nature neuroscience

ARTICLES

https://doi.org/10.1038/s41593-021-00840-6

Check for updates

Attractor dynamics gate cortical information flow during decision-making

Arseny Finkelstein[1,3], Lorenzo Fontolan[1,3], Michael N. Economo[1], Nuo Li[1,2], Sandro Romani[1]✉ and Karel Svoboda[1]✉

Decisions are held in memory until enacted, which makes them potentially vulnerable to distracting sensory input. Gating of information flow from sensory to motor areas could protect memory from interference during decision-making, but the underlying network mechanisms are not understood. Here, we trained mice to detect optogenetic stimulation of the somatosensory cortex, with a delay separating sensation and action. During the delay, distracting stimuli lost influence on behavior over time, even though distractor-evoked neural activity percolated through the cortex without attenuation. Instead, choice-encoding activity in the motor cortex became progressively less sensitive to the impact of distractors. Reverse engineering of neural networks trained to reproduce motor cortex activity revealed that the reduction in sensitivity to distractors was caused by a growing separation in the neural activity space between attractors that encode alternative decisions. Our results show that communication between brain regions can be gated via attractor dynamics, which control the degree of commitment to an action.

Information flow can be dynamically routed across the brain[1,2]. Routing of sensory inputs is crucial for behavior so that only relevant stimuli will drive an appropriate response. Studies of attention have shown that sensory inputs can be gated at the behavioral and neurophysiological level[3]. Classically, gating has been associated with increased neural responses to the attended sensory stimulus[4,5] and decreased responses to irrelevant or unattended stimuli (that is, distractors)[6,7].

Several cellular and network mechanisms have been proposed to underlie gating of sensory stimuli, including top-down modulation in the neocortex[8] or thalamic gain[9,10], changes in the balance between excitation and inhibition so that incoming inputs are canceled by local inhibition[11–13], and selective transmission of information via the synchronization of neural population activity across brain regions[2,14,15]. These gating mechanisms imply that as the inputs propagate from one population of neurons to the next, neural responses to distractors should be suppressed compared with responses elicited by relevant stimuli (that is, distractor suppression). However, specific neural network dynamics also allow for context-dependent computation[16,17] in which only certain network states are sensitive to incoming stimuli, whereas other network states are unaffected. According to this scheme, distractors are not suppressed altogether but rather lose influence on future neural dynamics that are relevant for behavior.

Here, we investigated the mechanisms controlling the information flow between and within brain areas during decision-making. When sensation and action are separated in time (by a memory or a delay epoch), perceptual decisions persist in the memory until movement execution and thus become vulnerable to sensory interference. Behavioral studies using monkeys have suggested that additional inputs arriving during the delay epoch can be effectively gated[18], which is critical for protecting the decision against interferences.

On the neural level, decisions about upcoming actions (motor plans) are manifested by preparatory activity observed in the motor cortex and related areas[19–22], which predicts future movements[23–27]. To probe the information flow during motor planning in the delay epoch, we studied the ability of calibrated stimuli to perturb the motor plan for the upcoming movement. We found that the communication channel between the sensory and the motor cortex remained open throughout the delay epoch, which allowed inputs from the sensory cortex to reach the motor cortex without suppression. Nevertheless, stimuli of identical strength became progressively less capable of affecting preparatory activity and the behavioral response. Our results show that temporal modulation of attractor dynamics gradually renders the motor plan insensitive to additional input, and reveal a novel network mechanism underlying the gating of information between cortical areas. More broadly, this work reveals how a dynamic modification of input–output mapping in cortical circuits allows for flexible regulation of information flow.

Results

Sensorimotor transformation evoked by direct cortical photostimulation. The anterior lateral motor cortex (ALM) is a critical node in a decision-making circuit that controls directional licking in response to sensory inputs[24,26,27]. To study communication principles in this circuit, we developed a delayed-response task in which mice learned to respond by directional licking to the optogenetic stimulation of genetically defined neurons in layer 4 of the vibrissal somatosensory cortex (vS1; Fig. 1a). A photostimulus applied to the left vS1 during the sample epoch (delineated by auditory cues) instructed the mice to lick right, whereas the absence of a photostimulus instructed the mice to lick left (Fig. 1b,c). The time of action was signaled by an auditory go cue played after a 2-s delay epoch. Training mice to respond to photostimulation of sensory cortex neurons[28] allowed us to apply calibrated photostimuli of different amplitudes at different times and to investigate cortical gating mechanisms isolated from upstream modulations (for example, in the sensory thalamus[9,10]). This paradigm also allowed us to measure the light-evoked responses at the source (vS1) and in downstream motor areas (ALM) during different stages of decision-making.

[1]Janelia Research Campus, Howard Hughes Medical Institute, Ashburn, VA, USA. [2]Department of Neuroscience, Baylor College of Medicine, Houston, TX, USA. [3]These authors contributed equally: Arseny Finkelstein, Lorenzo Fontolan. ✉e-mail: romanis@janelia.hhmi.org; svobodak@janelia.hhmi.org

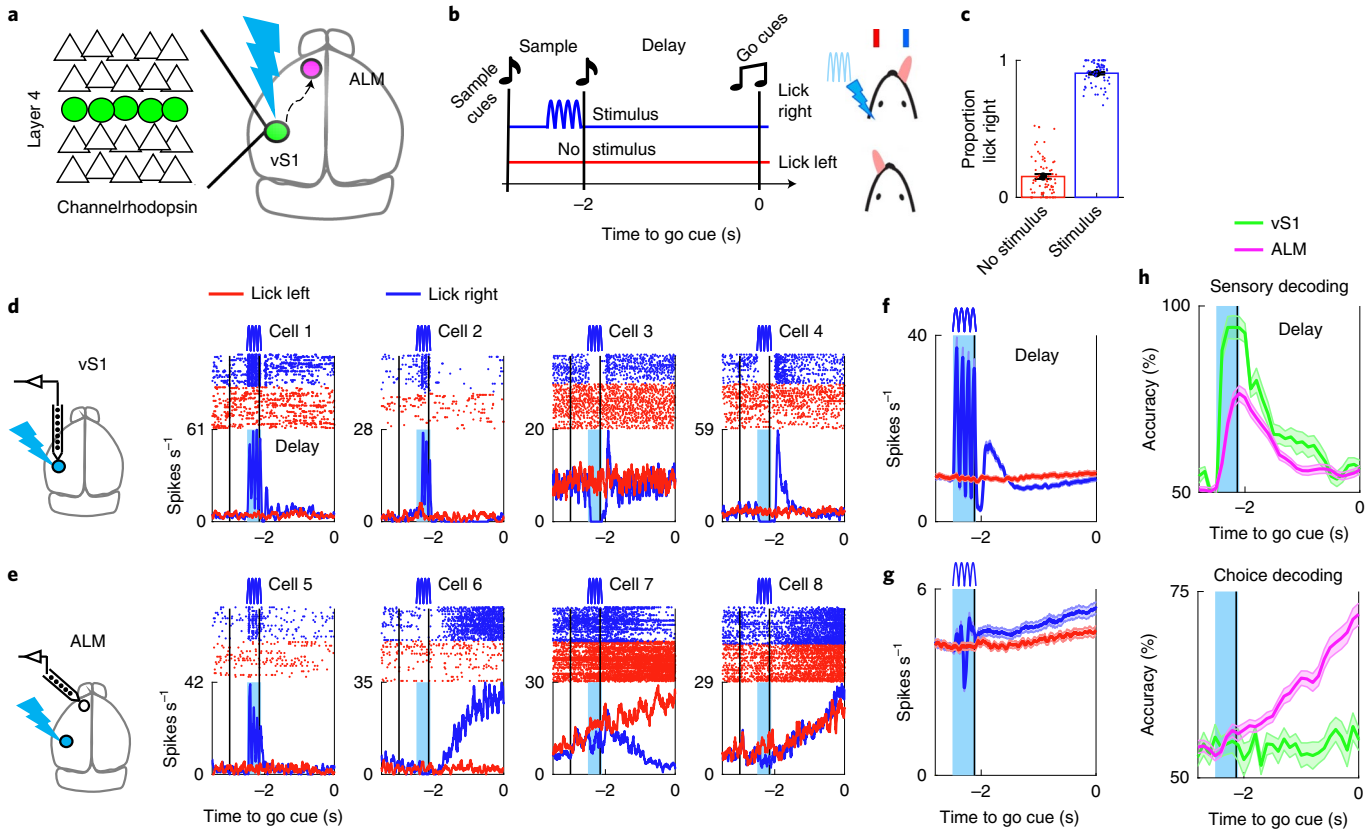

Fig. 1 | Sensorimotor transformations evoked by direct cortical photostimulation. a, Cortical regions involved in tactile decision-making. Left: schematics of channelrhodopsin-2–EYFP-expressing neurons in layer 4 of the vS1 (barrel cortex) of transgenic mice. **b**, Basic task. Mice had to lick right in response to vS1 photostimulation during the sample epoch (delineated by auditory cues). An absence of a stimulus instructed mice to lick left. A go cue after a 2-s delay epoch signaled the start of the response epoch. **c**, Proportion of lick-right responses ($n=5$ mice, 67 sessions). Each point corresponds to a behavioral session (mean ± s.e.m.). **d**,**e**, Silicon probe recordings showing example neurons with transient responses in the vS1 (**d**) and a mixture of transient and sustained responses in the ALM (**e**). Spike raster and trial-averaged spike rates on correct lick-right and lick-left trials are shown for each example neuron. The cyan vertical bar indicates photostimulation. The black vertical bars delineate the trial epochs. **f**,**g**, Grand-average population responses of putative pyramidal cells on lick-right (blue) and lick-left (red) trials in the left vS1 (**f**; $n=547$ cells) and the left ALM (**g**; $n=2{,}417$ cells). **h**, Decoding sensory-related (top) and choice-related (bottom) activity in the vS1 and the ALM at different times along the trial. Shaded area represents the s.e.m. computed across sessions.

We will refer to photostimulation-evoked neural activity in the vS1 as 'sensory input'.

We used silicon probes to record activity from individual putative pyramidal cells in the vS1 (the site of photostimulation) and the ALM (Fig. 1d,e and Extended Data Fig. 1). Photostimulation transiently modulated the spike rate of most vS1 neurons (55%, 300 out of 547 neurons; Fig. 1d). On the population level, responses of vS1 neurons peaked during the time of photostimulation and rapidly decayed thereafter (Fig. 1f). In the ALM, only a small proportion of neurons (13%, 316 out of 2,417 neurons) responded transiently to photostimulation (Fig. 1e, example cell 5). These neurons were more abundant on the side of the photostimulated vS1 (that is, left hemisphere ALM; Extended Data Fig. 1b, cluster 4). The majority of ALM neurons in both hemispheres showed preparatory activity during the delay epoch (Extended Data Fig. 1b, clusters 1–3 and 6) in the form of either trial-type selective (Fig. 1e, cells 6 and 7) or nonselective (Fig. 1e, cell 8) ramping. Putative fast-spiking interneurons in the vS1 and the ALM showed a diversity of response profiles (Extended Data Fig. 2a–c), which suggests that their activity might not simply reflect the average population activity of excitatory neurons (Fig. 1f,g).

Trial-type-selective neural activity in the vS1 and the ALM could encode sensory evidence or the chosen action[24,29]. In the latter case, the same sensory input would result in a different neural activity profile on correct versus error trials. We used vS1 and ALM neural activity on correct and error trials to decode sensory (photostimulus) information or the upcoming choice (lick direction) at different times during the task. Sensory-related information was transient and peaked first in the vS1 followed by the ALM (Fig. 1h, top). Choice-related information was nearly absent in the vS1, but emerged in the ALM early during the delay epoch and increased with time (Fig. 1h, bottom). This demonstrates that brief photostimulation causes transient sensory-related activity in both the sensory cortex and the motor cortex, whereas the motor cortex encodes the chosen action during the delay epoch.

Gating of sensory input during decision-making. We next asked how sensory input is processed as time unfolds during motor planning. We first trained mice on the basic task (Fig. 1b) and then probed their behavior in response to weak distractors, which consisted of photostimuli delivered at different times during lick-left trials (Fig. 2a, left, and Methods; for detailed behavior analyses, see Supplementary Fig. 1). These distractor-naive mice were not exposed to distractors during training. The presence of weak distractors during sample or early delay significantly increased lick-right responses compared with no-stimulus trials. In contrast, a weak distractor presented late in the delay epoch did not affect behavior. This demonstrates that a temporal gating effect is occurring (Fig. 2c, left).

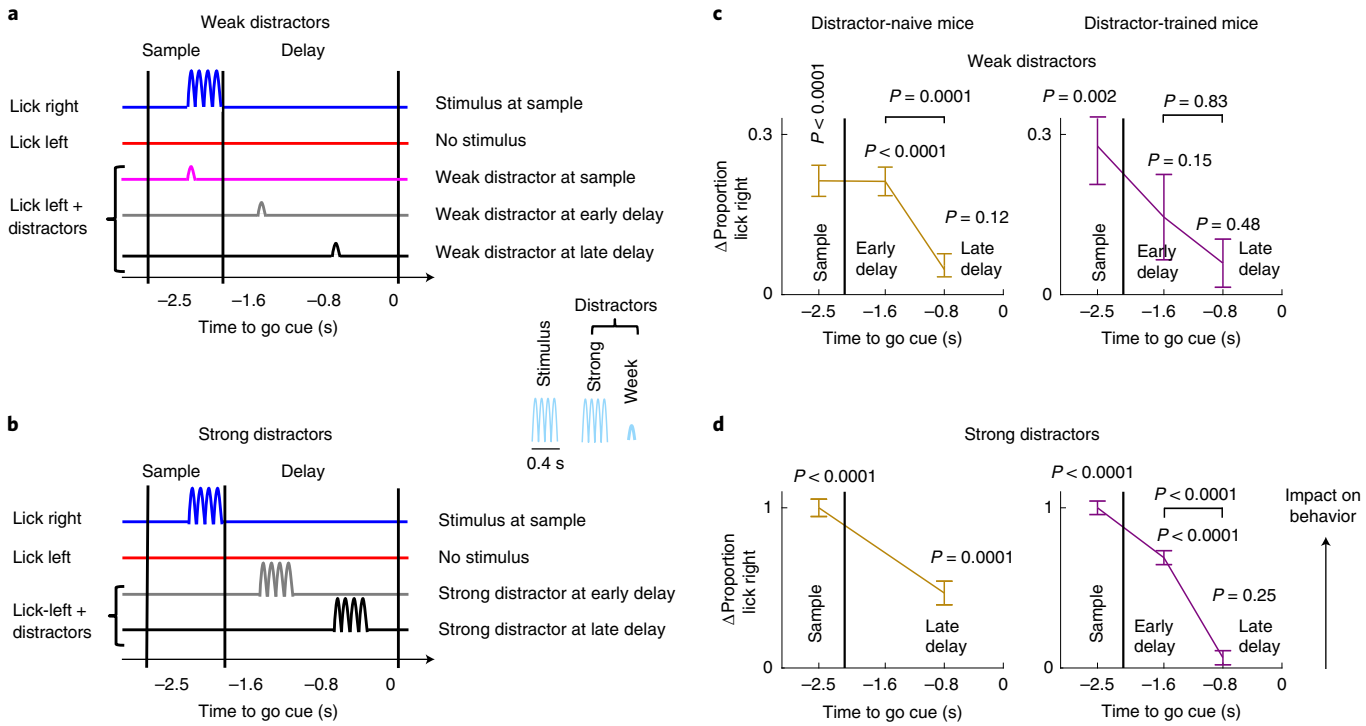

Fig. 2 | Gating of sensory inputs during decision-making. a,**b**, Schematics of the tasks with weak (**a**) and strong (**b**) distractors. **c**,**d**, Behavioral impact of weak (**c**) and strong (**d**) distractors delivered at different times in distractor-naive mice (gold) and distractor-trained mice (purple). Behavioral impact was quantified as the change in proportion of lick-right responses during photostimulation trials relative to no-stimulus trials. This value was normalized so that the change in proportion of lick-right trials in response to stimulus during the sample epoch was set to 1 (Methods). Data are shown as mean ± s.e.m. (error bars) across sessions; *P* values indicate statistical significance by paired Student's *t*-test (two-sided) for photostimulation trials versus non-stimulus trials and for early-delay versus late-delay distractor trials. The following numbers of mice and sessions were analyzed: distractor-naive mice, $n = 5$ mice, 67 sessions for weak distractors and $n = 2$, 17 sessions for strong distractors; distractor-trained mice, $n = 4$ mice, 19 sessions for weak distractors and $n = 4$ mice, 67 sessions for strong distractors.

A strong late-delay distractor (of the same amplitude and duration as the sample stimulus) was not fully gated and increased the proportion of lick-right responses (Fig. 2d, left), which suggests that gating depends on both the timing and the strength of the distractors.

A different group of mice was exposed to distractors of variable strength during training (distractor-trained; Methods) and without receiving a reward for licking right in response to distractors. We probed the response of distractor-trained mice to weak and strong distractors delivered during the delay epoch. Weak distractors delivered early or late in the delay did not significantly increase lick-right responses, which suggests that weak distractors were gated (Fig. 2c, right). A strong early-delay distractor often resulted in a right lick, whereas a strong late-delay distractor did not affect behavior (Fig. 2d, right). In other words, the same stimulus that was consistently acted on when presented during the sample epoch and, to some extent, early in the delay epoch was ignored when delivered late in the delay epoch. This demonstrates that temporal gating occurs even for strong distractors.

Gating without distractor suppression by population dynamics in the ALM. We considered two alternative mechanisms to explain gating. Hypothesis 1: stimulus- and distractor-evoked responses in the vS1 or in the ALM are suppressed when delivered later during the delay epoch (distractor suppression; Fig. 3a). Hypothesis 2: even though stimuli can reach the ALM unaltered throughout the delay epoch, network dynamics make the motor plan progressively more robust as the moment to act approaches (Fig. 3b). Therefore, the ability of distractors to switch the motor plan would diminish with time, which results in gating of the sensory inputs at the behavioral level.

To investigate potential gating mechanisms on the neural level, we first tested whether temporal gating operated early during sensory processing by analyzing neural responses in the vS1. We focused on distractor-trained mice because they exhibited gating even though the stimulus and the distractors were identical (Fig. 2d, right). Population spike-rate responses to photostimulation did not show distractor suppression at any time point during the task (Fig. 3c and Extended Data Fig. 3a–c, top), which indicates that gating is absent in the vS1. Similarly, analysis of the population spike-rate in the ALM showed that distractors reached the ALM without attenuation throughout the delay epoch (Fig. 3d and Extended Data Fig. 3a–c, bottom) despite the attenuation of lick-right behavioral responses to distractors presented during the delay epoch. Because the ALM contained both stimulus- and choice-related information (Fig. 1h), we performed dimensionality reduction on the population activity (Fig. 3e and Extended Data Fig. 4) to separate stimulus-related activity (stimulus mode) from the activity predicting the upcoming movement (choice mode). We examined the ALM population activity projected on the stimulus mode and found no suppression of stimulus-evoked responses during the delay epoch (Fig. 3f and Extended Data Fig. 3d,e), thereby ruling out hypothesis 1. Similar to distractor-trained mice, distractor-naive mice gated distractors in a time-dependent manner and displayed no suppression of distractors in the vS1 or the ALM during the delay epoch (Supplementary Fig. 2).

