance of an example rat, averaged over sessions and delay intervals ranging from 2 to 8 s. The colored boxes represent the set of [s_a, s_b] pairs used in a session. On each trial, a pair was selected at random and presented to the rat. Small purple squares show [s_a,s_b] pairs used in a subset

of sessions to assess performance at psychometric threshold. **c**, Contraction bias schematic. Presented stimuli (black boxes) are thought to produce a percept (grey boxes), and drive behavior, as if the initially presented s_a were closer to the average stimulus <s> (vertical midline) than its presented value. For some [s_a,s_b] pairs this will decrease the difference between s_a and s_b, and thus impair performance (Bias-, red) while for others it will have the opposite effect (Bias+, green). **d**, Performance, averaged across n = 21 rats and across all stimuli, as a function of delay duration. Error bars represent standard deviation over rat subjects. **e**, Psychometric curves for one example rat (fits to a four-parameter logistic function; see Methods), for three different delay durations. Dashed line shows the tangent to the fitted sigmoid at the midpoint, used to calculate the slopes shown in the panel (f). See **Extended Data Fig. 2** for other 14 rats. **f,** Midpoint slopes of psychometric curves for each animal, as a function of delay duration. Red is the average and standard deviation over rats. Magenta indicates the example rat in (e).

contained sufficient information to solve the task (**Fig. 1b**). As with any magnitude discrimination task, the smaller the difference between s_a and s_b, the harder the task (**Fig. 1c**). Classical contraction bias[5] argues that during the delay interval, the memory of the magnitude of s_a drifts towards the mean of all stimuli presented in the task (**Fig. 1c**, vertical line labeled "<s>"). Consequently, on those [s_a, s_b] pairs in which s_a drifts away from the high difficulty s_a=s_b diagonal, the memory of s_a would become more distinct from s_b and thus contraction bias would cause performance to improve ("Bias+" regions with green dashed outline in **Fig. 1c**). Whereas in pairs where s_a drifts towards the diagonal, performance would decrease (red "Bias-", **Fig. 1c**). This predicted pattern can be seen in our rat behavior (**Fig. 1b,** high % correct for Bias + stimuli [s_a, s_b] = [84, 92] and [68, 60], lower % correct for Bias - stimuli [s_a, s_b] = [60, 68] and [92, 84]). The same pattern has been observed in monkeys[17] and humans (Extended data **Fig 1d-e**, and[5,6,18]). History-dependent effects are likely adaptive in the natural world, where there are many long-term regularities. But in our laboratory task, in which each trial is generated independently, such history biases lead, on average, to suboptimal performance.

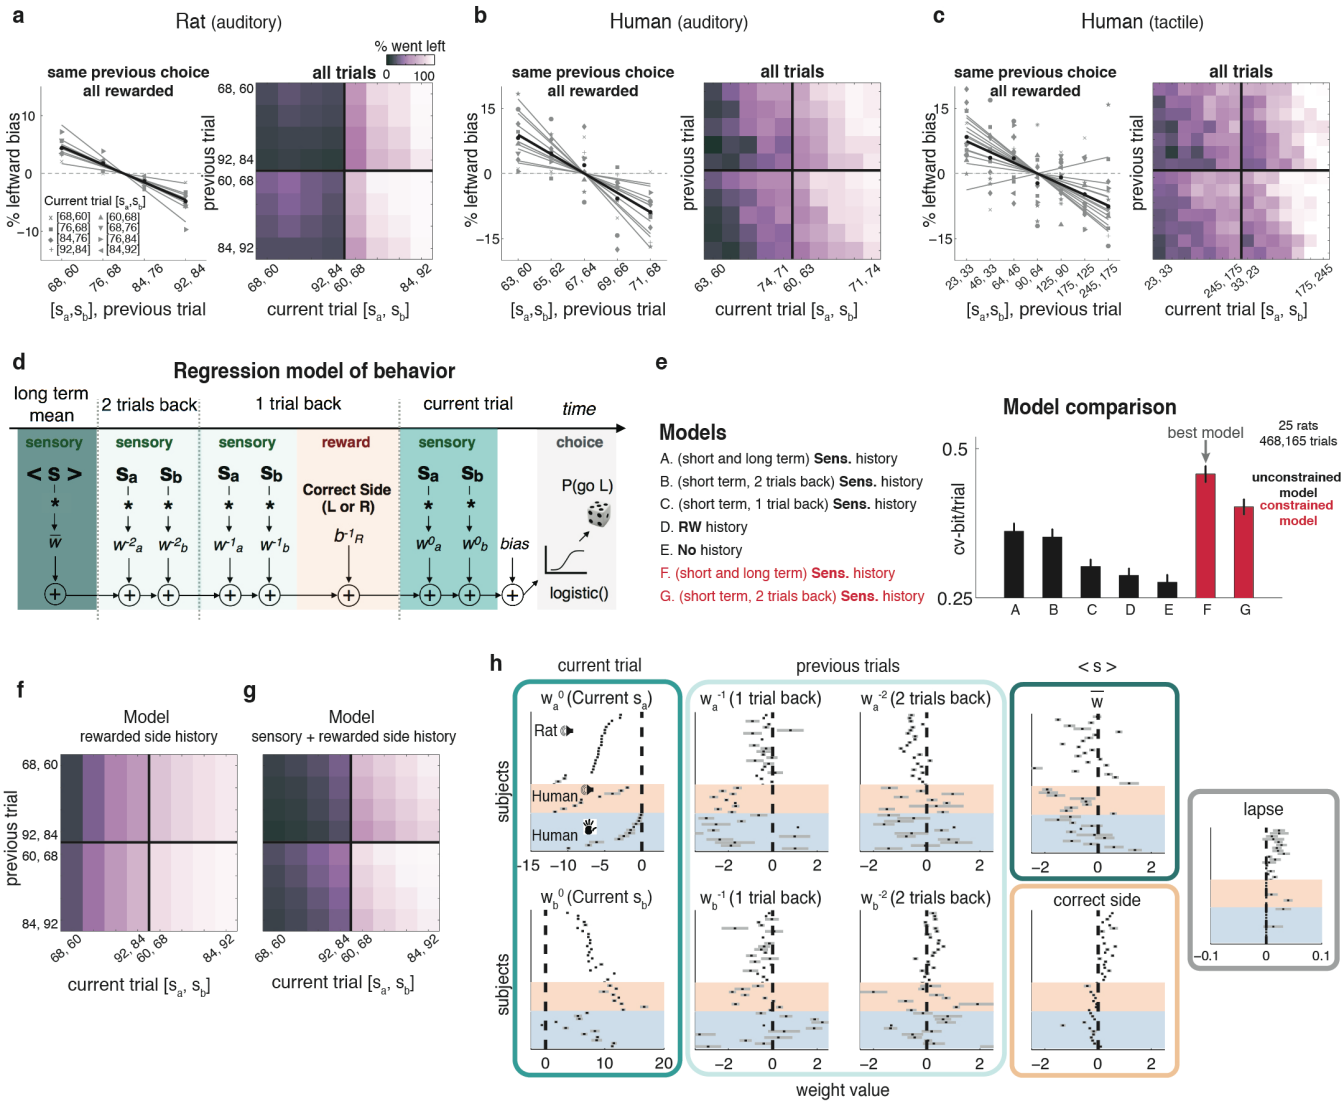

Figure 2. Sensory history biases behavior. a, Left, behavioral bias as a function of previous trial's stimuli, for fixed previous trial response choice and reward. Grey lines are different current trial [s_a,s_b] pairs, black line is average over pairs. Right, Stimulus History Matrix, showing percentage of trials with a left response for all combinations of current and previous stimuli. Modulation along the vertical indicates a previous trial effect. **b**, Same format as in (a), data from human subjects engaged in a similar auditory PWM task (n = 11). **c**, Same format as in (a), human tactile data (n = 14). **d**, The logistic regression model of behavior. Weights on a linear weighted sum of 9 regressors are used to predict the logarithmic probability ratio $\log(P_{go\ Left} / P_{go\ right})$, with weights fit to best match training data, and evaluated on left-out cross-validation data (Methods). Regressors are: average of sensory stimuli over the last few tens of trials (but excluding last two trials); each of stimuli s_a and s_b from each of the last two trials;

correct (i.e., "winning") side on the last trial, which is sufficient to drive a win-stay/lose-shift strategy; current trial s_a and s_b; and an overall side bias. **e**, Evaluation of a variety of model variants with different regressors on left-out cross-validation data (Methods; see Extended Data Fig. 6 for full set of models tested). Model performance is quantified in bits/trial (see Methods). Leftmost model: Model with regressors as in panel (d). Moving rightwards, regressors progressively removed from the models are: long-term sensory history; each of last two trials' sensory stimuli s_a and s_b (short-term sensory history); previous trial correct side. Next two models in crimson have the same regressors as the first two models, but the weights on the current trial's s_a plus all previous sensory stimuli weights are constrained to add to 1, thus removing one free parameter. Best-performing model is the constrained model with regressors shown in panel (d). **f**, Stimulus History Matrix predicted by model with current trial regressors and previous trial correct side only is a poor match to the data in (a). See **Extended Data Fig. 6c** for quantification of matrix similarity. **g,** As in (f), now for a model also including sensory history regressors, as in (d) is a much better match to panel (a). **h,** Summaries of values of best-fit parameter values over all subjects. Black ticks are best-fit values; gray bars span the 95% CIs. Each panel has been divided by task (pink highlight for human auditory task, blue highlight for human tactile task) and then sorted based on the parameter value of w^0_a (similar ordering is used for other panels). Note the comparatively small strength of previous correct side regression weight (see **Extended Data Fig. 6d** for comparison between sensory history versus correct side regression weights).

The overall performance of our rats was robust and similar across memory delay intervals ranging from 2 s to 12 s (**Fig. 1d**; see **Extended Data Fig.1 b-c** for performance over the 3 to 4-month long course of learning). In a subset of sessions, we included [s_a, s_b] stimulus pairs that were closely spaced along s_a (small purple squares in **Fig. 1b**), and used these to measure the psychometric discrimination threshold (**Fig. 1e**). This threshold did not vary significantly across delays intervals ranging from 2 to 6 seconds (P=0.7, **Fig. 1f**).

In contrast to the negligible effects induced by varying the memory delay interval (**Fig. 1d,f**), we found that a strong effect was induced by the history of sensory experience. As will be quantified below, this sensory history effect was stronger than the well-documented influence of previous rewards and choices [19–21] (see also **Extended Data Fig.1 g-i**). Examining only trials in which in the immediately previous trial the subject had gone to the right and had received a

reward, and therefore holding previous reward and choice fixed, **Fig. 2a, left** shows that the smaller the previous trial's s_a and s_b stimuli, the greater the percentage of leftward choices in the current trial (slope = -3.06 percent per decibel, $p < 0.0001$, see **Extended Data Fig. 5a** for slopes from leftward biases from n=1,..,7 trials back). This is consistent with a contraction bias in which the estimate of <s> is weighted towards recent stimuli[18], because small values of recent stimuli would then make the current s_a more likely to be perceived as small, increasing the likelihood of an "$s_a < s_b$ (go Left)" response. **Fig. 2a, right** shows the same effect across all combinations of current and previous trial stimuli from our standard stimulus set (see **Extended Data Fig. 3** for Sensory History Matrices for n=1,..,5 trials back, and **Extended Data Fig. 4** for full set of such trial history matrices, controlled for action and reward). Similar effects were found in human versions of the task using auditory (**Fig. 2b**) or tactile (**Fig. 2c**) sensory modalities. To simultaneously take into account effects across multiple previous trials of the history of rewards, choices, and sensory stimuli, we fit logistic regression models with these variables as regressors, and compared the performance of a variety of such models on cross-validation data (**Fig. 2d;** Methods and **Extended Data Fig. 6**). Consistent with human data[18], short-term (last two trials) sensory history had strong effects on behavior. In addition, our large dataset revealed a smaller but nevertheless important effect of longer-term (average of last few tens of trials) sensory history (**Fig. 2e**). In a weighted average, all weights sum to 1. Constraining the sum of regression weights on the current trial's s_a stimulus plus weights on previous sensory stimuli to be equal to 1, thus effectively removing one parameter from the model, led to the best performance on cross-validation data (**Fig. 2e**, red). This is consistent with the contraction bias proposal that sensory history does not add a behavioral bias that is independent of working memory, but instead produces a value of s_a held in working memory that is a weighted average of

the current stimulus and sensory history[22–27] (**Fig. 2e**, see **Extended Data Fig. 6** for full set of models and comparison between them). Sensory history was critical for accounting for behavior (**Fig. 2f,g**): examining the weights in the regression model shows that the weights for sensory history terms are larger than those for the rewarded side history term (**Fig. 2h**; see also **Extended Data Fig. 6d**).

