

Posterior parietal cortex represents sensory history and mediates its effects on behaviour

Athena Akrami^{1,2,3}, Charles D. Kopec^{1,2}, Mathew E. Diamond⁴ & Carlos D. Brody^{1,2,3}

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 not relevant to the current task at hand. Nevertheless, the neural mechanisms and regions of the brain that are necessary for computing and using such prior experience 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 task, and found that they displayed substantial and readily quantifiable behavioural effects of sensory-stimulus history, similar to those observed in humans^{5,6} and monkeys⁷. Earlier proposals that the PPC supports working memory^{8,9} predict that optogenetic silencing of this region would impair behaviour in our working memory task. Contrary to this prediction, we found that silencing the PPC significantly improved performance. Quantitative analyses of behaviour 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 the sensory stimuli of previous trials than about the stimuli of the current trial. Furthermore, for a given rat, the more information about previous trial sensory history in the neural firing rates of the PPC, the greater the behavioural effect of sensory history, suggesting a tight link between behaviour and PPC representations of stimulus history. Our results indicate that the PPC is a central component in the processing of sensory-stimulus history, and could enable further 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. However, little is known about the neural mechanisms involved in estimating these regularities or their impact on memory. The history of sensory stimuli affects working memory^{10,11} and many other tasks involving sensory percepts^{12,13}. One salient example, discovered over a century ago¹⁴ and repeatedly observed in human cognition^{5,14,15}, is contraction bias, in which the representation of a stimulus held in working memory shifts towards the centre 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.

On the basis of previous work using somatosensory stimuli⁶, and inspired by parametric working memory (PWM) tasks in primates⁷, we developed a computerized protocol to train rats, in high-throughput facilities, to perform a novel auditory PWM task (behavioural shaping code at <http://brodylab.org/auditory-pwm-task-code>). PWM tasks involve the sequential comparison of two graded (that is, analogue) stimuli separated by a delay of a few seconds. Here we used auditory

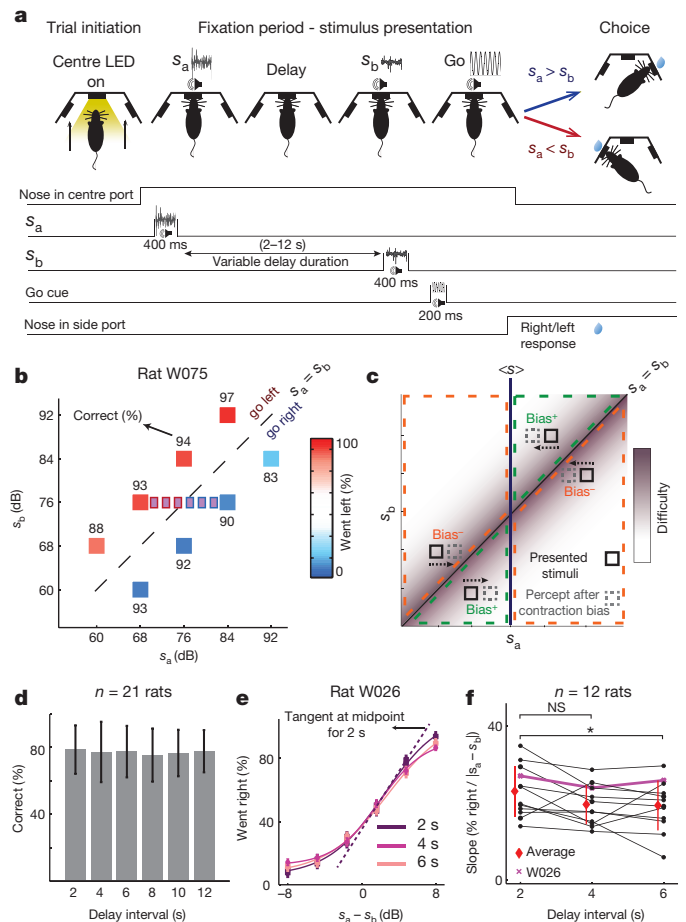


Figure 1 | Rat performance and contraction bias. **a**, Rats compared two sequentially presented auditory stimuli s_a and s_b , separated by a delay, and were rewarded for correctly reporting, through their choice of inserting their nose into the left or right ports, which stimulus was louder. **b**, Set of (s_a, s_b) pairs used in a session. In each trial, one randomly selected pair was presented. The small purple squares represent stimuli used in a subset of sessions to assess performance at the psychometric threshold. **c**, Contraction bias proposes that the presented stimuli (black boxes) drive behaviour as if s_a were closer (dashed boxes) to the average stimulus (\bar{s}) (vertical midline) than its actual value. For some (s_a, s_b) pairs, this decreases their difference, and thus impairs performance (bias⁻, red), whereas for others it has the opposite effect (bias⁺, green). **d**, Overall average performance as a function of delay duration ($n = 21$ rats, mean \pm s.d. over subjects). **e**, Psychometric curves for one example rat ($n = 120$ sessions, mean \pm s.e.m. over sessions), fits to a four-parameter logistic function; see Methods). **f**, Midpoint tangent line slopes for the psychometric curves for each of 12 rats. These are significantly greater (reflecting better performance) at delays of 2 s compared with 6 s (two sample t -test: 2 versus 4 s: $P = 0.051$; 2 versus 6 s: $P = 0.045$; 4 versus 6 s: $P = 0.86$). The mean and s.d. of the 12 rats are shown in red; the results for rat W026 are shown in magenta. NS, not significant.

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

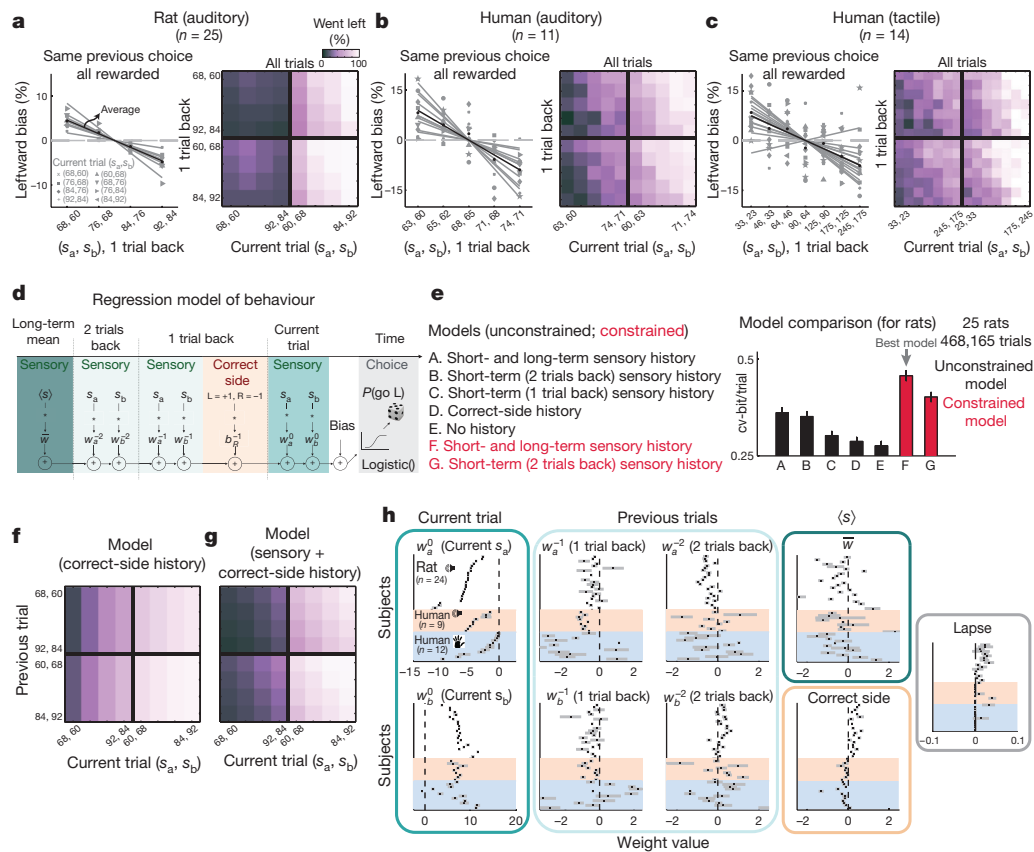


Figure 2 | Sensory history biases behaviour. **a**, Rat auditory trials. Left, percentage of trials in which rats went left minus the average value of left choices, as a function of the stimuli of the previous trial, for fixed previous trial response choice and reward. Right, percentage of trials in which rats went left for each combination of current and previous stimuli; vertical modulation indicates the previous-trial effect. **b**, As in **a**, but for human subjects in the auditory PWM task. **c**, As in **a**, but for human tactile trials. **d**, A linear weighted sum of 9 regressors is used to predict the probability ratio $\log(P_{\text{go left}}/P_{\text{go right}})$; weights are fit to best-match training data, and are evaluated on left-out cross-validation data (see Methods). Regressors: averaged stimuli over the last 20–50 trials (excluding the last two); stimuli from last two trials; correct side on last trial (that is, the side baited with a reward; when b_R^{-1} is positive, this increases the probability of going towards the previously baited side, that is, win-stay/lose-shift); current trial stimuli; overall side bias. **e**, Evaluation of model variants with different regressors

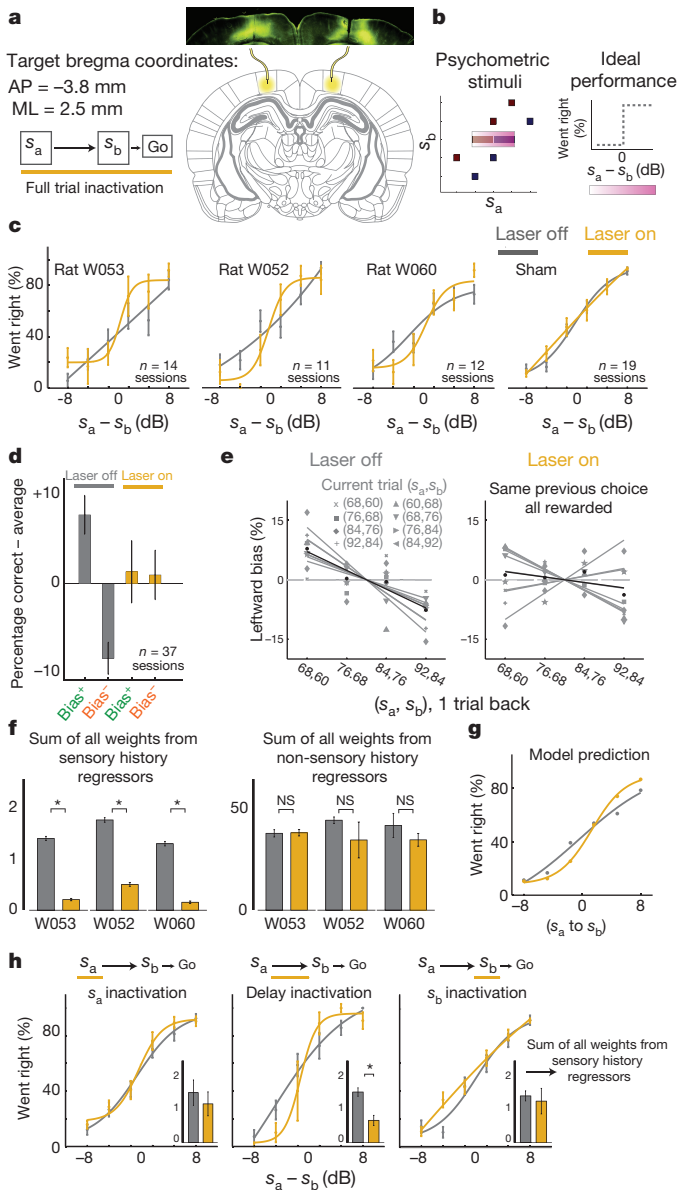
(see Methods; Extended Data Fig. 6). The best model has the regressors shown in **d**. Model A has all regressors shown in **d**; in models B to E, the following regressors are progressively removed: long-term sensory history (B); stimuli of the last two trials (short-term history, C and D); previous trial correct side (E)—model E therefore has no history information. The next two models (F and G, shown in red) have the same regressors as A and B, respectively, but the weights on the s_a of the current trial plus all previous sensory-stimuli weights are constrained to sum to 1, removing one free parameter. cv, cross validation. **f**, A poor match to the data in **a** is found from the predictions of a model with the current trial and previous trial's correct side regressors only. **g**, As in **f**, but now including sensory-history regressors, which greatly improves the match to the data. **h**, Summaries of best-fit parameters over all subjects, from model F in **e**, with additional 'lapse' term. Black ticks, best-fit parameter values, per subject; grey bars, 95% confidence intervals. All panels sorted based on value of w_a^0 .

pink noise stimuli, denoted s_a and s_b ; rats were rewarded for correctly reporting which of the two was louder (Fig. 1a). Following ref. 16, the set of (s_a, s_b) pairs used across trials in a session was chosen so that neither stimulus alone contained sufficient information to solve the task (Fig. 1b). As with any magnitude-discrimination task, the smaller the difference between the stimuli, the harder the task (Fig. 1c). Classical contraction bias⁵ argues that during the delay, the memory of the magnitude of s_a drifts towards the mean of all stimuli presented (Fig. 1c, vertical line $\langle s \rangle$). Consequently, for those pairs in which s_a drifts away from the high-difficulty $s_a = s_b$ diagonal, s_a becomes more distinct from s_b in memory, so contraction bias would improve performance (bias⁺, Fig. 1c). In pairs for which s_a drifts towards the diagonal, performance would decrease (bias⁻).