Decoding of upcoming actions from the choice mode in the presence of distractors (Fig. 3g) showed that changes in the choice mode closely predicted the behavioral responses to distractors at different times of the task. For trials in which the animal ignored the distractor and licked left, choice mode activity was initially perturbed

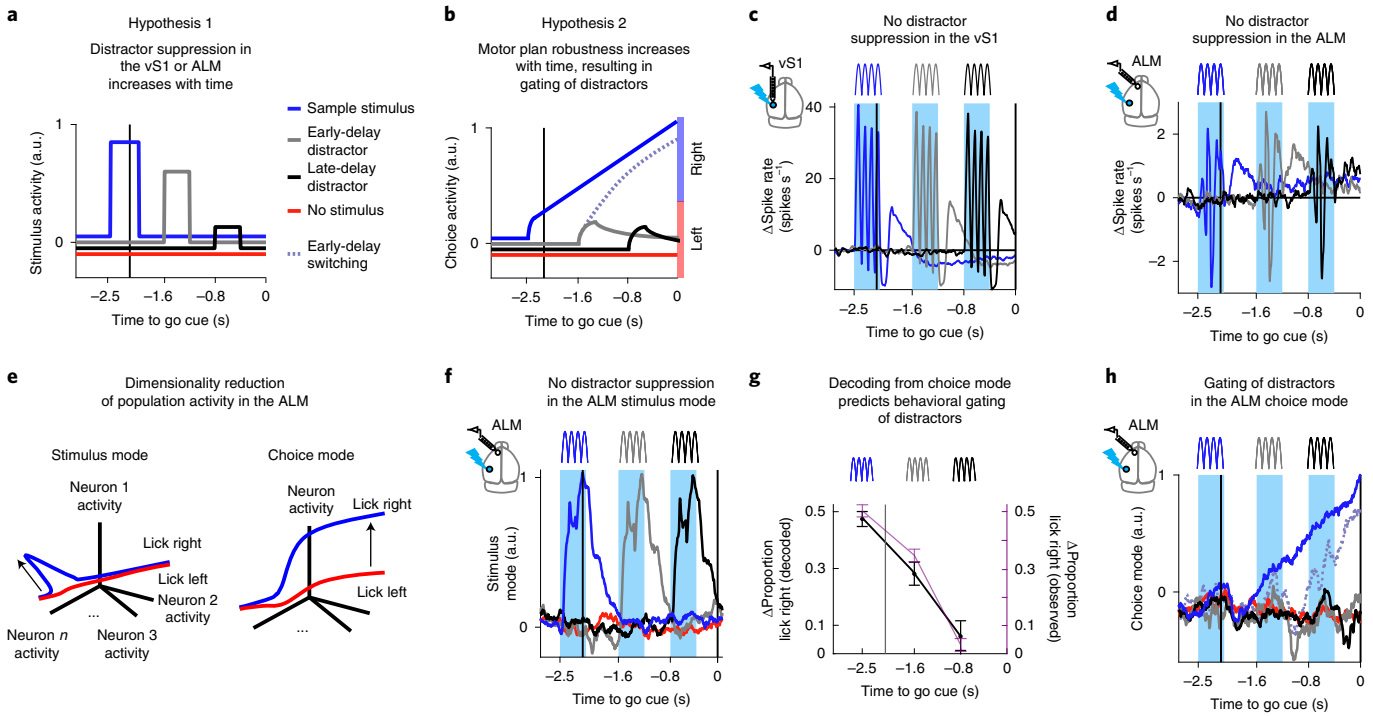

Fig. 3 | Gating without distractor suppression by population dynamics in the ALM. a,**b**, Potential mechanisms for gating. **a**, Distractors are suppressed when delivered closer to the time to act in the vS1 or at the input stages to the ALM. **b**, Distractors can reach the ALM unaltered throughout the delay epoch, but population dynamics increase motor plan robustness with time, leading to gating of the late-delay distractor. a.u., arbitrary units. **c**,**d**, Session-averaged population responses to the presentation of the stimulus (blue), the early-delay distractor (gray) and the late-delay distractor (black) in the vS1 (**c**) and the ALM (**d**). No attenuation of the response was observed during delay. **e**, Schematic of the targeted dimensionality reduction approach to define the stimulus mode (left) and the choice mode (right) in the neural activity space (Methods). **f**, Session-averaged neural activity of the left ALM projected on the stimulus mode. Lick-left trajectory without stimulation (red), early-distractor (gray) or late-distractor (black) trajectories. Lick-right trajectory during sample-epoch stimulation (blue). Average trajectories were computed using both correct and error trials. The black bars delineate the delay epoch. The timing of photostimulation on different trial types is indicated by stimulus-profile icons and cyan vertical bars. No suppression of the early- or late-delay distractor was observed. **g**, The proportion of lick-right responses relative to control trials based on behavior (purple, as in Fig. 2d; $n=67$ sessions) or on the decoding of choice from neural activity on single trials projected on the choice mode (black; Methods; $n=17$ sessions). Data shown as mean ± s.e.m. across sessions. **h**, Session-averaged neural activity of the left ALM projected on the choice mode in response to vS1 photostimulation. Correct trials lick-left trajectory without stimulation (red), correct trials early-distractor (gray) or late-distractor (black) trajectories. Correct lick-right trajectory during sample-epoch stimulation (blue). Trials with an early-delay distractor that resulted in a lick-right response (switching trials, dashed gray). The high correspondence between the temporal gating of distractors observed on the neural and behavioral level (**g**) confirmed that neural trajectories along the choice mode were more robust against the late-delay distractor compared with the early-distractor (**h**). Analyses in this figure were performed using data from distractor-trained mice.

by the distractor but then recovered to its original trajectory (Fig. 3h, 'robust trials' full gray and black, and Extended Data Fig. 5), which resulted in gating of the distracting inputs. This observation indicated that for robust trials, the choice mode was still affected by sensory inputs during the delay epoch. However, the impact of distractors on the choice mode was transient for robust trials, and the proportion of trials in which neural activity was robust was significantly increased for a late-delay distractor (Fig. 3g). In contrast, for trials with an early-delay distractor resulting in right licks (switching trials), neural activity exhibited switches from lick-left to lick-right trajectories (Fig. 3h, dashed gray), which is consistent with hypothesis 2. These results suggest that sensory inputs can reach the ALM with the same intensity at all times, whereas the selected motor plan becomes increasingly robust to sensory inputs during the delay epoch (hypothesis 2), which yields a temporal gating of information flow between sensory and motor areas.

Temporal modulation of attractor dynamics as a mechanism for sensory gating. To understand the circuit dynamics underlying temporal gating, we trained recurrent neural networks (RNNs) to model ALM activity (Fig. 4a,b). We required the activity in the RNNs to match the heterogeneous neural dynamics[30] observed in the ALM recordings (Extended Data Fig. 1). We trained RNN units to produce an output that matched the peristimulus time histogram (PSTH) of neurons recorded in the left ALM (Fig. 4b and Supplementary Fig. 3) computed during correct left or right trials without distractors. An analysis of ALM activity showed that a significant proportion of the variance during the delay epoch (Extended Data Fig. 4c,d) was explained by a nonselective ramping component of the neural activity (ramping mode; Methods), which likely reflects an external input[26,27]. Hence, the RNN was trained with an external input that mimicked the nonselective ramping component that we found in the ALM. Because this RNN was trained without using trials with distractors and received an external ramping input, we called it the 'distractor-free ramping network' (Fig. 4a). At the end of training, network dynamics were robust to both noise in the dynamics and small structural changes in network connectivity (Methods). We next tested the ability of the RNN to generalize to trial types that were not utilized to train the RNN by presenting it with distractors during the delay epoch. Even though the network was not trained to reproduce the neural activity in trials with distractors, its dynamics showed gating of the late-delay distractor (Fig. 4c,d) that resembled

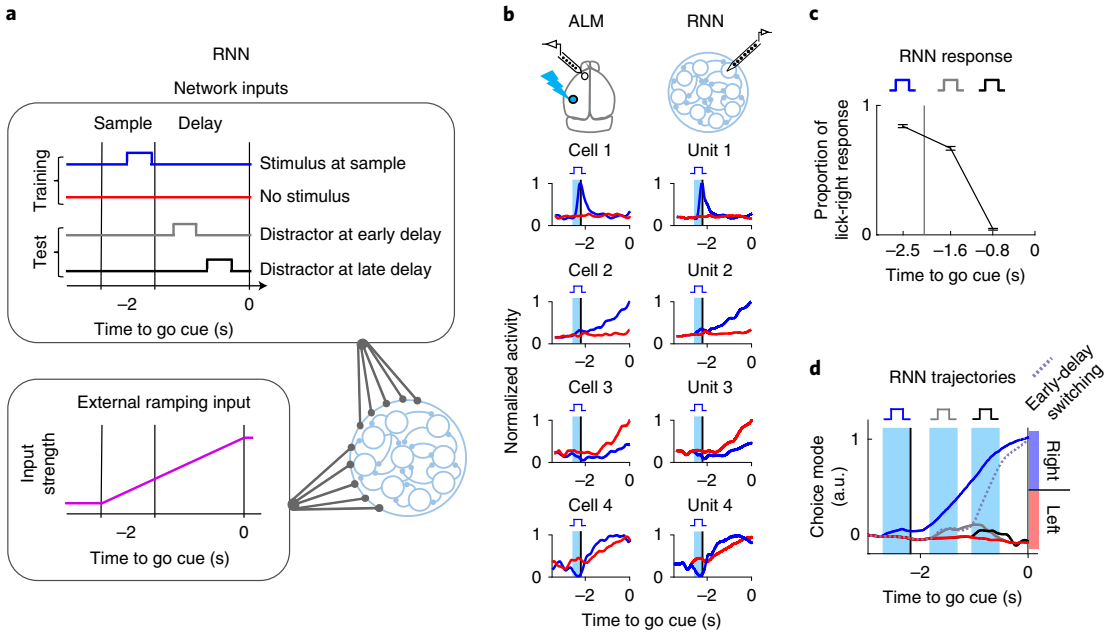

Fig. 4 | A heterogeneous RNN model of the ALM exhibits temporal gating. a, ALM activity was modeled as a network of recurrently connected nonlinear rate units (RNN; $n=668$ units). A pulse delivered to a subset of units mimicked the optogenetic stimulus instructing lick-right trials. The sample epoch was delimited by two additional inputs representing the auditory cues. A linear ramping input (black) was delivered to the network during the sample epoch and the delay epoch. **b**, Each unit (right, example RNN units) was successfully trained to reproduce the PSTH of a recorded neuron (left, example ALM cells) during correct right-lick (blue) and left-lick (red) trials. **c**, The proportion of right-lick outputs generated by the RNN model in response to the stimulus (blue) and distractors during early (gray) and late (black) delay. The network outcome (left or right) was defined by whether the activity along the choice mode (computed in the activity space of RNN units) crossed the choice boundary at the delay end (Methods). Data shown as mean ± s.e.m. across sessions ($n=10$, 100 trials each). **d**, RNN activity projected on the choice mode for correct trials (solid lines). The color scheme is as in Fig. 3h. Dashed gray line indicates trials with an early-delay distractor that resulted in a right-lick response (switching trials).

the gating effect observed in ALM recordings (Fig. 3g,h). Similar results were obtained when the network was explicitly trained to reproduce the activity obtained from early-delay distractor trials (called the 'distractor-trained ramping network'; Extended Data Fig. 6h,i) and when the network was trained on PSTHs of ALM neurons recorded from both hemispheres (Extended Data Fig. 7). We trained additional networks without an external ramping input ('autonomous' RNNs) in which the slow ramping dynamic was generated exclusively via recurrent dynamics (Extended Data Fig. 6). While all networks exhibited some form of temporal gating of stimuli, only the dynamics of RNNs trained with an external ramping stimulus were consistent with ALM recordings (Extended Data Fig. 6k).

We next reverse engineered[31] the RNNs to uncover the mechanism underlying temporal gating (for the distractor-free ramping RNN, see Fig. 5; for the other RNNs, see Extended Data Fig. 6 and Supplementary Fig. 4). During the delay epoch, small perturbations of the network state in the vicinity of left and right trajectories decayed to the original trajectories (Fig. 5a, right inset, green circles). These dynamics indicated the presence of two point-attractors, with attractor basins separated by a saddle point (Fig. 5a,b, purple diamond). As the ramping input increased during the delay epoch, it moved the attractors further apart from each other and increased the distance between the lick-left trajectory and the saddle point (Fig. 5c, right, regular ramping). This produced a progressive change in the spike rate of neurons encoding the choice made by the animal in a given trial (Fig. 4b, units 2 and 3). The mechanism by which an external ramping input can increase the separation between attractors and the depth of the attractor basins over time can be illustrated in a two-population model with mutual inhibition and an external ramping signal[27] (Supplementary Fig. 5 and Supplementary Math Note 1.3). The observed dynamics in the ramping RNN suggest that distractor stimuli became progressively less likely to overcome the distance between left and right basins of attraction (Fig. 5g), thereby resulting in temporal gating.

Ramping signal controls gating of distractors by separating attractors over time. The separation between attractors in the ramping RNN was controlled by the amplitude of the ramping input. That is, when the RNN received a weaker than average ramping input (Fig. 5c, left), the separation between the lick-left attractor and the saddle during the delay epoch was smaller (Fig. 5c, right). As a consequence, distractors were more likely to cause a switch during trials with a weaker ramping input compared with trials with a stronger ramping input (Fig. 5d and Extended Data Fig. 8a,b). To test this prediction from the model in ALM data, we analyzed single-trial activity of the left ALM projected along the nonselective ramping mode (Fig. 5e). To dissociate the effect of ramping from that of elapsed time on generating behavioral robustness to distractors, we looked at the across-trial variability of ramping at a single time point ($t=-1.6$ s), immediately before the onset of the early-delay distractor. We focused on early-delay distractor trials because they exhibited a highly variable behavioral response. In fact, approximately 50% of trials were robust to an early-delay distractor, whereas trials with a late-delay distractor were almost always robust (Fig. 2d). We therefore examined whether the degree of ramping in the ALM predicted the robustness to the upcoming distractor in single early-delay distractor trials. We grouped trials on the basis of the nonselective ramping amplitude before distractor onset. Trials that exhibited weaker ramping in the ALM before the distractor onset were more likely to result in a switch compared with trials with stronger ramping (Fig. 5f; $P=0.003$), which confirmed the prediction from the RNN (Fig. 5d). Furthermore, early-delay

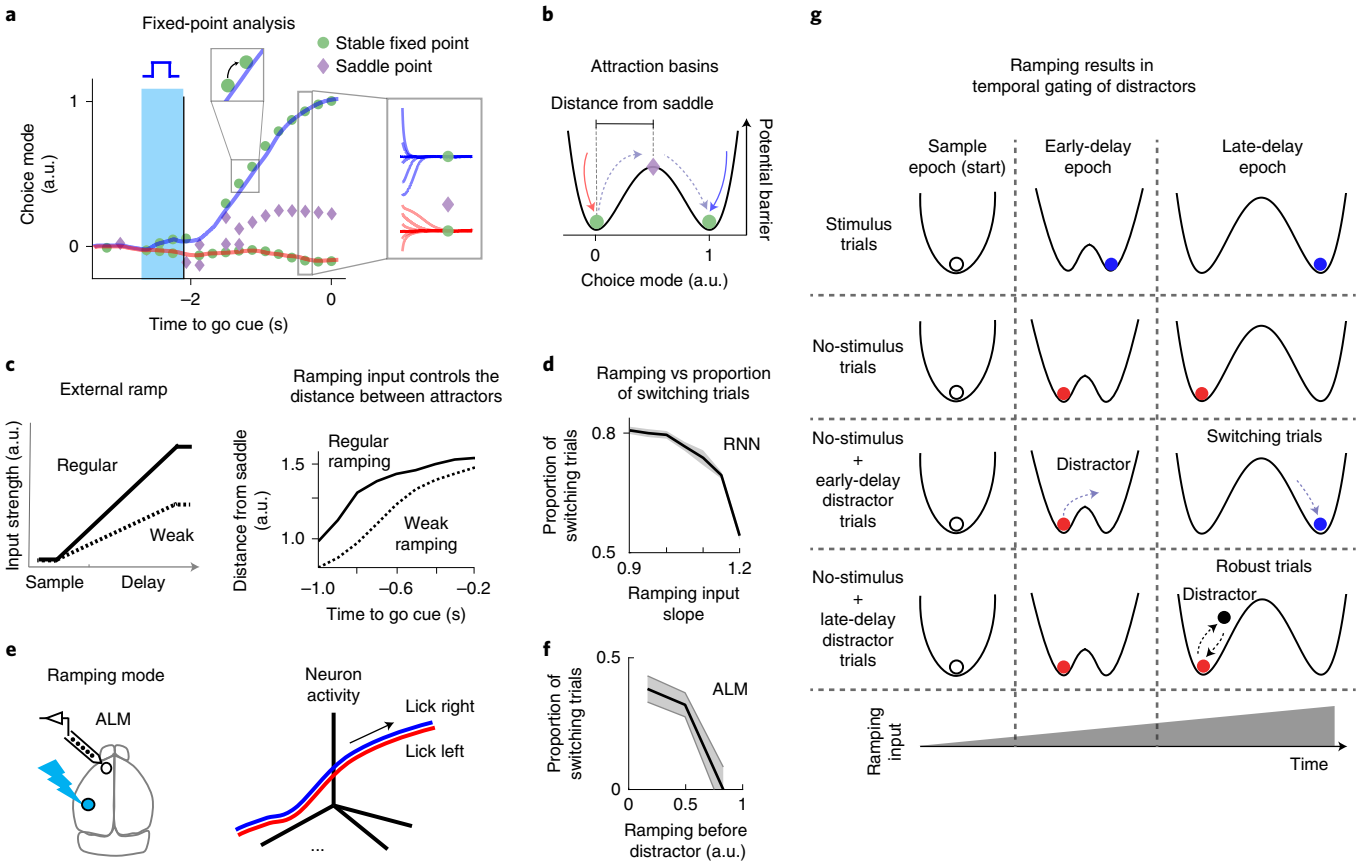

Fig. 5 | Temporal modulation of attractor dynamics as a mechanism for sensory gating. a, Results of a fixed-point search in a RNN model (distractor-free ramping RNN) of the ALM along the choice mode. The ramping input resulted in a gradual increase in the distance between fixed points during the delay. Right inset: example trajectories starting from different initial conditions are attracted towards the fixed points. To search for fixed points at time *t*, we set all network inputs to their value at *t* and searched for regions of the phase space where the dynamical flow was approximately zero (Supplementary Math Note 1.2). Top inset: fixed points moved within the phase space as the external ramping input increased with time. **b**, Snapshot of the RNN phase space along the choice mode during the delay epoch (cartoon). Trajectories converge toward one of the stable fixed points when inside the attraction basins (solid arrows). Trajectories crossing the saddle point will end in the opposite attraction basin (dashed arrows). **c**, Left: schematic of the external ramping inputs of two different slopes (regular $\mu_{ramp}=1$ and weak $\mu_{ramp}=0.9$) delivered to the RNN. Right: the distance between the left stable fixed point and the saddle point during delay for different ramping inputs. A stronger ramping input results in a larger distance to the saddle point during delay. The distance to the saddle point started increasing from $t=-1.2$ s, which corresponded to the offset time of the early-delay distractor. **d**, The proportion of RNN right-output (switching trials) following an early-delay distractor as a function of ramping-input slope. A stronger ramping input reduced the distractor propensity to induce a switch in the RNN output (mean ± s.d. across sessions, $n=10$, 100 trials each). **e**, Schematics of the nonselective ramping mode in the ALM. **f**, Single-trial analysis of the proportion of switching trials as a function of nonselective ramping mode amplitude in the left ALM of distractor-trained mice (ramping was computed before the early-delay distractor; Methods). Data shown as the mean ± s.e.m. across sessions, $n=16$ sessions. Distractors were less capable of inducing a behavioral switch (lick-right response) for trials with a stronger ramping before distractors. **g**, The gating mechanism in the ALM as revealed by RNN analysis (cartoon). Only one stable fixed point (or attractor), corresponding to baseline activity, was present at the start of the sample epoch (left column). The external ramping input caused a bifurcation at the beginning of the delay epoch that led to the appearance of two attractors near the lick-left and lick-right trajectories. As time progressed during the delay epoch, the external ramping increased the separation between the attractor basins. Early-delay distractor trajectories were often able to cross the saddle point, resulting in frequent switches to the right basin (middle column). Late-delay distractor trajectories (right column) did not reach the saddle point, and thus converged back to left basin (bottom), resulting in distractor gating.

distractor trials with stronger ramping exhibited nearly complete robustness (Fig. 5f), and a proportion of switching trials were comparable to the one observed in response to the late-delay distractor (Fig. 2d). This suggests that similar to the ramping RNN (Extended Data Fig. 8a–c), choice robustness in the ALM is controlled by the ramping amplitude rather than by elapsed time per se.

Additional stimuli during lick-right trials caused an increase in the lick-right probability if the additional stimulus was delivered early, but not late, in the delay epoch (Extended Data Fig. 9a,b). A potential explanation for this could be that for some error trials that would have resulted in a left lick, neural dynamics were 'rescued' by the second stimulus and eventually ended up in a right lick. Such rescued trials would be more likely to occur if the second stimulus was presented early in the delay epoch, when the attractors were still close to each other, rather than late. We also found that an additional stimulation delivered during the pre-sample epoch led to a higher level of activity along the choice mode compared with additional stimulation during the delay epoch (Extended Data Fig. 9c,d), which suggests that ALM dynamics are influenced by the time elapsed from the initiation of preparatory activity but not by the number of stimuli. We tested the effect of delivering two stimuli (one during the pre-sample epoch and one during the sample epoch) to the distractor-free ramping RNN (Supplementary Fig. 6). Results were consistent with the experimental data (Extended

Data Fig. 9d), which again suggests that the level of activity along the choice mode is controlled by the nonspecific ramping input that moved the attractors apart over time.