The PPC has been proposed as critical for working memory ([8,9,28] but see [29,30]), and we therefore examined its role in our working memory task. We injected bilaterally an AAV virus that drives expression of the light-activated inhibitory opsin halorhodopsin eNpHR3.0, under the CaMKIIa promoter (center of injection located at AP -3.8mm and ML 2.5mm from Bregma, **Fig 3a, Extended Data Fig. 7a**). Optical fibers were inserted at the centers of the injection sites to deliver laser illumination, and we inactivated PPC during a randomly-chosen 20% of trials. To best probe for any small effects, we included psychometric stimuli (purple squares, **Figs. 1b** and **3b**). Expecting a performance impairment[8], we were surprised to observe instead an improvement in psychometric performance in all the animals tested (**Fig. 3c**). However, the effect was not simply an overall performance improvement: looking beyond the psychometric stimuli,

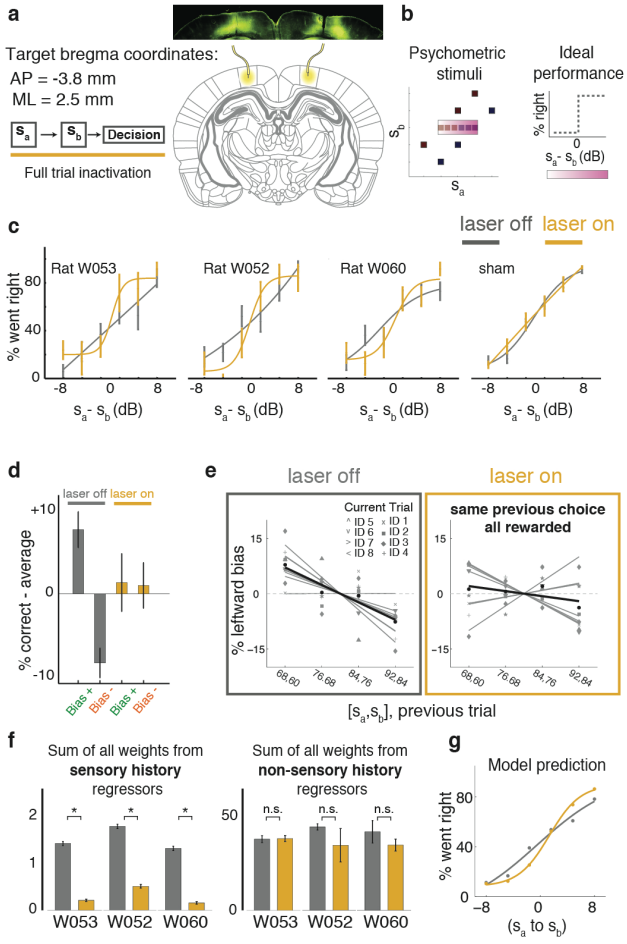

Figure 3. PPC is specifically necessary for behavioral effect of previous sensory stimuli. a, Schematic indicating injection of virus expressing the inhibitory opsin eNpHR3.0 in bilateral PPC and laser stimulation on 20% of randomly chosen trials. **b**. Stimuli used in these optogenetic silencing experiments included both our standard set of $[s_a,s_b]$ pairs (black) as well as stimuli used to measure psychometric performance (purple). Inset at right: psychometric performance of an ideal agent. **c**, Psychometric curves for individual animals for control trials (cyan) and trials with PPC inactivated (yellow). In all animals tested, PPC inactivation led to a psychometric curve closer to the ideal performance. Right, sham inactivation in rats implanted with optic fibers with no virus expressing eNpHR3.0 had no effect (n=2). **d**, Percent correct averaged across all Bias+ trials or all Bias- trials (see **Fig. 1c**), relative to overall average performance. PPC inactivation eliminates the difference between performance in Bias+ trials, versus Bias- trials, or versus the overall average (laser off: "Bias +" - "Bias -" = 14.29, $p < 0.00001$; laser on: ("Bias +" - "Bias -" = 1.79, $p = 0.706$; laser off vs laser on: $p = 0.0027$) **e**, The bias induced by previous trial sensory stimuli is reduced under PPC inactivation (laser off trials: slope of -4.74, $p = 0.0017$; laser on trials: slope of -1.36

p = 0.42; laser on vs. laser off: p=0.044; equal number of trials for two conditions) **f**, PPC inactivation selectively reduces the sensory history weights (left panel) in the regression model of **Fig. 2d.** Error bars show 95% confidence interval. Right panel includes sum of all the other weights for non-sensory-history regressors, which are not significantly affected by the inhibition. See **Extended Data Fig. 9a** for comparison of all individual weights. **g**, Reducing the sensory history weights in the model is sufficient to improve psychometric performance in a manner comparable to the experimental data of panel (c).

while performance with respect to control on Bias- trials was indeed improved, performance on Bias+ trials was impaired by PPC silencing (**Fig. 3d, Extended Data Fig. 9b**). Moreover, the effect was such that the difference between Bias+ and Bias- trials was eliminated, as was their difference with respect to the average performance (**Fig. 3d**). Similarly, the bias as a function of the previous trial's sensory stimuli was markedly reduced (**Fig 3e,** laser off: p=0.42; laser on: p=0.0017; laser on vs. laser off: p=0.044, see **Extended Data Fig. 8** for the impact of inhibition on the sensory history matrices). Fitting our regression model separately to the set of laser on versus laser off trials, we found that sensory history regression weights were significantly reduced when PPC was inactivated (**Fig. 3f, Extended Data Fig. 9c**), and no other regression terms were significantly affected (each individual weight shown in **Extended Data Fig. 9a**). A model in which sensory history effects were reduced to the levels shown in **Fig. 3f** was sufficient to reproduce the improvement in psychometric trials observed experimentally (**Fig. 3g**). Thus PPC silencing appeared to have no impact on working memory but instead induced a specific and substantial reduction of sensory history effects.

To examine whether signatures of sensory history are present in the region whose inactivation appears to cancel the effect of history, we conducted extracellular recordings during task performance, from a total of 936 units in five animals implanted with microwire arrays, targeting PPC. Neurons with a mean firing rate of less than 2Hz were discarded from analyses

(see Methods), resulting in total of 361 units analyzed. Again in contrast to expectations of a role for PPC in working memory, most cells were similar to the example cell shown in **Fig. 4a**, in that their firing rates during the working memory delay period did not distinguish between different values of s_a, the stimulus held in memory, and therefore did not carry information about it. Instead, robust information about the stimulus pair appeared approximately 1 second after the trial had terminated, during the inter-trial-interval (ITI). We used mutual information (MI; see Methods) to quantify the amount of information carried, in neuronal firing rates, about which [s_a, s_b] sensory stimulus pair was presented (**Fig. 4b,c**; see **Extended Data Fig. 10** for MI about other task components such as current and previous previous choices or rewards). During the ITI prior to the start of a new trial, a large fraction of PPC neurons carried significant information about the stimuli presented in the previous trial (22% of analyzed neurons, red curve in **Fig. 4c**). A smaller fraction of cells continued to code the previous trial's stimuli into the start of the new trial (blue curve in **Fig. 4c**). We computed the fraction of neurons with significant MI about the previous trial's stimuli, both during the ITI (**Fig 4d**) as well as during the current trial (**Fig 4e**), and compared this to the strength of the rat's sensory history behavioral bias (**Fig 4f, Extended Data Fig. 10e**). During the new trial, but not during the ITI, these two measures were perfectly correlated (**Fig 4g**, Spearman's rank correlation $r=1$, $p<0.01$ during current trial; $r=0.3$, $p=0.68$ during ITI; $p<0.00001$ ITI vs. full current trial from Steiger's Z-test). This suggests first, a tight link between sensory history representations in PPC and sensory history behavioral biases, and second, that the PPC sensory history representation is used during or shortly after the presentation of the new trial's s_a (**Fig. 4h**), consistent with the idea that contraction bias is an effect on the representation of s_a.

Parametric working memory tasks, with their readily quantifiable behavior, are well-suited to investigating the effect of sensory history on perception and behavior. Rodent versions of these tasks, combined with semi-automated training as used here, are an efficient platform for causal and cellular-resolution investigation of the neural mechanisms involved. Our work provides this platform, and using it,

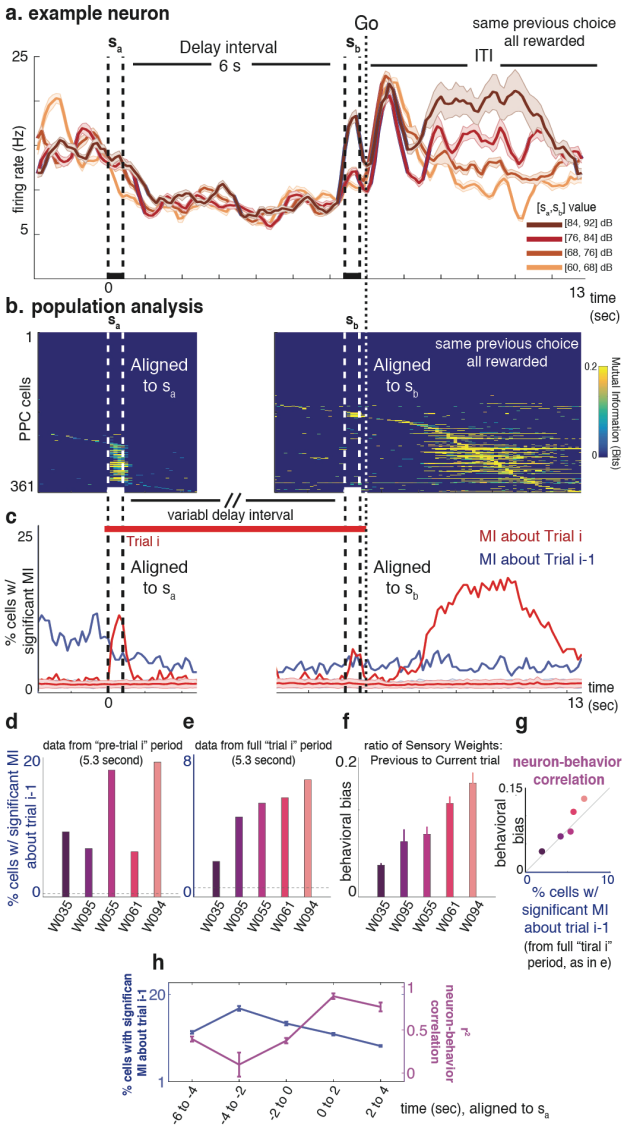

Fig. 4. PPC neurons carry more information about previous trial sensory stimuli than about current trial sensory stimuli, and it predicts behavioral bias. a, Firing rate of an example neuron in response to different

