The frontal cortex adjusts striatal integrator dynamics for flexible motor timing

Zidan Yang[1], Miho Inagaki[1], Charles Gerfen[2], Lorenzo Fontolan[3], Hidehiko K. Inagaki[1*]

1 Max Planck Florida Institute for Neuroscience, Jupiter, FL 33458, USA.
2 National Institute of Mental Health, Bethesda, MD 20814, USA.
3 Turing Centre for Living Systems, Aix- Marseille University, INSERM, INMED U1249, Marseille, France

*Correspondence: Hidehiko.inagaki@mpfi.org

Abstract

Flexible control of motor timing is crucial for behavior. Before movement begins, the frontal cortex and striatum exhibit ramping spiking activity, with variable ramp slopes anticipating movement onsets. This activity may function as an adjustable 'timer,' triggering actions at the desired timing. However, because the frontal cortex and striatum share similar ramping dynamics and are both necessary for timing behaviors, distinguishing their individual roles in this timer function remains challenging. To address this, we conducted perturbation experiments combined with multi-regional electrophysiology in mice performing a lick-timing task. Following transient silencing of the frontal cortex, cortical and striatal activity swiftly returned to pre-silencing levels and resumed ramping, leading to a shift in lick timing close to the silencing duration. Conversely, briefly inhibiting the striatum caused a gradual decrease in ramping activity in both regions, with ramping resuming from post-inhibition levels, shifting lick timing beyond the inhibition duration. Thus, inhibiting the frontal cortex and striatum effectively paused and rewound the timer, respectively. Additionally, the frontal cortex, but not the striatum, encodes trial-history information guiding lick timing. These findings suggest specialized functional allocations within the forebrain: the striatum temporally integrates input from the frontal cortex to generate ramping activity that regulates motor timing.

Introduction

Whenever we decide on an action, the brain needs to determine the timing of that action and execute it accordingly. Flexible control of movement onset, or motor timing, is crucial for a wide range of behaviors, including vocal communication, driving, playing sports, and performing music. In natural settings, many animals adjust the timing of their vigilance and foraging behaviors in response to anticipated predation threats and food availability[1–3]. In laboratory settings, most vertebrate species can adjust their motor timing to obtain reward or avoid punishment[4–11], implying the ancient origins and significance of flexible motor timing. Without these abilities, our behaviors would be restricted to immediate reactions to stimuli.

To execute timed actions, the brain tracks the passage of time over seconds and then triggers actions, akin to how a timer beeps at the end of a preset duration. To perform actions at flexible timings, this 'timer' in the brain should be adjustable to the desired durations[4–8,10,11]. Neurons in the frontal cortex and striatum exhibit neural correlates of such a flexible timer. Before voluntary movements begin, many neurons in these areas demonstrate a gradual change in spiking activity, often characterized by ramping activity that peaks at the onset of movement[6,11–25]. When animals perform actions at various timings, the slope of this ramping activity varies accordingly, leading to the attainment of a hypothetical threshold level that triggers action at different timings[5,11,22]. Hence, the slow spiking dynamics in the forebrain may function as a timer, and the alteration in the speed of these dynamics could be analogous to adjusting the timer (Fig. 1a).

Because isolated neurons can sustain activity only for tens of milliseconds, the seconds-long slow dynamics underlying motor timing likely arise from interactions among populations of neurons[12,26]. From a dynamical systems perspective, the population activity traces a trajectory within a high-dimensional state space, with each dimension corresponding to the activity of individual neurons[27]. Network interactions, for example via feedback connections, govern the evolution of trajectories and can stabilize certain states, known as attractors[12,26,28]. Within this framework, the slow dynamics can be generated by groups of quasi-stable states, forming a continuous attractor or point attractors with shallow basins. Such slow dynamics allow for the temporal integration of inputs to the network (when the input has a component aligned with the attractor), enabling the network to operate as an 'integrator'[28–34]. Integrator networks can generate ramping activity by temporally integrating a step input[30,31](or other types of inputs[28]), with the amplitude of the input adjusting the speed of the ramp (Fig. 1a). Therefore, the integrator has been proposed as a mechanism to generate ramping activity for flexible motor timing[20,22,30].

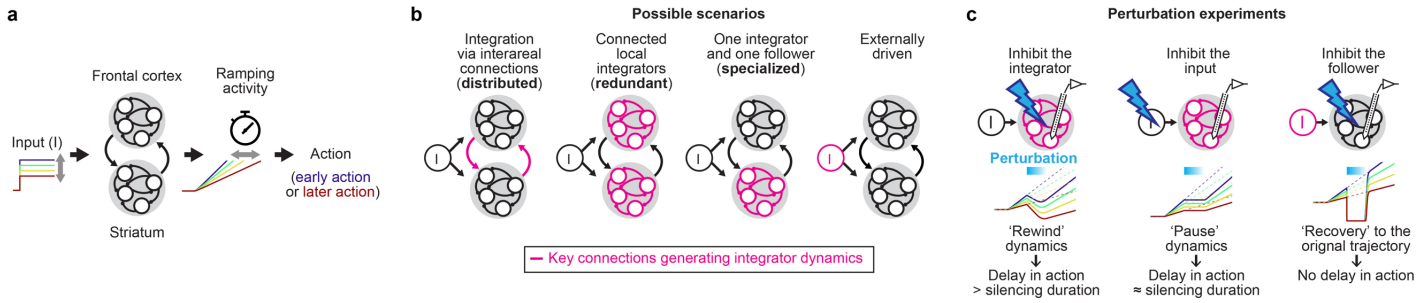

Figure 1. Multi-regional models of flexible motor timing

a. Schema of the computation performed by the forebrain for motor timing. The network may integrate step input to generate ramping dynamics with varying slopes, resulting in different lick timings. For simplicity, other areas connecting the striatum and frontal cortex are omitted.

b. Possible configuration of the cortico-striatal network to implement integrator dynamics. The key connections to generate integrator dynamics are shown in pink.

c. Schema of perturbation experiments and expected results.