In distractor-naive mice, which were less robust to weak distractors compared with distractor-trained mice (Fig. 2c), the nonselective ramping mode had a consistently smaller contribution to neural activity (Extended Data Fig. 4c,d). For weaker ramping input, the RNN model predicted that distractors could more easily push the activity beyond the saddle point and toward the lick-right fixed point (Fig. 5b, dotted arrow), which resulted in a higher proportion of switches from one attractor to the other (Fig. 5c,d and Extended Data Fig. 8). Indeed, a strong late-delay distractor, which was systematically ignored in distractor-trained mice, induced a significant number of switches in distractor-naive mice (Fig. 2d). Furthermore, the impact of weak distractors on choice mode in distractor-naive mice was more prominent and persisted for a longer time than in distractor-trained mice (Extended Data Fig. 10a–c). These results suggest that the attractor basins in distractor-naive mice are shallower compared with those in distractor-trained mice (Extended Data Fig. 10d). Taken together, our results indicate that sensory inputs are gated by attractor dynamics in the ALM in a manner that is shaped by experience.

Discussion

We found that distracting stimuli arriving in the delay epoch between the sample epoch and the motor response became less capable of influencing decisions as the time to act approached, thereby demonstrating a form of temporal gating. Temporal gating was previously proposed to operate via distractor suppression in the prefrontal cortex[7]. Our results show that temporal gating does not require the suppression of distracting inputs but can instead arise from a gradual increase in robustness of the chosen motor plan, which results in gating of the distractors. This novel gating mechanism complements previously proposed mechanisms by which temporal expectations can modulate attention during stimulus selection[5,17,32,33].

To infer the mechanism underlying temporal gating, we trained RNNs[16,17,34] to reproduce the heterogeneous dynamics of individual neurons[30] in the ALM. We found that the RNN models were able to generalize to conditions they were not trained to reproduce and predicted temporal gating as observed in the data. Reverse engineering of the trained RNN models[31] revealed a novel mechanism for temporal gating, which operated by combining attractor dynamics with a ramping signal. Attractor dynamics can underlie short-term memory[35–41], preparatory activity[26,27] and context-dependent stimulus selection[16], and were hypothesized to support distractor-resistant memory[26,27,40]. Here, we found that attractor dynamics in the ALM enable flexible motor responses to incoming inputs via the local regulation of information flow. Temporal gating was achieved by moving the attractors apart and by a continuous, nonspecific ramping input that reflects task-relevant timing[22,23,27,42,43]. Attractor dynamics were observed in similar memory-guided decision-making paradigms that relied on somatosensory[26] or auditory sensory inputs[27] instead of the optogenetic inputs that we used here. Therefore, attractor dynamics are expected to play a role in the filtering of stimuli incoming to the ALM irrespective of the type of stimuli (that is, somatosensory, auditory or optogenetic) used to form the decision. However, in the case of natural sensory stimuli, additional gating mechanisms may regulate sensory information flow on its path from the sensory periphery to the cortex[9].

We found that putative fast-spiking interneurons had trial-type-selective heterogenous responses[44] (Extended Data Fig. 2a–c). Population dynamics of fast-spiking interneurons in the ALM showed no distractor suppression along the stimulus mode, but did show gating of distractors along the choice mode (Extended Data Fig. 2d–g), which is similar to putative pyramidal cells.

Previous behavioral studies showed that direct perturbations of decision-making circuits had a stronger influence on behavior when applied early during the trial[18,40]. Our results provide a mechanistic interpretation of these behavioral findings, and our RNN model may constitute a more biologically constrained alternative to the static two-population attractor model to explain behavior[40]. By combining electrophysiological recordings, targeted optogenetic stimulations and training RNNs to reproduce the heterogeneous dynamics of individual neurons, we were able to find a distinct mechanism whereby attractors were moved by an external ramping input and then verified its predictions using ALM data. Specifically, we showed that differences in behavior are due to a progressive reconfiguration of circuit dynamics in the ALM, mediated by a ramping signal that increases the separations between attractors, and hence produces the gating of sensory inputs over time. This gating mechanism may also explain how decisions can become resistant to additional sensory evidence in evidence accumulation[45–47] or in oculomotor-memory[7] tasks.

We found that gating of cortical inputs can be shaped by experience (Fig. 2 and Extended Data Fig. 10), which was reflected in changes in the dynamical structure in the ALM of distractor-naive versus distractor-trained mice, with attractors becoming stronger and more separated with experience. This could be caused by either changes in the strength of inputs to the ALM (for example, the ramping input) or modifications of local connectivity within the ALM.

Our results also suggest a new computational role for the ramping signal, which has been observed in many decision-making tasks and is related to time estimation and a sense of urgency[22,23,27,42,43]. Ramping may act to gradually diminish the impact of incoming sensory inputs and thus increase the commitment to the selected action[48] by moving attractors apart over time. Previous studies suggested a relationship between ramping of neural activity and reaction time[23,49]. In our task, however, movement was instructed by a go cue, and the ramping mode had only weak relationship with reaction time (Pearson correlation coefficient $r=-0.05$ and $r=-0.14$, for distractor-naive and distractor-trained mice, respectively). Our model framework, in which a timing-varying ramping signal controls the flexibility of action selection, may also be relevant for countermanding (stop signal) tasks in which the go and the stop signals are modeled as competing ramping processes[50]. In summary, our work uncovered a novel gating mechanism that controls the impact of sensory inputs on decisions about future actions.

Online content

Any methods, additional references, Nature Research reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at https://doi.org/10.1038/s41593-021-00840-6.

Received: 29 April 2020; Accepted: 12 March 2021;
Published online: 19 April 2021

References

1. Cole, M. W. et al. Multi-task connectivity reveals flexible hubs for adaptive task control. *Nat. Neurosci.* **16**, 1348–1355 (2013).
2. Akam, T. & Kullmann, D. M. Oscillatory multiplexing of population codes for selective communication in the mammalian brain. *Nat. Rev. Neurosci.* **15**, 111–122 (2014).
3. Desimone, R. & Duncan, J. Neural mechanisms of selective visual attention. *Annu. Rev. Neurosci.* **18**, 193–222 (1995).
4. Moran, J. & Desimone, R. Selective attention gates visual processing in the extrastriate cortex. *Science* **229**, 782–784 (1985).
5. Ghose, G. M. & Maunsell, J. H. R. Attentional modulation in visual cortex depends on task timing. *Nature* **419**, 616–620 (2002).

6. Lennert, T. & Martinez-Trujillo, J. Strength of response suppression to distracter stimuli determines attentional-filtering performance in primate prefrontal neurons. *Neuron* **70**, 141–152 (2011).
7. Suzuki, M. & Gottlieb, J. Distinct neural mechanisms of distractor suppression in the frontal and parietal lobe. *Nat. Neurosci.* **16**, 98–104 (2013).
8. Moore, T. & Armstrong, K. M. Selective gating of visual signals by microstimulation of frontal cortex. *Nature* **421**, 370–373 (2003).
9. McAlonan, K., Cavanaugh, J. & Wurtz, R. H. Guarding the gateway to cortex with attention in visual thalamus. *Nature* **456**, 391–394 (2008).
10. Wimmer, R. D. et al. Thalamic control of sensory selection in divided attention. *Nature* **526**, 705–709 (2015).
11. Vogels, T. P. & Abbott, L. F. Gating multiple signals through detailed balance of excitation and inhibition in spiking networks. *Nat. Neurosci.* **12**, 483–491 (2009).
12. Yu, J., Gutnisky, D. A., Hires, S. A. & Svoboda, K. Layer 4 fast-spiking interneurons filter thalamocortical signals during active somatosensation. *Nat. Neurosci.* **19**, 1647–1657 (2016).
13. Schneider, D. M., Sundararajan, J. & Mooney, R. A cortical filter that learns to suppress the acoustic consequences of movement. *Nature* https://doi.org/10.1038/s41586-018-0520-5 (2018).
14. Fries, P., Reynolds, J. H., Rorie, A. E. & Desimone, R. Modulation of oscillatory neuronal synchronization by selective visual attention. *Science* **291**, 1560–1563 (2001).
15. Cardin, J. A. et al. Driving fast-spiking cells induces gamma rhythm and controls sensory responses. *Nature* **459**, 663–667 (2009).
16. Mante, V., Sussillo, D., Shenoy, K. V. & Newsome, W. T. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* **503**, 78–84 (2013).
17. Carnevale, F., de Lafuente, V., Romo, R., Barak, O. & Parga, N. Dynamic control of response criterion in premotor cortex during perceptual detection under temporal uncertainty. *Neuron* **86**, 1067–1077 (2015).
18. Seidemann, E., Zohary, E. & Newsome, W. T. Temporal gating of neural signals during performance of a visual discrimination task. *Nature* **394**, 72–75 (1998).
19. Fuster, J. M. & Alexander, G. E. Neuron activity related to short-term memory. *Science* **173**, 652–654 (1971).
20. Tanji, J. & Evarts, E. V. Anticipatory activity of motor cortex neurons in relation to direction of an intended movement. *J. Neurophysiol.* **39**, 1062–1068 (1976).
21. Erlich, J. C., Bialek, M. & Brody, C. D. A cortical substrate for memory-guided orienting in the rat. *Neuron* **72**, 330–343 (2011).
22. Fried, I., Mukamel, R. & Kreiman, G. Internally generated preactivation of single neurons in human medial frontal cortex predicts volition. *Neuron* **69**, 548–562 (2011).
23. Churchland, A. K., Kiani, R. & Shadlen, M. N. Decision-making with multiple alternatives. *Nat. Neurosci.* **11**, 693–702 (2008).
24. Guo, Z. V. et al. Flow of cortical activity underlying a tactile decision in mice. *Neuron* **81**, 179–194 (2014).
25. Kaufman, M. T., Churchland, M. M., Ryu, S. I. & Shenoy, K. V. Cortical activity in the null space: permitting preparation without movement. *Nat. Neurosci.* **17**, 440–448 (2014).
26. Li, N., Daie, K., Svoboda, K. & Druckmann, S. Robust neuronal dynamics in premotor cortex during motor planning. *Nature* **532**, 459–464 (2016).
27. Inagaki, H. K., Fontolan, L., Romani, S. & Svoboda, K. Discrete attractor dynamics underlies persistent activity in the frontal cortex. *Nature* **566**, 212–217 (2019).
28. O'Connor, D. H. et al. Neural coding during active somatosensation revealed using illusory touch. *Nat. Neurosci.* **16**, 958–965 (2013).
29. Hernández, A. et al. Decoding a perceptual decision process across cortex. *Neuron* **66**, 300–314 (2010).
30. Rajan, K., Harvey, C. D. & Tank, D. W. Recurrent network models of sequence generation and memory. *Neuron* **90**, 128–142 (2016).
31. Sussillo, D. & Barak, O. Opening the black box: low-dimensional dynamics in high-dimensional recurrent neural networks. *Neural Comput.* **25**, 626–649 (2012).
32. Ede, F., van, Chekroud, S. R., Stokes, M. G. & Nobre, A. C. Decoding the influence of anticipatory states on visual perception in the presence of temporal distractors. *Nat. Commun.* **9**, 1449 (2018).
33. Nobre, A., Correa, A. & Coull, J. The hazards of time. *Curr. Opin. Neurobiol.* **17**, 465–470 (2007).
34. Sussillo, D. & Abbott, L. F. Generating coherent patterns of activity from chaotic neural networks. *Neuron* **63**, 544–557 (2009).
35. Wang, X.-J. Probabilistic decision making by slow reverberation in cortical circuits. *Neuron* **36**, 955–968 (2002).
36. Machens, C. K., Romo, R. & Brody, C. D. Flexible control of mutual inhibition: a neural model of two-interval discrimination. *Science* **307**, 1121–1124 (2005).
37. Lim, S. & Goldman, M. S. Balanced cortical microcircuitry for maintaining information in working memory. *Nat. Neurosci.* **16**, 1306–1314 (2013).
38. Barak, O. & Tsodyks, M. Working models of working memory. *Curr. Opin. Neurobiol.* **25**, 20–24 (2014).
39. Wimmer, K., Nykamp, D. Q., Constantinidis, C. & Compte, A. Bump attractor dynamics in prefrontal cortex explains behavioral precision in spatial working memory. *Nat. Neurosci.* **17**, 431–439 (2014).
40. Kopec, C. D., Erlich, J. C., Brunton, B. W., Deisseroth, K. & Brody, C. D. Cortical and subcortical contributions to short-term memory for orienting movements. *Neuron* **88**, 367–377 (2015).
41. Kamiński, J. et al. Persistently active neurons in human medial frontal and medial temporal lobe support working memory. *Nat. Neurosci.* **20**, 590–601 (2017).
42. Maimon, G. & Assad, J. A. A cognitive signal for the proactive timing of action in macaque LIP. *Nat. Neurosci.* **9**, 948–955 (2006).
43. Finnerty, G. T., Shadlen, M. N., Jazayeri, M., Nobre, A. C. & Buonomano, D. V. Time in cortical circuits. *J. Neurosci.* **35**, 13912–13916 (2015).
44. Najafi, F. et al. Excitatory and inhibitory subnetworks are equally selective during decision-making and emerge simultaneously during learning. *Neuron* **105**, 165–179.e8 (2020).
45. Kiani, R., Hanks, T. D. & Shadlen, M. N. Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment. *J. Neurosci.* **28**, 3017–3029 (2008).
46. Zuo, Y. & Diamond, M. E. Rats generate vibrissal sensory evidence until boundary crossing triggers a decision. *Curr. Biol.* **29**, 1415–1424.e5 (2019).
47. Peixoto, D. et al. Decoding and perturbing decision states in real time. *Nature* https://doi.org/10.1038/s41586-020-03181-9 (2021).
48. Thura, D. & Cisek, P. Deliberation and commitment in the premotor and primary motor cortex during dynamic decision making. *Neuron* **81**, 1401–1416 (2014).
49. Roitman, J. D. & Shadlen, M. N. Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J. Neurosci.* **22**, 9475–9489 (2002).
50. Boucher, L., Palmeri, T. J., Logan, G. D. & Schall, J. D. Inhibitory control in mind and brain: an interactive race model of countermanding saccades. *Psychological Rev.* **114**, 376–397 (2007).

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

© The Author(s), under exclusive licence to Springer Nature America, Inc. 2021

Methods

Animals. We used nine adult Scnn1a-TG3-Cre × Ai32 transgenic mice, which express channelrhodopsin-2 (ChR2H134R–EYFP) in layer 4 stellate cells in the barrel cortex (Ai32 mice, The Jackson Laboratory, 012569; Scnn1a-TG3-Cre mice, The Jackson Laboratory, 009613)[28,51–53]. All experimental procedures were approved by the Janelia Institutional Animal Care and Use Committee. Detailed procedures for water restriction and surgical procedures are described elsewhere[24,27].

Mouse behavior. Water-restricted mice were trained to report detection of photostimulation by directional licking for a water reward. Mice were implanted with a head bar[54] before water restriction and behavioral training.

Basic task. Mice were head-fixed in front of two lick ports placed to the right and left of their snout and were instructed to lick right or left for a water reward (Fig. 1b). Each trial began with a pre-sample epoch (1.2 s in duration) followed by a sample epoch (1 s in duration) that was demarcated by auditory cues (0.15 s in duration, sample cues). Photostimulation during the sample epoch instructed mice to lick right. Absence of photostimulation during the sample epoch instructed mice to lick left. Mice were trained to withhold licking for an additional 2 s after the sample epoch (delay epoch) until an auditory go cue (pure tone, 3.4 kHz, 0.1 s in duration) signaled the beginning of the response epoch (1.5 s in duration). Early licks (occurring before the go cue) were punished by restarting the delay epoch.

Task with distractors. Distractors consisted of photostimuli presented outside the sample epoch (Fig. 2a,b). Mice were not rewarded for licking right in response to distractors. Lick-left trials with distractors resulting in right licks (which were considered error trials) were referred to as switching trials (for example, Fig. 3h, dashed gray, switching trials). We used two groups of mice: distractor-naive ($n=5$) and distractor-trained ($n=4$). Initially all mice were trained on the basic task without distractors until reaching criterion (75% correct trials). Distractor-naive mice were not trained beyond the basic task, whereas distractor-trained mice were further trained to explicitly ignore distractors. During training of the distractor-trained mice, the amplitude of distracting photostimuli was gradually increased (across sessions) until it was identical to that of the stimulus. In half of the trials, distractors were also delivered on lick-right trials (Extended Data Fig. 9). In each session ~75% of the trials included distractors. During recording sessions, both distractor-naive and distractor-trained mice were tested in the presence of distractors of different amplitudes and timings (see the section 'Photostimulation' for details).

Photostimulation. Photostimuli from a 473-nm laser (Laser Quantum) were controlled by an acousto-optical modulator (Quanta Tech). Photostimuli were delivered to the vS1 (anterior–posterior (AP): −1.3 mm; medial–lateral (ML): 3.5 mm, relative to bregma) in the left hemisphere through a clear skull cup[24]. A guide tube (1.6-mm inner diameter) was glued on top of the stimulation site, and an optical fiber was inserted into the guide tube at the beginning of each behavioral session. During sessions in which we combined electrophysiological recordings from the vS1 with photostimulation, the laser beam was positioned directly above the recording site.

To prevent mice from using visual cues to distinguish photostimulation trials from control trials, a masking flash (10 Hz) was delivered using 470-nm light-emitting diodes (LEDs) throughout the trial. The masking flash was synchronized to the photostimulation in the recording rig, but not in the training rig. We did not see a change in performance after transitioning from the training rig to the recording rig, which indicated that the mice did not use visual signals to detect photostimulation. In a handful of sessions, we also moved the laser beam ~1-mm away from the usual photostimulation site (vS1). This resulted in a marked drop in performance on lick-right trials, which indicated that the mice did not detect the photostimulation when a neighboring part of the cortex was stimulated. This further shows that mouse behavior was guided by direct cortical stimulation of the vS1 rather than by unintended visual signals from the photostimulation.

In the recording rig, photostimulation was delivered by a laser whose shutter generated a brief clicking sound. To prevent mice from using the laser shutter sound as an auditory cue to detect photostimulation, we set the shutter to open on all trials, including for those in which there was no photostimulation. For such no-photostimulation trials, the shutter was still opening but the laser power was set to 0. Importantly, for all trials, the shutter was opening during the pre-sample epoch and then stayed open until after the go cue. These two measures ensured that mice could not use the shutter sound to predict the presence of photostimulation or the time of photostimulation. In the training rig, photostimulation was delivered via a LED that did not have a shutter and thus did not create detectable sounds.

We used three types of photostimuli (Fig. 2a,b): stimulus, strong distractors and weak distractors.

1. Stimulus was delivered during the sample epoch (at −2.5 s before the go cue) and instructed mice to lick right. It contained 4 pulses lasting 0.4 s (10 Hz, sinusoidal temporal profile). The peak power for the stimulus was typically 2.2 mW, but in a few sessions, power was adjusted to correct for performance biases (mean power = 2.2 mW, s.d. = 0.2 mW, range (1.2–3.2 mW), across sessions).
2. Strong distractors had the same photostimulation profile as the stimulus, but were delivered outside the sample epoch (onsets = −3.8, −1.6 and −0.8 s relative to the go cue).
3. Weak distractors contained one pulse that lasted for 0.1 s, with the peak power typically set at 33.3% of the stimulus peak power (mean power = 0.8 mW, s.d. = 0.2 mW, range (0.4–1.1 mW), across sessions). The weak-distractor photostimulus power was calibrated before the first recording session to ensure that it elicited a change in behavioral performance when delivered during the sample epoch. Weak distractors were delivered during the sample or outside of it (onsets = −3.8, −2.5, −1.6 and −0.8 s relative to the go cue).

Electrophysiology. We recorded extracellular spikes using silicon probes (2 shanks × 32 channels) with 250-µm spacing between shanks and 25-µm spacing between channels (H2; Cambridge Neurotech). The 64-channel voltage signals were multiplexed, recorded on a PCI6133 board (National Instrument) and digitized at 14 bits. The signals were demultiplexed into 64 voltage traces, sampled at 25 kHz and stored for offline analyses. Data acquisition was performed using SpikeGL (https://github.com/cculianu/SpikeGL) and Wavesurfer (https://www.janelia.org/open-science/wavesurfer) software. Recordings were done through a small craniotomy (diameter of 0.5–1.0 mm) made 1 day before the recording. The craniotomy was centered over the left or right ALM (AP: 2.5 mm; ML: 1.5 mm, relative to bregma) or the left vS1 (AP: −1.3 mm; ML: 3.5 mm). The brain was allowed to settle for at least 10 min following penetration with the silicon probe before the recording started. Depths of recorded neurons in the ALM (ranging between 250 and 1,350 µm) and the vS1 (between 175 and 1,150 µm) were inferred from manipulator readings. We performed three to ten recording sessions from each craniotomy on consecutive days. Distractor-naive mice: left ALM $n=26$ sessions; right ALM $n=11$; and vS1 $n=10$. Distractor-trained mice: left ALM $n=23$; right ALM $n=11$; and vS1 $n=16$.