values of s_a. Only trials in which the animal responded to the left after the Go cue and was rewarded are shown in this panel, and for clarity only trials with 6 s delay interval are shown. **b**, Population analysis. Each row represents one neuron, and shows, as a function of time, the Mutual Information (MI) between the cell's firing rate and the sensory stimulus pair [s_a,s_b]. Data from trials with delay intervals ranging from 2 to 6 s were used. The left and right parts of the plot are aligned to the start of s_a and the start of s_b, respectively. Only MI values that are significantly larger than the shuffled distribution ($p < 0.005$) are included, non-significant values are blanked out in dark blue. To control for reward and side choice, MI values were calculated using only trials with fixed choice and reward, and only then averaged across the different, separately calculated reward and choice groups. **c**, Summary of the population analysis, showing percentage of cells with significant coding of stimuli presented on trial i (in red), or on trial i-1 (in blue). **d**, Percent of cells with significant MI about the previous trial, for each of the five recorded animal, calculated over a 5.3 seconds period during the ITI prior to the start of the new trial. Horizontal dashed line shows % of cells expected by chance from the shuffled data. **e**, As in d, but data is from the 5.3 seconds duration of the new trial. **f**, Behavioral bias for individual animals, calculated as the relative sensory history weights (to weights for s_a and s_b regressors of trial i) from the best model fit. **g**, Data from panel (f) plotted against data from panel (e), Spearman's rank correlation $r = 1.0$, $p<0.01$. **h.** Calculation as in panels d-g, but using narrower 2 s-long time windows. The high neuron-behavior correlation appears concurrently with the presentation of the new trial's s_a.

revealed the PPC as an essential node in both the representation and causal effects of sensory stimulus history. This opens a critical window towards a cellular resolution understanding of long-standing questions about how sensory stimulus history affects working memory and perception. Important issues that can now be addressed include precisely how sensory history representations in PPC interact with current stimulus representations so as to modulate perception, how sensory history information reaches PPC, and which other brain regions connected to PPC are also essential nodes of the relevant circuit.

Author contributions. AA and CDB conceived the project. AA carried out all experiments and analyzed the data, with the optogenetic inactivations carried out together with CDK. Human tactile data was gathered in MED's laboratory. AA and CDB wrote the manuscript, based on a first draft by AA, with extensive comments by CDK and MED.

Methods

Rat Subjects. A total of 30 male Long–Evans rats (*Rattus norvegicus*) between the ages of 6 and 24 months were used for this study. Of these, 25 were used for behavioral assessments, 5 rats were used for neural recordings, and 5 for optogenetic inactivations. All statistical tests were made between groups with similar sample sizes. Investigators were not blinded to experimental groups during data collection or analysis. Animal use procedures were approved by the Princeton University Institutional Animal Care and Use Committee and carried out in accordance with National Institutes of Health standards.

Human Subjects (auditory). 11 human subjects (8 males and 3 females, ages 22-40) were tested and all gave their informed consent. Participants were paid to be part of the study and were naive to the main conclusions of the study. The consent procedure and the rest of the protocol were approved by the Princeton University Institutional Review Board.

Human Subjects (tactile). 14 human subjects (8 males and 6 females, ages 22–35) were tested. Protocols conformed to international norms and were approved by the Ethics Committee of the International School for Advanced Studies. Subjects signed informed consent.

Rat Behavior. We developed a computerized protocol to train rats, in high-throughput facilities, to perform an auditory delayed comparison task, adapted from a tactile version[6]. All trainings happen in three-port operant conditioning chambers, where ports are arranged side-by-side along one wall, and with two speakers, placed above the right and left nose ports. Figure 1a shows the task structure. A visible LED in the center port signals the availability of each trial. Rat subjects initiate a trial by inserting their nose into the center port that causes the center LED to turn off. Rats must keep their nose in the center port ("fixation" period) until an auditory "go" cue signals the end of fixation. Only after the "go" cue, subjects can withdraw and orient to one of the side pokes in order to receive water reward. During the fixation period, two auditory stimuli, 's_a' and 's_b', separated by a variable delay, are played for 400 ms. There are short delay periods of 250 ms inserted before 's_a' and after 's_b'. Stimuli consist of broadband noise (2K-20K Hz), generated as a series of Sound Pressure Level (SPL) values sampled from a zero-mean normal distribution. Overall mean intensity of sounds vary from 60-92 dB. Rats should judge which of s_a and s_b had greater SPL standard deviation. If $s_a > s_b$ then the correct action is to poke in the right side poke in order to collect reward, and if $s_a < s_b$ rats have to orient to the left side poke. Trial durations are independently varied on a trial-by-trial basis, by varying the delay interval between the two stimuli, that can be as short as 2 s or as long as 12 s. Rats progressed through a series of shaping stages, before the final version of the delayed comparison task, in which they learned to

1) associate light in center poke with availability of trials 2) associate sounds with reward 3) maintain their nose in the center poke until they hear an auditory "go" signal and 4) compare the two s_a and s_b stimuli.

Human Auditory Behavior. Similar auditory stimuli to those used for rats were used in the human version of the task. In this experiment, subjects received on each trial a pair of sounds played from the earphone. First sound was presented together with a green square on the left side. Then there was a delay period, indicated by "WAIT!" on the screen, then the second sound was presented together with a red square on the right side. Subjects, at the end of the second stimulus after the go cue, were required to compare the two sounds and decide which one was louder, then indicate their choice by pressing a key with their right or left hand. Written feedback about the correctness of their response was provided on the screen.

Human Tactile Behavior. Human subjects performed the tactile version of the task. The details of this task have been previously described and the behavior has been characterized[6]. Briefly, at each trial two noisy vibration stimuli, interleaved with a variable delay interval, were delivered to the subject's fingertip. Subjects viewed a computer monitor and wore headphones that presented acoustic noise and eliminated ambient sounds. To start a trial, the subject pressed the keyboard up arrow with the right hand. This triggered presentation of the two stimuli. After a post-stimulus delay, a blue panel illuminated on the monitor, and the subject pressed the left or right arrow on the keyboard, signifying selection of the first or the second stimulus, respectively. They received feedback (correct/incorrect) on each trial through the monitor. Human experiments were controlled using LabVIEW software (National Instruments, Austin, Texas).

Stimulus Set. If the first stimulus, s_a were fixed across all trials and only the second stimulus s_b changed, subjects might solve the task by ignoring the first stimulus and applying a constant threshold to the second stimulus. Likewise, if the second stimulus were fixed, subjects might apply a constant threshold on the first stimulus. To prevent such alternative strategies, it is necessary to vary both s_a and s_b, and use a set of stimuli composed of pairs of s_a and s_b which guarantees that across trials the same value of SPL standard deviation is randomly presented for the first stimulus or the second stimulus. Figure 1b represents such stimulus set. A minimum of 8 pairs of stimuli span a wide range of SPL standard deviation values (Figure 1b). Using this stimulus set, if the subject were to ignore either of "s_a" or "s_b", then the maximum performance would be 63%. The mean amplitudes of stimuli were evenly distributed in a logarithmic scale (linear in dB). The diagonal line represents $s_a = s_b$; all stimulus pairs on one side of the diagonal were associated with the same action, and all have the same ratio of s_a to s_b. For each trial, one of these 8 pairs of stimuli is randomly selected to determine s_a and s_b.

Psychometric curves

Psychometric plots (as shown in Fig. 1e and Fig. 3c) show the probability of the subject responding leftward as a function of the difference between s_a and s_b when s_b is fixed. The fits were to a 4-parameter logistic function of the form

$$y(x) = y_0 + \frac{a}{1 + e(\frac{-(x-x_0)}{b})}$$

Where $y_0$is the left endpoint, $y_0 + a$us the right endpoint, $x_0$is the bias, and $a/4b$ is the slope. Fits were non-linear least square regression done using the nlinfit.m function from Matlab2013.

Regression Model of Behavior

Semi-automated training facilitated the generation of a behavioral data set comprising 468,165 trials from 25 animals, which in turn enabled statistical characterization of the decision-making process. In order to quantify rats' behavior we carried out an analysis to "weigh" the contributions of s_a , s_b, of the current trial and several trials in the past, as well as the history of choice and reward to the animal's choice, as follows. From the data originating in a single training session, for each [s_a , s_b] stimulus pair, we fit animal choice with a logistic regression model. This model allows linear combinations of s_a and s_b and other desired factors. The linear combination is then mapped nonlinearly into the animal choice i.e. probability of trials in which the subject judged $s_a > s_b$, through a logistic function as:

$$P(Left) = lapse + \frac{1 - 2lapse}{1 + e^{(-A)}}$$

Where

$$A = \sum_{t=n}^{0} (s_a^{-t} W_a^{-t} + s_b^{-t} W_b^{-t} + Rw^{-t} + \beta)$$

Where W_a^{-t}andW_b^{-t} are the s_b and s_b regressors, respectively, from t trials back. Rw^{-t} is the correct side on trial $-t$: left = +1, right = -1. This regressor captures win-stay/lose-switch strategy. β is the baseline regressor that captures the overall (stimulus-independent) bias of the subject in calling $s_a > s_b$ (for instance, a bias against turning right, the side associated with the

judgment $s_a > s_a$). The absolute values of all the regressors are normalized between 0 and 1. We used the loglikelihood as the cost function:

$$Cost function = -(yLog(p)+(1-y)Log(1-p) = -\frac{1}{m}[\sum_{i}^{m} y_i Log(p_i) + (1-y_i)Log(1-p_i)]$$

The model was fit using a gradient descent algorithm to minimize the negative log likelihood function cost function. We used the "sqp" algorithm of fmincon function from Matlab 2013. Weights were calculated using L2-regularization to prevent over-fitting. The hyperparameter value (lambda) was selected independently for each rat using evidence optimization, based on 5-fold cross-validation. Different variants of the model, that systematically studies relevance of various sensory and reward history factors, are discussed in the Extended Data Fig. 6.

Model comparison and cross-validation

All models were fit separately for each individual rat (n = 25), using 200 runs of 5-fold cross-validation. For each run we calculated the log likelihood of the test dataset given the best-fit parameters on the training set (Logl). We also calculated the log likelihood of the test dataset for the mean value of %Left (the experimentally-measured fraction of trials in which the animal went left). This gives us a null log likelihood reference value ($Logl_0$). In order to quantify the efficiency of each model we defined the "cross-validated bit/trial" (CV-bit/trial) as the trial-averaged excess likelihood of the model compared to the null model[31]:

$$CV - bit/trial = \frac{(Logl - Logl_0)/n_{trials}}{Log(2)}$$

For each model, we first chose the right regularization value (lambda) that would maximize the CV-bit/trial. To compare different models we calculated the median value of CV-bit/trial across 10000 fits for each subject.

To compare models with different number of parameters, we used two common metrics known as the Bayesian and Akaike Information Criterion (BIC and AIC). They are defined as:

$$BIC = -2Logl + kln(n)$$

$$AIC = -2Logl + 2k$$

Where Logl is the maximum log likelihood of the model with k parameters on n data points.