This predicted pattern is observed in the behaviour of the rats during our experiments (Fig. 1b, high percentage correct for bias⁺ stimulus pairs ($s_a = 84, s_b = 92$) and ($s_a = 68, s_b = 60$), lower percentage correct for bias⁻ stimulus pairs ($s_a = 60, s_b = 68$) and ($s_a = 92, s_b = 84$)). The same pattern has been observed in monkeys⁷ and humans (Extended Data Fig. 1d, e, and refs 5, 6, 17). History-dependent effects are probably adaptive in the natural world, in which there are many long-term

regularities. But in our laboratory task, in which each trial is generated independently, such biases, on average, produce suboptimal performance. The overall performance of the rats was robust and similar across delay intervals from 2–12 s (Fig. 1d; see Extended Data Fig. 1b, c for performance over learning). In some sessions, we used stimulus pairs that were closely spaced along s_a (Fig. 1b) to measure the psychometric discrimination threshold. This worsened slightly, but significantly, at longer delay periods (Fig. 1e, f).

Whereas variation of the delay interval resulted in small effects on overall performance (Fig. 1d), sensory history, by contrast, had a strong effect on performance. As quantified below, this effect was greater than the well-documented influence of previous rewards and choices^{18–20} (see also Extended Data Fig. 1g–i). In the trials that followed a rewarded, rightwards choice—thus holding previous choice and reward fixed—it was found that the smaller the stimuli of the previous trial, the greater the percentage of leftwards choices in the current trial (Fig. 2a, left; slope = -3.06% per decibel, linear fits to percentage leftwards minus average, $P < 0.0001$; see Extended Data Fig. 5a for slopes from $n = 1–7$ trials back). This is consistent with a contraction bias in which the estimate of $\langle s \rangle$ is weighted towards recent stimuli¹⁷, because recent



small values make it more likely that the current s_a is perceived as small, increasing the likelihood of an $s_a < s_b$ (leftwards) response. Figure 2a, right shows that the same effect occurred across all combinations of current and previous trial stimuli from our standard stimulus set ($|s_a - s_b|$ fixed at 8 dB; see Extended Data Fig. 3 for $n = 1-5$ trials back, and Extended Data Fig. 4 for controlling for action and reward). Similar effects were found in human auditory (Fig. 2b) or tactile (Fig. 2c) versions of the task, and increased for larger delay intervals^{21,22} (Extended Data Fig. 2). To simultaneously take into account effects across several previous trials of the history of rewards, choices and 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, Extended Data Fig. 6). Consistent with human data¹⁷, short-term (previous two trials) sensory history had a strong effect on behaviour. In addition, our large dataset revealed a smaller but nevertheless important effect of longer-term (average of previous few tens of trials) sensory history (Fig. 2e, Extended Data Fig. 5b). It has been suggested^{21,23-27} that sensory history does not add a behavioural bias independent of working memory, but instead produces a value of s_a in working memory that is a weighted average of the current stimulus and sensory history. In weighted averages, the weights sum to one. Consistent with that suggestion, we found

Figure 3 | The PPC is specifically necessary for the behavioural effect of previous sensory stimuli. **a**, Schematic of virus injection and full trial inactivations (delay = 2 s). The atlas panel is taken from ref. 31. AP, anteroposterior; ML, mediolateral. **b**, Left, the stimuli included our standard (s_a, s_b) set (black) plus psychometric pairs (purple); right, the ideal psychometric performance. **c**, Psychometric curves for all rats were closer to ideal during PPC inactivation (yellow) trials than during control (grey) trials. Far right, sham inactivation in rats with optic fibres but without virus had no effect ($n = 2$). Error bars show s.e.m. **d**, Percentage correct averaged across all bias⁺ or all bias⁻ trials (Fig. 1c), relative to overall average performance. PPC inactivation eliminates the difference between bias⁺ and bias⁻ trials, or versus overall average (t -test, bias⁺ - bias⁻ significantly different from zero, laser off, $P < 0.00001$; laser on, $P = 0.706$; laser off versus on, $P < 0.01$; two-way analysis of variance, interaction of laser on/off and bias⁺/bias⁻: $P < 0.01$). Error bars show s.e.m. **e**, The bias induced by previous stimuli is reduced under PPC inactivation (laser off: slope of -4.74 , $P = 0.0017$; laser on: slope of -1.36 , $P = 0.42$; laser on versus off: $P = 0.044$). **f**, Left, PPC inactivation selectively reduces sensory-history weights in the regression model shown in Fig. 2d. Error bars show 95% confidence interval ($n = 600$, 200 iterations of threefold cross-validation; $*P < 0.01$, one-sided t -test). Right, sum of all other weights. See Extended Data Fig. 9 for individual weights. **g**, Reducing sensory-history weights in the model is sufficient to improve psychometric performance, comparable to experimental data (c). **h**, Similar to c, for inactivation during either s_a (left), delay (middle) or s_b (right). The insets show the sum of all sensory-history regressors, as in **f** ($n = 600$). Only inactivation of the PPC during the delay produces a significant effect (permutation test, s_a , laser off versus on: $P = 0.17$; delay: $P < 0.0001$; s_b : $P = 0.18$; s_a versus delay: $P = 0.03$, s_b versus delay: $P = 0.02$).

that constraining the sum of regression weights on the s_a stimulus of the current trial plus weights on previous sensory stimuli to equal one—thus removing one free parameter from the model—produced the best performance on cross-validation data (Fig. 2e, red; see Extended Data Fig. 6 for all models and comparisons; the best-fitting model, for every individual rat, had the regressors in Fig. 2d). Sensory history was essential in accounting for behaviour (Fig. 2f, g): examining the weights in the regression model (Fig. 2h) shows that those for sensory history are significantly larger than those for correct side history ($P < 0.001$), and have a greater impact on behaviour (Extended Data Fig. 6g).

It has been proposed that the PPC is critical for working memory (refs 8, 9, 28, but see also refs 29, 30), and therefore we examined its role in our task. We injected bilaterally an AAV virus that drives expression of the light-activated inhibitory opsin halorhodopsin eNpHR3.0, under the CaMKII α promoter (centre of injection, anteroposterior -3.8 mm and mediolateral 2.5 mm from the bregma, Fig. 3a, Extended Data Fig. 7a). Sharp optical fibres were inserted at the centres of injection sites to deliver laser illumination, and we inactivated the PPC during a randomly chosen 20% of trials. To best probe for any small effects, we included psychometric stimuli (Figs 1b and 3b).

Expecting a performance impairment⁸, we were surprised to instead observe an improvement in psychometric performance in all rats tested (Fig. 3c). However, the effect was not simply an overall performance improvement: looking beyond the psychometric stimuli, although performance was indeed improved with respect to control on bias⁻ trials, PPC silencing instead impaired performance on bias⁺ trials (Fig. 3d). Moreover, the difference between bias⁺ and bias⁻ trials was eliminated, as was their difference from control average performance (Fig. 3d, Extended Data Fig. 9a). Similarly, bias as a function of the stimuli of the previous trial was markedly reduced (Fig. 3e, laser off: $P = 0.42$; laser on: $P = 0.0017$; laser on versus off: $P = 0.044$, see Extended Data Fig. 8 for effect on history matrices). Fitting our regression model separately to the set of laser on versus off trials, we found that PPC inactivation significantly reduced sensory-history regression weights (Fig. 3f), and no other regression terms were significantly affected (individual weights in Extended Data Fig. 9). A model with reduced sensory-history effects as in Fig. 3f was sufficient to reproduce the experimental data (Fig. 3g). Therefore, PPC silencing appeared to have no impact on working memory, but specifically and substantially

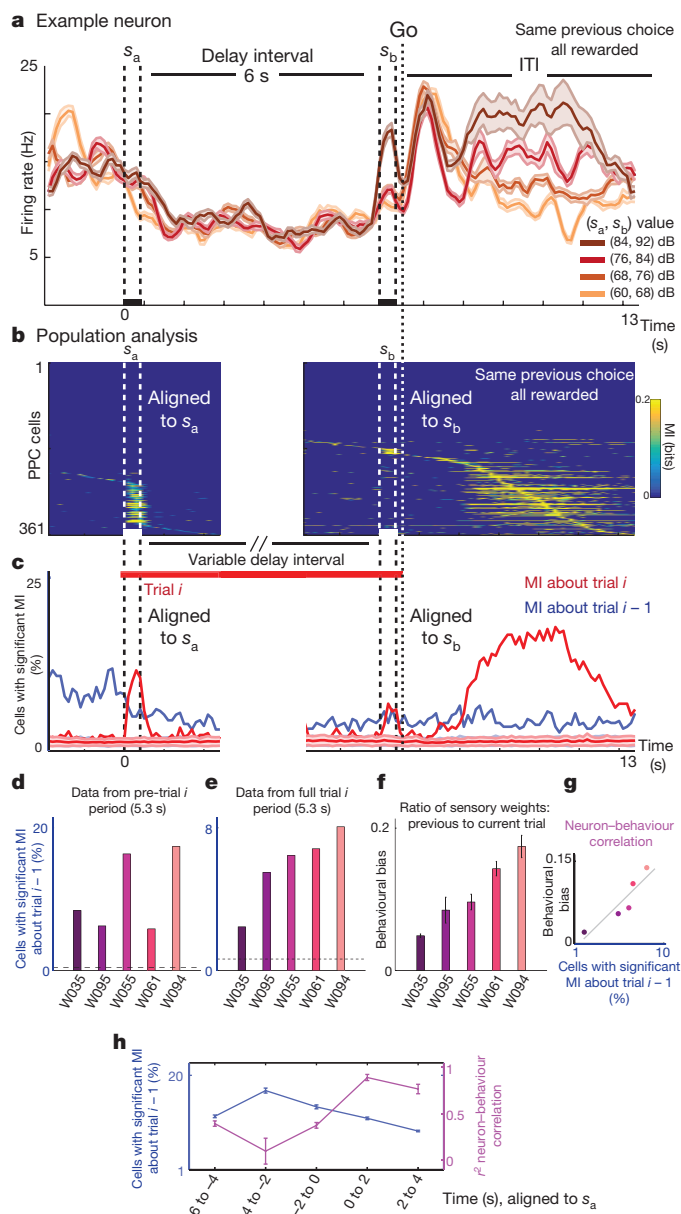


Figure 4 | PPC neurons carry more information about previous-trial sensory stimuli than about current trial stimuli, which predicts behavioural bias. **a**, Firing rate of example neuron in response to different s_a values. For clarity, the graph shows only trials in which the rat responded left after the 'go' cue, was rewarded, and the delay interval was 6 s. **b**, Mutual information (MI) between the firing rate of each neuron and stimulus pair (s_a, s_b) as a function of time. Data are from trials with delay intervals of 2–6 s. Left, aligned to the start of s_a ; right, aligned to the start of s_b . Only mutual information values significantly larger than the shuffled distribution ($P < 0.005$) are included; non-significant values are dark blue. To control for reward and choice, mutual information values were calculated using only trials with fixed choice and reward (see Methods). **c**, Summary of population analysis, showing the percentage of cells with significant coding of stimuli from trial i (red), or trial $i - 1$ (blue). **d**, Percentage of cells with significant mutual information about the previous trial, for each of five recorded rats, calculated over 5.3 s during the ITI before a new trial. The horizontal dashed line indicates the percentage of cells expected by chance from shuffled data (see Methods). **e**, As in **d**; data from 5.3 s of the new trial. **f**, Behavioural bias for individual rats, calculated as relative sensory-history weights (to weights for s_a and s_b regressors of trial i) from the best model fit. **g**, Neuron–behaviour correlation; data from **f** plotted against data from **e**, Spearman's rank correlation $r = 1.0$, $P < 0.01$, $n = 5$ rats. **h**, Calculation as in **d–g**, using 2-s windows. High neuron–behaviour correlation appears concurrently with the s_a of the new trial. Error bars show s.d. of the mean over $n = 1,000$ bootstrap samples with replacement.

reduced sensory-history effects. This reduction did not persist into future trials (Extended Data Fig. 8), and occurred with PPC inhibition during the working memory delay, but not during s_a or s_b (Fig. 3h), suggesting a focused role for the PPC in the interaction between the previous stimuli and the working memory of the current trial.