Behavioral and electrophysiological data were stored and analyzed in custom pipelines in the DataJoint framework[55].

Behavioral data analyses. We excluded early-lick trials (trials during which the animal started licking before the go cue) and no-response trials (trials that did not elicit any licking until 1.5 s after the go cue) from behavioral analyses. These trials were also excluded from analyses of neural activity. The proportion of lick-right responses in Fig. 1c is shown for distractor-naive mice. For distractor-trained mice, the proportion of lick-right responses was as follows: 0.79 ± 0.02% for stimulus (lick-right) trials and 0.25 ± 0.01% for no-stimulus (lick-left) trials ($n=4$ mice, 87 sessions). The behavioral performance for each session was calculated as the proportion of correct responses for each trial type and is shown in Supplementary Fig. 1a.

To assess the effect of stimulus/distractors on performance (Fig. 2c,d), we calculated for each session the change in proportion of lick-right responses for trials with stimulus or distractors relative to control trials (that is, trials with instruction to lick left without distractors; Fig. 2a,b, red) as follows:

$$\Delta\text{Proportion lick}-\text{right} = \frac{\text{Right response trials}_{\text{Stimulus/distractor}}}{\text{All trials}_{\text{Stimulus/distractor}}} - \frac{\text{Right response trials}_{\text{Control}}}{\text{All trials}_{\text{Control}}} \quad (1)$$

Control trials accounted for the proportion of spontaneously occurring lick-right trials in the absence of stimulus or distractors. The data in Fig. 2 are based on the following sessions: distractor-naive mice with weak distractors ($n=5$ mice, 67 sessions) and strong distractors ($n=2$, 17 sessions); and distractor-trained mice with weak distractors ($n=4$ mice, 19 sessions) and strong distractors ($n=4$ mice, 67 sessions). To account for variations in baseline performance (for trials without distractors) between different group of mice, the value of Δproportion lick-right (equation (1) and Fig. 2c,d) was normalized so that the change in proportion of lick-right trials in response to the full stimulus during the sample epoch was set to 1. Raw behavioral data (the proportion of lick-right responses for each trial-type) for individual mice is shown in Supplementary Fig. 1b–e.

Electrophysiological data analyses. Spike sorting was done using JRClust[56]. A total of 3,385 well-isolated single units were recorded, of which we included 3,011 based on the following criteria: (1) the cell had a mean spike rate ≥0.5 Hz; (2) the cell was recorded for at least ten lick-left and ten lick-right correct trials without distractors. In addition, we recorded 1,799 multiunits, of which we included 1,611 units based on the above criteria.

Putative pyramidal neurons and fast-spiking interneurons were differentiated on the basis of spike waveforms[24,27]. Putative pyramidal cells were recorded in the following areas (numbers in parentheses relate to well-isolated neurons and multiunits): the left ALM (1,648; 769); the right ALM (679; 365); and the left vS1 (256; 291). Putative fast-spiking interneurons recorded in the following areas (well-isolated neurons; multiunits): the left ALM (184; 21); the right ALM (96; 13); and the left vS1 (92; 116).

All analyses were performed on putative pyramidal neurons, except for the analyses of putative fast-spiking interneurons presented in Extended Data Fig. 2. For all population analyses, we combined single units with multiunits, except for the analyses presented in Extended Data Figs. 1 and 2. We only analyzed trials with ≥5 simultaneously recorded units (mean number of units per trial = 45.3 ± 0.1).

All analyses in Fig. 1 (except panel c) and Extended Data Figs. 1 and 2 (except panels d–g) were performed on data from distractor-naive and distractor-trained mice combined. Analyses in Figs. 3–5 and Extended Data Figs. 3, 6, 7 and 9 were performed on distractor-trained mice.

Population-average spike rate and selectivity. Spike rates were computed in 5-ms time bins, with the spike time in each trial defined relative to the go cue. Trial-average spike rates were smoothed with a 50-ms causal boxcar filter. For plotting PSTHs of example cells we used a 50-ms filter (Fig. 1d–e, Extended Data Figs. 2a and 3a and Supplementary Fig. 2b) or a 200-ms filter (Extended Data Figs. 5a and 10a).

Grand-average population responses (Fig. 1f,g and Extended Data Fig. 2a) were computed by averaging the trial-average spike rates across cells using correct lick-left or lick-right trials without distractors.

Trial-type selectivity $S_i(t)$ at time t of cell i was defined as follows:

$$S_i(t) = \bar{r}_i^{\mathrm{R}}(t) - \bar{r}_i^{\mathrm{L}}(t) \quad (2)$$

where $\bar{r}_i^K$ are the trial-average spike rates at time t, computed using correct right ($K=\mathrm{R}$) or left trials ($K=\mathrm{L}$).

Neurons were identified as trial-type selective during the sample or delay epochs if their epoch-averaged selectivity was significantly different from zero (Wilcoxon rank-sum test, $P \leq 0.001$, two-sided). To determine the selectivity of neuron i during the sample epoch, we averaged $S_i(t)$ in a 0.5-s long window following the stimulus onset. To test for significant selectivity during the delay epoch, we averaged $S_i(t)$ during the entire 2 s of the delay.

Hierarchical clustering of neural activity profiles in response to photostimulation. We used hierarchical clustering (implemented using Matlab functions pdist, linkage and cluster, where the distance between observations was based on pairwise correlations) to categorize trial-averaged neural activity responses across recorded cells (Extended Data Figs. 1a,b and 2b,c). For each cell, we concatenated the trial-average spike rates for both lick-left and lick-right trial types computed using correct trials. Trial-average spike rates were smoothed with a 200-ms causal boxcar filter. Clustering was done on the basis of the correlation between trial-average responses across cells. For clustering we used a time window of $t=[-3.5, 0]$ s, which included the sample and delay epochs. Putative pyramidal cells and fast-spiking interneurons were clustered separately. In the ALM, clustering was done using well-isolated cells recorded from the left and right ALM. In the vS1, we combined well-isolated cells and multiunits together.

For display purposes (Extended Data Figs. 1a,b and 2b,c), we averaged the spike rate across all neurons that belonged to the same cluster separately for correct and error trials. For correct trials, we used only trials without distractors. For error trials (which were less frequent), we combined error trials from all trial types, including trials with distractors. Only cells with at least five trials of each trial type were included. Only clusters with more than 2% of cells are shown.

Decoding of sensory- and choice-related neural activity. To decode sensory- and choice-related neural activity (Fig. 1h), we trained a support vector machine classifier (implemented using the Matlab function fitcsvm with the linear kernel function). To decode sensory-related activity, we labeled the trials based on the instruction (that is, whether a stimulus was present during sample epoch). To decode choice-related activity, we labeled the trials based on whether the mouse licked left or right irrespective of the instruction. We trained the classifier on the matched number of correct and error trials in both lick directions to ensure that the instruction is decorrelated from the outcome. Specifically, each training set was composed of four categories of trials (correct left, error left, correct right and error right), with the same number of trials in each category. Matching was performed by subsampling, whereby we randomly selected the same number of trials from all categories (subsampling was repeated 20 times) and the number of trials selected was equal to the size of the smallest category. We included sessions that had at least five simultaneously recorded cells, including multiunits, and with at least ten trials in the smallest category, including trials with distractors. In each session, for each subsample, we trained the classifier based on 80% of the selected correct and error trials and tested it on the remaining 20% of trials. This process was repeated 100 times. The decoder performance was taken as the average decoder performance across all repetitions. Decoding was done in 0.1-s time bins, and the classifier was retrained for each time bin. For decoding, we used only stable cells; that is, cells that were recorded on at least 80% of the trials within each session.

Population dynamics in the neural activity space. We analyzed the population dynamics of n neurons simultaneously recorded in a session (Fig. 3, Extended Data Figs. 2d,f, 3d,e, 4, 5b, 7e,f, 9c,d and 10b, and Supplementary Fig. 2c, bottom). During each trial, the population activity of n neurons drew a trajectory in the n-dimensional activity space, where each dimension represents the spike rate of one neuron. We identified directions (modes) in the activity space that maximally separated the neural trajectories for different trial conditions and times during the trial[26,27] (Extended Data Fig. 4a). When analyzing responses to distractors in the activity space, we restricted the analysis to trial types that had at least five neurons recorded simultaneously with at least five trials per condition.

The stimulus mode was computed as a $n \times 1$ vector of trial-averaged spike rate differences of n neurons during trials with lick-right and lick-left instructions, averaged within a 0.5-s window following the stimulus onset during the sample epoch as follows:

$$\Delta\bar{\mathbf{r}} = \bar{\mathbf{r}}^{\mathrm{R}} - \bar{\mathbf{r}}^{\mathrm{L}} \quad (3)$$

The resulting vector (equation (3)) was normalized by its l_2 norm, resulting in a vector of weights (one weight per neuron) as follows:

$$\mathbf{m} = \frac{\Delta\bar{\mathbf{r}}}{\sqrt{\sum_i^N |\Delta\bar{r}_i|^2}} \quad (4)$$

Projections of the neural activity along the mode over time were calculated as follows:

$$\mathbf{p}^k = \left(X^k\right)^{\mathsf{T}} \mathbf{m} \quad (5)$$

Where X^k is the $N \times T$ matrix of spike rates of neurons over time for trial k. Normalizing each mode by its l_2 norm (equation (4)) ensured that projected activity $\mathbf{p}^k$ would not scale with the number of neurons recorded simultaneously. If an individual neuron was not recorded during a particular trial, its weight in equations (4) and (5) was set to zero when computing the projected neural activity for that trial.

In Extended Data Fig. 3d,e, we also computed the stimulus mode in the same way, but used trials with and without distractors during early or late delay. We computed the trial-averaged spike rates within a 0.5-s window following the distractor onset.

The choice mode was defined as a $n \times 1$ vector of trial-averaged spike rate differences of n neurons during trials with lick-right and lick-left outcomes, averaged within a 0.2-s window at the end of the delay epoch before the go cue (analogously to equations (3) and (4)). The number of correct and error trials in both lick directions was matched to ensure that the instruction is decorrelated from the outcome. Specifically, we used four categories of trials (correct left and error right trials (both resulting in a lick-left outcome), and correct right and error left trials (both resulting in a lick-right outcome)), with the same number of trials in each category. Matching was performed by subsampling, whereby we randomly selected the same number of trials from all categories (subsampling was repeated 20 times) and the number of trials selected was equal to the size of the smallest category. If a session did not have at least ten trials in the smallest category (including trials with distractors), we matched the number of trials with the left and right outcome irrespective of whether they came from correct or error trials. To compute the choice mode using left-preferring cells (Extended Data Fig. 5b, right), we set the positive weights of the choice mode to zero and then projected the neural activity on this mode as described above. Positive weights corresponded to right-preferring cells, whereas negative weights corresponded to left-preferring cells (equations (3) and (4)).

We also identified a nonselective ramping mode (Fig. 5e and Extended Data Figs. 2d,g and 4a,b), which captured the difference in neural activity (regardless of the trial type) between the beginning of the trial and the end of the delay epoch. We computed the ramping mode in a similar way to the choice mode and stimulus mode by taking the difference in trial-averaged spike rates (but pooling together both lick-left and lick-right trials) during the last 0.5 s of the delay epoch versus the 0.5 s that preceded the sample epoch as follows:

$$\Delta\bar{\mathbf{r}} = \bar{\mathbf{r}}_{\mathrm{Pre-sample}} - \bar{\mathbf{r}}_{\mathrm{Delay}} \quad (6)$$

The modes were orthogonalized to each other using a Gram–Schmidt process in the following order: choice, ramping, stimulus.

To compare trial-averaged projections across sessions, we first subtracted from each trial-averaged projection the baseline activity computed in the same session. Baseline activity was defined as the trial-averaged activity computed on correct lick-left trials without photostimulation, averaged over the 1-s window preceding the sample epoch. Trial-average projections were smoothed with a 100-ms causal boxcar filter for display. After averaging the projected activity across sessions, the resulting averaged projection was normalized to its maximum value reached at any time before the go cue on lick-right correct trials. Projections along the stimulus mode and the ramping mode were computed by combining correct and error trials, with at least 15 trials per trial type. Projections along the choice mode were computed separately for correct trials and error trials, with at least five trials per trial type and outcome.

Variance explained by different modes. To calculate the fraction of variance explained by the neural activity projected on each mode for each session (Extended Data Fig. 4c,d), we first performed baseline subtraction from the spike rate of each neuron. Specifically, we subtracted the trial-averaged baseline spike rate of each neuron from the spike rate on each trial. The baseline spike rate of each neuron was computed by averaging the spike rate across all trials over 1-s window that preceded the sample epoch. We used the same baseline-subtracted spike rates to project the neural activity on the different modes. The projections to each mode

were smoothed with a causal sliding boxcar window that was 0.5-s long with 0.1-s steps. For the stimulus mode and the ramping mode, we used correct and error trials without distractors. For the choice mode, we used only correct trials without distractors.

To calculate the fraction of trial-to-trial variance explained (Extended Data Fig. 4c) by each mode, we calculated the trial-to-trial variance in each time bin as the sum (across neurons $i = 1,..., N$, and trials $k = 1,..., M$) of the square of the spike rate averaged within each time bin. We divided this value by the trial-to-trial variance of the neural activity projected on a given mode, calculated as the sum (across trials) of the square of the projection averaged within each time bin as follows:

$$\text{Trial}-\text{to}-\text{trial variance explained}(t) = \frac{\sum_k^M \sum_i^N (r_i^k(t))^2}{\sum_k^M (p^k(t))^2} \quad (7)$$

To calculate the fraction of trial-averaged variance explained (Extended Data Fig. 4d), we calculated the total trial-averaged variance in each time bin as a sum (across neurons) of squares of the trial-average spike rate left and right trial types averaged in each time bin. We divided this value by the square of the trial-average projection in left and right trial types averaged in each time bin as follows:

$$\text{Trial}-\text{averaged variance explained}(t) = \frac{\sum_i^N (\bar{r}_i^R(t))^2 + \sum_i^N (\bar{r}_i^L(t))^2}{(\bar{p}^R(t))^2 + (\bar{p}^L(t))^2} \quad (8)$$

Distractor response-size quantification. We computed the changes in spike rate of neurons in response to photostimulation (Δspike rate; Fig. 3c,d, Extended Data Fig. 3b and Supplementary Fig. 2c, top) as follows. For each session, we first computed the grand-average population spike rates (see the definition in section "Population-average spike rate and selectivity") separately for each trial type, with correct and error trials pooled together. To obtain the Δspike rate in each time bin, we took the difference between the grand-average population spike rate for each trial type and that for lick-left trials without distractors.

The response size was defined as the number of spikes added during photostimulation (Δspikes; Extended Data Fig. 3c and Supplementary Fig. 2d, top). For each session, we averaged the Δspike rate during the time window of the photostimulation (strong distractors, 0.4-s time window; weak distractors 0.1-s time window).

Because the number of responsive cells was smaller in the ALM compared to the vS1, we also quantified the change in ALM activity during photostimulation using the stimulus mode (stimulus mode response, Extended Data Fig. 2e, left, Extended Data Fig. 3d and Supplementary Fig. 2d, bottom), which allowed us to detect changes in the neural activity during photostimulation. For each session and for each trial type with photostimulation, we projected trial-averaged neural activity on the stimulus mode, combining correct and error trials. From this, we subtracted the trial-averaged projection on lick-left trials without distractors. The resulting change in projection activity was averaged over a given time window during photostimulation (stimulus and strong distractors, 0.4-s time window; weak distractors, 0.1-s time window). The stimulus mode response to distractors at different times and strengths was normalized to the stimulus mode response to the sample (that is, to the photostimulation presented during the sample epoch).

We compared photostimulation response sizes in the vS1 and the ALM at different time points (Extended Data Fig. 2e, right, Extended Data Fig. 3c and Supplementary Fig. 2d) using repeated-measures analysis of variance (ANOVA) and found no significant differences in responses to photostimulation at different time points during the task ($P > 0.05$). We also analyzed these data by fitting a linear regression. Regression analysis supported our original conclusions based on repeated-measures ANOVA; there were no significant changes in responses to distractors along the stimulus mode over time ($P > 0.05$). Statistical significance between response sizes to early- and late-delay distractors was done using Student's *t*-test with Tukey–Kramer correction for multiple comparison. Taken together, these result show absences of distractor suppression during delay.

Distractor impact on the choice mode. To assess the impact of distractors on the choice mode (Extended Data Figs. 5c and 10c), for each session we computed the trial-averaged projections of neural activity on the choice mode for trial types with distractors. We aligned the trajectories to the distractor onset and subtracted the projection for control trials (correct lick-left trials without distractors). This was done separately for trials with distractors delivered during early delay or late delay. In Extended Data Fig. 10c, we then averaged the activity across trial-averaged projections corresponding to different distractor times. We limited this analysis to 0.8 s following the distractor onset, which was the minimal common duration for each distractor trajectory until the go cue (the time interval from the late-delay distractor onset and the go cue). The resulting trajectory was defined as the distractor impact on the choice mode (Extended Data Figs. 5c and 10c).

Single-trial analyses of the choice mode. To analyze the separation of the neural activity projected on the choice mode (Extended Data Fig. 5d,e) on lick-left versus lick-right trials, we analyzed the distribution of single-trial projections on the choice mode, averaged over the last 0.2 s of the delay epoch (end-delay points). To compare across sessions, the distribution of end-delay points for each session was normalized[27] so that the median of end-delay points on lick-left trials was set to 0, and the median on lick-right trials was scaled to 1. To compute the medians, we used all response trials (including distractors) with lick-left and lick-right trials referring to the trial outcome. Extended Data Fig. 5d shows the end-delay points distribution (± s.e.m. across sessions) on correct lick-left versus lick-right trials without distractors. In Extended Data Fig. 5e (right), we computed the end-delay points distribution (± s.e.m. across sessions) for trials with an early-delay distractor, in which the animal ignored the distractor and licked left ('robust trials', gray), and trials in which the animal licked-right ('switching trials', light blue). Analogously, in Extended Data Fig. 5e (left), we computed the distribution of trajectories for robust and switching trials during mid-delay ([−0.8, −0.6] s relative to the go cue). This time window corresponded to the middle of the interval between the onset of the early-delay distractor (−1.6 s relative to the go cue) and the go cue.

Decoding of behavioral outcome from the neural activity projected along the choice mode. For decoder analyses shown in Fig. 3g, for each session, we first projected the neural activity on the choice mode. We then computed the optimal classification threshold for single-trial projections of lick-right versus lick-left trials via the receiver operating characteristic method. We averaged the projections over the last 0.2 s of the delay epoch. Lick-left or lick-right class labels were assigned to each trial according to the corresponding behavioral outcome (including distractor trials). Then, for each session, we counted the proportion of projected trials that could be classified as resulting in a lick-right response using the optimal threshold, after subtracting the proportion of no-stimulus trials that spontaneously resulted in a right lick (as we did for the behavioral analyses in Fig. 2c,d).

Single-trial analyses of the ramping mode. To analyze whether ramping predicted the degree of robustness to distractors (Fig. 5f), we analyzed the population activity projected on the nonselective ramping mode (equation (6) and Extended Data Fig. 4a,b) on single trials. For each trial, we smoothed the projected neural activity with a 200-ms causal boxcar filter. We limited our analysis to early-delay distractor trials to ensure that there were enough switching trials. For a given trial, we computed the ramping amplitude before an early-delay distractor as the ramping-mode projection averaged in a 0.4-s time window ending before the onset of the distractor (−1.6 s). For each session, we binned the distribution of ramping amplitudes across trials into three equally sized bins, corresponding to weak, intermediate and strong ramping amplitudes. We then computed the proportion of distractor trials with a lick-right outcome in each bin. We repeated the same procedure using control trials (that is, trials with instruction to lick left without distractors) to obtain the proportion of spontaneously occurring lick-right trials. To compute the ramping values for control trials, we used the same time window that we used for computing ramping values before the distractor onset for early-distractor trials. For each session, we subtracted the proportion of lick-right responses on control (no distractor) trials from the proportion of lick-right responses on distractor-trials computed for the corresponding ramping bins. The resulting difference was defined as proportion of switching trials for each level of non-selective ramping (x).

$$\text{Proportion of switching trials } (x) = \frac{\text{Right response trials } (x)_{\text{Distractor}}}{\text{All trials } (x)_{\text{Distractor}}} - \frac{\text{Right response trials } (x)_{\text{Control}}}{\text{All trials } (x)_{\text{Control}}} \quad (9)$$

Subtracting the control distribution allowed us to estimate the proportion of switches caused by distractors by accounting for the proportion of spontaneous lick-right responses that occurred in trials without distractors. Statistical significance of the effect of ramping amplitude on the proportion of switching trials was assessed using repeated-measures ANOVA ($F(2, 24) = 7.2$, $P = 0.003$; Fig. 5f).