Optogenetic virus injection and fibre implantation

For optogenetic perturbation experiments, the general surgery techniques and fiber etching follow previous reports [32], except that we began construction with a standard off the shelf 50/125 μm LC-LC duplex fibre cable (http://www.fibercables.com), instead of the usual FC-FC duplex fiber cables. The cable jacket, strengthening fibres, and outer plastic coating (typically white or orange) were fully removed leaving 1 cm of fibre optic cable and inner plastic coating (typically clear) intact. Then 2 mm of the fibre tip (with final layer of plastic coating still attached) was submerged in 48% hydrofluoric acid topped with mineral oil for 85 min, followed by water for 5 min (submerging 5 mm), and acetone for 2 min (to soften the plastic). The plastic coating was then gently cut with a razor and pulled off with tweezers to reveal a 1 mm sharp-etched fibre tip. Enough plastic was removed, depending on the depth of the targeted site, to ensure that only the glass fibre optic would be inserted into the brain.

For viral injection, 2 µl of adeno-associated virus (AAV) (AAV5-CaMKIIα-eNpHR3.0-eYFP), that drives expression of the light-activated inhibitory opsin halorhodopsin eNpHR3.0, under the CaMKIIa promoter, coupled to eYFP, was lightly dyed with fast green powder and front loaded into a glass pipette mounted to a Nanoject (Drummond Scientific) prefilled with mineral oil. The pipette tip was manually cut to ~30 µm diameter. Five closely spaced injection tracts were used with each animal. For the central injection tract, one injections of 23 nl were made every 100 µm in depth starting 100 µm below brain surface for PPC for 1.5 mm. Four additional injection tracts were completed, using procedures identical to the central tract, one 500 µm anterior, posterior, medial and lateral from the central tract. Each injection was followed by a 10 s pause, with 1 min following the final injection in a tract before the pipette was removed. A total of 1.5 µl of virus was injected over a 30-min period consisting of ~160 separate injections. A chemically sharpened fibre optic (50 µm core, 125 µm cladding) was then lowered down the central injection tract to a depth of 1 mm. The craniotomy was filled with kwik-sil (World Precision Instruments), allowed to set for 10 min, and the fibre optic was secured to the skull with C&B Metabond and dental acrylic. Dental acrylic covered all the incision site and allowed only the LC connector to protrude. Halorhodopsin expression was allowed to develop for 6 weeks before behavioural testing began.

Optogenetic perturbation

The animal's implant was connected to a 1-m patch cable attached to a single fibre rotary joint (Princetel) mounted on the ceiling of the behavioural chamber. This was connected to a 200 mW, 532 nm laser (OEM Laser Systems) operating at 25 mW, which was triggered with a 5 V transistor–transistor logic (TTL) pulse. Laser illumination occurred on a random 20% of trials.

See **Extended Data Fig. 7** for physiological confirmation of optogenetic inactivation effects in an anesthetized animal.

Recordings

5 animals were operated to implant microwire arrays in their left or right PPC (n=2 in rPPC, n=3 in lPPC, see **Extended Data Fig. 7** for histological localization of electrodes). The target region was accessed by craniotomy, using standard stereotaxic technique (centered 3.8 mm posterior to bregma and 2.5 mm lateral to the midline). Dura mater was removed over the entire craniotomy with a small syringe needle. The remaining pia mater, even if usually not considered to be resistant to penetration, nevertheless presents a barrier to the entry of the microelectrode arrays, due to high-density arrangement of electrodes in multi-channel electrode arrays. The term dimpling is commonly used to describe the situation where the electrodes are pushing the brain cortex in without penetrating. Obviously, the more electrodes there are in an array, the more pronounced this effect becomes. In addition to potentially injuring the brain tissue, dimpling is obviously a source of error in the determination of depth measurements. Ideally, if dimpling could be eliminated, the electrodes would move in relation to the pial surface, allowing effective and accurate electrode placement. After the craniotomy is made, and the dura is carefully removed over the entire craniotomy, a petroleum-based ointment such as bacitracin ointment or sterile petroleum jelly is applied to the exact site of electrode implantation. The cyanoacrylate adhesive is then applied to the zone of the pia surrounding the penetration area. This procedure fastens the pia mater to the overlying bone and the resulting surface tension prevents the brain from compressing under the advancing electrodes. Once the polymerization of cyanoacrylate adhesive happened, over a period of few minutes, the petroleum ointment at the target site is

removed, and the 32 electrode microwire array (Tucker-Davis Technologies (TDT), Alachua FL) was inserted by slowly advancing a Narashige micromanipulator. After inserting the array(s), the remaining exposed cortex was covered with biocompatible silicon (kwik-sil, World Precision Instruments), and the microwire array was secured to the skull with C&B Metabond and dental acrylic.

During the 10 days of recovery time, rats had unlimited access to water and food. Recording sessions in the apparatus began thereafter. Extracellular activity of PPC was manually sorted into single units and multiunits, based on the spike waveform and the refractory period observed in interspike interval histogram, using spike3 software, and verified later using a MATLAB-based software, UltraMegaSort 2000. In total 936 single or multiunits were recorded in PPC of 5 rats. Only neurons with overall firing rate within the session was at least 2 Hz were included in the analysis. These neurons summed to total of 361.

Neural Analysis

Mutual Information. In order to quantify the type and amount of information that PPC neurons carry about various task parameters we computed Shannon's Mutual Information [33]. In this formulation, the amount of information which can be extracted from the firing rate of a neuron R, about the task-related parameter X can be computed as:

$$I(X,R) = \sum_{x} P(x) \sum_{r} P(r,x) Log_2 \frac{P(r|x)}{P(r)}.$$

Where *P(r|x)* is the conditional probability of observing a neuronal response r given the presentation of the task parameter x, *P(r)* is the marginal probability of occurrence of neuronal

response r among all possible responses, and *P(x)* is the probability of task parameter x. Information measured in this way quantifies how well an ideal observer can discriminate between members of a stimulus set based on the neuronal responses of a single trial [34]. For each trial, neuronal response was defined as the rate of spiking during time windows of 100 ms. The conditional probability in the above formula is not known a priori and must be estimated empirically from a limited number, *N*, of experimental trials for each stimulus. Limited sampling of response probabilities can lead to an upward bias in the estimate of mutual information [35]. In order to correct for the bias, we used combination of two techniques. First we estimated and corrected the bias based on Quadratic Extrapolation (QE) method [36], that assumes the bias can be accurately approximated as second order expansions in 1/N. Then we used bootstrap procedure that consists of many rounds of pairing stimuli and responses at random in order to destroy all the information that the responses carry about the stimulus. Due to limited data sampling, the information computed using the bootstrapped responses may still be positive. The average value of the bootstrapped information was then used to estimate the residual bias of the information calculation, and was subtracted out. Moreover, the distribution of bootstrapped information values were used to build a non-parametric test of whether the corrected information computed using QE method is significantly different from zero[37].

Code and data availability.

All software used for behavioral training is available on the Brody lab website at http://brodylab.org/auditory-pwm-task-code. Software used for data analysis, as well as raw and processed data, are available from the authors upon reasonable request.

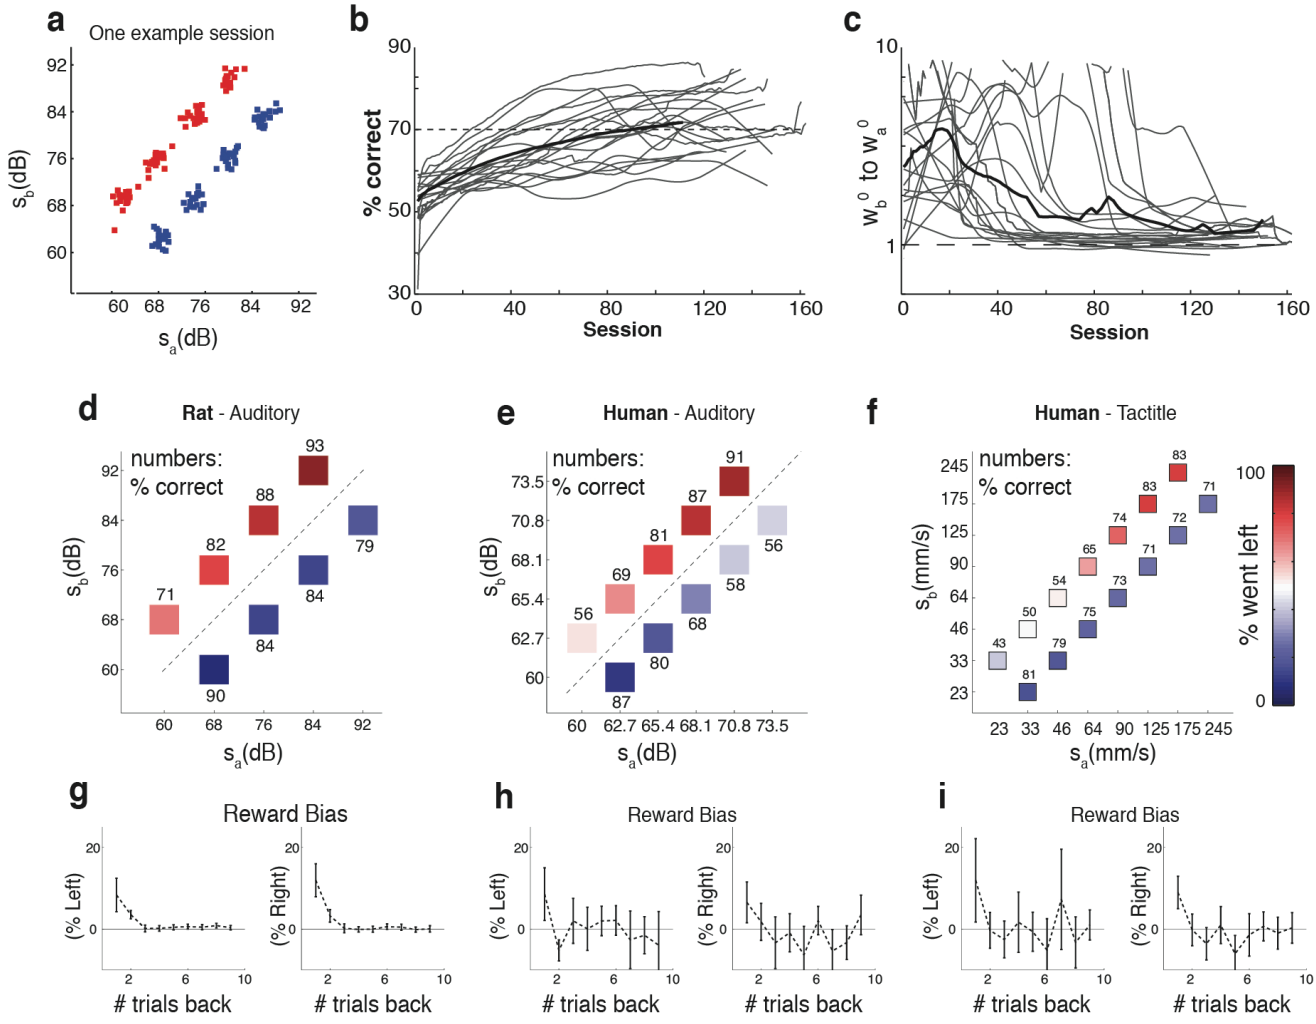

Extended Data Fig. 1, Full stimulus set - Learning curve - Mean performance - Reward bias

a, Each stimulus is made of a series of Sound Pressure Level (SPL) values sampled from a zero-mean normal distribution, and standard-deviation value of s. For each trial, SPL values are randomly drawn and therefore due to sampling statistics, the actual standard deviation value of the stimulus always differed slightly from its designated value. The coordinates of each small box represent the actual joint values of [s_a,s_b] for one sample training session. **b,** Individual gray lines shows learning curves presented as the change in % correct over months of training, for N = 25 rats. Average rat (black line) reaches 70% of performance after 90 sessions. **c**, Learning curve presented as the ratio of the best fit weights for the second stimulus, s_b, to the first stimulus, s_a, using the model described in Figure 2e (3-parameter, "No history" version). **d,** Rat auditory working memory performance, data from 21 rat subjects (total of 468,165 trials) are separated by [s_a,s_b] pair but averaged across subjects and over different delay durations

(2-8 sec). **e**, Human auditory working memory performance. For human, interstimulus delay varied randomly from 2 s to 6 s. (11 subjects, 12623 trials). **f**, Human tactile working memory performance. Similar to panel "e" but for humans engaged in the tactile version of the task. In this task, interstimulus delay varied randomly from 2 s to 8 s. Data from 14 human subjects (total of 4694 trials) are pooled together. **g**, Reward history bias. Y-axis shows the difference of performance on "turn right" trials, when the "k" trials back was rewarded either on the left or right. Data from n = 21. Errorbars show 95% confidence interval. **h-i**, similar to (g) for human auditory (h) and tactile (i) PWM task.