To examine whether signatures of sensory history are present in the region in which inactivation appears to cancel history effects, we made extracellular recordings that targeted the PPC during task performance, recording from 936 units in 5 rats implanted with microwire arrays. Neurons with a mean firing rate of below 2 Hz were discarded (Methods), leaving 361 units for the analysis. Most cells gave results similar to those of the example cell in Fig. 4a; their firing rates during the working memory delay did not distinguish between values of s_a held in memory, and therefore did not carry information about them. Instead, robust information about the sensory stimulus pair appeared approximately 1 s after the trial had terminated, during the inter-trial interval (ITI). We used mutual information (Methods) to quantify the amount of information, in neuronal firing rates, about which (s_a, s_b) sensory-stimulus pair was presented (Fig. 4b, c; see Extended Data Fig. 10 for mutual information about other components, including current and previous choices, rewards and s_a alone). During the ITI before a new trial, a large fraction of PPC neurons carried significant information about the previous trial stimuli (22% of analysed neurons; Fig. 4c). A smaller fraction of cells continued to code the stimuli of the previous trial both into the start of the new trial and throughout the working memory delay of the new trial (Fig. 4c, Extended Data Fig. 10b). We computed the fraction of neurons with significant mutual information about the stimuli of the previous trial, during the ITI (Fig. 4d) or the current trial (Fig. 4e), and compared this to the strength of the sensory history behavioural bias of the rat (Fig. 4f). 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 versus full current trial from Steiger's Z -test). This suggests two things: first, a tight link between sensory history representations in the PPC and history biases, and second, that the PPC history representation is used during or shortly after the presentation of the new s_a (Fig. 4h), consistent with the idea that contraction bias affects the working memory representation of s_a .

Parametric working memory tasks, with their quantifiable behaviour, are well suited to investigating the effect of sensory history on perception and behaviour. Rodent versions of these tasks, with semi-automated training, are an efficient platform for causal and cellular-resolution investigation of neural mechanisms. By using this platform, we identified the PPC as an essential node in both the representation and causal effects of sensory-stimulus history. This represents a step towards a cellular-resolution understanding of long-standing questions about how sensory-stimulus history affects working memory and perception. Important issues that may now be addressed include how history representations in the PPC interact with current stimulus representations to modulate perception, how history information reaches the PPC, and which brain regions connected to the PPC are also essential nodes of the circuit.

Online Content Methods, along with any additional Extended Data display items and Source Data, are available in the online version of the paper; references unique to these sections appear only in the online paper.

Received 10 August 2017; accepted 9 January 2018.

Published online 7 February 2018.

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. in *Brain Mapping: An Encyclopaedic Reference* Vol. 3, (ed. Toga, A. W.) 501–506 (Academic, 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., Esmaili, V. & Diamond, M. E. Tactile perception and working memory in rats and humans. *Proc. Natl Acad. Sci. USA* **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. *J. Philos. Psychol. Sci. Methods* **7**, 461–469 (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. Hernández, A., Salinas, E., García, R. & Romo, R. Discrimination in the sense of flutter: new psychophysical measurements in monkeys. *J. Neurosci.* **17**, 6391–6400 (1997).
17. Raviv, O., Ahissar, M. & Loewenstein, Y. How recent history affects perception: the normative approach and its heuristic approximation. *PLoS Comput. Biol.* **8**, e1002731 (2012).
18. Busse, L. *et al.* The detection of visual contrast in the behaving mouse. *J. Neurosci.* **31**, 11351–11361 (2011).
19. Abrahamyan, A., Silva, L. L., Dakin, S. C., Carandini, M. & Gardner, J. L. Adaptable history biases in human perceptual decisions. *Proc. Natl Acad. Sci. USA* **113**, E3548–E3557 (2016).
20. Hwang, E. J., Dahlen, J. E., Mukundan, M. & Komiyama, T. History-based action selection bias in posterior parietal cortex. *Nat. Commun.* **8**, 1242 (2017).
21. Olkkonen, M., McCarthy, P. F. & Allred, S. R. The central tendency bias in color perception: effects of internal and external noise. *J. Vis.* **14**, 5 (2014).
22. Fritsche, M., Mostert, P. & de Lange, F. P. Opposite effects of recent history on perception and decision. *Curr. Biol.* **27**, 590–595 (2017).
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. Papadimitriou, C., Ferdoash, A. & Snyder, L. H. Ghosts in the machine: memory interference from the previous trial. *J. Neurophysiol.* **113**, 567–577 (2015).
27. Fischer, J. & Whitney, D. Serial dependence in visual perception. *Nat. Neurosci.* **17**, 738–743 (2014).
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**, e05457 (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. Paxinos, G. & Watson, C. *The Rat Brain in Stereotaxic Coordinates* 5th edn (Elsevier, 2004).

Supplementary Information is available in the online version of the paper.

Acknowledgements We thank C. Duan, R. Low, A. Piet, L. Pinto, B. Scott and I. Witten for their comments on the manuscript. We thank K. Osorio and J. Teran for animal and laboratory support.

Author Contributions A.A. and C.D.B. conceived the project. A.A. carried out all experiments and analysed the data, with the optogenetic inactivations carried out with assistance from C.D.K. A.A. gathered human tactile data in M.E.D.'s laboratory. A.A. and C.D.B. wrote the manuscript, based on a first draft by A.A., with extensive comments from C.D.K. and M.E.D.

Author Information Reprints and permissions information is available at www.nature.com/reprints. The authors declare no competing financial interests. Readers are welcome to comment on the online version of the paper. Publisher's note: Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations. Correspondence and requests for materials should be addressed to A.A. (aakrami@princeton.edu) or C.D.B. (brody@princeton.edu).

Reviewer Information *Nature* thanks L. Busse and J. de la Rocha for their contribution to the peer review of this work.

METHODS

Rat subjects. A total of 33 male Long–Evans rats (*Rattus norvegicus*) between the ages of 6 and 24 months were used for this study. Of these, 25 rats were used for behavioural assessments, 6 were used for neural recordings, and 7 for optogenetic inactivations. All rats were assigned randomly to different experimental conditions. All statistical tests were performed between groups with similar sample sizes. No statistical methods were used to predetermine sample size. 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, aged 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, aged 22–35) were tested. Protocols conformed to international norms and were approved by the Ethics Committee of the International School for Advanced Studies (Trieste, Italy). Subjects signed informed consent.

Rat behaviour. We developed a computerized protocol to train rats, in high-throughput facilities, to perform an auditory delayed comparison task, adapted from a tactile version⁶. All training takes place in three-port operant conditioning chambers, in which ports are arranged side-by-side along one wall, with two speakers placed above the left and right nose ports. Figure 1a shows the task structure. A visible light-emitting diode in the centre port signals the availability of each trial. Rat subjects initiate a trial by inserting their nose into the centre port, which causes the centre light to turn off. Rats must keep their nose in the centre port (fixation period), until an auditory ‘go’ cue (a 6-kHz pure tone for 200 ms) signals the end of fixation. Only after the ‘go’ cue can subjects withdraw and orient to one of the side ports in order to receive a reward of water. During the fixation period two auditory stimuli, s_a and s_b , separated by a variable delay, are played for 400 ms, with short delay periods of 250 ms inserted before s_a and after s_b . The stimuli consist of broadband noise (2,000–20,000 Hz), generated as a series of sound pressure level (SPL) values sampled from a zero-mean normal distribution. The overall mean intensity of sounds varies from 60–92 dB. Rats should judge which out of the two stimuli, s_a and s_b , had the greater SPL standard deviation. If $s_a > s_b$, the correct action is to poke the nose into the right-hand nose port in order to collect the reward, and if $s_a < s_b$, rats should orient to the left-hand nose port. Trial durations are independently varied on a trial-by-trial basis, by varying the delay interval between the two stimuli, which 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: associate light in the centre poke with the availability of trials; associate special sounds from the side pokes with reward; maintain their nose in the centre poke until they hear an auditory ‘go’ signal; and compare the two s_a and s_b stimuli.

Although a substantial amount of data has been collected on all delay intervals from 2 to 12 s, in this manuscript we focus on delay durations of 2, 4 and 6 s, as most of the rats were consistently trained on these values. Training began when rats were two months old, and typically required three to four months for rats to display stable performance on the complete version of the task.

Human auditory behaviour. Similar auditory stimuli to those used for rats were used in the human version of the task. In this experiment, subjects received, in each trial, a pair of sounds played from ear-surrounding noise-cancelling headphones (brand 233621-H501). The subject self-initiated each trial by pressing the space bar on the keyboard. The first sound was then presented together with a green square on the left side of a computer monitor in front of the subject. This was followed by a delay period, indicated by ‘WAIT!’ on the screen, then the second sound was presented together with a red square on the right side of the screen. At the end of the second stimulus and after the go cue, subjects were required to compare the two sounds and decide which one was louder, then indicate their choice by pressing the ‘k’ key with their right hand (second was louder) or the ‘s’ key with their left hand (first was louder). Written feedback about the correctness of their response was provided on the screen, for individual trials as well as the average performance updated every ten trials.

Human tactile behaviour. In a separate set of experiments, run at the International School For Advanced Studies (SISSA), human subjects performed the tactile version of the task. The details of this task have been previously described and the behaviour has been characterized⁶. In brief, at each trial two noisy vibration stimuli, interleaved with a variable delay interval, were delivered to the subject’s fingertip on their left hand. 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 their right hand. This triggered presentation

of the two stimuli. After a post-stimulus delay, a blue panel was 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 as greater, respectively. They received feedback (correct/incorrect) on each trial via the monitor. Human experiments were controlled using LabVIEW software (National Instruments).

Stimulus set. If the first stimulus, s_a , was 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 was 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 that guarantee that, across trials, the same value of SPL standard deviation is randomly presented for the first stimulus or the second stimulus. The stimuli along the diagonal in Fig. 1b represent such a stimulus set. A minimum of eight pairs of stimuli span a wide range of SPL standard deviation values (Fig. 1b). Using this stimulus set, if the subject were to ignore either 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 of the stimulus pairs on one side of the diagonal are associated with the same action, and all have the same ratio of s_a to s_b . For each trial, one of these eight pairs of stimuli is randomly selected to determine s_a and s_b .

Psychometric curves. In some sessions, we used stimulus pairs that were closely spaced along s_a (Fig. 1b) to measure the psychometric discrimination threshold. 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 four-parameter logistic function of the form

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

where y_0 is the left endpoint, $y_0 + a$ is the right endpoint, x_0 is the bias, and $a/4b$ is the slope. Fits were performed using the nonlinear least square regression method (nlinfit.m function) in MATLAB 2013.

Regression model of behaviour. Our semi-automated training protocol facilitated the generation of a behavioural dataset comprising 468,165 trials from 25 trained rats, which in turn enabled statistical characterization of the decision-making process. To quantify the behaviour of the rats, we carried out an analysis to weight the contributions of s_a and s_b on the current trial and several trials in the past, as well as the contributions of the history of choice and reward on the rat’s choice in the current trial. Using the data generated by concatenating several training sessions, we fit the rat’s choice with a logistic regression model that allows for the linear combinations of s_a and s_b and other desired factors. The linear combination is then mapped non-linearly into the rat’s choice; that is, the probability of trials in which the subject judged $s_a > s_b$, through a logistic function as:

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

where

$$A = \sum_{t=0}^n (s_a^{-t} W_a^{-t} + s_b^{-t} W_b^{-t} + b_R^{-1} + \beta)$$

where W_a^{-t} and W_b^{-t} are coefficients of the s_a and s_b regressors, respectively, from t trials back. b_R^{-1} is the correct side on the previous trial: left = +1, right = -1. This regressor captures the 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 of $s_a > s_b$). The absolute values of all of the regressors are normalized between 0 and 1. We used the log-likelihood as the cost function C :

$$C = -y \log(p) - (1 - y) \log(1 - p) = -\frac{1}{m} \left[\sum_i^m y_i \log(p_i) + (1 - y_i) \log(1 - p_i) \right]$$

The model was fit using a gradient-descent algorithm to minimize the negative log-likelihood cost function. We used the sqp algorithm in the fmincon function from MATLAB 2013. Weights were calculated using L2-regularization to prevent overfitting. The hyperparameter value (λ) was selected independently for each rat using evidence optimization, on the basis of fivefold cross-validation. Different variants of the model, which systematically study the relevance of various sensory and reward history factors^{18,19,32}, capturing not only win-stay/lose-switch but also perseveration, are discussed in Extended Data Fig. 6.