To compare the neural activity projected on the ramping mode for simultaneously recorded putative pyramidal cells and fast-spiking interneurons (Extended Data Fig. 2g), we first computed the ramping mode separately for these two populations. For this analysis, we only included sessions that contained more than five putative fast-spiking interneurons. We then projected the neural activity on single trials, on the ramping mode, and averaged the projection over the last 0.1 s of the delay epoch. To compare across sessions, the distribution of end-delay points for each session was normalized.

Choice-mode trajectory slope on correct and switching trials. To compute the choice mode projection slope (Extended Data Fig. 6k) on correct lick-right trials versus switching early-delay distractor trials, we computed the trial-averaged projection of neural activity along the choice mode for each of the trial types. We then aligned the trajectory to stimulus/distractor onset. The slope was computed as the difference in the trial-averaged trajectory averaged over a time window $t_1 = [-0.2, 0]$ s and time window $t_2 = [1.4, 1.6]$ s divided by the difference in time window midpoints ($t_2 - t_1 = 1.6$ s). Time was defined relative to the stimulus/distractor onset. These time windows were chosen to cover the entire period between the onset of early-delay distractor and the end of delay epoch. We used the same procedure to determine the slope of the trajectories generated by the RNN (see below).

RNN models. Below is a summary of the RNN models presented in this study. For detailed descriptions of the models, see the Supplementary Math Note.

We built a rate-based RNN in which the output of each unit in the network matched the PSTH of an experimentally recorded neuron from the ALM of distractor-trained mice, pooled across sessions[30]. We trained four classes of RNNs (Extended Data Fig. 6a) defined by the presence or absence of an external ramping input and by whether the training included early-delay distractor PSTHs or not. RNNs that did not receive a ramping input (Extended Data Fig. 6a, green) were called 'autonomous'; note, however, that these networks received the same transient inputs as the ramping ones. In these networks, the slow ramping dynamics observed in the data were generated solely via recurrent internal dynamics. RNNs that received an additional, linearly ramping, external input were called 'ramping' RNNs; in these networks, in addition to the internal recurrent dynamics, the time-varying external input could be integrated to generate the slow ramping timescales seen in the ALM data[27] (Fig. 1e). We trained four networks to reproduce neural recordings from the left-hemisphere ALM: a distractor-free autonomous RNN (Extended Data Fig. 6b–d and Supplementary Fig. 4); a distractor-trained autonomous RNN (Extended Data Fig. 6e–g); a distractor-trained ramping RNN (Extended Data Fig. 6h–j); and a distractor-free ramping RNN (Figs. 4 and 5, Extended Data Figs. 6 and 8, and Supplementary Figs. 3, 4 and 6). We also trained a distractor-free ramping RNN to reproduce neural recordings from the ALM in both hemispheres (Extended Data Fig. 7).

Network dynamics, inputs to the network and connectivity. The networks were composed of N rate units whose membrane currents, $\mathbf{x}$, were modeled with the following system of N-coupled first-order differential equations as follows:

$$\tau \dot{\mathbf{x}}(t) = -\mathbf{x}(t) + W_{\text{rec}}\,\mathbf{r}(t) + \mathbf{w}_{\text{ramp}}\,I_{\text{ramp}}(t) + \mathbf{w}_{\text{stim}}\,I^{K}_{\text{stim}}(t) + \mathbf{w}_{\text{cue}}\,I_{\text{cue}}(t) + \sigma_{\eta}\boldsymbol{\eta}(t) \quad (10)$$

The spike rate was computed from the membrane currents by applying a sigmoidal transduction function elementwise as follows:

$$r_i = \phi(x_i) = \frac{1}{1+e^{-\beta(x_i-\theta)}} \quad (11)$$

The parameters of $\phi(x)$ were set to $\beta=0.8$ and $\theta=3$ to generate a Gaussian distribution of currents from the skewed log-normal distribution of spike rates (Supplementary Fig. 3a). The integration time constant τ was 10 ms for all units. The white noise term $\eta_i(t)$ was independently drawn at each time step for each neuron. Synaptic weights from a given unit in the network could be excitatory to some postsynaptic units and inhibitory to other postsynaptic units. We used the following three weight vectors for the different task-dependent inputs supplied to the network: $\mathbf{w}_{\text{stim}}$ for stimulus and distractors, $\mathbf{w}_{\text{ramp}}$ for external ramping and $\mathbf{w}_{\text{sample cue}}$ for the external cues delimiting the sample epoch. The weight vectors were drawn from the following three Gaussian distributions: $\mathbf{w}_{\text{stim}} \sim \mathcal{N}(0,1)$, $\mathbf{w}_{\text{ramp}} \sim \mathcal{N}(0,1)$ and $\{\mathbf{w}_{\text{cue}} \sim \mathcal{N}(0,0.1)$. $I^{K}_{\text{stim}}(t)$ was zero during lick-left trials (K=L) and had a square pulse profile (duration $T=400$ ms) during lick-right trials (K=R). The peak amplitude of the stimulus current $I^{\text{R}}_{\text{stim}}(t)$ in lick-right trials was drawn at each trial from a Gaussian distribution with $\mu_{\text{stim}}=1$ and $\sigma_{\text{stim}}=0.1$ during training and $\sigma_{\text{stim}}=0.4$ during testing. $I^{\text{L}}_{\text{stim}}(t)$ was always set to zero. The stimulus started at $t^{\text{start}}_{\text{stim}} = -2.5$ s for the regular stimulus delivered during the sample epoch right trials, at $t^{\text{start}}_{\text{early}} = -1.6$ s for early-delay distractor trials and at $t^{\text{start}}_{\text{late}} = -0.8$ s for late-delay distractor trials. $I_{\text{sample cue}}(t)$ had a square pulse profile of amplitude max ($I_{\text{sample cue}}$)=1 and duration $T=150$ ms. The first auditory cue was delivered at $t^{1}_{\text{cue}} = -3$ s before the go cue, the second at $t^{2}_{\text{cue}} = -2.15$ s. In the RNN of Figs. 4 and 5 (distractor-free ramping network), the external input $I_{\text{ramp}}(t)$ was a linearly ramping function of time, which started to ramp at the beginning of the sample epoch $\left(t^{\text{start}}_{\text{ramp}} = -3\text{ s}\right)$ until the go cue ($t_{\text{go cue}}=0$ s). The slope of $I_{\text{ramp}}(t)$ was drawn from a Gaussian distribution of mean $\mu_{\text{ramp}}=1$ and standard deviation $\sigma_{\text{ramp}}=0.1$. Instead, in RNNs with autonomous dynamics, $I_{\text{ramp}}(t)$ was set to zero for the entire duration of a trial. The auditory cues and the nonspecific external input were identical during lick-right and lick-left trials. All inputs were smoothed with a 400-ms square window. All parameters were identical across networks trained on the left ALM, with some exceptions. In particular, for distractor-trained RNN the following parameters were applied: (1) the slope of the ramping input current $I_{\text{ramp}}(t)$ during training was drawn from a Gaussian distribution with $\mu_{\text{ramp}}=1$ and $\sigma_{\text{ramp}}=0.05$; (2) during training, the peak amplitude of the stimulus/early-delay distractor current $I^{\text{R}}_{\text{stim}}(t)$ was drawn at each trial from a Gaussian distribution with $\mu_{\text{ramp}}=1$ and $\sigma_{\text{ramp}}=0.05$; (3) the amplitude of fast noise σ_{η} was set to 0.092 during both training and testing of the RNN models; and (4) in the testing phase, the stimulus/distractor amplitude was drawn from a Gaussian distribution with mean $\mu_{\text{stim}}=1$ and standard deviation $\sigma_{\text{stim}}=0.2$. Numerical solutions of network dynamics were obtained using the first-order Euler–Maruyama method with a time step of $\Delta t=0.1$ ms.

Network training and testing. We used the FORCE[34] algorithm to train networks to reproduce the neural activity of single cells[30] (Supplementary Math Note 1.1). Each unit in distractor-free RNNs was trained to reproduce the PSTHs of recorded neurons averaged across correct lick-right (K=R) or lick-left trials (K=L) without distractors. The training phase of distractor-trained networks also included the PSTHs averaged across correct or switching early-delay distractor trials (Extended Data Fig. 6a). To construct the PSTHs for each recorded neuron, spike trains were smoothed with a 400-ms boxcar window. To train the RNNs, we used only neurons recorded from distractor-trained mice, including both well-isolated neurons and multiunits. Except for the two-hemisphere network (Extended Data Fig. 7), we only included neurons recorded in the left ALM. The total number of units in networks trained to reproduce the neural activity recorded from the left-hemisphere ALM was $N=668$ for distractor-free networks and $N=545$ for distractor-trained networks due to the further exclusion of units that displayed an average spike rate lower than 0.01 Hz in at least one time-bin (resolution of 1 ms) in the early-delay distractor PSTH types (correct or error). The two-hemisphere network was composed of $N=668$ left-hemisphere units and $N=351$ right-hemisphere units.

It may be theoretically possible to conceive networks that display faster slopes by finely tuning the weights so that small differences in initial conditions (before stimulus presentation versus before early-delay distractor presentation) may generate different dynamics. However, fine-tuned solutions are unstable to both fast noise and structural changes and are sensitive to stimulus timing. The training procedure we implemented aimed at reducing the risk of convergence to fine-tuned solutions by introducing fast noise in the dynamics of single units (see equation (10)) and variability in both the amplitude and the presentation schedule of the network inputs. Therefore, independently for each trial, all time-dependent inputs were subject to a random jitter drawn from a uniform distribution within the interval [−10, 10] ms to potentially account for temporal variability of information transfer from other brain areas to the ALM. The training phase lasted for 1,000 trials for the distractor-trained ramping network and 2,000 trials for the distractor-trained autonomous network.

At the beginning of a given trial, a label that defined the trial type (and, consequently, which inputs were delivered to the RNN in that trial; Extended Data Fig. 6a) was randomly assigned to the trial. For each trial, the initial conditions, the jitters and the fast noise realizations were randomized. We excluded aberrant trials (<10% of all trials from all analyses, that is, trials that exhibited fluctuations six or more times bigger than the standard deviation of lick-right or lick-left trials at any time point during the delay epoch. The choice mode was computed from a subset of trials as the vector of spike rate differences during correct lick-right and lick-left trials, averaged during the last 400 ms of the delay epoch ($\mathbf{r}_{\text{choice}} = \bar{\mathbf{r}}^{\text{R}} - \bar{\mathbf{r}}^{\text{L}}$, the external ramping input was drawn from a Gaussian distribution of mean $\mu_{\text{ramp}}=1$ and standard deviation $\sigma_{\text{ramp}}=0.1$). To investigate whether the networks were able to generalize to conditions that were not presented during training, we probed the RNN both with trial types that were presented during the training phase and trial types that were not presented during the training phase (Fig. 4a and Extended Data Fig. 6a). Hence, after we estimated the choice mode, we simulated 20 sessions, each containing 100 trials of each type (lick-right, lick-left, early-delay and late-delay distractor). A lick-right trial k was considered to be correct if the network trajectory, projected onto the choice mode $r^{k}_{\text{proj}}(t) = (\mathbf{r}^{k}(t))^{\intercal}\,\mathbf{r}_{\text{choice}}$, lay above the halfway point between $\bar{r}^{\text{R}}_{\text{proj}}$ and $\bar{r}^{\text{L}}_{\text{proj}}$ at the end of the delay epoch (3.5 s into the trial corresponding to the time of the go cue; Fig. 4d). Conversely, for the remaining trial types, a trial was considered to be correct if its projection lay below the halfway point. A lick-right trial was defined as an error trial if its projected trajectory lay below the halfway point (and vice versa for the other trial types). We furthermore tested the ability of the ramping distractor-free RNN to generalize to multiple stimuli presentation (Supplementary Fig. 6). We presented the distractor-free ramping network with two stimuli: the standard sample epoch stimulus (onset at $t=-2.5$ s; Supplementary Fig. 6, blue) and an additional stimulus presented during the pre-sample epoch (−3.3 s; Supplementary Fig. 6, dark blue). According to our hypothesis, the external ramping input reflects task-relevant timing acquired through learning. For this reason, in this test, the external input began ramping at the onset of the pre-sample stimulus instead of at the beginning of the sample epoch.

Robustness to quenched noise in recurrent weights. The networks displayed some degree of robustness to multiplicative synaptic noise. We applied multiplicative Gaussian noise with zero mean and $\sigma_W=0.03$ to all recurrent weights. We simulated 100 realizations of the network with 200 stimulus trials and 200 no-stimulus trials in each realization. More than 50% of the realizations showed similar behavior (proportion of correct/switching trials) and dynamics (comparable end-delay points range) to the original network, both in the external ramping and in the autonomous case. With $\sigma_W=0.05$, 20% of network realizations would exhibit dynamics compatible with the ALM recordings.

RNN trained to reproduce ALM data from the left and right hemispheres. A number of experiments have demonstrated that the ALM is bilaterally coupled[26]. In Figs. 4 and 5, we trained RNNs on data from the left ALM since the optogenetic perturbation was performed in the left vS1. To show that excluding the right ALM hemisphere does not qualitatively affect the results of Figs. 4 and 5, we trained an additional RNN to reproduce the PSTHs of recorded putative excitatory neurons from both the left ($N=668$ units) and right ($N=351$ units) ALM during correct lick-right trials and correct lick-left trials (Extended Data Fig. 7). A linear ramping input (magenta) was delivered to units belonging to both hemispheres during the sample epoch and the delay epoch. Stimuli and distractors were delivered to units belonging to left ALM, but not to units belonging to the right ALM, to conform

with the experimental design according to which photostimulation was delivered to the left hemisphere. Initial recurrent weights and external input weights were drawn from the same statistical distribution as for the original non-autonomous RNN. All other parameters were also identical, with the following exceptions: (1) the slope of the ramping input current $I_{ramp}(t)$ during training was drawn from a Gaussian distribution with $\mu_{ramp}=1$; (2) the peak amplitude of the stimulus current $I^{R}_{stim}(t)$ was drawn at each trial from a Gaussian distribution with $\mu_{stim}=1$ and $\sigma_{stim}=0.01$; (3) the amplitude of fast noise σ_{η} was set to 0.05 during both training and testing of the RNN models; and (4) in the testing phase, the stimulus/distractor amplitude was drawn from a Gaussian distribution with mean $\mu_{stim}=1$ and standard deviation $\sigma_{stim}=0.1$ and the slope of the external ramping input was drawn from a Gaussian distribution of mean $\mu_{ramp}=1$ and standard deviation $\sigma_{ramp}=0.1$.

Analysis of RNN dynamics. We used a previously published method[31] to reverse engineer the RNNs (Supplementary Math Note 1.2). To search for slow points of the dynamics, at each time point, for example $t=-1.2$ s, we set all network inputs to their value at that time, for example, $I_{ramp}(t=-1.2\,s)$, and ran an optimization procedure in the vicinity of the network trajectories (Supplementary Math Note 1.2). The analysis of the autonomous RNN revealed the presence of two stable fixed points (that is, attractors) near the end-delay points of correct lick-right and lick-left trajectories along the choice mode. In this network, the two fixed points were separated by a series of slow points (that is, points where the dynamics was slow but not still; Extended Data Fig. 6d). Near a slow point, the network dynamic evolved on a time scale that was significantly slower than the single neuron decay time ($\tau=10$ ms), thereby generating the ramping observed in the data. In the ramping network, lick-left and lick-right correct trajectories remained close to a pair of stable fixed points during the entire delay epoch and were separated by a single saddle point (Fig. 5a,b). The distractor-trained networks (Extended Data Fig. 6g,j) displayed a phase space qualitatively similar to their distractor-free counterparts (Extended Data Fig. 6d and Fig. 5a). The slopes of switching trials (Extended Data Fig. 6k) generated by the RNNs were computed using the same procedure applied to the experimental data (see the section "Choice-mode trajectory slope on correct and switching trials").

Statistics. No statistical methods were used to predetermine sample sizes, but our sample sizes were similar to those used in other similar studies[27]. We performed recordings for multiple behavior sessions and from multiple animals to confirm reproducibility. All results were reproducible across behavior sessions and animals. We did not exclude any mouse for data analyses. During experiments, trial types were randomly determined by a computer program. During spike sorting, experimenters could not tell the trial type, so experimenters were blinded to the experimental conditions. Analyses of neural and behavioral data were conducted regardless of the identity of the animal from which the data were collected. All comparisons using Student's *t*-test and Wilcoxon rank-sum test were two-sided. Data distribution was assumed to be normal, but this was not formally tested.

Reporting Summary. Further information on research design is available in the Nature Research Reporting Summary linked to this article.

Data availability

Data in NWB format are available for download at https://dandiarchive.org/dandiset/000060/draft.

Code availability

The Matlab code for data analyses is available at https://github.com/arsenyf/FinkelsteinFontolan_2021NN. The Matlab code for network models is available at https://github.com/fontolanl/RNN_ALM_gating.

References

51. Madisen, L. et al. A robust and high-throughput Cre reporting and characterization system for the whole mouse brain. *Nat. Neurosci.* **13**, 133–140 (2010).
52. Madisen, L. et al. A toolbox of Cre-dependent optogenetic transgenic mice for light-induced activation and silencing. *Nat. Neurosci.* **15**, 793–802 (2012).
53. Pluta, S. et al. A direct translaminar inhibitory circuit tunes cortical output. *Nat. Neurosci.* **18**, 1631–1640 (2015).
54. Liu, L., Finkelstein, A., Chen, S., Li, N. & Svoboda, K. Headbar implantation. *Protocols.io* https://doi.org/10.17504/protocols.io.bcrsiv6e (2020).
55. Yatsenko, D. et al. DataJoint: managing big scientific data using MATLAB or Python. Preprint at *bioRxiv* https://doi.org/10.1101/031658 (2015).
56. Jun, J. J. et al. Real-time spike sorting platform for high-density extracellular probes with ground-truth validation and drift correction. Preprint at *bioRxiv* https://doi.org/10.1101/101030 (2017).

Acknowledgements

We thank D. Hansel and S. Druckmann for discussions; J. Aljadeff, K. Daie, R. Darshan, J. Drugowitsch, H. Inagaki, M. Rozsa and T. Wang for comments on the manuscript; T. Nguyen and other Vathes members for help with DataJoint, and T. Pluntke for animal training. This work was funded by the Howard Hughes Medical Institute and the Simons Foundation. A.F. is a Rothschild Foundation and EMBO Long-Term postdoctoral fellow (ALTF 869-2015).

Author contributions

A.F., M.N.E. and K.S. designed the experiments, with help from N.L. A.F. and M.N.E. collected the experimental data with help from N.L. A.F. analyzed the behavioral and neural data, with input from M.N.E., L.F., S.R. and K.S. L.F. and S.R. conceived the computational models with input from A.F. and K.S. L.F. built the models and performed model simulations. A.F., L.F., S.R. and K.S. wrote the paper, with input from all the authors.

Competing interests

The authors declare no competing interests.

Additional information

Extended data is available for this paper at https://doi.org/10.1038/s41593-021-00840-6.

Supplementary information The online version contains supplementary material available at https://doi.org/10.1038/s41593-021-00840-6.

Correspondence and requests for materials should be addressed to S.R. or K.S.

Peer review information *Nature Neuroscience* thanks Athena Akrami, Anne Churchland, and Albert Compte for their contribution to the peer review of this work.

Reprints and permissions information is available at www.nature.com/reprints.