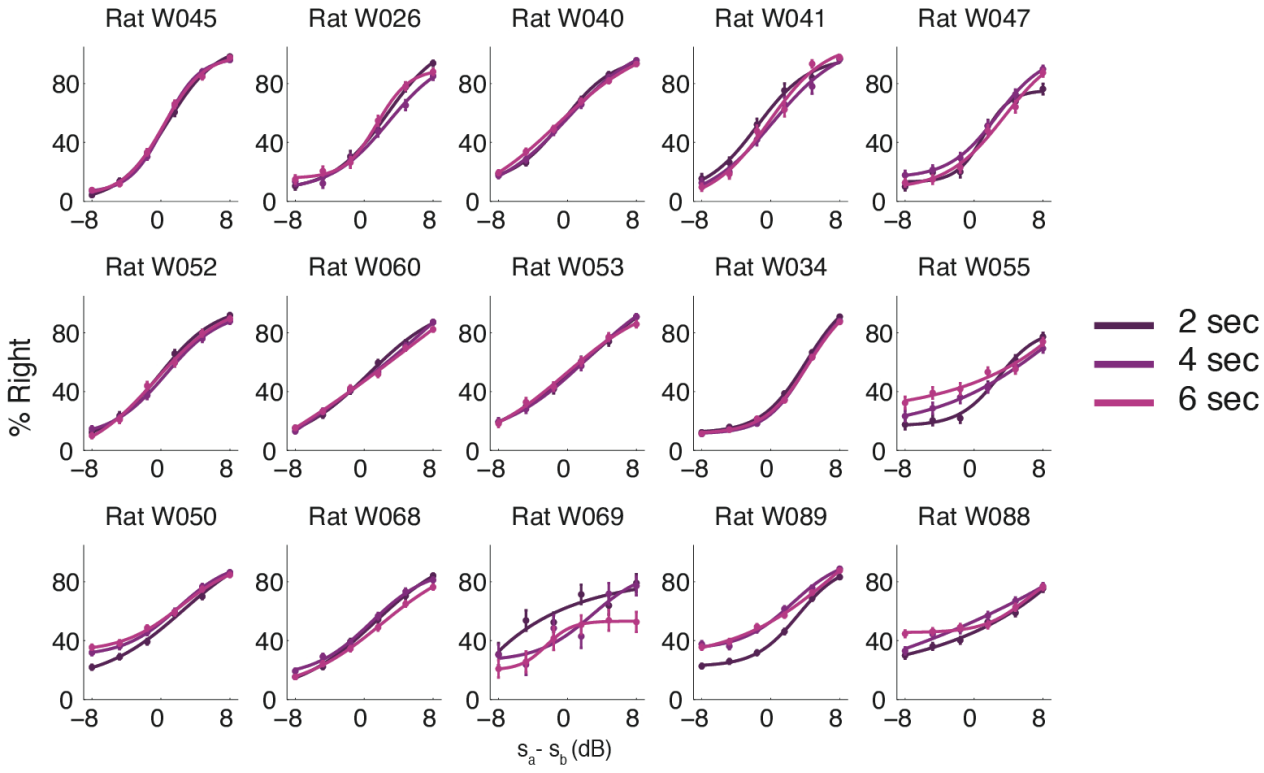

Extended Data Fig. 2, Psychometric figures, individual rats

Psychometric curves (fits to a four-parameter logistic function for one 15 rats; see Methods), for three different delay durations.

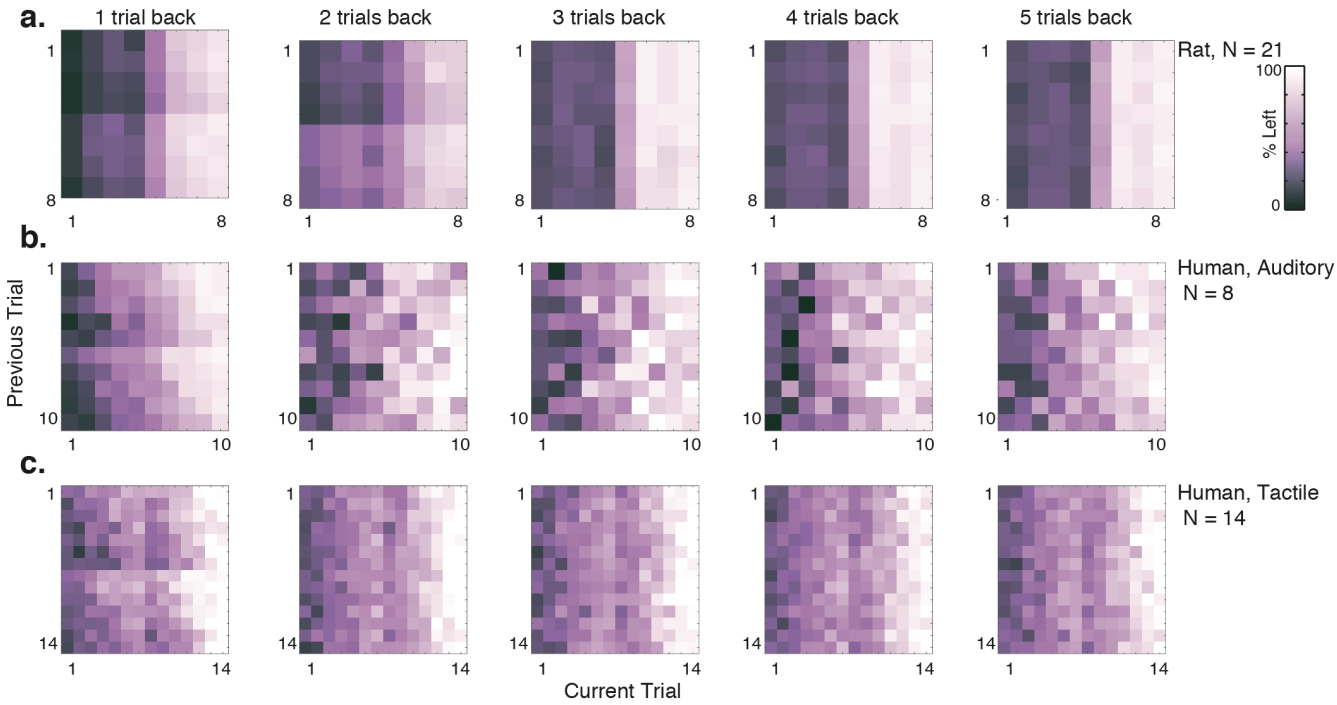

Extended Data Fig. 3, Sensory History matrix, from 1 to 5 trials back

a, Stimulus history matrix, as described in Figure 2b, when %left is shown given any combination of the stimuli in current trial (x-axis) and n-trials back (y-axis), n = 1,2,3,4,5. Data from N = 21 rats, composing total of 381,612 trials is used in this analysis. **b**, Similar to panel "a", for human auditory task. **c**, Similar to panel "a", for human tactile task.

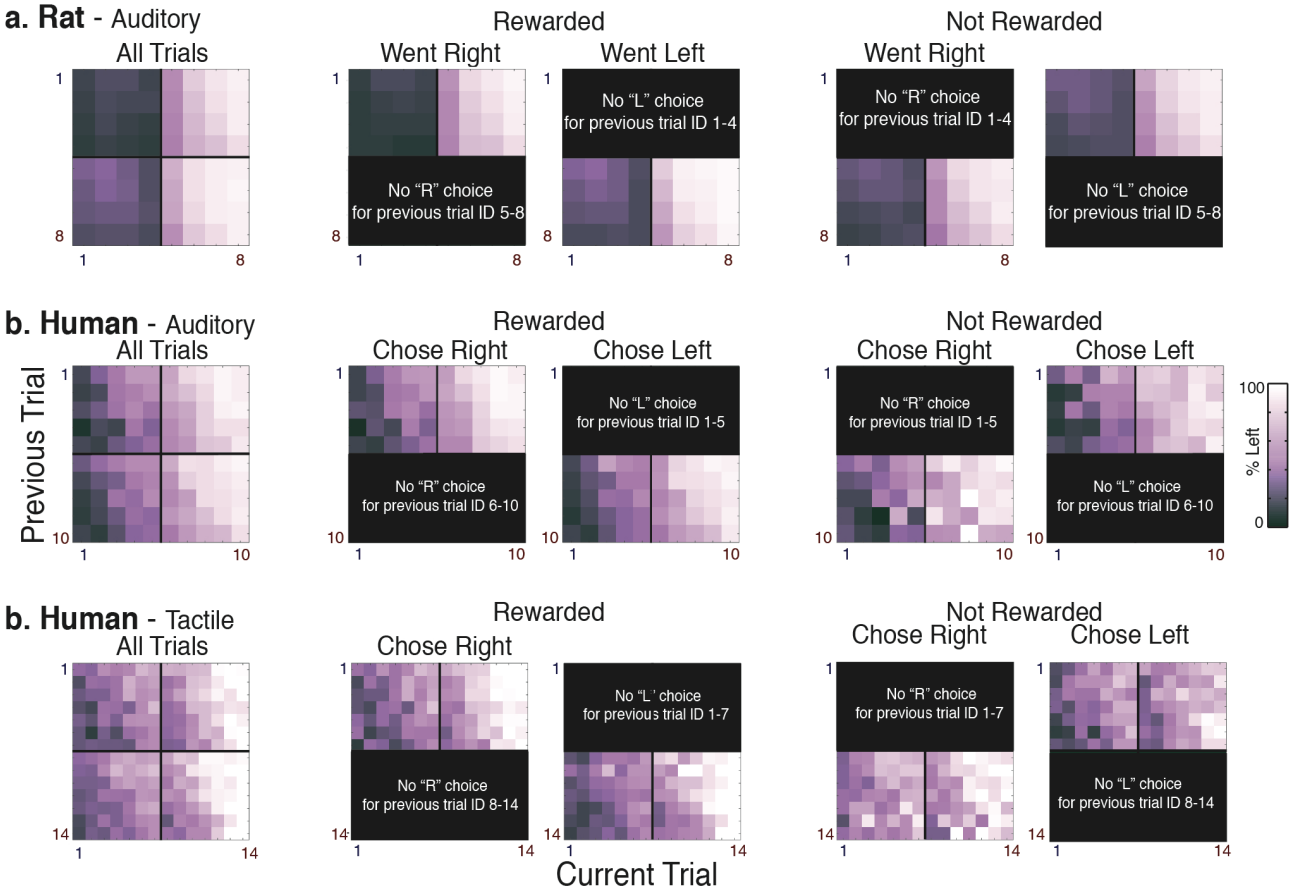

Extended Data Fig. 4, Sensory History matrix, controlled for Reward and Choice

Similar to Extended Figure 2, except that in this plot only trials for which the previous trial was a "turn right" trial and the animal was rewarded are included. Therefore, modulation by previous trial cannot be due to "action" or "reward" history.