Model comparison and cross-validation. All models were fit separately for each individual rat ($n = 25$), using 200 runs of fivefold cross-validation. For each run

we calculated the log-likelihood of the test dataset given the best-fit parameters on the training set ($\log l$). 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 rat went left). This gives us a null log-likelihood reference value ($\log l_0$). 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 model³³:

$$\frac{(\log l - \log l_0)/n_{\text{trials}}}{\log(2)}$$

For each model, we first chose the optimal regularization value (λ) that would maximize the CV-bit/trial. To compare different models, we calculated the median value of CV-bit/trial across 10,000 fits for each subject. Because in this method we measure the log-likelihood using the cross-validated data, it automatically addresses the overfitting problem, such that if additional parameters of one model result in overfitting in the training set, it would penalize it in the cross-validated test set.

Optogenetic virus injection and fibre implantation. For optogenetic perturbation experiments, the general surgery techniques and fibre etching follow previous reports³⁴, except that we began the 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 fibre cables. The cable jacket, strengthening fibres and outer plastic coating (typically white or orange) were fully removed, leaving 1 cm of the fibre-optic cable and the inner plastic coating (typically clear) intact. Then 2 mm of the fibre tip (with the 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 removed with tweezers to reveal a 1-mm sharp-etched fibre tip. Sufficient 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 injections, 2 μl of adeno-associated virus (AAV) (AAV5-CaMKII α -eNpHR3.0-eYFP), which drives expression of the light-activated inhibitory opsin halorhodopsin eNpHR3.0, under the CaMKII α 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 a diameter of approximately 30 μm . Five closely spaced injection tracts were used for each rat. For the central injection tract (anteroposterior -3.8 mm and mediolateral 2.5 mm from the bregma, Fig. 3a; brain image from Paxinos and Watson³¹), one injection of 23 nl was administered every 100 μm in depth, starting 100 μm below the brain surface of the PPC, to a total depth of 1.5 mm. Four additional injection tracts were completed, using procedures identical to those for the central tract, with one injection at 500 μm anterior, posterior, medial and lateral from the central tract. Each injection was followed by a 10-s pause, with 1 min after 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 approximately 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 the entire incision site and allowed only the LC connector to protrude. Halorhodopsin expression was allowed to develop for six weeks before the behavioural testing began.

Optogenetic perturbation. The implant in the rat was connected to a 1-m patch cable attached to a single-fibre rotary joint (Princeton) 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 pulse. Laser illumination occurred on 20% of randomly selected trials. See Extended Data Fig. 7 for physiological confirmation of optogenetic inactivation effects in an anesthetized rat. On the basis of our previous quantifications of optogenetic effects³⁴, we estimate that using eNpHR3.0 we can inhibit, almost entirely, neurons in a radius of approximately 750 μm from the tip of the optic fibre, amounting to a sphere of around 1,500 μm in diameter.

Recordings. Six rats were implanted with microwire arrays in their left or right PPC ($n = 3$ in lPPC, $n = 3$ in rPPC, see Extended Data Fig. 7 for histological localization of electrodes). The target region was accessed by craniotomy, using standard stereotaxic techniques (centred 3.8 mm posterior to the 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 not usually considered to be resistant to penetration, nevertheless presents a barrier to the entry of the micro-electrode arrays, owing to the high-density arrangement of electrodes in the multi-channel electrode arrays. This dimpling phenomenon, when the electrodes are

pushing the brain cortex down without penetrating, is more pronounced for arrays with larger numbers of electrodes. In addition to potentially injuring the brain tissue, dimpling is 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 for effective and accurate electrode placement. To overcome the dimpling problem, we implemented the following procedure. After the craniotomy was made, and the dura was carefully removed over the entire craniotomy, a petroleum-based ointment (such as bacitracin ointment or sterile petroleum jelly (Puralube Vet Ointment)) was applied to the exact site of electrode implantation. The cyanoacrylate adhesive (Vetbond Tissue Adhesive) was 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 was complete, over a period of few minutes, the petroleum ointment at the target site was removed, and the 32-electrode microwire array (Tucker-Davis Technologies) was inserted by slowly advancing a Narishige hydraulic micromanipulator. After inserting the array(s), the remaining exposed cortex was covered with biocompatible silicone (kwik-sil), and the microwire array was secured to the skull with C&B Metabond and dental acrylic.

During the ten-day recovery period, rats had unlimited access to water and food. Recording sessions in the apparatus began thereafter, using Neuralynx acquisition systems. Extracellular activity of the PPC neurons was manually sorted into single units and multi units, on the basis of the spike waveform and the refractory period observed in the interspike interval histogram, using SpikeSort3D software. In total 1,081 single or multi units were recorded in the PPC of 6 rats. Only neurons for which the overall mean firing rate within the session was at least 2 Hz were included in the analysis, giving a total of 456.

Neural analysis. Mutual information. To quantify the type and amount of information that the PPC neurons carry about various task parameters, we computed Shannon's mutual information³⁵. In this formulation, the amount of information that can be extracted from the firing rate of a neuron R , about the task-related parameter X can be computed as:

$$I(X, R) = \sum_x P(x) \sum_r P(r, x) \log_2 \frac{P(r|x)}{P(r)}$$

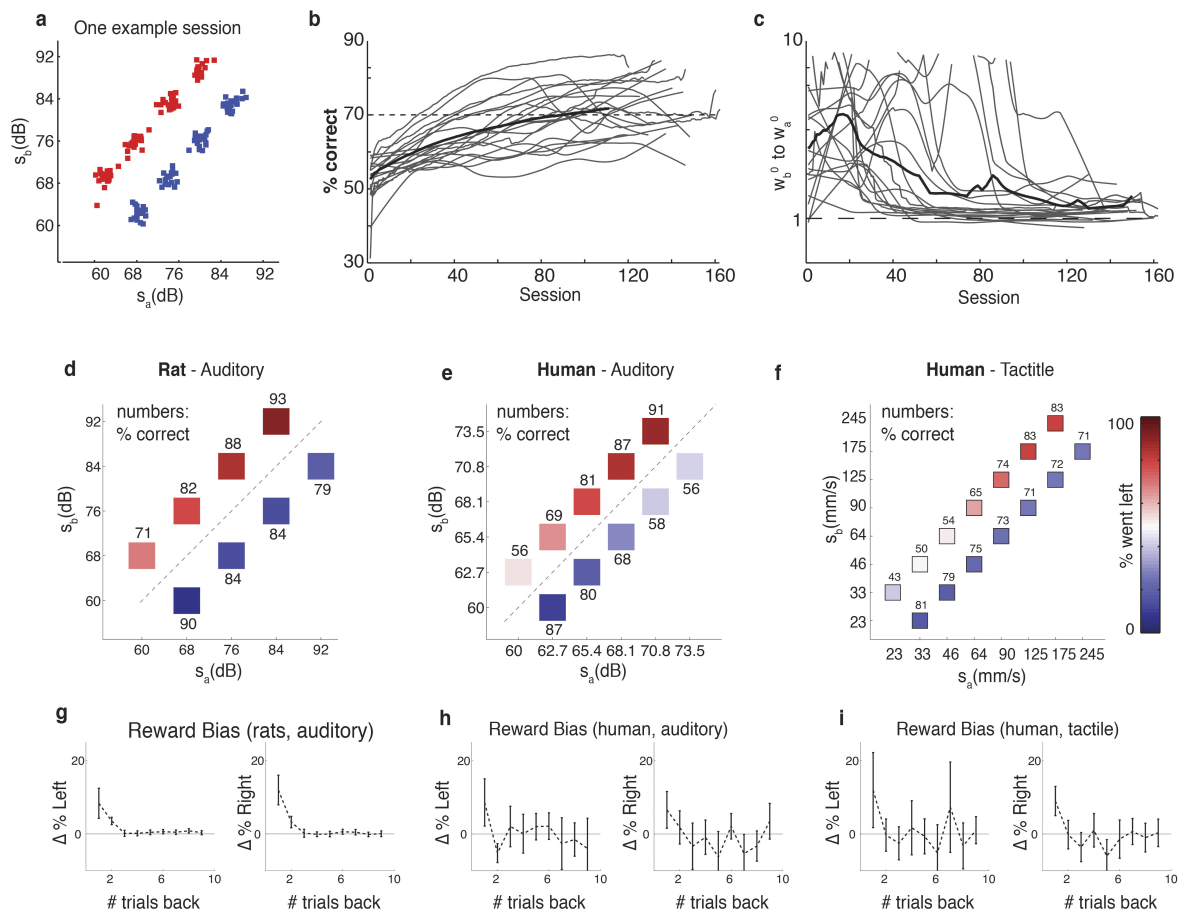
where $P(r|x)$ is the conditional probability of observing a neuronal response r given the presentation of the task parameter x , $P(r)$ is the marginal probability of occurrence of neuronal response r among all possible responses, and $P(x)$ is the probability of task parameter x . Information measured in this way quantifies how well an ideal observer can discriminate between members of a stimulus set on the basis of the neuronal responses of a single trial³⁶. For each trial, neuronal response was defined as the rate of spiking during a time window 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 information³⁷. In order to correct for this bias, we used a combination of two techniques. First we estimated and corrected the bias based on the quadratic extrapolation method³⁸, which assumes that the bias can be accurately approximated as second order expansions in $1/N$. Then we used a bootstrap procedure (shuffling) 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. Owing 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 was used to build a non-parametric test of whether the corrected information computed using quadratic extrapolation method is significantly different from zero³⁹.

Using the mutual information distribution from a shuffled dataset, at each time bin, in which trials are randomly labelled, we first calculated the bin-by-bin estimate of the percentage of cells with significant value of mutual information expected by chance (Fig. 4c, shuffled data). We then computed the average over the ITI or the duration of trial i to find the mean values depicted by dashed lines in Fig. 4d, e. To control for reward and choice, mutual information values were calculated using only trials with fixed choice and reward, and only then averaged across different, separately calculated reward and choice groups.

Code availability. All software used for behavioural training is available on the Brody laboratory website at <http://brodylab.org/auditory-pwm-task-code>. Software used for data analysis is available from the corresponding authors upon reasonable request.