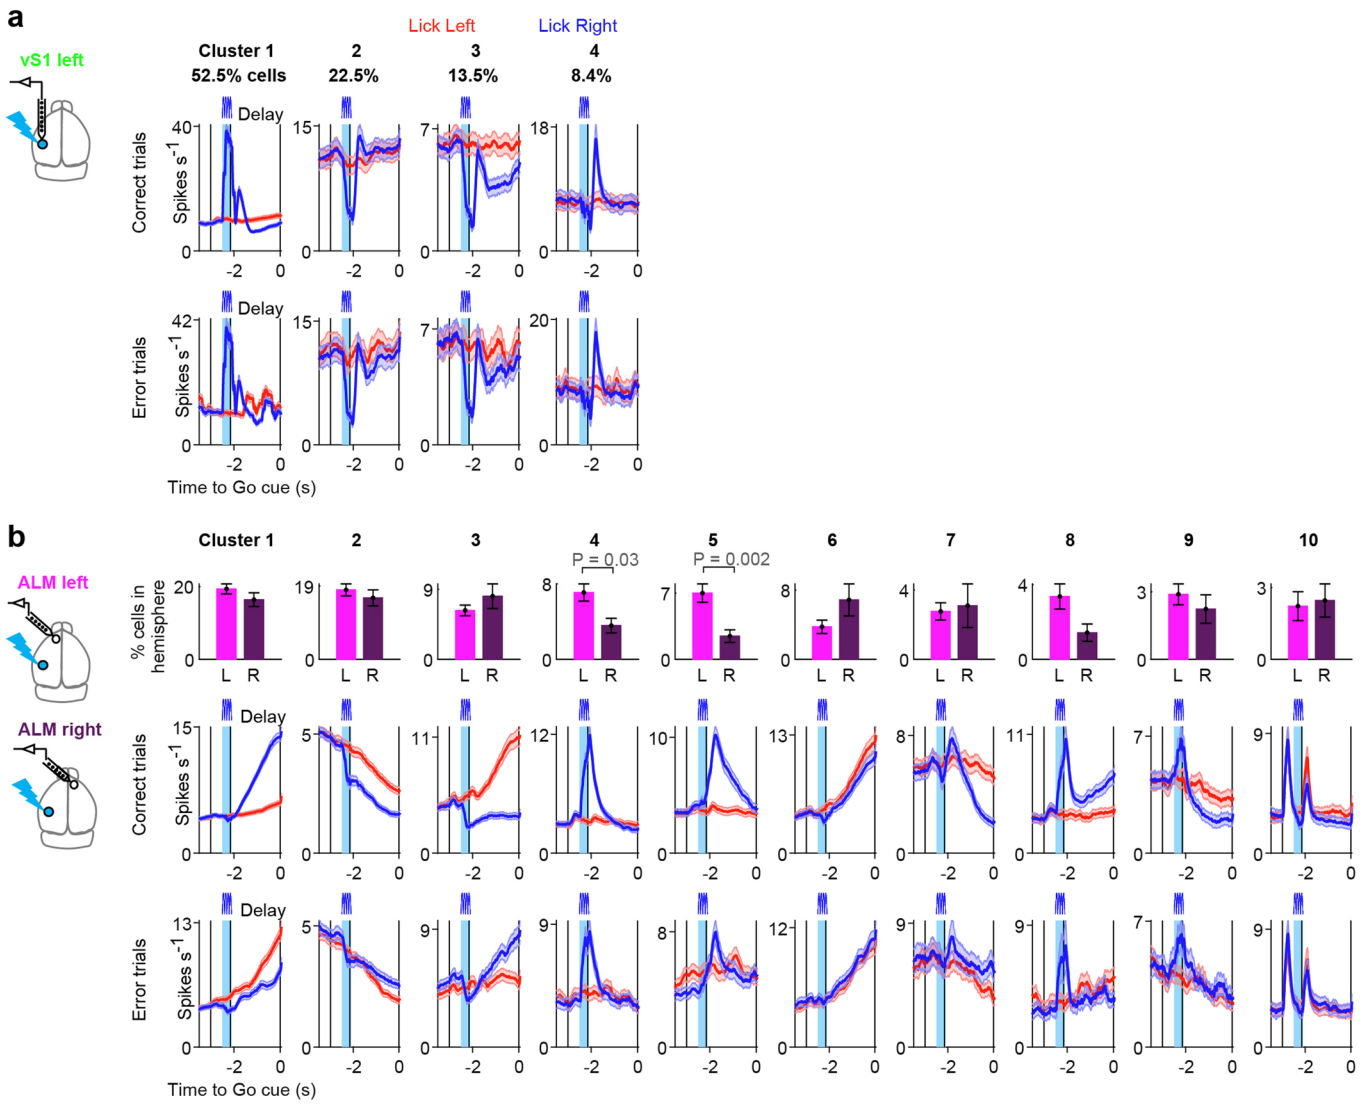

Extended Data Fig. 1 | Dynamics of putative pyramidal neurons in the basic task. Hierarchical clustering of trial-averaged PSTHs of putative pyramidal cells recorded in vS1 (**a**, $n = 256$ cells) and ALM (**b**, left ALM $n = 1648$; right ALM $n = 679$ cells). **a**, Average spike rate of vS1 cells in each cluster and percentage of cells in each cluster. Red: trials without photostimulation, (lick-left). Blue: trials with photostimulation (lick-right); Cyan bar indicates the photostimulus. Top, correct trials. Bottom, error trials. Data is presented as trial-averaged activity, then averaged across cells belonging to each cluster ± s.e.m. (across cells, shaded). Cells in vS1 (clusters 1–4) did not switch selectivity on error trials, indicating that they tracked stimulus-related information. **b**, Same for ALM neurons. Neurons recorded from both hemispheres were clustered together. Bars show the percentage of putative pyramidal cells in left and right ALM that fall within each cluster, mean ± s.e.m across cells. Statistical significance was assessed by Wilcoxon rank sum test (two-sided). Cells in ALM often showed opposite selectivity on error trials (for example, Cluster 1–3), indicating that they tracked choice-related information. Despite the laterality of the movement, trial type selective preparatory activity was similar in left and right ALM (see distribution of cells in clusters 1–3 across hemispheres). Only a minority of cells in ALM did not switch selectivity on error trials – indicating sensory responses (for example, Cluster 4; note that cells in this cluster were more prevalent in left hemisphere). In addition, there was a cluster with neurons exhibiting non-selective ramping responses (Cluster 6), and a cluster with neurons responding to auditory cues that delineated the sample epoch (Cluster 10).

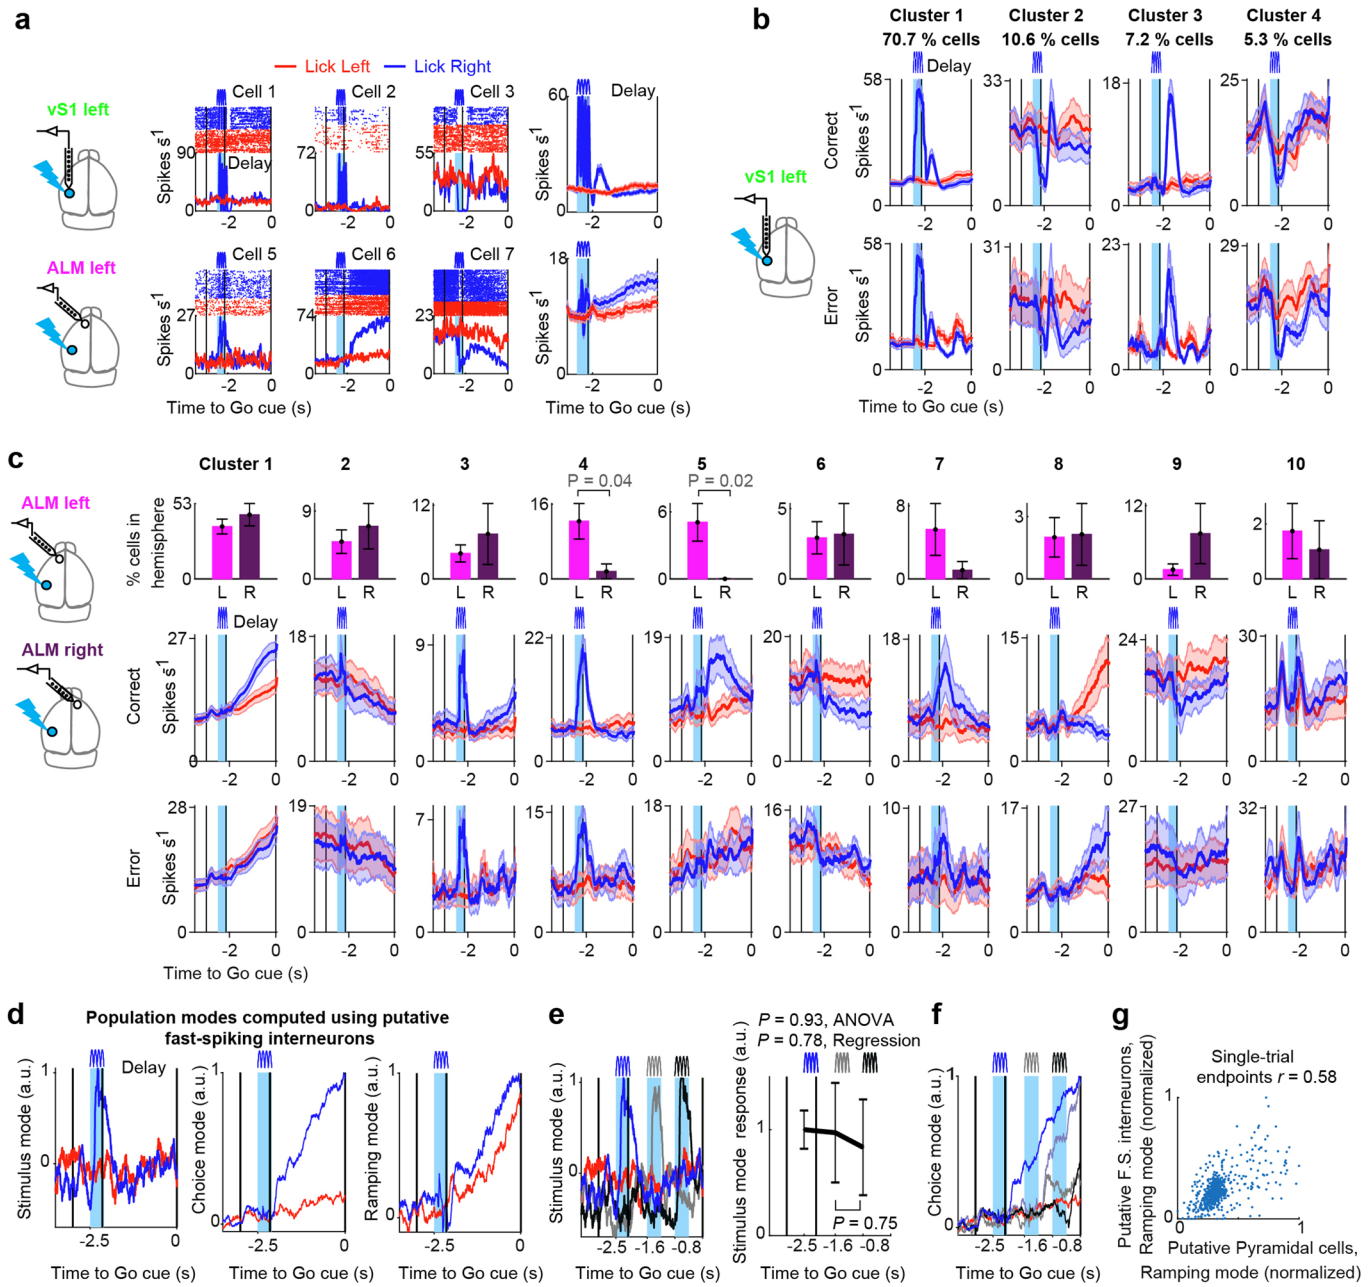

Extended Data Fig. 2 | Dynamics of putative fast-spiking interneurons in the basic task and in the presence of distractors. Activity of putative fast-spiking interneurons (Methods) recorded in vS1 ($n = 92$ cells) and ALM (left ALM $n = 184$; right ALM $n = 96$ cells). **a**, Example cells and grand population average (labels as in Fig. 1d-g). **b-c**, Hierarchical clustering of trial-averaged PSTHs of putative interneurons (labels as in Extended Data. Figure 1a-b, bars indicate mean ± s.e.m across cells). Statistical significance was assessed by Wilcoxon rank sum test (two-sided). Putative interneurons in vS1 and ALM show a diversity of response profiles, suggesting that their activity does not simply reflect average activity of excitatory neurons. **d**, Session-averaged projections of neural activity of putative fast-spiking interneurons in left ALM on Stimulus, Choice, and Ramping modes ($n = 104$ cells, distractor trained mice). Red, lick-left trials; blue: lick-right trials. **e-f**, Trials with distractors projected on Stimulus mode (**e**, left) and Choice mode (**f**). Color-code as in Fig. 3 f,h. There was no suppression of distractors in the Stimulus mode during delay (e, right), whereas Choice mode exhibited robustness to distractors at late delay. Error bars, mean ± s.e.m., across sessions. Response size was assessed by repeated measures ANOVA, and also verified by fitting linear regression model. Statistical significance between response size to early- and late-delay distractors was done using two-sided Student t-test with Tukey-Kramer correction for multiple comparison. **g**, Endpoints of ramping mode of simultaneously recorded putative pyramidal cells versus putative fast-spiking interneurons. Each dot represents the Ramping mode endpoint on a single trial computed for the last 0.2 s of the delay epoch. Endpoints of ramping mode of simultaneously recorded putative fast-spiking (F.S.) interneurons and putative pyramidal neurons were correlated on single trials (r represents the Pearson correlation coefficient).

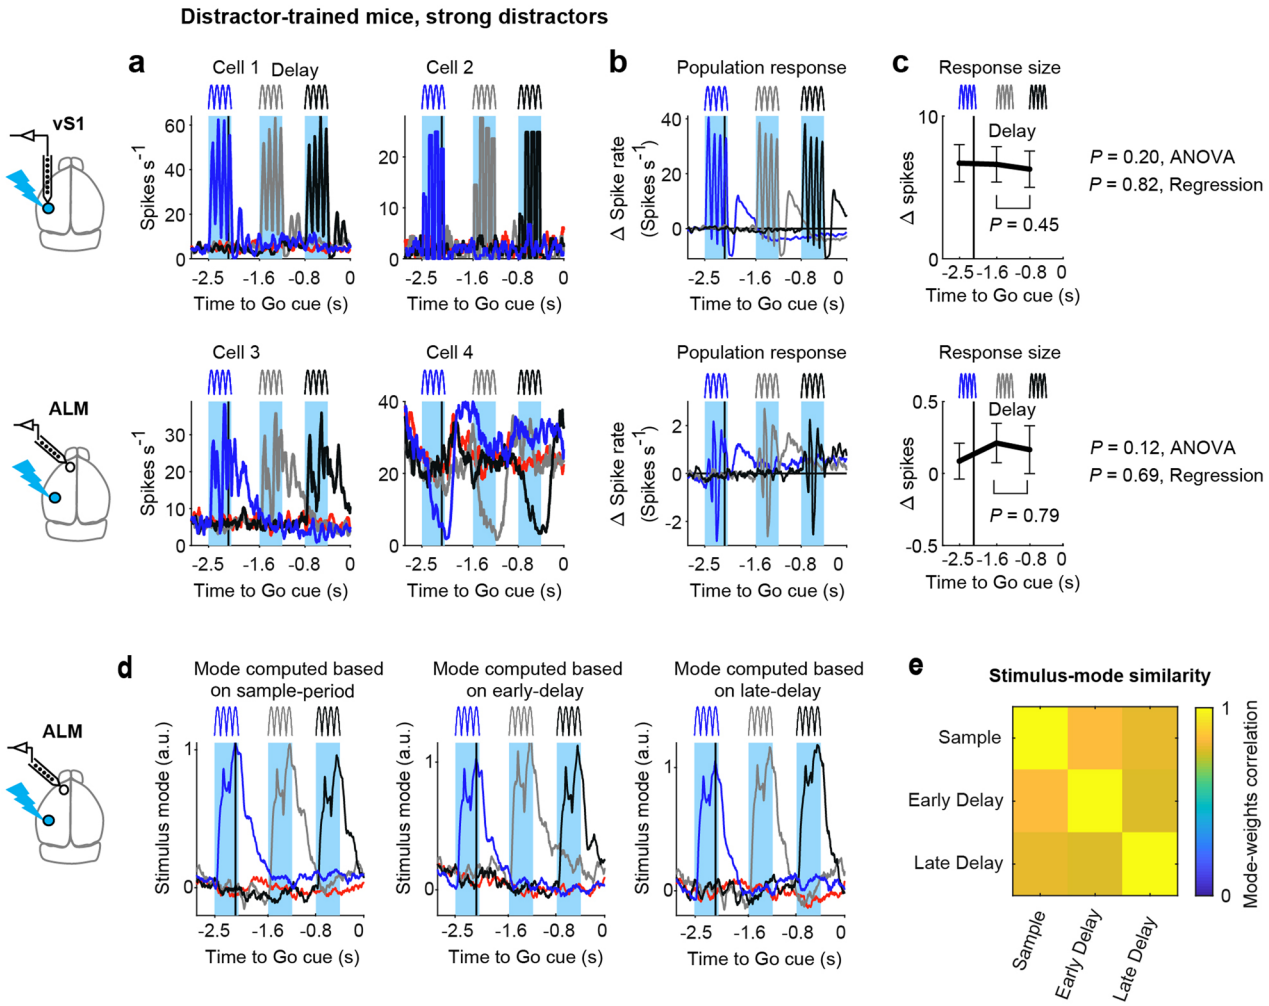

Extended Data Fig. 3 | Transient activity in vS1 and ALM in response to photostimulations in distractor-trained mice. Spike rate modulation in vS1 (n = 317 cells) and left ALM (n = 1061 cells) of distractor-trained mice in response to strong distractors at different times, computed using all response trials, regardless of the behavioral outcome. Lick-left trials without stimulation (red), with distractors during early-delay (gray), or late-delay (black). Lick-right trials during sample-epoch stimulation (blue). **a**, Spike rates of example cells. **b**, Population-averaged responses. **c**, Quantification of the response size based on the population average (**b**). There was no reduction of distractor size in vS1 or in ALM of distractor-trained mice as assessed by repeated measures ANOVA, and also verified by fitting linear regression model, despite differences in the effect of distractors on behavior (Fig. 2d). Statistical significance between response size to early- and late-delay distractors was done using two-sided Student t-test with Tukey-Kramer correction for multiple comparison. **d**, For ALM recordings, we also used targeted dimensionality reduction to extract stimulus-related activity ('Stimulus mode', Methods). Throughout the paper, we defined the weights of Stimulus mode based on the difference in firing-rate between trials with and without stimulus, computed during the 0.5 s time-window following stimulus presentation at the Sample epoch (left panel). Here we also calculated the Stimulus mode based on the difference in neural activity following either early- (middle panel) or late-delay distractor (right panel). Neural activity following stimulus or distractors projected on these three modes was rather similar, regardless of what time epoch was used to compute the modes. This indicates that there was no distractor suppression during late-delay regardless of the way the Stimulus mode was computed. **f**, Correlation between the vectors of neuronal weights for the different definitions of Stimulus modes, showing that these modes were very similar (average Pearson correlation coefficient r = 0.80). Taken together, this suggests that the input channel to ALM (that is the Stimulus mode) does not change substantially through the sample and delay epochs. These results further support our observations that distractors were not suppressed during delay at the level of inputs to ALM. Error bars, mean ± s.e.m., across sessions (n = 17 sessions).

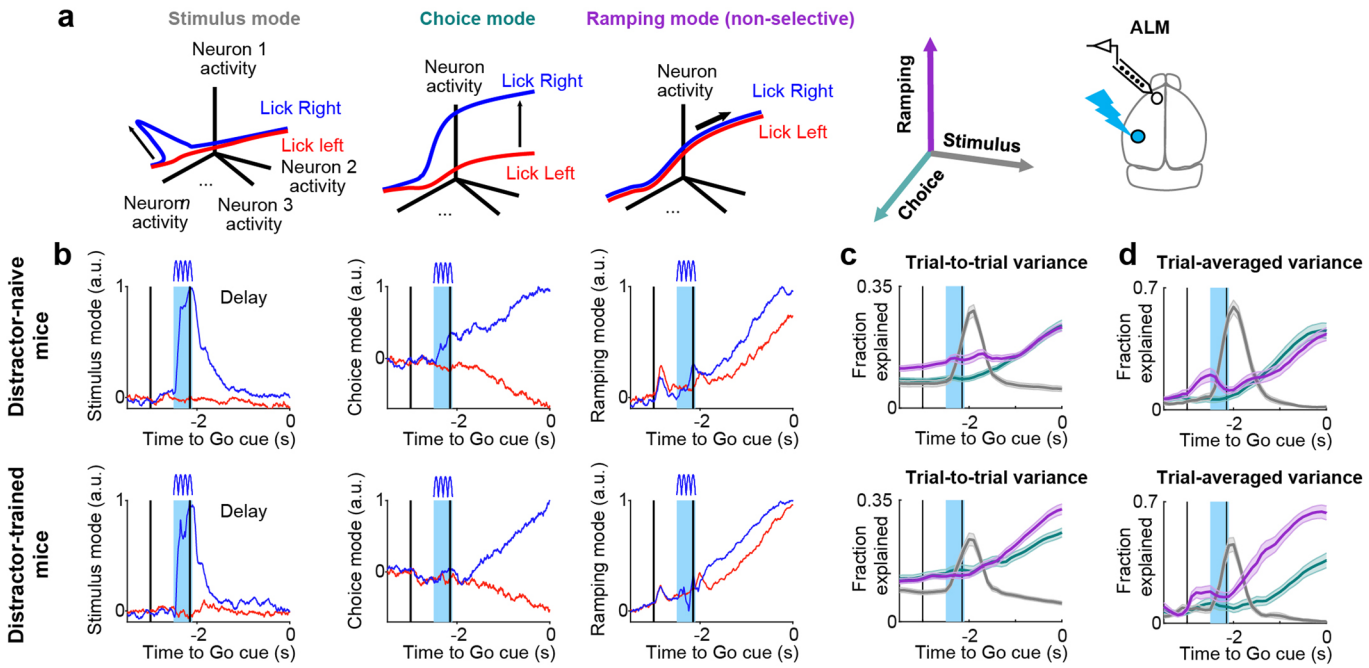

Extended Data Fig. 4 | Dimensionality-reduction on population dynamics in ALM. a, Targeted dimensionality-reduction to define Stimulus, Choice (trial type selective), and Ramping (non-selective) modes in neural activity space. Modes are defined as one-dimensional subspaces in activity space (arrows) and are orthogonalized with respect to each other (Methods). **b**, Session-averaged projections of neural activity in left ALM on different modes in distractor-naive (top, $n = 1356$ cells) and distractor-trained (bottom, $n = 1061$ cells) mice. Red, lick-left trials; blue: lick-right trials. **c**, Proportion of trial-to-trial variance explained by different modes in distractor-naive (top) and distractor-trained (bottom) mice. Variance was computed at different time-points along the trial. **d**, Same as in c, for proportion of trial-averaged variance explained. Trial-averaged variance was computed separately for lick-left and lick-right trials. In all plots, for Choice mode analyses we used correct trials, excluding trials with distractors; for Stimulus and Ramping modes analyses we used correct and error trials, excluding trials with distractors. Shaded area represents the standard error of the mean (s.e.m.), across sessions.