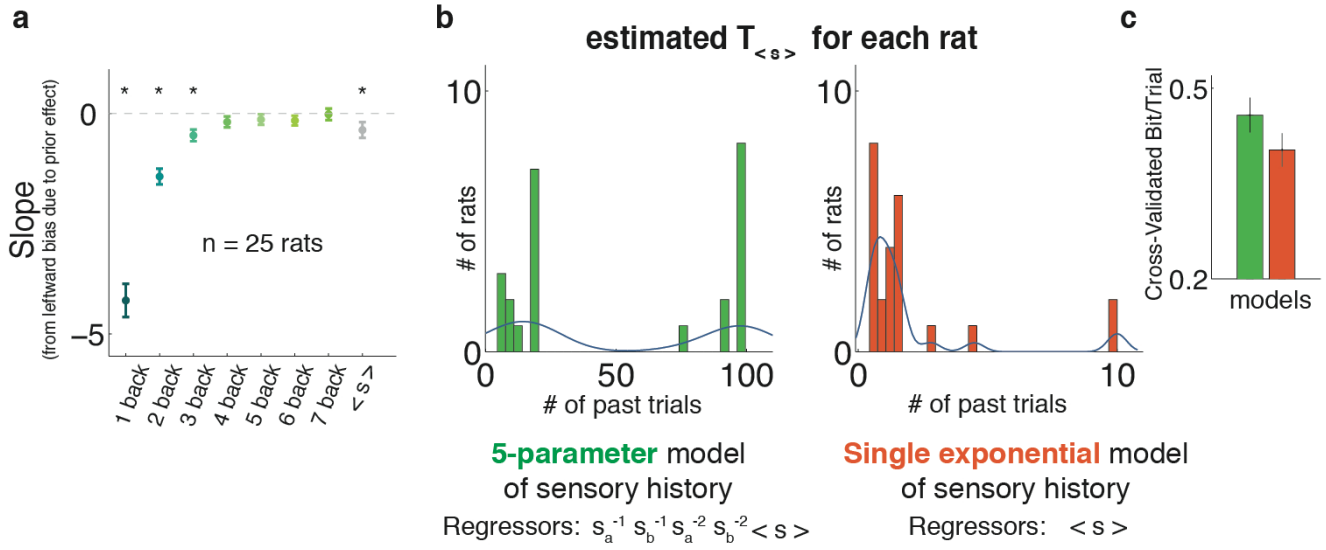

Extended Data Fig. 5, Estimating optimal window of < s >

a, Slopes from linear fit to the %leftward bias from n-back trials (n=1:7, as in **Fig. 2a** where n=1 was used), and also < s > which is a window of 17 trials, from n=4 to n=20, in gray. Errorbars show 95% confidence intervals. Data from n = 25 rats. **b**, For each rat the optimal exponential window over the past trials was estimated such that it would maximize the cross-validation bit/trial measurement. Two models are compared here: green shows the distribution of taus from a model that has 5 regressors to account for the sensory history - first and second stimulus from the 2 trials back and a separate exponential window over the remaining past trials (Fig 2d). In orange, instead, only one regressor which is a single exponential window over all the past trials accounts for the sensory history. In the single-exponential model, the best fit value of tau comes out very small, practically as if only past 1 or 2 trials back are inducing most of the effect. **c**, The 5-parameter model of sensory history outperforms the single-exponential model.

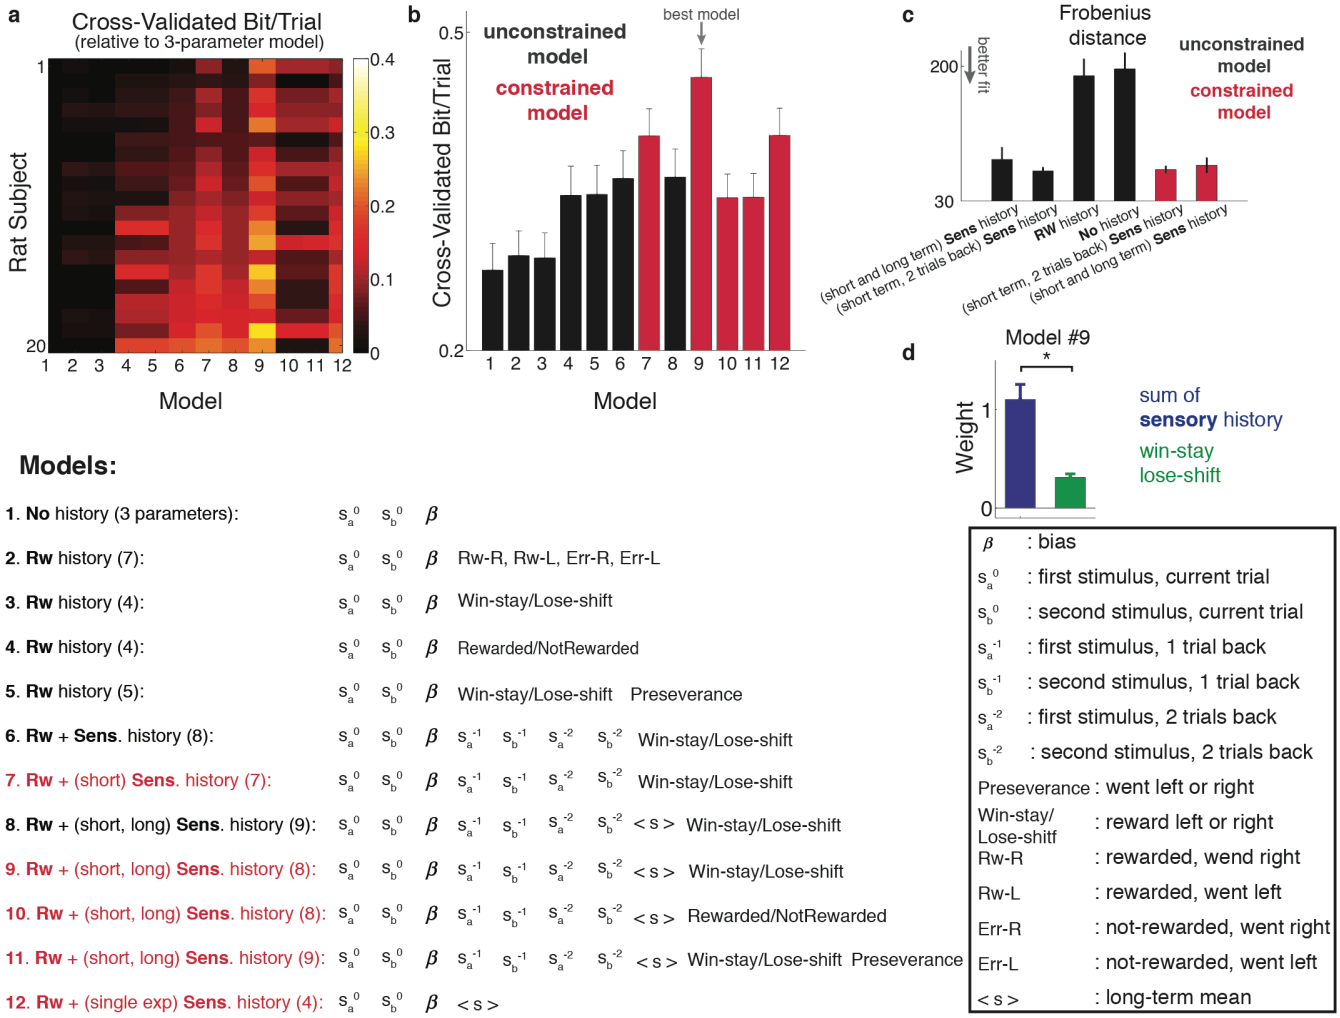

Models:

1. **No** history (3 parameters):	s_a^0	s_b^0	β						
2. **Rw** history (7):	s_a^0	s_b^0	β	Rw-R, Rw-L, Err-R, Err-L					
3. **Rw** history (4):	s_a^0	s_b^0	β	Win-stay/Lose-shift					
4. **Rw** history (4):	s_a^0	s_b^0	β	Rewarded/NotRewarded					
5. **Rw** history (5):	s_a^0	s_b^0	β	Win-stay/Lose-shift	Preseverance				
6. **Rw** + **Sens**. history (8):	s_a^0	s_b^0	β	s_a^{-1}	s_b^{-1}	s_a^{-2}	s_b^{-2}	Win-stay/Lose-shift	
7. **Rw** + (short) **Sens**. history (7):	s_a^0	s_b^0	β	s_a^{-1}	s_b^{-1}	s_a^{-2}	s_b^{-2}	Win-stay/Lose-shift	
8. **Rw** + (short, long) **Sens**. history (9):	s_a^0	s_b^0	β	s_a^{-1}	s_b^{-1}	s_a^{-2}	s_b^{-2}	$< s >$	Win-stay/Lose-shift
9. **Rw** + (short, long) **Sens**. history (8):	s_a^0	s_b^0	β	s_a^{-1}	s_b^{-1}	s_a^{-2}	s_b^{-2}	$< s >$	Win-stay/Lose-shift
10. **Rw** + (short, long) **Sens**. history (8):	s_a^0	s_b^0	β	s_a^{-1}	s_b^{-1}	s_a^{-2}	s_b^{-2}	$< s >$	Rewarded/NotRewarded
11. **Rw** + (short, long) **Sens**. history (9):	s_a^0	s_b^0	β	s_a^{-1}	s_b^{-1}	s_a^{-2}	s_b^{-2}	$< s >$	Win-stay/Lose-shift Preseverance
12. **Rw** + (single exp) **Sens**. history (4):	s_a^0	s_b^0	β	$< s >$					

β	: bias
s_a^0	: first stimulus, current trial
s_b^0	: second stimulus, current trial
s_a^{-1}	: first stimulus, 1 trial back
s_b^{-1}	: second stimulus, 1 trial back
s_a^{-2}	: first stimulus, 2 trials back
s_b^{-2}	: second stimulus, 2 trials back
Preseverance	: went left or right
Win-stay/ Lose-shitf	: reward left or right
Rw-R	: rewarded, wend right
Rw-L	: rewarded, went left
Err-R	: not-rewarded, went right
Err-L	: not-rewarded, went left
$< s >$	: long-term mean

Extended Fig. 6, Model comparison

a, Model comparisons, 200 runs of 5-fold cross validation were done, on data from each rat, in order to find the best fit parameters and to compare different model fits using the "Cross-Validated Bit/Trial" quantity defined as the relative value of the log likelihood of each model, to the null log likelihood, normalized in log2. Removing one parameter by constraining the regression weights on the current trial's s_a stimulus plus the weights on previous sensory stimuli to add to 1 (constrained model, in crimson) improved performance on cross-validated data compared to the unconstrained model (in black). **c**, to compare the Sensory History Matrix from the real data to the ones predicted from the best model fits, Frobenious distance norm was used, defined as the square root of the sum of the absolute squares of the difference between elements of two matrices. Frobenious distance is a measure of similarity and the smaller the value, the more similar the two matrices. **d**, examining the weights in the regression model #9,

which is determined to be the best model, shows that the weights for sensory history terms are larger than those for the rewarded side history term.