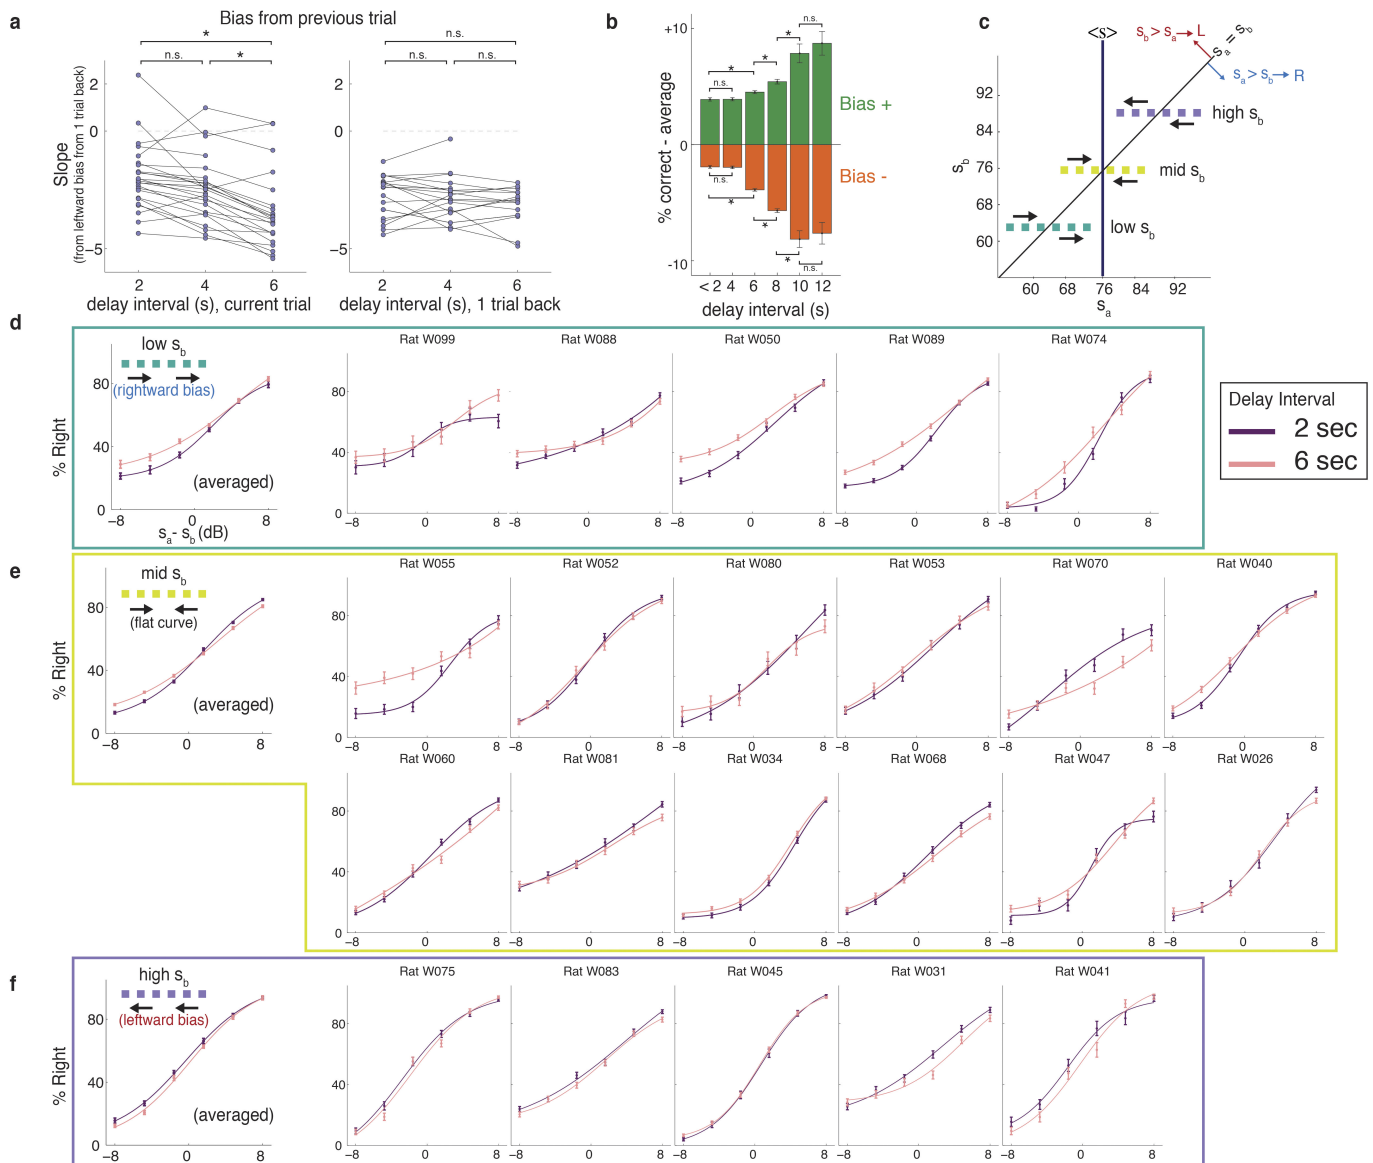
Data availability. Raw and processed data are available from the corresponding authors upon reasonable request.

32. Nogueira, R. *et al.* Lateral orbitofrontal cortex anticipates choices and integrates prior with current information. *Nat. Commun.* **8**, 14823 (2017).
33. 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).
34. Hanks, T. D. *et al.* Distinct relationships of parietal and prefrontal cortices to evidence accumulation. *Nature* **520**, 220–223 (2015).
35. Shannon, C. E. A mathematical theory of communication. *Bell Syst. Tech. J.* **27**, 379–423 (1948).
36. Rieke, F., Warland, D. & Bialek, W. Coding efficiency and information rates in sensory neurons. *Europhys. Lett.* **22**, 151–156 (1993).
37. Treves, A. & Panzeri, S. The upward bias in measures of information derived from limited data samples. *Neural Comput.* **7**, 399–407 (1995).
38. Nemenman, I., Bialek, W. & de Ruyter van Steveninck, R. Entropy and information in neural spike trains: progress on the sampling problem. *Phys. Rev. E* **69**, 056111 (2004).
39. 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**, 62–70 (2010).



Extended Data Figure 1 | Full stimulus set, learning curves, mean performance and reward bias. **a**, Each stimulus is composed of a series of SPL values sampled from a zero-mean normal distribution, and standard deviation of s . For each trial, SPL values are randomly drawn and therefore, owing 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 grey lines show learning curves presented as the change in percentage correct over months of training, for $n = 25$ rats. An 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 Fig. 2e (three-parameter, no-history version). **d**, Rat auditory working memory performance, data from 21 rat subjects (total of 468,165 trials) are grouped according to (s_a, s_b) pair but averaged across subjects and over different delay durations (2–8 s). **e**, Human auditory working memory


performance. For humans, the interstimulus delay varied randomly from 2 s to 6 s. (11 subjects, 12,623 trials). **f**, Human tactile working memory performance; similar to **e** but for humans engaged in the tactile version of the task. In this task, the interstimulus delay varied randomly from 2 s to 8 s. Data from 14 human subjects (total of 4,694 trials) are pooled together. **g**, Reward history bias. Left, the y axis shows, for turn-left trials and as a function of k , the percentage of subjects that went left when the k^{th} trial back was rewarded on the left, minus the percentage that went left when the k^{th} trial back was rewarded on the right. Right, the complementary plot for turn-right trials: the percentage that went right when the k^{th} trial back was rewarded on the right, minus the percentage that went right when the k^{th} trial back was rewarded on the left. Data from $n = 21$ rats. Each point shows the mean value of the bias over subjects. Error bars show 95% confidence intervals. **h**, **i**, Similar to **g** for human auditory (**h**, $n = 11$ subjects) and tactile (**i**, $n = 14$ subjects) PWM tasks.



Extended Data Figure 2 | Contraction bias grows as a function of the working memory delay interval of the current trial. **a**, Slopes from linear fits to the percentage leftward bias (as in Fig. 2a), for rats that were each trained on delay intervals of 2, 4, and 6 s ($n = 21$). The plot on the left shows the behavioural bias (percentage that went left minus the average) as a function of working memory delay interval of the current trial. The plot on the right shows the behavioural bias as a function of working memory delay interval from one trial back. Each dot represents a rat; lines connect the different delay intervals for each rat. Left: from a one-sided paired t -test, 2 versus 4 s: $P = 0.012$, 2 versus 6 s: $P < 0.001$; 4 versus 6 s: $P < 0.001$ * $P < 0.001$, one-sided paired t -test. Right: 2 versus 4 s: $P = 0.76$, 2 versus 6 s: $P = 0.37$; 4 versus 6 s: $P = 0.65$. The behavioural bias increases with greater current working memory delay period, but no significant dependence on the working memory delay period of the previous trial is found²⁶. **b**, Percentage correct averaged across all bias⁺ trials or all bias⁻ trials, relative to overall average performance, as a function of working memory delay interval on the current trial. Data are pooled from a dataset in which different rats were trained on different sets of delay intervals; data for each delay interval may therefore contain different rats than data for other delay intervals ($n = 25$ rats total). Error bars show s.d. As in **a**, behavioural effect grows as a function of the current working memory

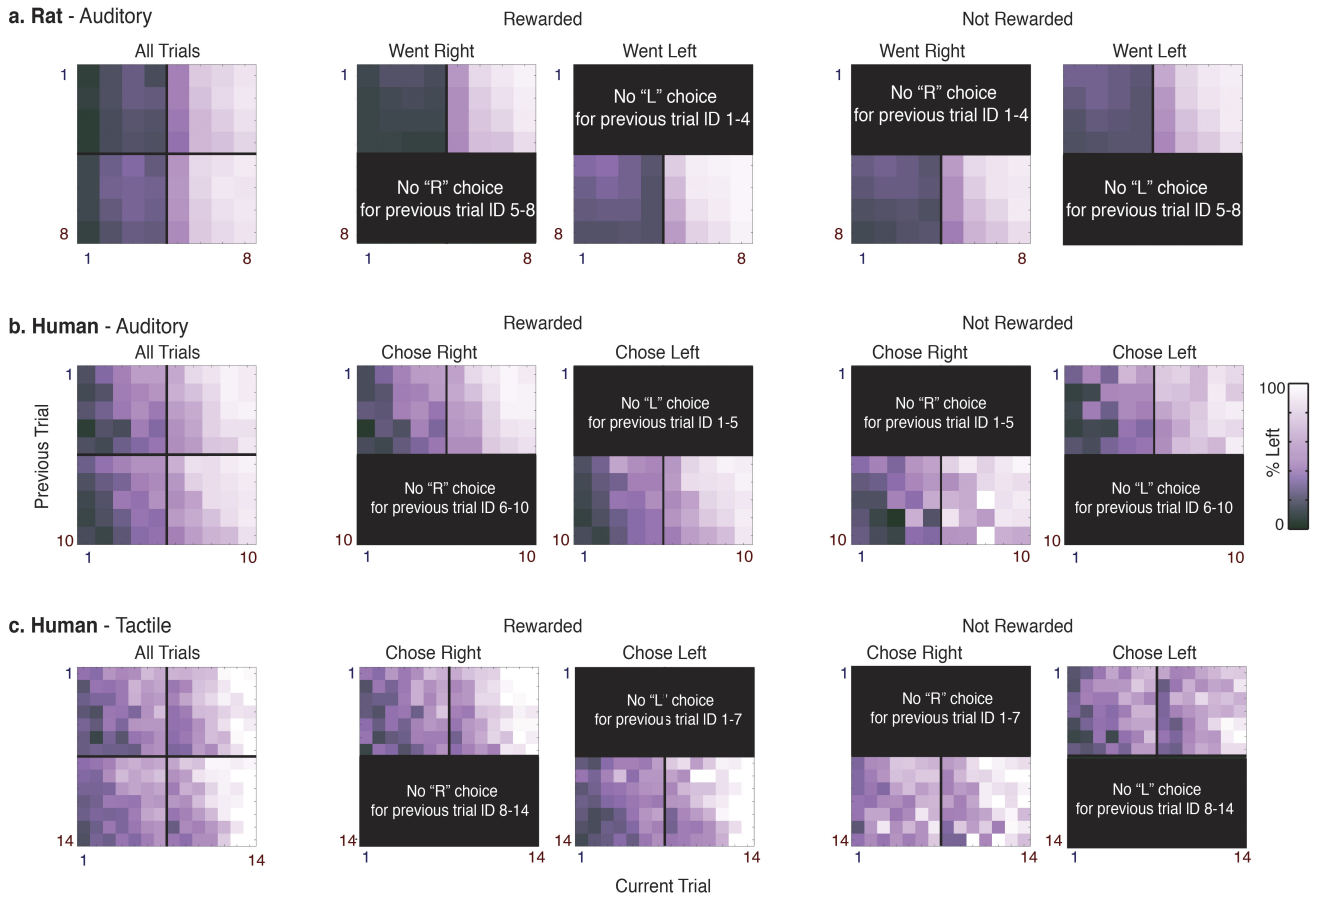
delay period. * $P < 0.001$, one-sided t -test.

c, Schematics of stimuli used for three different psychometric curves: high s_b , in which contraction bias would lead all the s_a stimuli to be treated as lower than they actually were (indicated by the leftward arrows), producing a rightward shift of the psychometric curve; mid s_b , in which contraction bias would lead all the s_a stimuli to be treated as closer to s_b than they actually were, producing a flattening of the psychometric curve; and low s_b , in which contraction bias would lead all the s_a stimuli to be treated as higher than they actually were, producing a leftward shift of the psychometric curve. **d**, Psychometric curves for low- s_b trials, averaged across rats and separately for each individual rat, for trials with a 2-s working memory delay interval, and for trials with a 6-s working memory delay interval. Curves are fits to a four-parameter logistic function (see Methods). As the working memory delay interval grows, the leftward shift predicted by contraction bias shift is more pronounced. For each individual rat, $n = 120$ sessions of data were used. Error bars show the s.e.m. over sessions. **e**, as in **d** but for the mid- s_b trials. As the working memory delay interval grows, the flattening predicted by contraction bias is more pronounced. **f**, as in **d** but for the high- s_b trials. As the working memory delay interval grows, the rightward shift predicted by contraction bias is more pronounced.