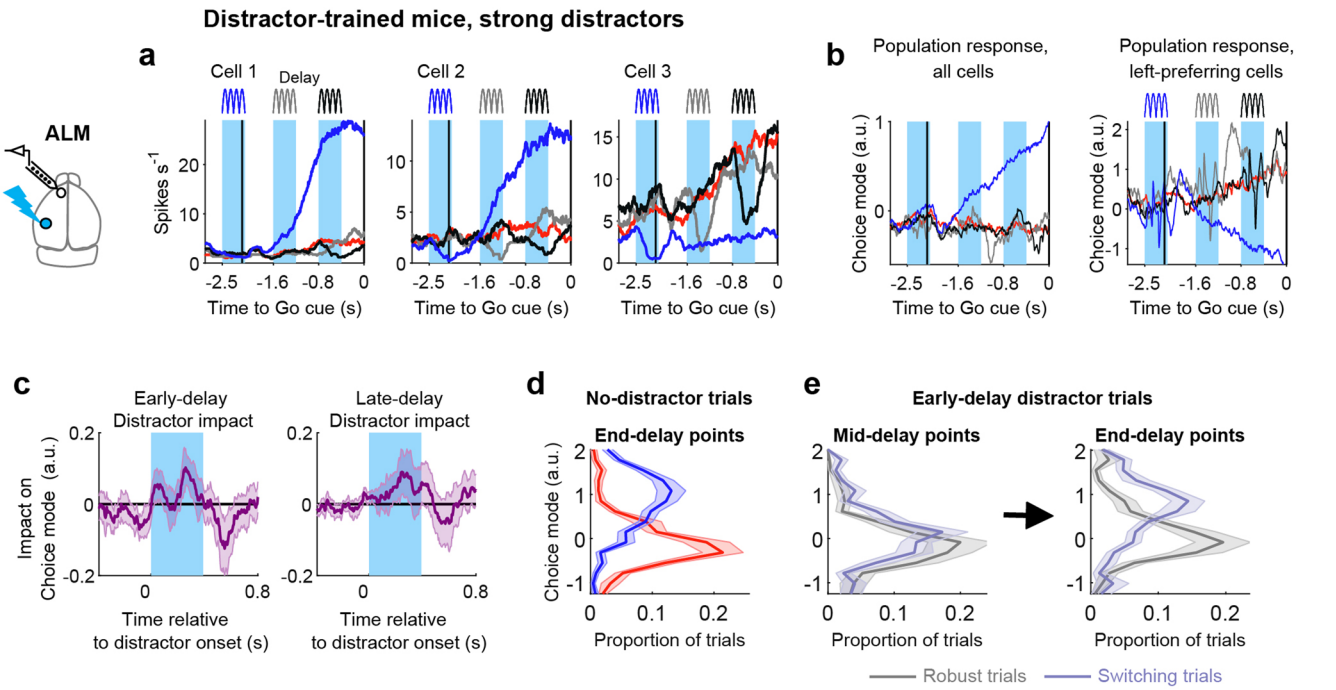

Extended Data Fig. 5 | Robustness of the Choice mode in response to distractors. Spike rate modulation in left ALM ($n = 1061$ cells) in response to strong distractors at different trial epochs, in distractor-trained mice. **a**, Example cells. **b**, Projection of neural activity on the Choice mode using all cells (left panel) or using only left-preferring cells (right panel). Lick-left trajectory without stimulation (red), with distractors during early-delay (gray), or late-delay (black), computed using correct trials. Lick-right trajectory during sample-epoch stimulation (blue). **c**, Impact of distractors on the Choice mode. Trajectories were aligned to the onset of each distractor. Data is shown as average across sessions ± s.e.m. (shaded). Data in a-c was computed using correct trials. The effect of distractors on the Choice mode on correct trials was temporary: distractors often resulted in transient change in spike rate of individual cells that contributed to the Choice mode, but their activity later recovered to the unperturbed (red) trajectory, indicating robustness. **d**, Distribution of Choice mode projections on single trials without distractors, averaged over the last 0.2 s of the delay epoch (*end-delay points)* for lick-left (red) versus lick-right (blue) trials. **e**, Distribution of Choice mode projections on trials with early-delay distractor (gray, robust trials; light blue, switching trials). Projections were averaged during mid-delay (left panel) or end-delay (right panel, Methods). Data is expressed as mean ± s.e.m. (shaded) across sessions.

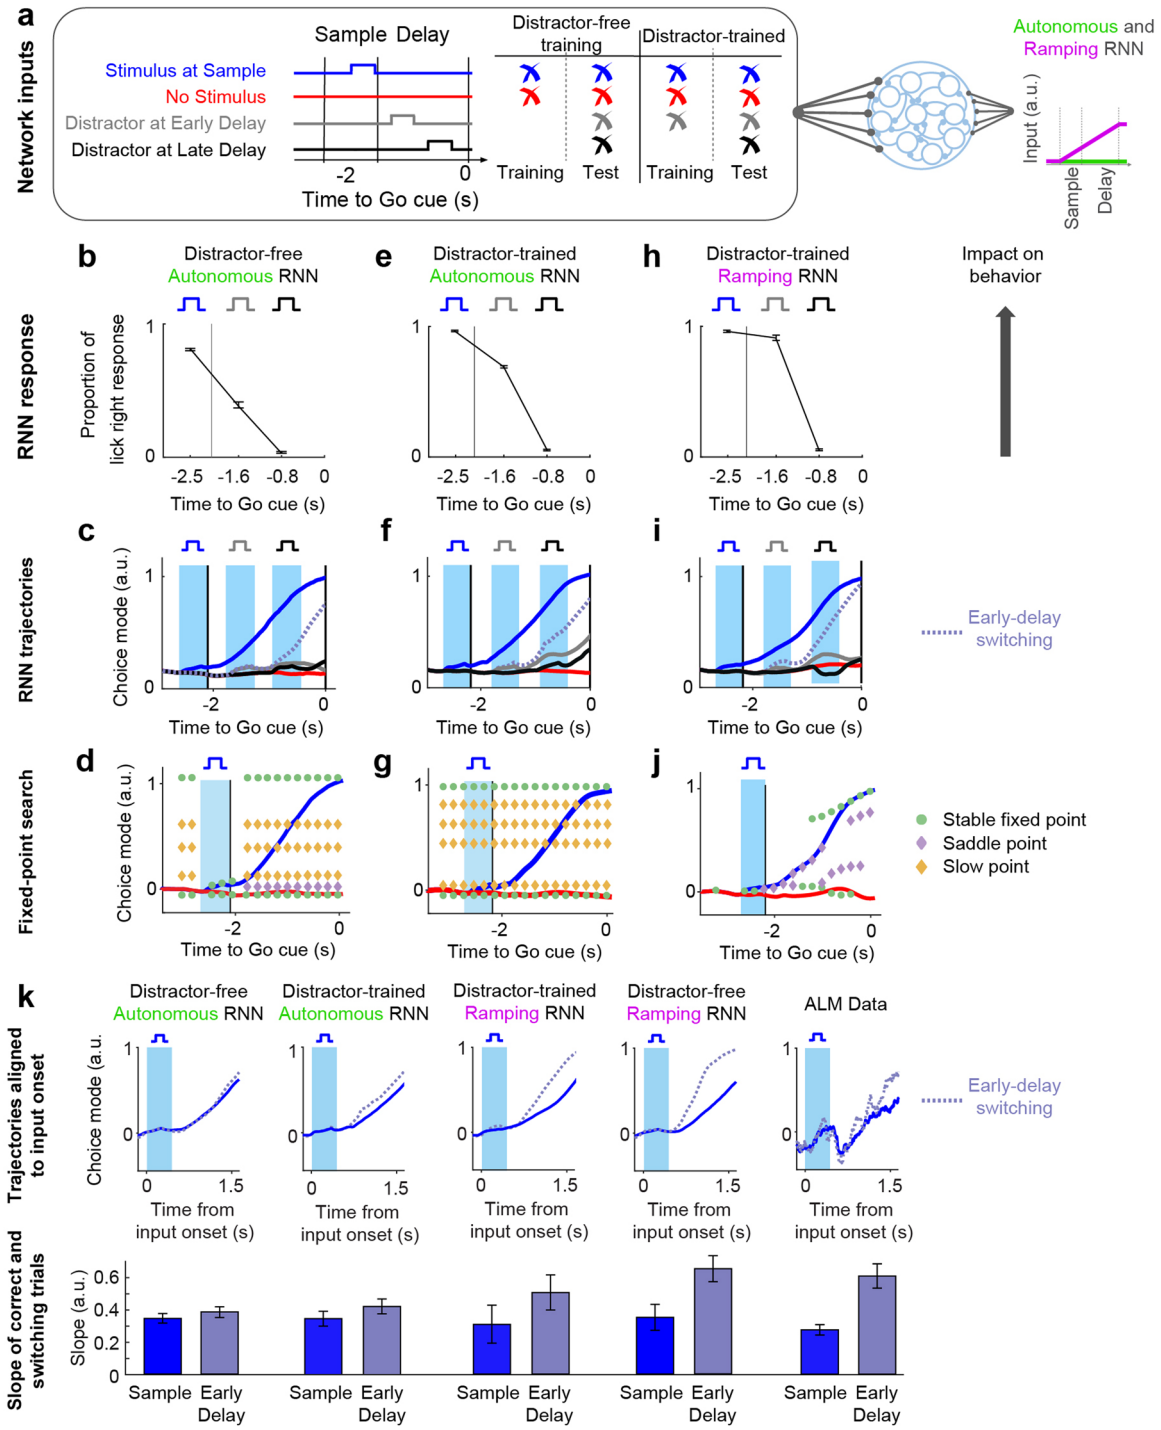

Extended Data Fig. 6 | See next page for caption.

Extended Data Fig. 6 | Distractor-free and distractor-trained RNN models with and without external ramping, comparison with ALM data. a, Schematics of inputs to RNN models. Distractor-free RNNs were trained to reproduce stimulus and no-stimulus trials only. Distractor-trained RNNs were trained to additionally reproduce the PSTH of correct early-delay distractor trials. In addition, RNNs could receive a linearly ramping external input (purple, 'ramping RNN') or a constant zero input (green, 'autonomous RNN'). Autonomous RNNs relied solely on internal recurrent (autonomous) dynamics to reproduce the slow ramping observed in the data, while ramping RNNs could make use of both the recurrent dynamics and the ramping input. We trained four networks based on the activity of neurons in left ALM: a distractor-free autonomous RNN (**b-d**), a distractor-trained autonomous RNN (**e-g**), a distractor-trained ramping RNN (**h-j**) and a distractor-free ramping RNN (Figs. 4–5). **b,e,h**, Proportion of lick right responses. In the 3 networks shown, the proportion of lick-right distractor-evoked trials was qualitatively similar to those observed in the data (Fig. 2d) and in the distractor-free ramping RNN (Fig. 4c), implying that all RNNs displayed a form of temporal gating. Error bars, mean ± s.e.m across sessions ($n = 20$, 100 trials per condition). **c,f,i**, RNNs activity projected on Choice mode for correct stimulus trials (blue), correct no-stimulus trials (red), correct and switching trials evoked by early-delay distractor (full and dotted gray, respectively), correct late-delay distractor trials (black). **d,g,j**, Fixed-point search results projected onto Choice mode. Green circles, stable fixed points. Purple diamonds, saddle points. Yellow diamonds, slow points (that is points where the network flow was significantly different from zero but much slower than the typical timescale of a single neuron, see Methods). In the autonomous RNNs (**d**,**g**), trajectories moved through a number of slow points before reaching either of the two stable-fixed points, in contrast to ramping networks in which, during most of the delay epoch, trajectories remained close to either of the two stable fixed points that were moved by the external ramping input (j, Fig. 5a). Hence, ramping and autonomous networks relied on different mechanism to achieve the slow ramping observed in the data. The fixed-point search revealed a partial difference in the dynamics of distractor-free and distractor-trained ramping networks. In the distractor-free ramping network, trajectories closely followed the moving stable fixed-point during the delay epoch (Fig. 5a). In the distractor-trained ramping network (**j**), lick-right trajectories (blue) were close to a slow saddle point during the first half of the delay epoch (purple diamonds) and then approached the moving stable fixed points during the second part of the delay. From the offset of the early-delay distractor until the end of the delay epoch the dynamic**s** of the distractor-trained ramping network was qualitatively similar to that of the distractor-free ramping network. In addition, both the saddle point (present in the first part of the delay epoch) and the stable fixed-points (present in the second part of the delay epoch) were moving because of the external ramping stimulus, confirming that the ramping stimulus drives the dynamics in the distractor-trained ramping network. **k** (top), Trial-averaged trajectories of stimulus-evoked correct trials (blue) and early-distractor evoked (gray) switching trials, aligned to stimulus/distractor onset for all trained RNNs and for ALM data. **k** (bottom), Slopes of stimulus-evoked correct trials (blue) and early-distractor-evoked (gray) switching trials trajectories. Error bars, mean ± s.d. across trials (RNNs, 400 trials per condition) or mean ± s.e.m. across sessions (ALM data from distractor-trained mice only, $n = 17$ sessions). Statistical significance between stimulus-evoked and early-delay distractor-evoked slopes was assessed by Student t-test (two-sided, no corrections were made for multiple comparisons); from left to right: P=0.0074, P <0.0001, P <0.0001, P<0.0001, P=0.0003. Switching trajectories evoked by early-delay distractor in ALM data (rightmost panels) exhibited a steeper slope compatible with trajectories generated by the RNN models with external ramping input, but not with trajectories generated by the autonomous RNNs.

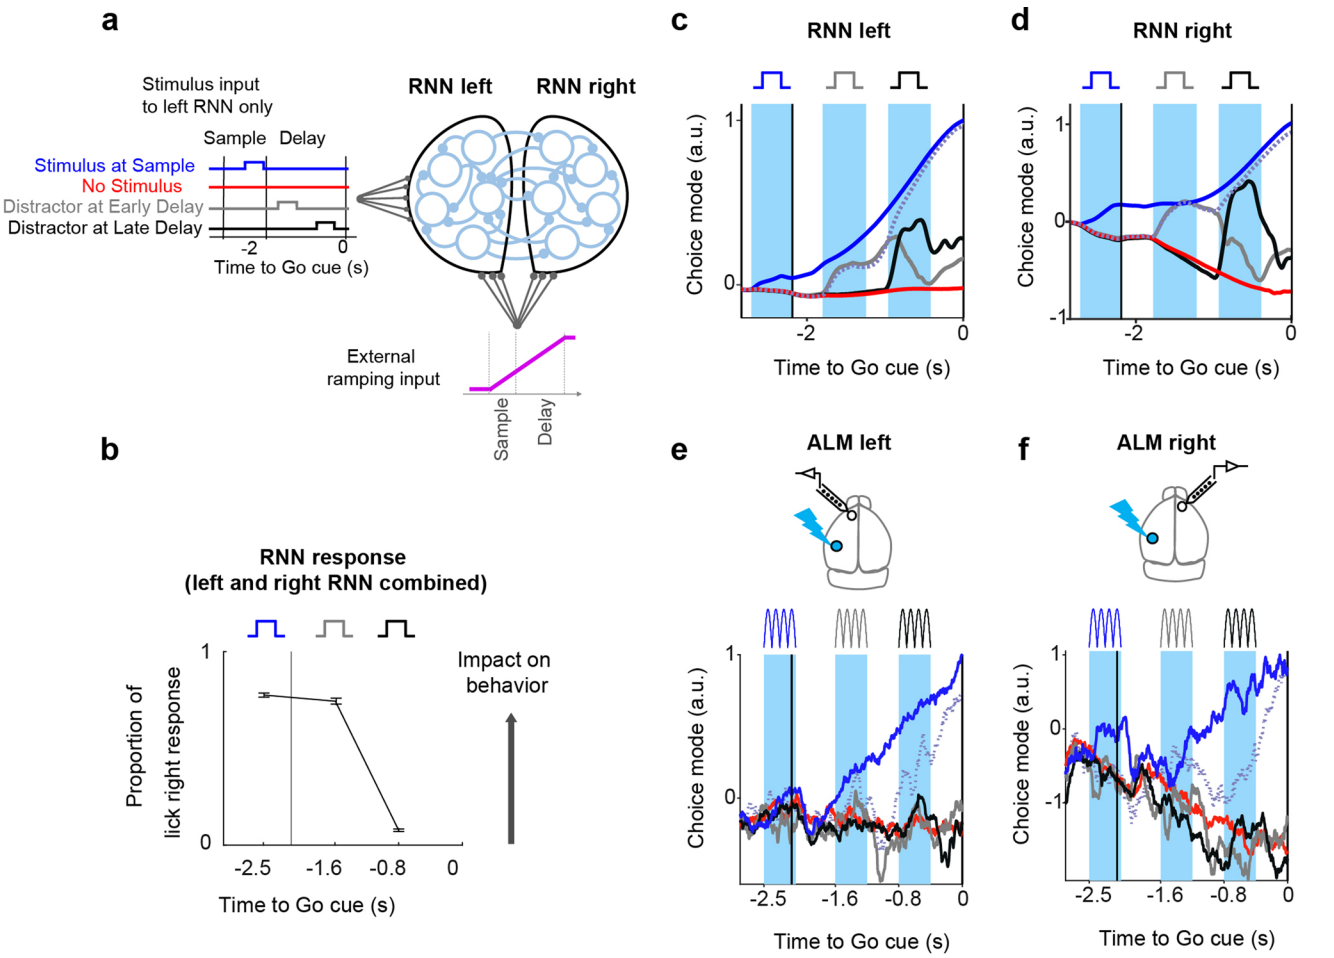

Extended Data Fig. 7 | Recurrent neural network trained on ALM data from the left and right hemispheres. a, RNN trained to reproduce the PSTHs of recorded putative excitatory neurons from left ALM ($n = 668$ units) and right ALM ($n = 351$ units), during correct lick-right trials and correct lick-left trials. A linear ramping input (black) was delivered to units belonging to both hemispheres during sample and delay epochs. Stimuli and distractors were delivered to units belonging to left ALM, but not to units belonging to the right ALM, to conform with the experimental design according to which photostimulation was delivered to the left hemisphere. **b**, Proportion of right outputs generated by the two-hemisphere RNN model with ramping input, in response to stimulus (blue), and distractors during early (gray) and late (black) delay. The network outcome (left/right) was defined by whether the activity along the Choice mode (computed in activity space of left hemisphere RNN units) crossed the choice boundary at the end of the delay (Methods). Error bars, mean ± s.e.m. across sessions ($n = 20$, 100 trials per condition). **c-d**, Two-hemisphere RNN activity projected on Choice mode (**c**, computed using units belonging to left ALM (**c**) or right ALM (**d**) for correct trials (solid lines) and trials with early-delay distractor that resulted in right response ('switching trials', dashed lines); color scheme as in Fig. 3h. **e-f**, Neural activity recorded from ALM on left (**e**) and right hemisphere (**f**), projected onto Choice mode; color scheme as in panel c-d. The dynamics of the two-hemisphere network exhibited gating of distractors during late-delay, similarly to the dynamics observed in ALM in both hemispheres.