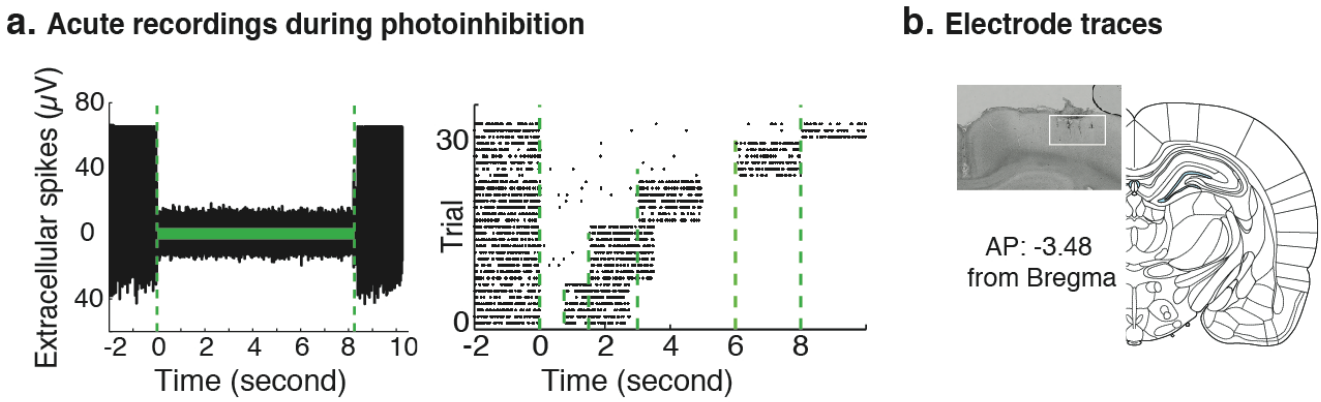

Extended Data Fig. 7

a, Physiological confirmation of optogenetic inactivation effect in an anesthetized animal. Left, acute extracellular activity of an example cell in the PPC, expressing eNpHR3.0, is shown in response to light stimulation. Laser illumination period (8 s) is marked by the light green bar. Right, raster-plot for 32 trials, for variable durations of light stimulation. **b**, histological localization of electrodes targeting PPC. The inset shows example of electrode locations in a coronal slice at AP=3.48 from Bregma. In all cases, the electrode and fiber placements in PPC were within between 2.8 and 4 mm anterior to Bregma and between 2 and 3.5 mm lateral to the midline.

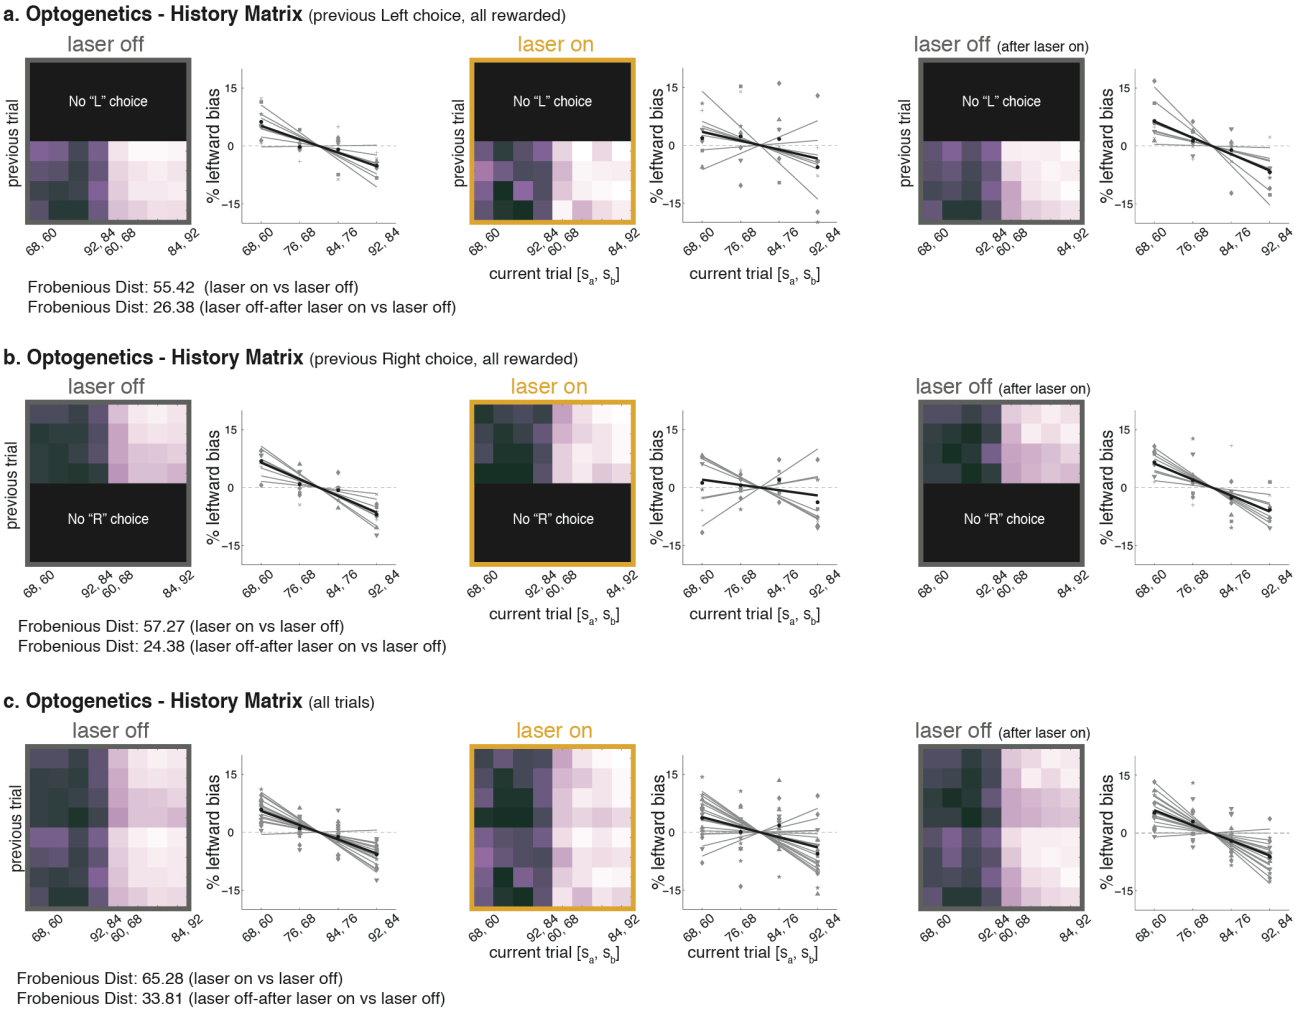

Extended Data Fig. 8, Optogenetics - PPC inhibition reduces leftward bias due to past sensory stimuli

a, Sensory History Matrix, and leftward biases due to past sensory stimuli, similar to **Figure 2a-c**, but now for three types of trials: "laser off" trials (two leftmost panels) which consist of trials with no PPC inactivation on either "current" or "previous" trial; "laser on" trials (two middle panels) which consist of trials with PPC inactivation on "current" trial; "laser off after laser on" trials (two rightmost panels) which consist of trials immediately after the "laser on" trials. This last set controls for number of trials, as it contains equal number of trials to "laser on" condition. Modulation along the vertical indicates a previous trial effect behavioral bias as a function of previous trial's stimuli, for trials for which animals went Left, and were rewarded, therefore history of reward and choice is held fixed. Grey lines are different current trial [s_a,s_b] pairs, black line is average over pairs. **b**, similar to (a) for trials for which animals went Right and were rewarded. **c**, similar to (a) for all combinations of current and previous stimuli.

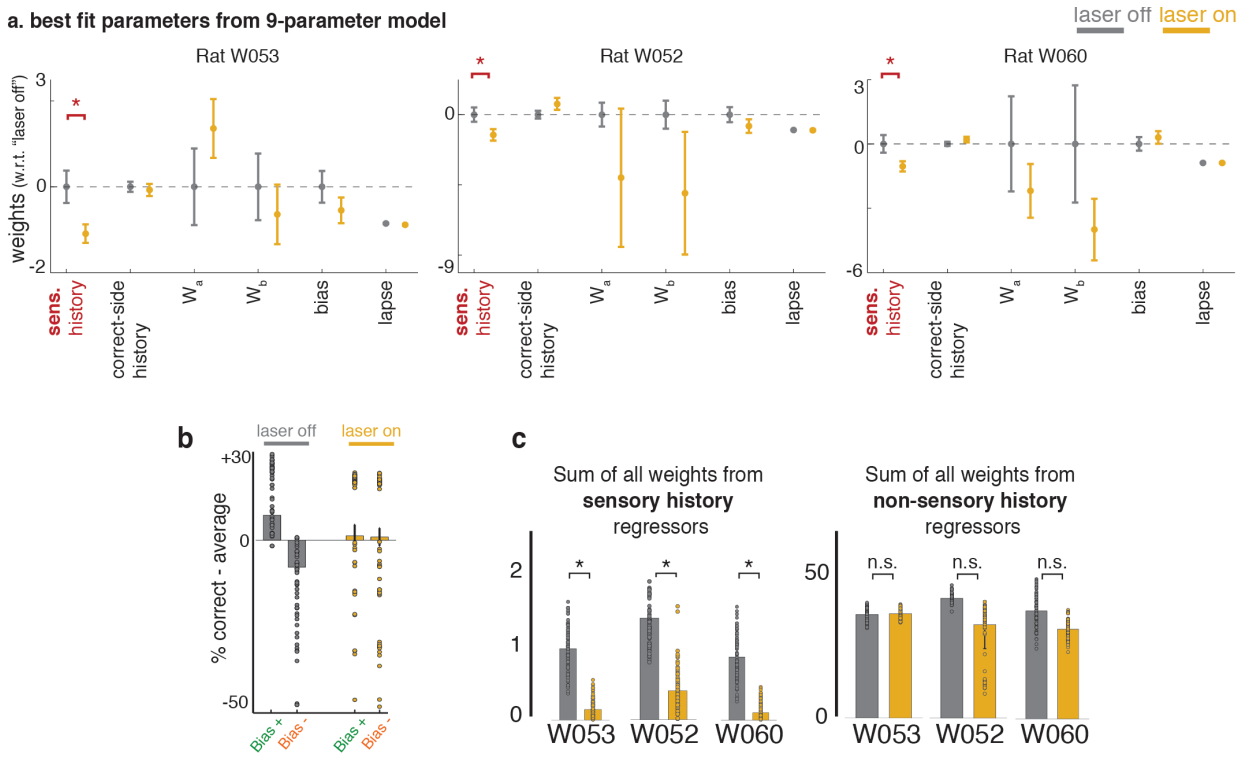

Extended Data Fig. 9, Optogenetics - Best fit parameters for non-sensory-history weights, Data distribution for Fig3d-f

a, Best fit parameter values for non-"sensory history" weights, from the 9-parameter model (short-term, **Sens** history model, constrained version, figure 2d-e). Except for the “sensory history”, none of the other weights were significantly affected by optogenetic inactivation of PPC. Errorbars show standard deviation of the mean. **b-c**, similar to Fig. 3d-f, with all data points overlaid on the bar-plots.