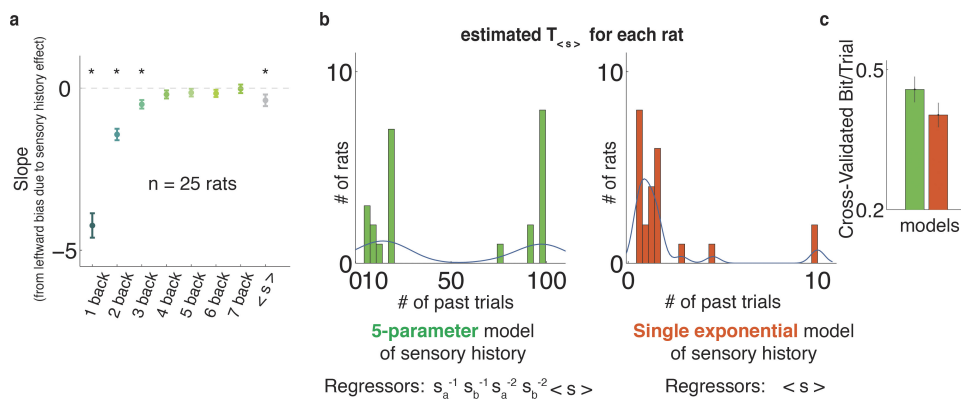
Extended Data Figure 3 | Sensory-history matrix, from one to five trials back. **a**, Stimulus-history matrix, as described in Fig. 2a, when percentage left is shown given any combination of the stimuli in the current trial (x axis) and n -trials back (y axis), $n = 1, 2, 3, 4, 5$. Trial numbers indicate pairs of (s_a, s_b), values in dB. 1: (68, 60); 2: (76, 68); 3: (84, 76); 4: (92, 84); 5: (60, 68); 6: (68, 76); 7: (76, 84); 8: (92, 84). Data from $n = 21$ rats, comprising a total of 468,165 trials used in this analysis. **b**, Similar to **a**,

for the human auditory task. Trial numbers, with values in dB: 1: (62.7, 60); 2: (65.4, 62.7); 3: (68.1, 65.4); 4: (70.8, 68.1); 5: (73.5, 70.8); 6: (60, 62.7); 7: (62.7, 65.4); 8: (65.4, 68.1); 9: (68.1, 70.8); 10: (70.8, 73.5). **c**, Similar to **a**, for the human tactile task. Trial numbers, in mm s^{-1} : 1: (33, 23); 2: (46, 33); 3: (64, 46); 4: (90, 64); 5: (125, 90); 6: (175, 125); 7: (245, 175); 8: (23, 33); 9: (33, 46); 10: (46, 64); 11: (64, 90); 12: (90, 125); 13: (125, 175); 14: (245, 175).



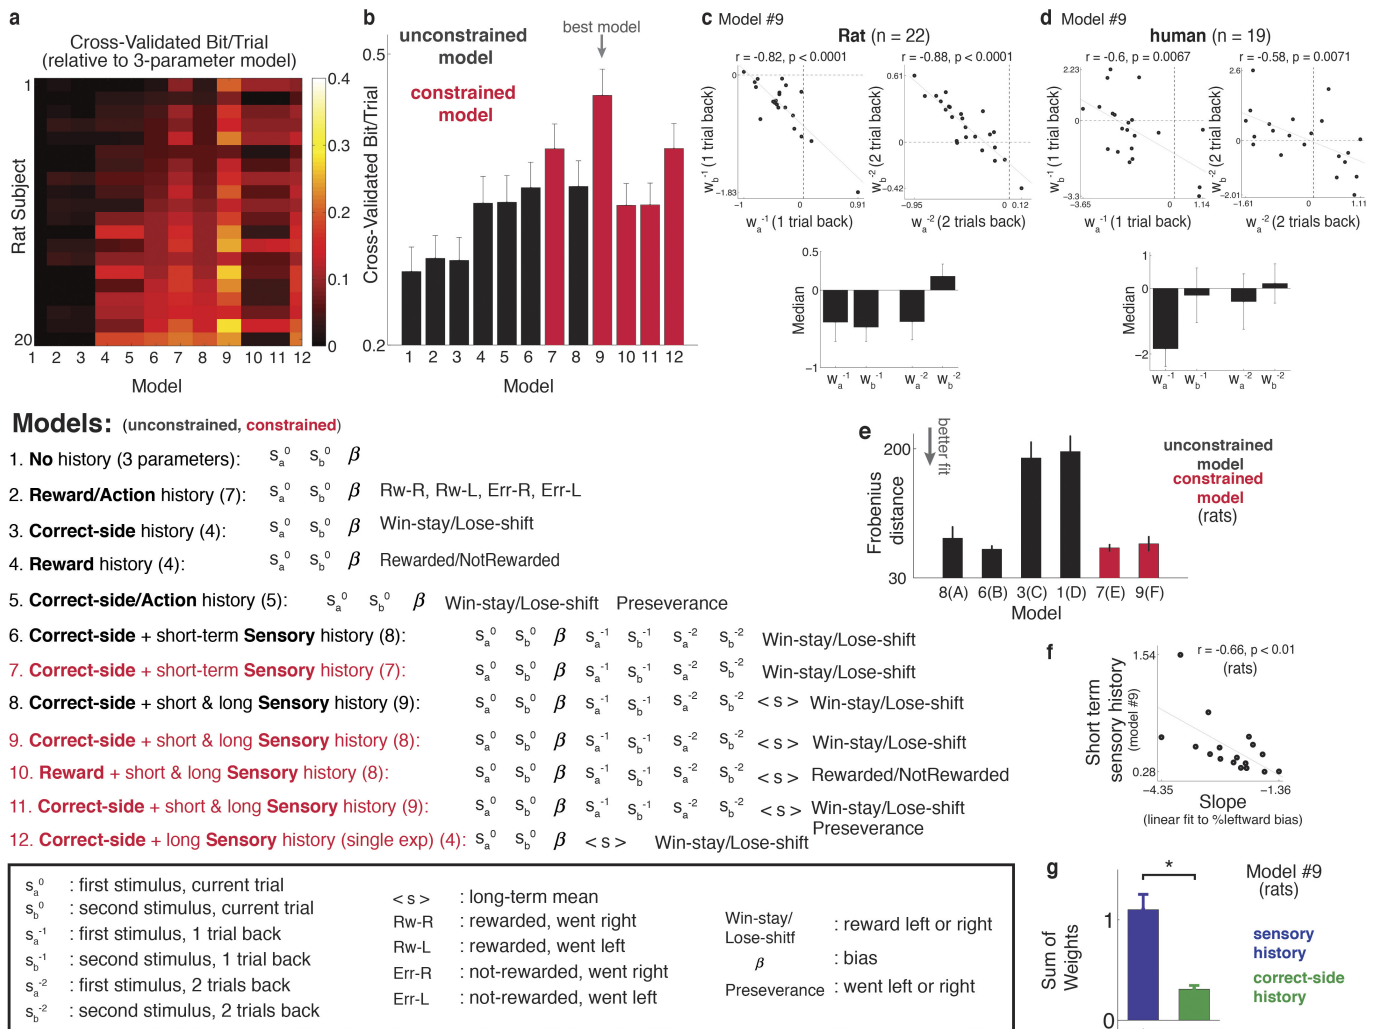
Extended Data Figure 4 | Sensory-history matrix, controlled for reward and choice. Similar to Extended Data Fig. 2, except that in this plot only trials for which the previous trial resulted in the same action and reward

status are included. Therefore, modulation by previous trial cannot be due to action history or reward history. Trial numbers are similar to those in Extended Data Fig. 3.





Extended Data Figure 5 | Short-term and long-term sensory history, and estimating the optimal window of $\langle s \rangle$. **a**, Slopes from linear fits to the percentage leftward bias from n -back trials ($n = 1-7$, as in Fig. 2a where $n = 1$ was used), and also $\langle s \rangle$ which is a window of 17 trials, from $n = 4$ to $n = 20$, in grey. Each point shows the mean of the slope values over $n = 25$ rats. Error bars show 95% confidence intervals. **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 τ values from a model that has five regressors to account for the sensory history—the first and second

stimulus from the two trials back and a separate exponential window over the remaining past trials (Fig. 2d). The results shown in orange are from a model containing only one regressor: a single exponential window over all the past trials accounts for the sensory history. In the single-exponential model, the best-fit value of τ is very small, practically as if only past one or two trials back are inducing most of the effect. **c**, The five-parameter model of sensory history outperforms the single-exponential model. Two hundred iterations of fivefold cross validation were used to calculate the cross-validated bit/trial (see Methods). Accordingly, each bar shows the mean of $n = 1,000$ data points. Error bars denote s.d.



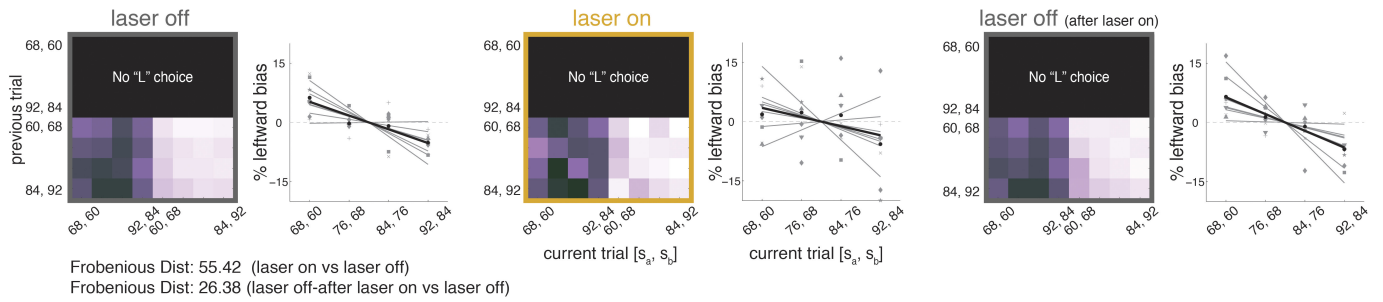
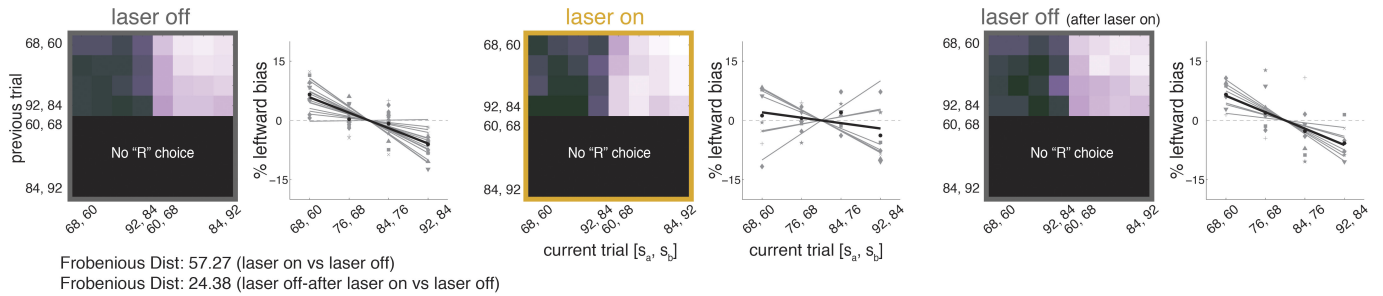
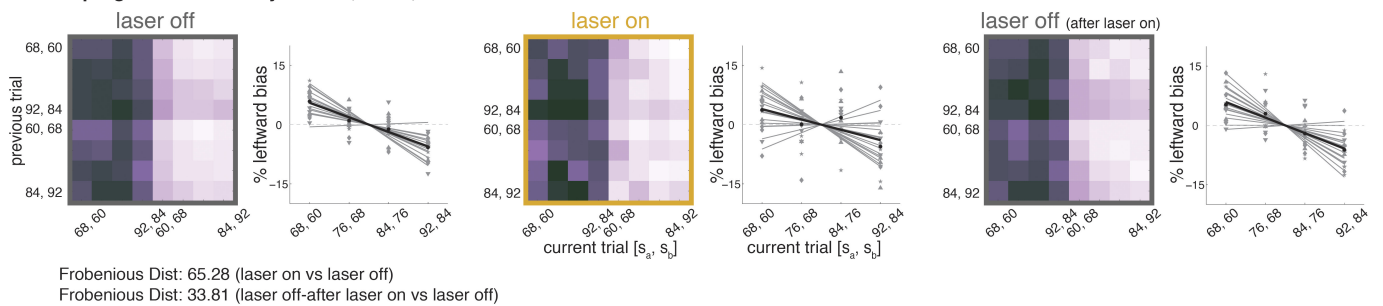
Extended Data Figure 6 | Model comparison. **a**, Model comparisons, 200 runs of fivefold cross validation were performed, on data from each rat, in order to find the best fit parameters and to compare different model fits using the cross-validated bit/trial quantity defined as the relative value of the log likelihood of each model, to the null log likelihood, normalized in log2. Removing one parameter by constraining the regression weights on the s_a stimulus of the current trial plus the weights on previous sensory stimuli to add to 1 (constrained model, in red) improved performance on cross-validated data compared to the unconstrained model (in black). A total of 12 different variants of the model are compared. Regressors are described in the box. **b**, Mean value of cross-validated bit/trial for different variants of the model as in **a**, over $n = 20$ rats. Error bars show s.e.m. Unconstrained models are shown in black, constrained models are shown in red. **c**, Top, raster plots of W_a^{-t} versus W_b^{-t} ($t = 1, 2$, from model 9). Each dot represents a subject. Pearson correlation values (r), and corresponding two-sided P values are shown for each plot. Bottom, median value of W_a^{-t} and W_b^{-t} ($t = 1, 2$), across rats. Error bars show median absolute deviation. **d**, Similar to **c**, for human subjects (auditory and tactile tasks are pooled together). Similar to rat subjects, model 9 shows the best performance for human subjects as well (data not shown). **e**, To compare the sensory-history matrix from the real data to the ones predicted from the best model fits (Fig. 2f, g), Frobenius distance norm was used, defined as the