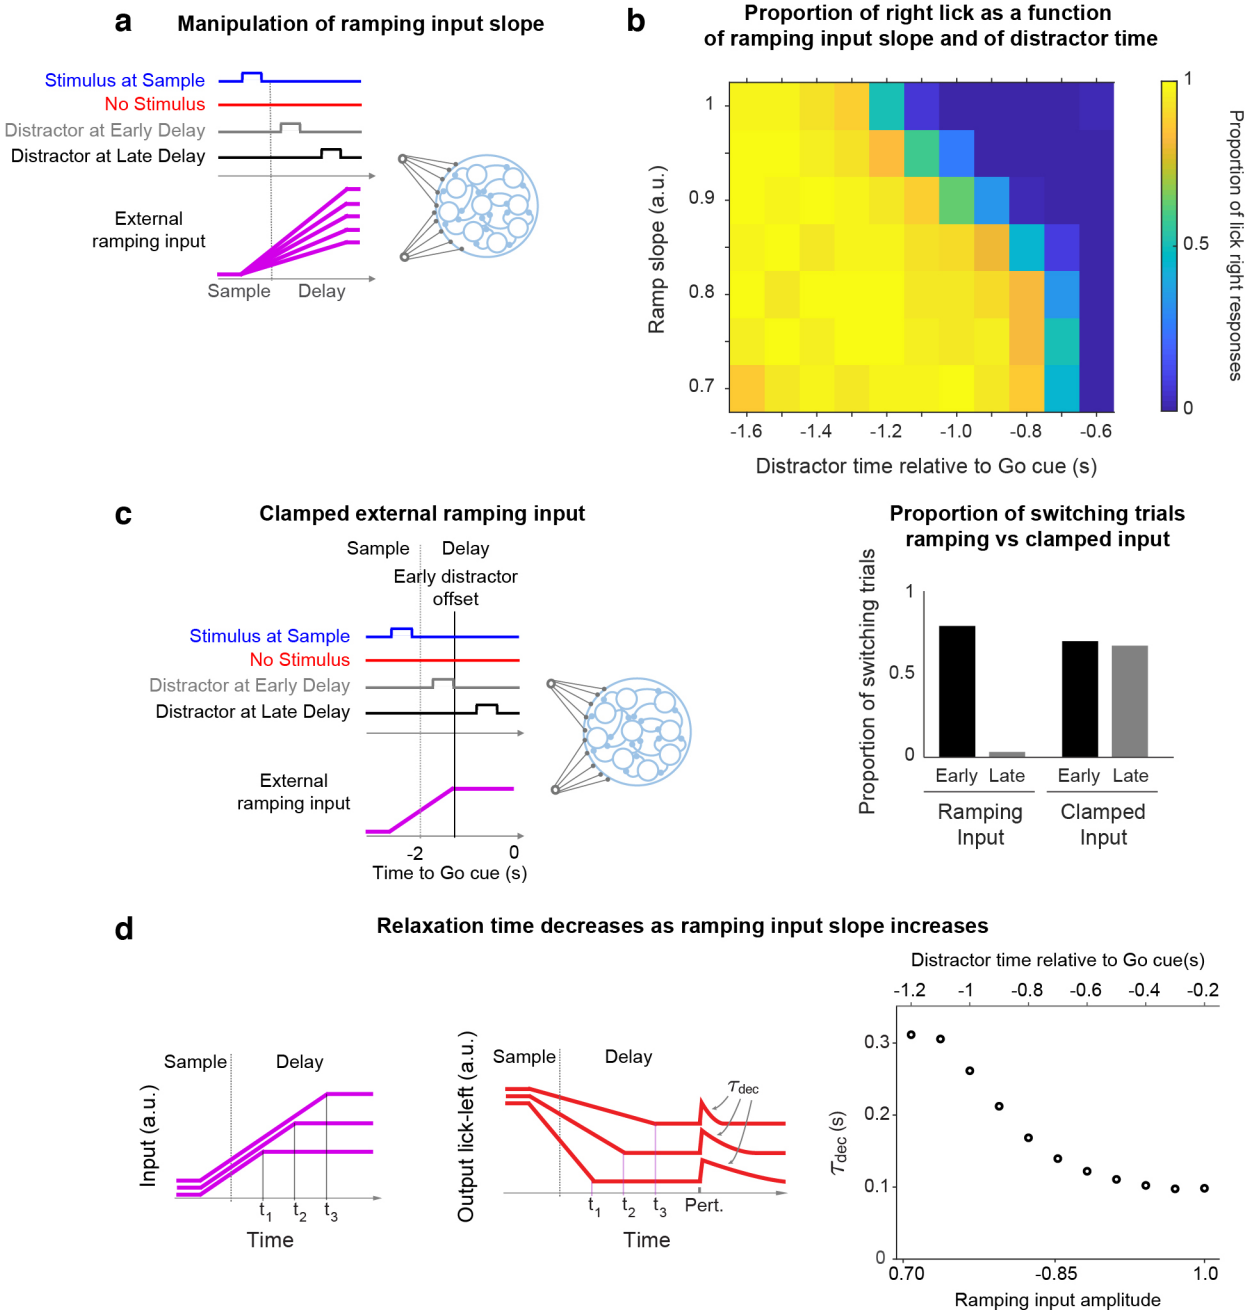

Extended Data Fig. 8 | External ramping input level determines robustness to distractors. a, Schematics of distractor-free ramping RNN when an external ramping input of variable slope (purple) was delivered to the network. **b**, Proportion of distractor-evoked switching trials as a function of the time at which the distractor was delivered and of ramping input strength. Note that i) weaker ramping input induce more errors and ii) distractors delivered late in the delay are gated more efficiently. This suggests that elapsed time and amplitude of the ramping input were directly related: the gating of late distractors could be explained by the increased amplitude of the external ramping compared to earlier times in the delay. **c**, Assessing the contribution of external ramping input to temporal gating of distractors in the distractor-free ramping network. Left panel: To dissociate the impact on RNN dynamics of the external ramping input (which progressively increased during sample and delay epochs) from the effects of passage of time *per se* (that is regardless of the ramping), we clamped the external ramping input at the level it reached at the offset of the early-delay distractor (t = –1.2 s, black vertical line). Right panel: Proportion of switching trials for early- (gray) and late-delay (black) distractors: unlike the case examined in Fig. 4 (in which, for the same RNN, the external input kept ramping throughout the delay epoch), the proportion of switching trials evoked by late-delay distractor when the ramping input was clamped was similar to that of early-delay distractor. This indicates that temporal gating is regulated by the external ramp level and not by the passage of time *per se*. **d**, Relaxation time constants of distractor-free ramping network as a function of time at which the perturbation was delivered. In this test, the ramping input was set to the value it had reached at time t_i in the delay (left panel), and, for each ramping input level, a perturbation of amplitude 0.2 a.u. along the Choice mode was delivered to the RNN during lick-left trials (middle panel). We then measured the time it took the network activity to drift back towards the lick left trajectory, that is the relaxation constant τ, by fitting a single exponential decay function $f(t) = A \exp\left(-\frac{t}{\tau}\right) + B$ to the decaying trajectories. The relaxation time constant decreased as ramping slope increased (right panel), suggesting a deepening of the lick-left attraction basin with increased ramping (which corresponds, equivalently, to later times in the delay).

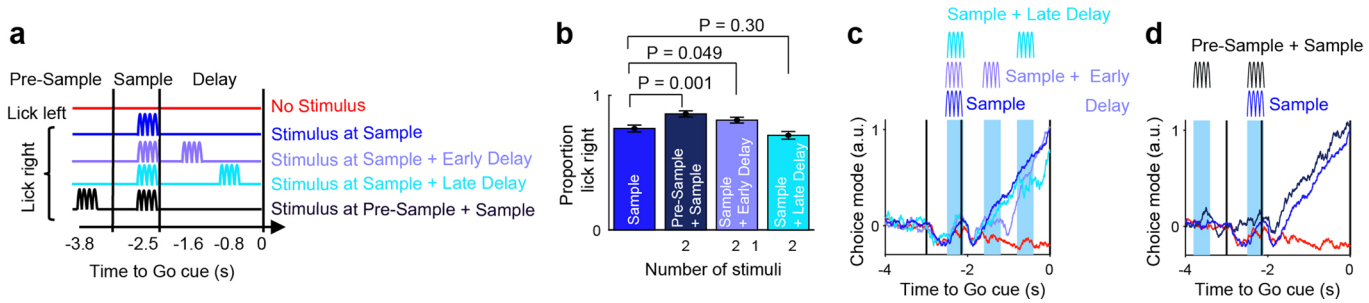

Extended Data Fig. 9 | Effects of multiple stimulations on behavior and ALM dynamics. a, Task schematics with additional stimulation on lick-right trials. **b**, Probability to lick-right after single or multiple vS1 photostimulations. Error bars, Mean ± s.e.m. across sessions ($n = 37$ sessions). *P* values indicate statistical significance by a two-sided paired Student t-test. **c**, ALM neural activity projected on Choice mode on correct lick-right trials with additional stimulation during early delay (light-blue) or late delay (cyan); activity on trials without stimulation (red) or with single stimulation during sample epoch (blue) are shown for comparison. **d**, Same as in **c**, for additional stimulation during pre-sample epoch (dark blue). Data was computed based on activity of neurons recorded in left ALM ($n = 1061$ cells) of distractor-trained mice. To study the effect of additional stimuli on lick-right trials, we first compared a single stimulus during sample-epoch ('Sample') to additional identical stimuli during delay ('Sample + Early Delay', and 'Sample + Late Delay'; **a**). We considered three possible outcomes: 1) Additional stimuli during delay would not affect performance because a lick-right instruction was already provided during sample epoch; 2) Additional stimuli provide additional evidence in favor of lick-right decision, which can be encoded by larger activity along the Choice mode in ALM. In this case, additional stimuli would increase the probability to lick-right regardless of their timing (early or late in the delay); 3) Additional stimuli on lick-right trials could 'rescue' neural dynamics on error trials, which would otherwise follow lick-left neural trajectory. Such rescue by an additional stimulus would consist of a switch from the erroneous lick-left trajectory to the correct lick-right trajectory. Stimulus-induced rescue can be expected to occur more often for early- rather than late-delay stimulation, because later in time the two attractors would be further apart from each other and thus less amenable to switching. Also, because in our attractor model selectivity (distance between attractors) is controlled by the non-specific ramping input, trials with multiple stimulations are not expected to result in higher selectivity at the end of the delay epoch. Our results supported the third possibility because: i) Additional stimulation during early delay ('Sample + Early Delay') but not late delay ('Sample + Late Delay') increased lick-right probability compared to single stimulation during sample ('Sample'; **b**); ii) Choice mode trajectories were perturbed by the second stimulation during early or late delay, but then recovered to same level as the trajectory for single stimulation (**c**). These transient perturbations can be attributed to feedforward inhibition, as we observed that putative fast-spiking interneurons in ALM also responded to stimulation (Extended Data Fig. 2). The eventual recovery of the neural activity to the original lick-right trajectory, further supports the notion that lick-right trajectory was stabilized by attractor dynamics – which made the neural activity robust to transient perturbations. We also tested an experimental condition in which the additional stimulus was given during pre-sample epoch on lick-right trials ('Pre-sample + Sample', **a**). This increased the probability to lick right, and the Choice mode trajectory on these trials started earlier and ramped to a higher level compared to 'Sample' trajectory (**b,d**). The fact that both Pre-Sample stimulation and Early-Delay stimulation increased behavioral performance, but only Pre-Sample stimulation resulted in higher Choice mode activity by the end of the delay, is also consistent with a moving attractor model in which the fixed point corresponding to the lick-right attractor is moved in state-space by the ramping input. In this scenario, an earlier onset of the ramping input (for example during pre-sample) would result in larger separation between attractors by the end of the delay – which would be manifested by larger selectivity along the Choice mode (Supplementary Fig. 6). Taken together, these results showed that ALM dynamic: 1) was robust on lick-right trials, 2) was influenced by the time elapsed from the initiation of preparatory activity, and 3) did not encode the total number of stimuli. More broadly, these results suggest that similarly to frontal cortical regions in rats (Hanks, T. D. *et al. 2015, Nature* **520**, 220–223), ALM activity in mice in this task encodes the categorical choice rather than graded strength of evidence.

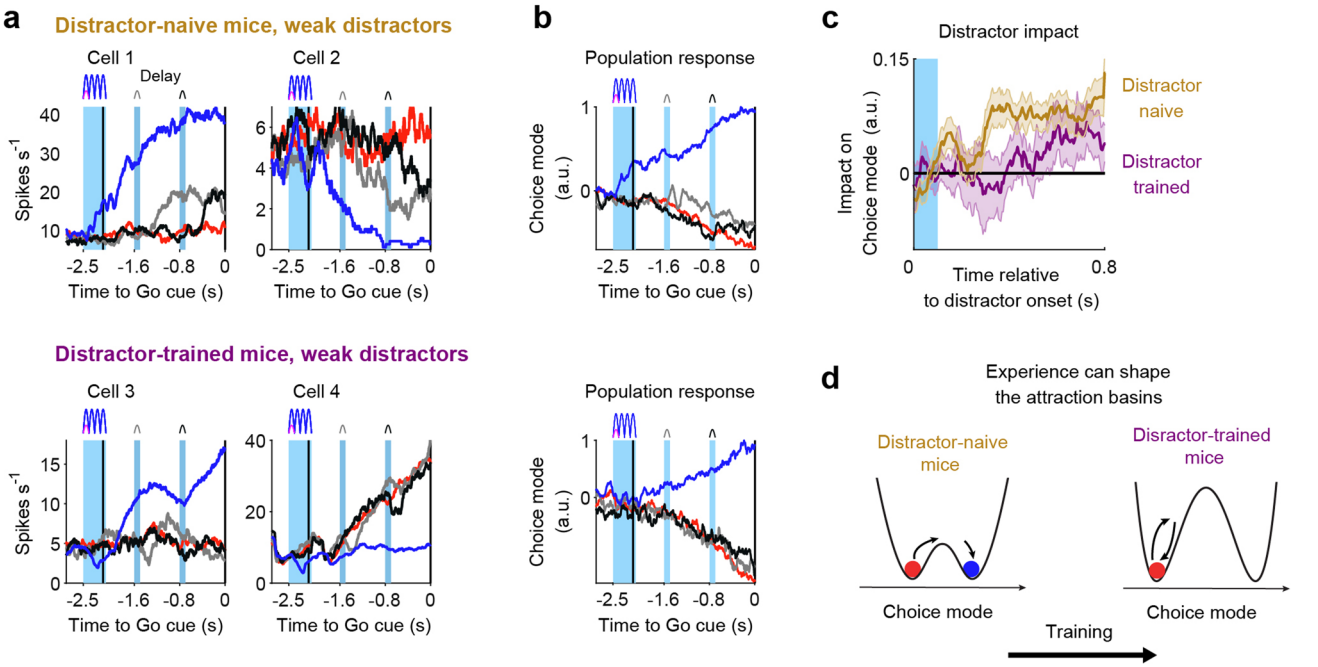

Extended Data Fig. 10 | Distractor-impact on persistent activity in ALM is learning-dependent. Spike rate modulations in left ALM in distractor-naive versus distractor-trained mice, in response to weak-distractors computed using correct trials. **a**, Spike rates of example cells and **b**, Activity along the Choice mode in the presence of weak distractors in distractor-naive (top, $n = 1356$ cells) and distractor trained mice (bottom, $n = 1061$ cells). Lick-left trajectory without stimulation (red), and with weak distractors during early-delay (gray), or late-delay (black). Lick-right trajectory during sample-epoch stimulation (blue). Note that weak distractors had a persistent effect in ALM of distractor-naive mice. Specifically, distractors shifted the activity from lick-left trajectory towards lick-right trajectory. This was evident from shifts in projections on Choice mode (**b**) and from activity of individual cells that contributed to this mode (**a**). In contrast, in distractor-trained mice, the effect of distractors was transient. **c**, Impact of weak distractors on Choice mode in distractor-naive (gold) and distractor-trained mice (purple). Trajectories were aligned to the onset of each distractor. Data is shown as average across distractors with different onset times and sessions ± s.e.m. (shaded). In distractor-naive mice the effect of distractor persisted for at least 0.8 s (the time interval from the late-delay distractor onset and the Go cue), whereas in distractor-trained mice the activity recovered to the unperturbed trajectory. **d**, Schematics of putative attraction basins in distractor-naive (left) and distractor-trained (right) mice (cartoon). Shallow basin of attraction in distractor-mice mice allowed sufficiently strong stimuli to switch the neural activity from one basin to another.

nature research

Corresponding author(s): Karel Svoboda, Sandro Romani

Last updated by author(s): Mar 11, 2021

Reporting Summary

Nature Research wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Research policies, see Authors & Referees and the Editorial Policy Checklist.

Statistics

For all statistical analyses, confirm that the following items are present in the figure legend, table legend, main text, or Methods section.

n/a	Confirmed	
☐	☒	The exact sample size (*n*) for each experimental group/condition, given as a discrete number and unit of measurement
☐	☒	A statement on whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
☐	☒	The statistical test(s) used AND whether they are one- or two-sided *Only common tests should be described solely by name; describe more complex techniques in the Methods section.*
☒	☐	A description of all covariates tested
☐	☒	A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons
☐	☒	A full description of the statistical parameters including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)
☐	☒	For null hypothesis testing, the test statistic (e.g. *F*, *t*, *r*) with confidence intervals, effect sizes, degrees of freedom and *P* value noted *Give P values as exact values whenever suitable.*
☒	☐	For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings
☒	☐	For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes
☐	☒	Estimates of effect sizes (e.g. Cohen's *d*, Pearson's *r*), indicating how they were calculated

Our web collection on statistics for biologists contains articles on many of the points above.

Software and code

Policy information about availability of computer code

Data collection	Data collection was done using: SpikeGL (v 20170315) https://github.com/cculianu/SpikeGL Wavesurfer (v 0.787) https://www.janelia.org/open-science/wavesurfer JRClust (2016)
Data analysis	Data analysis was done using: DataJoint (v 3.2.0) https://datajoint.io Custom MATLAB code. The code will be made available for download.

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors/reviewers. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Research guidelines for submitting code & software for further information.

Data

Policy information about availability of data

All manuscripts must include a data availability statement. This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A list of figures that have associated raw data
- A description of any restrictions on data availability

Data in the NWB format is available for download here: https://dandiarchive.org/dandiset/000060/draft
MATLAB code for data analysis is available here: https://github.com/arsenyf/FinkelsteinFontolan_2021NN
MATLAB code for network models is available here: https://github.com/fontolanl/RNN_ALM_gating

Field-specific reporting

Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

☒ Life sciences ☐ Behavioural & social sciences ☐ Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see nature.com/documents/nr-reporting-summary-flat.pdf

Life sciences study design

All studies must disclose on these points even when the disclosure is negative.

Sample size	The sample sizes are similar to sample sizes used in the field (Inagaki et al., Nature 2019). No statistical methods were used to determine sample size.
Data exclusions	We did not exclude any animal from data analysis.
Replication	We performed recordings for multiple behavior sessions (n = 97 sessions) and from multiple animals (n= 9 mice) to confirm reproducibility. All presented results were reproducible across experimental replicates.
Randomization	All animals were used for experiments. Trial types were randomly determined by a computer program.
Blinding	Analysis of neural and behavior data was conducted regardless of the identity of the animal from which the data was collected. During experiments, trial types were randomly determined by a computer program. During spike sorting, experimenters cannot tell the trial type, so experimenters were blind to conditions.

Reporting for specific materials, systems and methods

We require information from authors about some types of materials, experimental systems and methods used in many studies. Here, indicate whether each material, system or method listed is relevant to your study. If you are not sure if a list item applies to your research, read the appropriate section before selecting a response.

Materials & experimental systems

n/a	Involved in the study
☒	☐ Antibodies
☒	☐ Eukaryotic cell lines
☒	☐ Palaeontology
☐	☒ Animals and other organisms
☒	☐ Human research participants
☒	☐ Clinical data

Methods

n/a	Involved in the study
☒	☐ ChIP-seq
☒	☐ Flow cytometry
☒	☐ MRI-based neuroimaging

Animals and other organisms

Policy information about studies involving animals; ARRIVE guidelines recommended for reporting animal research

Laboratory animals	We used 9 adult male mice (age P>60). We used Scnn1a-TG3-Cre × Ai32 transgenic mouse line (See section 'Animals' in the Methods for details). Mice were maintained on a reverse 12:12-h light:dark cycle. The holding room temperature was maintained at 21±1°C with a relative humidity of 30% to 70%.
Wild animals	This study did not involve wild animals
Field-collected samples	This study did not involve field-collected samples.
Ethics oversight	All experimental procedures were approved by Janelia Institutional Animal Care and Use Committee.

Note that full information on the approval of the study protocol must also be provided in the manuscript.