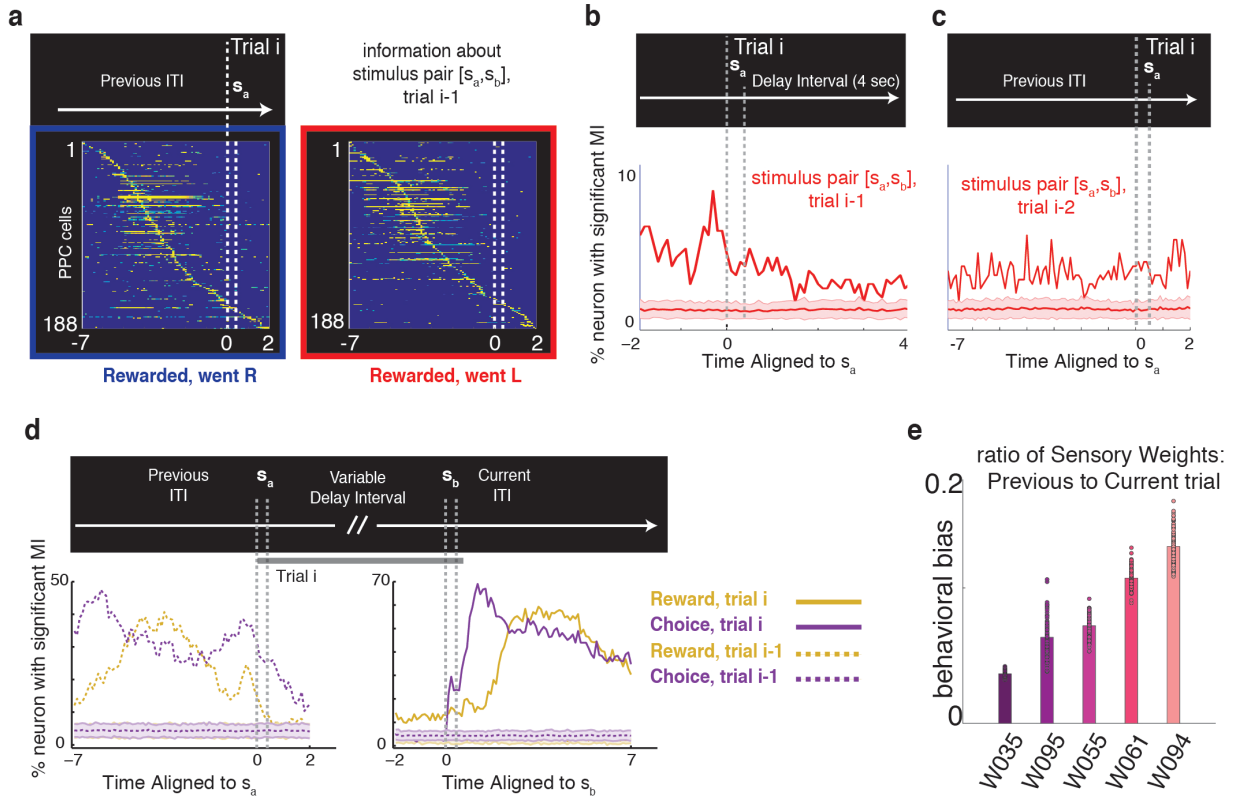

Extended Data Fig. 10, Mutual Information

a, **Sensory-history coding, 1 trial back, population analysis**, each row represents the time course of significant values of Mutual Information (MI) between cell's firing rate and the stimulus pair [s_a,s_b] presented on the previous trial. Data from all trials with variable delay duration (minimum of 2 sec) was pooled and plots are aligned to the beginning of s_a. Data from n=5 animals, and only cells with significant values of MI values are included. When estimating the MI, spurious information values can be caused by the inherent correlations between task parameters, like sensory stimuli and choice. To overcome this, conditional MI was calculated when only trials with same previous choice and reward status were considered, and sensory inputs were the only variable: Left panel, on the previous trial animals went right and were rewarded. Right panel, on the previous trial animals went left and were rewarded. **b**, **Sensory-history coding, 1 trial back**, % of cells with significant coding of stimuli presented on the previous trial (trial i-1), aligned to the start of trial i. Only trials with delay interval of larger 4 seconds are included in this analysis. **c**, **Sensory-history coding, two trials back**, % of cells with significant coding of stimuli presented two trials in the past trial (trial i-2), aligned to the start of trial i. **d**, % of cells with significant coding of animal's

choice and reward status, on both current trial (solid lines) and previous trial (dashed lines), when time is aligned to the current trial. **e**, Similar to Fig. 4f, with all data points overlaid on the bar-plots.

1. Steyvers, M., Griffiths, T. L. & Dennis, S. Probabilistic inference in human semantic

memory. *Trends Cogn. Sci.* **10,** 327–334 (2006).

2. Chater, N., Tenenbaum, J. B. & Yuille, A. Probabilistic models of cognition: conceptual foundations. *Trends Cogn. Sci.* **10,** 287–291 (2006).
3. Pouget, A., Beck, J. M., Ma, W. J. & Latham, P. E. Probabilistic brains: knowns and unknowns. *Nat. Neurosci.* **16,** 1170–1178 (2013).
4. Schapiro, A. & Turk-Browne, N. Statistical Learning. in *Brain Mapping* 501–506 (2015).
5. Ashourian, P. & Loewenstein, Y. Bayesian inference underlies the contraction bias in delayed comparison tasks. *PLoS One* **6,** e19551 (2011).
6. Fassihi, A., Akrami, A., Esmaeili, V. & Diamond, M. E. Tactile perception and working memory in rats and humans. *Proc. Natl. Acad. Sci. U. S. A.* **111,** 2331–2336 (2014).
7. Romo, R. & Salinas, E. Flutter discrimination: neural codes, perception, memory and decision making. *Nat. Rev. Neurosci.* **4,** 203–218 (2003).
8. Harvey, C. D., Coen, P. & Tank, D. W. Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature* **484,** 62–68 (2012).
9. Carandini, M. & Churchland, A. K. Probing perceptual decisions in rodents. *Nat. Neurosci.* **16,** 824–831 (2013).
10. Visscher, K. M., Kahana, M. J. & Sekuler, R. Trial-to-trial carryover in auditory short-term memory. *J. Exp. Psychol. Learn. Mem. Cogn.* **35,** 46–56 (2009).
11. Lockhead, G. R. & King, M. C. A memory model of sequential effects in scaling tasks. *J. Exp. Psychol. Hum. Percept. Perform.* **9,** 461–473 (1983).
12. Ernst, M. O. & Banks, M. S. Humans integrate visual and haptic information in a statistically optimal fashion. *Nature* **415,** 429–433 (2002).
13. Körding, K. P. & Wolpert, D. M. Bayesian integration in sensorimotor learning. *Nature* **427,**

244–247 (2004).

14. Hollingworth, H. L. The Central Tendency of Judgment. *The Journal of Philosophy, Psychology and Scientific Methods* **7,** 461 (1910).
15. Karim, M., Harris, J. A., Langdon, A. & Breakspear, M. The influence of prior experience and expected timing on vibrotactile discrimination. *Front. Neurosci.* **7,** 255 (2013).
16. Hernandez, A., Salinas, E., Garcıa, R. & Romo, R. Discrimination in the sense of flutter: new psychophysical measurements in monkeys. *Journal of Neuroscience* **17,** 6391–6400 (1997).
17. Romo, R., Brody, C. D., Hernández, A. & Lemus, L. Neuronal correlates of parametric working memory in the prefrontal cortex. *Nature* **399,** 470–473 (1999).
18. Raviv, O., Ahissar, M. & Loewenstein, Y. How Recent History Affects Perception: The Normative Approach and Its Heuristic Approximation. *PLoS Comput. Biol.* **8,** e1002731 (2012).
19. Busse, L. *et al.* The detection of visual contrast in the behaving mouse. *J. Neurosci.* **31,** 11351–11361 (2011).
20. Abrahamyan, A., Silva, L. L., Dakin, S. C., Carandini, M. & Gardner, J. L. Adaptable history biases in human perceptual decisions. *Proc. Natl. Acad. Sci. U. S. A.* **113,** E3548–57 (2016).
21. Nogueira, R. *et al.* Lateral orbitofrontal cortex anticipates choices and integrates prior with current information. *Nat. Commun.* **8,** ncomms14823 (2017).
22. Hellström, Å. The time-order error and its relatives: Mirrors of cognitive processes in comparing. *Psychol. Bull.* **97,** 35–61 (1985).
23. Lu, Z. L., Williamson, S. J. & Kaufman, L. Behavioral lifetime of human auditory sensory

memory predicted by physiological measures. *Science* **258,** 1668–1670 (1992).

24. Sinclair, R. J. & Burton, H. Discrimination of vibrotactile frequencies in a delayed pair comparison task. *Percept. Psychophys.* **58,** 680–692 (1996).
25. Preuschhof, C., Schubert, T., Villringer, A. & Heekeren, H. R. Prior Information biases stimulus representations during vibrotactile decision making. *J. Cogn. Neurosci.* **22,** 875–887 (2010).
26. Olkkonen, M., McCarthy, P. F. & Allred, S. R. The central tendency bias in color perception: effects of internal and external noise. *J. Vis.* **14,** (2014).
27. Papadimitriou, C., Ferdoash, A. & Snyder, L. H. Ghosts in the machine: memory interference from the previous trial. *J. Neurophysiol.* **113,** 567–577 (2015).
28. Pasternak, T. & Greenlee, M. W. Working memory in primate sensory systems. *Nat. Rev. Neurosci.* **6,** 97–107 (2005).
29. Erlich, J. C., Brunton, B. W., Duan, C. A., Hanks, T. D. & Brody, C. D. Distinct effects of prefrontal and parietal cortex inactivations on an accumulation of evidence task in the rat. *Elife* **4,** (2015).
30. Raposo, D., Kaufman, M. T. & Churchland, A. K. A category-free neural population supports evolving demands during decision-making. *Nat. Neurosci.* **17,** 1784–1792 (2014).
31. Paninski, L., Shoham, S., Fellows, M. R., Hatsopoulos, N. G. & Donoghue, J. P. Superlinear population encoding of dynamic hand trajectory in primary motor cortex. *J. Neurosci.* **24,** 8551–8561 (2004).
32. Hanks, T. D. *et al.* Distinct relationships of parietal and prefrontal cortices to evidence accumulation. *Nature* **520,** 220–223 (2015).
33. Shannon, C. E. A Mathematical Theory of Communication. *Bell System Technical Journal*

27, 379–423 (1948).

34. Rieke, F., Warland, D. & Bialek, W. Coding Efficiency and Information Rates in Sensory Neurons. *Europhys. Lett.* **22,** 151–156 (1993).
35. Treves, A. & Panzeri, S. The Upward Bias in Measures of Information Derived from Limited Data Samples. *Neural Comput.* **7,** 399–407 (1995).
36. Nemenman, I., Bialek, W. & de Ruyter van Steveninck, R. Entropy and information in neural spike trains: progress on the sampling problem. *Phys. Rev. E Stat. Nonlin. Soft Matter Phys.* **69,** 056111 (2004).
37. Ince, R. A. A., Mazzoni, A., Petersen, R. S. & Panzeri, S. Open source tools for the information theoretic analysis of neural data. *Front. Neurosci.* **4,** (2010).