square root of the sum of the absolute squares of the difference between elements of two matrices. Frobenius distance is a measure of similarity, and the smaller the value, the more similar the two matrices. Frobenius distance is calculated separately for individual rats and each bar shows its mean value over $n = 20$ rats. Error bars show s.e.m. Models are models A–F from Fig. 2e. **f**, Scatter plot of slopes from linear fits to percentage leftward bias (Fig. 2a) versus short-term sensory history (that is, sum of weights for s_a^{-1} , s_b^{-1} , s_a^{-2} and s_b^{-2}) from model 9. This plot shows significant correlation between the two measurements (Pearson correlation, $r = -0.66$, two-sided $P = 0.0084$, $n = 17$ rats), suggesting that when our logistic fit coefficients are particularly large, the subjects also have a particularly large contraction bias. **g**, Examining the weights in regression model 9, which is determined to be the best model, shows that the weights for sensory-history terms are significantly larger than those for the correct-side history term (paired-sample t -test, $P < 0.0001$, $n = 22$ rats). Data from individual rats are used to fit the model and bars show the mean value of sensory-history weights (in blue), and correct-side history weight (in green), over fit values from $n = 22$ rats. Error bars show s.e.m. Moreover, the sensory history regressor term, that is, sum of sensory-history weights \times regressors produces larger variance over trials (0.38) compared to the correct-side regressor (0.11), indicating a bigger impact on trial-by-trial behaviour.

a. Acute recordings during photoinhibition**b. Electrode traces****Extended Data Figure 7 | Physiological and histological confirmations.**

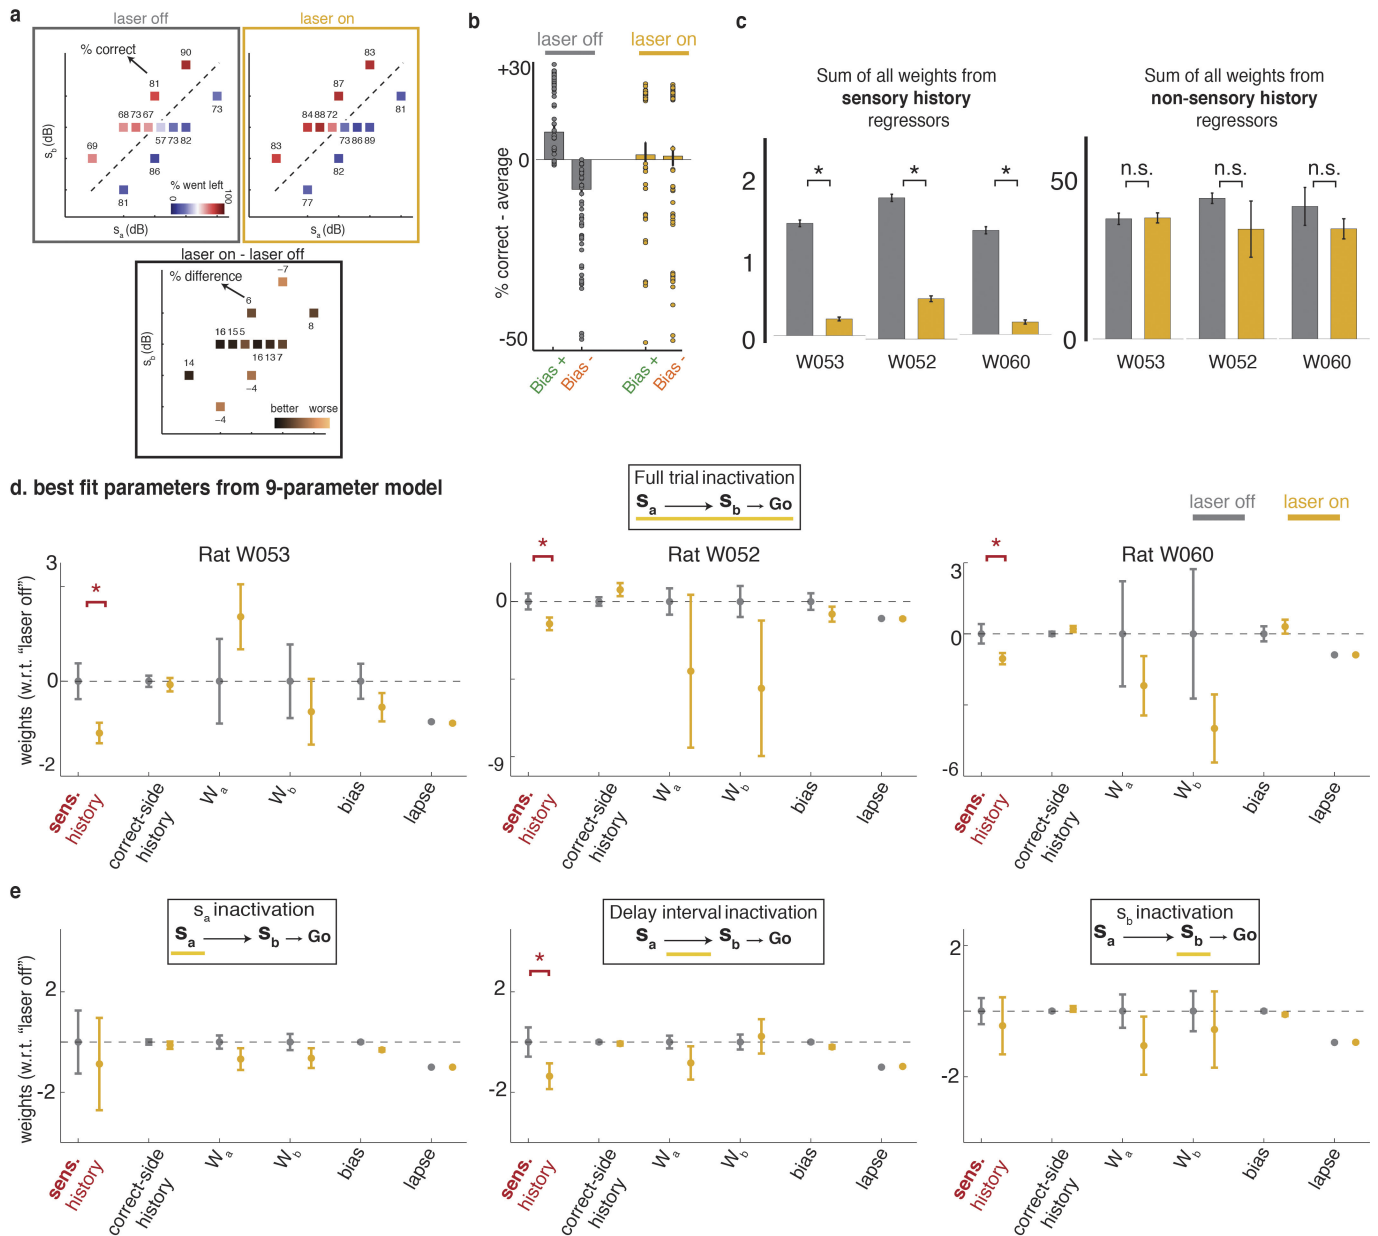
a, Physiological confirmation of optogenetic inactivation effect in an anesthetized rat. Left, single trace of acute extracellular activity of an example cell in the PPC, expressing eNpHR3.0, is shown in response to light stimulation. Laser illumination period (8 s) is marked by the light green bar. Right, raster-plot for 32 trials, for variable durations of light stimulation. The green vertical dashed line indicates the start of the laser illumination. The laser was on for variable durations of 750,

1,500, 3,000, 6,000 or 8,000 ms. The laser turning off is indicated by the vertical red dashed line. Recordings continued for 2 s after the laser was turned off. **b**, Histological localization of electrodes targeting the PPC. The inset shows an example of electrode locations in a coronal slice at anteroposterior = 3.48 from the bregma. In all cases, the electrode and fibre placements in the PPC were within between 2.8 and 4 mm posterior the bregma and between 2 and 3.5 mm lateral to the midline. Atlas panel is taken from Paxinos and Watson, 2004 (ref. 31).

a. Optogenetics - History Matrix (previous Left choice, all rewarded)**b. Optogenetics - History Matrix** (previous Right choice, all rewarded)**c. Optogenetics - History Matrix** (all trials)

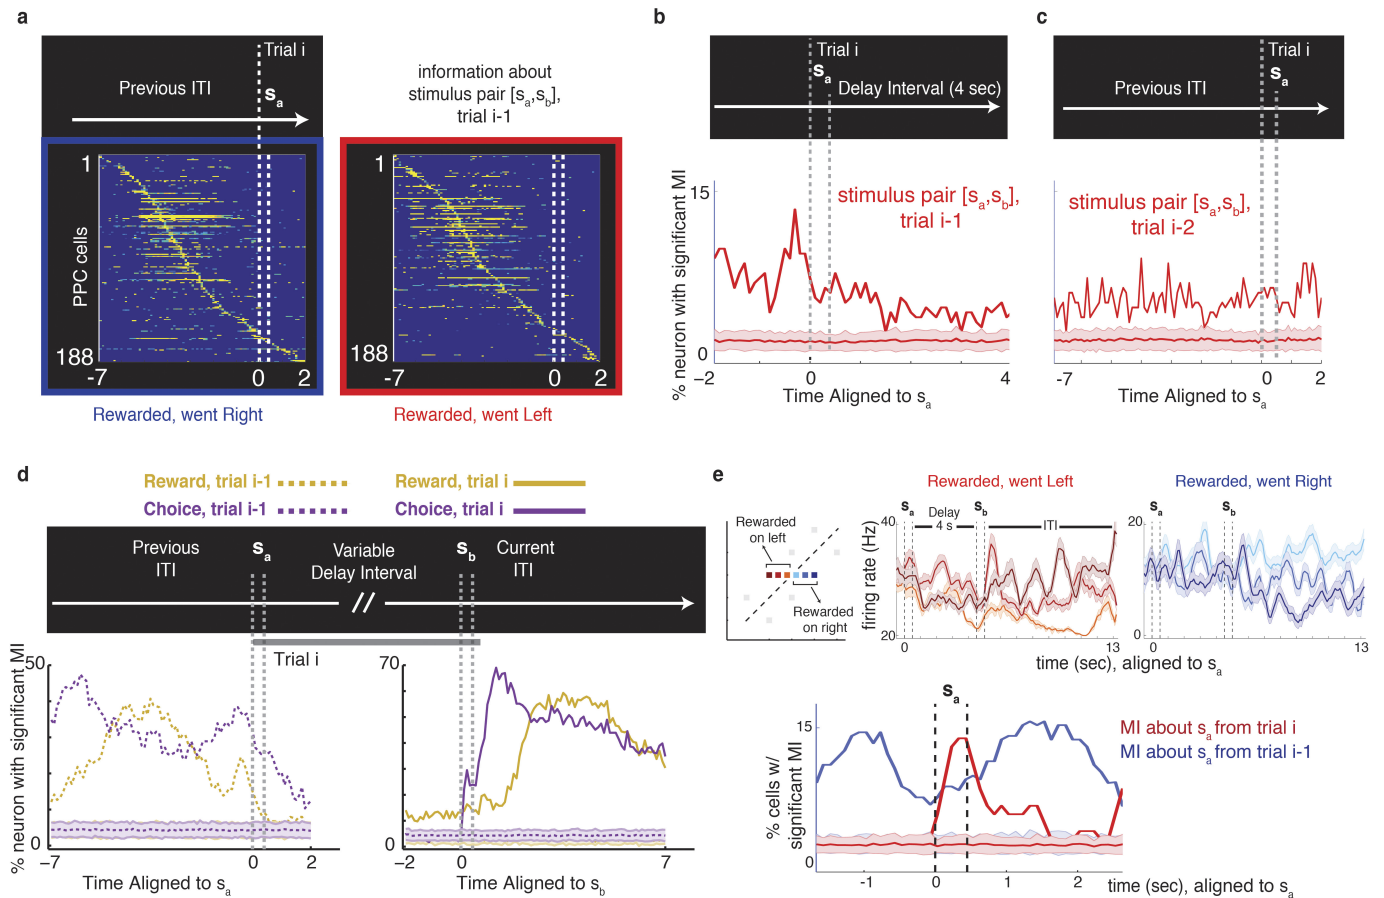
Extended Data Figure 8 | Optogenetics: PPC inhibition reduces leftward bias owing to past sensory stimuli. **a**, Sensory-history matrix and leftward biases due to past sensory stimuli, similar to Fig. 2a–c, but now for three types of trials: laser-off trials (two leftmost panels) that consist of trials with no PPC inactivation on either the current or the previous trial; laser-on trials (two middle panels) that consist of trials with PPC inactivation on the current trial; and laser-off-after-laser-on trials (two rightmost panels) that consist of trials immediately after the laser-on

trials. This last set controls for number of trials, as it contains equal numbers of trials to the laser-on condition. Modulation along the vertical indicates a previous trial effect behavioural bias as a function of the stimuli of the previous trial, for trials for which rats 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, the black line is the average over pairs. **b**, Similar to **a**, for trials for which rats went right and were rewarded. **c**, Similar to **a** for all combinations of current and previous stimuli.



Extended Data Figure 9 | Optogenetics: impact on contraction bias on the full stimulus set, individual data points and best fit parameters for non-sensory-history weights. **a**, Stimulus set and performance during optogenetic inhibition sessions, averaged over 37 sessions from 3 rats (delay interval of 2 s). Trials are grouped based on laser-off (left) and laser-on (right) conditions. The boxes represent the set of (s_a , s_b) pairs used in a session, with the colour representing the percentage that went left and the numbers above each box indicating the percentage correct. The plot in the bottom shows the difference between laser-off and laser-on conditions, with positive values indicating improved performance in laser-on conditions and negative values indicating impaired performance. **b**, **c**, Similar to Fig. 3d–f, with all data points overlaid on the bar plots.

For **b**, $n = 37$ for each bar plot (equal to the total number of inactivation sessions); for **c**, $n = 600$, from 200 iterations of threefold cross-validation data; $*P < 0.01$ from one sided t -test. **d**, Best-fit parameter values for all weights from the nine-parameter model (short-term sensory-history model, constrained version, Fig. 2d, e). Values are plotted as their mean once the average value from the laser-off condition is subtracted. Except for the sensory history, none of the other weights were significantly affected by optogenetic inactivation of the PPC. Error bars show s.d. ($n = 600$, 200 iterations of threefold cross-validation; $*P < 0.01$ from one sided t -test). **e**, Similar to **d**, for period-selective optogenetic inhibitions, in which the PPC is selectively inhibited during the first stimulus s_a (left), delay interval (middle) or second stimulus s_b (right).



Extended Data Figure 10 | Mutual information. **a**, Sensory-history coding, one trial back, population analysis, each row represents the time course of significant values of mutual information between the firing rate of a cell and the stimulus pair (s_a , s_b) presented on the previous trial. Data from all trials with variable delay duration (minimum of 2 s) were pooled and plots are aligned to the beginning of s_a . Data from $n = 5$ rats, and only cells with significant values of mutual information values are included. When estimating the mutual information, spurious information values can be attributed to the inherent correlations between task parameters, such as sensory stimuli and choice. To overcome this, conditional mutual information was calculated only when trials with same previous choice and reward status were considered, and sensory inputs were the only variable. Left, on the previous trial rats went right and were rewarded. Right, on the previous trial rats went left and were rewarded. **b**, Sensory-history coding, one trial back, percentage 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 a delay interval larger than 4 s are included in this analysis. **c**, Sensory-history coding, two trials back, percentage of cells with significant coding of stimuli presented two trials in the past (trial $i - 2$), aligned to the start of trial i . Shaded horizontal areas show the mean \pm s.d. of the percentage of neurons with significant mutual information (MI), calculated from random sets built by shuffling the firing rates of neurons and conditions. **d**, Percentage of cells with significant coding of a rat's choice and reward status, on both the current trial (solid lines) and previous trial (dashed lines), when time is aligned to the

current trial, either s_a (left), or s_b (right). Shaded horizontal areas show the mean \pm s.d. of the percentage of neurons with significant mutual information, calculated from random sets built by shuffling the firing rates of neurons and conditions. **e**, In the standard stimulus set (Fig. 1b, (s_a , s_b) pairs along the diagonal lines), knowledge of the rat's choice of side, whether it was rewarded or not, and one of either s_a or s_b enables unique identification of the other stimulus (s_b or s_a). Therefore, in order to probe whether neurons carried information for different values of s_a itself (as opposed to a combination of choice, reward and s_b), we ran recording sessions with psychometric stimuli added to the standard stimulus set (top left). In this way, three different values of s_a are assigned to one fixed value of s_b and one fixed action (left in different shades of red, and right in different shades of blue). The firing rate of an example neuron is shown in response to different values of s_a , only for trials in which the rat responded by going left (middle graph) or right (right graph) after the 'go' cue, was rewarded, and the delay interval was 4 s. Even though choice, reward and s_b are fixed, firing rates clearly differentiate values of s_a . The bottom graph shows a summary of population analysis from psychometric recording sessions (as in the examples in the graphs above), showing the percentage of cells with significant coding of s_a from trial i (red) or trial $i - 1$ (blue, $n = 142$ cells). Shaded horizontal areas show the mean \pm s.d. of the percentage of neurons with significant mutual information, calculated from random sets built by shuffling the firing rates of neurons and conditions.

Life Sciences Reporting Summary

Nature Research wishes to improve the reproducibility of the work that we publish. This form is intended for publication with all accepted life science papers and provides structure for consistency and transparency in reporting. Every life science submission will use this form; some list items might not apply to an individual manuscript, but all fields must be completed for clarity.

For further information on the points included in this form, see [Reporting Life Sciences Research](#). For further information on Nature Research policies, including our [data availability policy](#), see [Authors & Referees](#) and the [Editorial Policy Checklist](#).

► Experimental design

1. Sample size

Describe how sample size was determined.

Sample sizes were determined by the limits of what data could be collected within a reasonable timeframe and standards of the field. The effects we report were all reproducible across individual animals within our sample. This is documented by showing data and analysis from each individual animals, either in the main text, or in the Extended Data figures.

A total of 33 male Long–Evans rats were used for this study. Of these, 25 were used for behavioral assessments (total of 468,165 trials), 6 rats were used for neural recordings (total of 1081 neurons), and 7 for optogenetic inactivations (5 experimental animals, two controls, for a total of 30,156 trials).

Moreover, total of 25 human subjects participated in our research. Each human subject provided minimum number of 300 trials, summing to the total of 17,317 trials.

All statistical tests were made between groups with similar sample sizes.

2. Data exclusions

Describe any data exclusions.

No data has been excluded, except for the electrophysiological analysis, in which from the total of 936 single or multiunits that were recorded in 5 rats, only neurons with overall firing rate of at least 2 Hz were included in the analysis. These neurons summed to total of 361.

3. Replication

Describe whether the experimental findings were reliably reproduced.

No attempt at replication of the exact results has been done outside of our study. However, we are reporting results for all individual subjects, as well as the average effect, in behavioral assessments, optogenetic inactivations, and neural recordings, and show how consistent are the results across different subjects.

Moreover, our human psychophysics experiments consist of two different setups, one in Princeton University (Princeton, NJ, USA) and another one in International School for Advance Studies (SISSA, Trieste, Italy). We show the behavioral effect of sensory history can be replicated reliably in both setups, in two different modalities, and results from both human settings are remarkably similar to the result found in our rat experiments.

4. Randomization

Describe how samples/organisms/participants were allocated into experimental groups.

All subjects were randomly allocated into experimental groups.

5. Blinding

Describe whether the investigators were blinded to group allocation during data collection and/or analysis.

Investigators were not blinded to experimental groups during data collection or analysis.

Note: all studies involving animals and/or human research participants must disclose whether blinding and randomization were used.

6. Statistical parameters

For all figures and tables that use statistical methods, confirm that the following items are present in relevant figure legends (or in the Methods section if additional space is needed).

- n/a Confirmed
- The exact sample size (n) for each experimental group/condition, given as a discrete number and unit of measurement (animals, litters, cultures, etc.)
 - A description of how samples were collected, noting whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
 - A statement indicating how many times each experiment was replicated
 - The statistical test(s) used and whether they are one- or two-sided (note: only common tests should be described solely by name; more complex techniques should be described in the Methods section)
 - A description of any assumptions or corrections, such as an adjustment for multiple comparisons
 - The test results (e.g. P values) given as exact values whenever possible and with confidence intervals noted
 - A clear description of statistics including central tendency (e.g. median, mean) and variation (e.g. standard deviation, interquartile range)
 - Clearly defined error bars

See the web collection on [statistics for biologists](#) for further resources and guidance.

► Software

Policy information about [availability of computer code](#)

7. Software

Describe the software used to analyze the data in this study.

All codes used in this study were custom codes written in MATLAB2013. 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 request.

For manuscripts utilizing custom algorithms or software that are central to the paper but not yet described in the published literature, software must be made available to editors and reviewers upon request. We strongly encourage code deposition in a community repository (e.g. GitHub). *Nature Methods* [guidance for providing algorithms and software for publication](#) provides further information on this topic.

► Materials and reagents

Policy information about [availability of materials](#)

8. Materials availability

Indicate whether there are restrictions on availability of unique materials or if these materials are only available for distribution by a for-profit company.

no unique material was used.

9. Antibodies

Describe the antibodies used and how they were validated for use in the system under study (i.e. assay and species).

For all antibodies, as applicable, provide supplier name, catalog number, clone name, and lot number. Also describe the validation of each primary antibody for the species and application, noting any validation statements on the manufacturer's website, relevant citations, antibody profiles in online databases, or data provided in the manuscript OR state that no antibodies were used.

10. Eukaryotic cell lines

a. State the source of each eukaryotic cell line used.

Provide information on cell line source(s) OR state that no eukaryotic cell lines were used.

b. Describe the method of cell line authentication used.

Describe the authentication procedures for each cell line used OR declare that none of the cell lines used have been authenticated OR state that no eukaryotic cell lines were used.

c. Report whether the cell lines were tested for mycoplasma contamination.

Confirm that all cell lines tested negative for mycoplasma contamination OR describe the results of the testing for mycoplasma contamination OR declare that the cell lines were not tested for mycoplasma contamination OR state that no eukaryotic cell lines were used.

d. If any of the cell lines used are listed in the database of commonly misidentified cell lines maintained by [ICLAC](#), provide a scientific rationale for their use.

Provide a rationale for the use of commonly misidentified cell lines OR state that no commonly misidentified cell lines were used.

► Animals and human research participants

Policy information about [studies involving animals](#); when reporting animal research, follow the [ARRIVE guidelines](#)

11. Description of research animals

Provide details on animals and/or animal-derived materials used in the study.

Long–Evans rats (*Rattus norvegicus*) between the ages of 6 and 24 months were used for this study.

Policy information about [studies involving human research participants](#)

12. Description of human research participants

Describe the covariate-relevant population characteristics of the human research participants.

Two different groups of human subjects participated in our research:
1) 11 human subjects (8 males and 3 females, ages 22–40) were tested in Princeton University, in the "auditory" version of our task, and all gave their informed consent. The consent procedure and the rest of the protocol were approved by the Princeton University Institutional Review Board.

2) 14 human subjects (8 males and 6 females, ages 22–35) were tested in International School for Advance Studies (SISSA, Trieste, Italy), in the "tactile" version of our task. Protocols conformed to international norms and were approved by the Ethics Committee of the International School for Advanced Studies.

Participants were paid to be part of the study and were naive to the main conclusions of the study.