While neural correlates of an integrator have been observed in the frontal cortex and striatum[15,18,22,35–38]— two forebrain regions forming a long-range loop — the precise roles of these areas in implementing the integrator dynamics remain unclear. Manipulations of the frontal cortex and striatum affect timing behavior, supporting their causal roles[11,39–54]. However, the presence of neural correlates and necessity cannot determine their computational roles[8,55,56]. First, neural correlates may be generated locally (internally generated) or inherited from other brain regions (externally driven). Second, a brain area that is 'causal' for timing behavior might 1) house the attractor landscape that generates integrator dynamics, 2) supply essential inputs to another brain area that implements integrator dynamics, or 3) affect behavior through mechanisms independent of integrator dynamics, for example by controlling movement execution. This presents a general challenge in pinpointing the precise loci of computations within a multi-regional recurrent network. In the context of motor timing, the integrator generating the neural representation of time can either be: distributed across the frontal cortex and striatum; redundantly present in both areas; present in one area (specialized); or located in upstream areas (dynamics in the forebrain externally-driven) (Fig. 1b).

We addressed these knowledge gaps by performing a series of transient perturbations with simultaneous multi-regional electrophysiology and testing the results against computational models of ALM-striatum interactions. Models of multi-regional networks predict how the dynamics respond during and after a transient perturbation depending on their functional allocations[12,26,28,57] (Fig. 1c; see Extended Data Fig. 1 for simulation results). For instance, when inhibiting a brain area that acts as an integrator, transient perturbations can be integrated into the ramping dynamics (provided that the inhibition has a component along the attractor landscape). This leads to the rewinding of the representation of time, causing a delay in motor timing beyond the silencing duration (Fig. 1c, left). In contrast, silencing a brain area that provides input to the integrator (in another area) will temporarily pause the representation of time in the recipient area, thereby delaying the action by the duration of the silencing (Fig. 1c, middle; in a leaky integrator network, activity decays during the silencing of the input. See Extended Data Fig. 1e for the distinction of this scenario). Conversely, silencing a brain area with externally driven dynamics will result in a rapid return to the original trajectory after the silencing[58,59], with no impact on motor timing (Fig. 1c, right). To systematically dissect multi-regional dynamics following this model-driven approach, we developed a flexible lick-timing task in mice, where mice explore various lick times over 600 trials per session (652 ± 9 trials; mean ± SEM; 224 sessions, 48 mice; Fig. 2a). This enabled a series of perturbations within individual sessions. Large-scale electrophysiology in the frontal cortex and striatum allowed decoding of planned lick time in individual trials, providing an ideal testbed to quantify how the dynamics respond to perturbations. Leveraging this system, we identified specialized functional allocation within the forebrain that implements the integrator dynamics, generating ramping activity that serves as an adjustable timer.

Results

The frontal cortex is required for lick timing control

We developed a flexible lick-timing task for mice. In this task, an auditory cue (3 kHz, 0.6 s) signals trial onset followed by a delay epoch with unsignaled duration (1 - 3 s; Fig. 2a and Extended Data Fig. 2ab; Methods). If mice lick after the delay, they receive a water reward. In contrast, a premature lick during the delay terminates the trial without reward. We varied the delay duration in blocks of trials within a session (with the number of trials within each block randomly selected from 30 - 70 trials; Fig. 2b and Extended Data Fig. 2ab). Despite the absence of a sensory cue instructing the delay duration or block transitions, mice dynamically adjusted their lick time distribution according to the delay within 10 trials after the delay switch (Fig. 2bc).

The only information available for mice to infer the appropriate lick time in this task is their prior lick times and the outcomes. To investigate whether such 'trial history' guides trial-by-trial lick timing, we exploited a linear regression model to predict future lick times based on trial history[39,60] (n = 30 mice; Methods). This analysis revealed a positive correlation between upcoming lick times and previous lick times, alongside a negative correlation between upcoming lick times and previous reward outcomes: mice tended to lick earlier after a reward and later after a lack of reward (Fig. 2d and Extended Data Fig. 2l). Since an absence of reward indicates a premature lick in this task, delaying licking after an unrewarded attempt is an adaptive strategy. In contrast, when we keep the delay duration constant across trials and sessions ('constant delay condition'; n= 13 mice; Extended Data Fig. 2c), there was no significant influence of former trials on lick timing (Fig. 2d). Thus, mice use trial history to strategically adjust lick timings only when the delay duration is variable. We employed this behavior as a model system to examine how the brain strategically and flexibly adjusts the timing of actions (see Extended Data Figs. 2 and 3 for further quantifications, models explaining lick time distribution, and body part tracking).

To identify dorsal cortical areas that control lick timing in an unbiased manner, we performed optogenetic loss-of-function screening using transgenic mice expressing ChR2 in GABAergic neurons (Vgat-ChR2-EYFP mice[61]) with a clear skull preparation[62] (Methods). We bilaterally silenced individual dorsal cortical areas during the delay epoch (delay duration: 1.5 s) by scanning a blue laser in randomly interleaved trials (488 nm, 1.5 mW). Photostimulation started 0.6 s after the cue and lasted for 1.2 s including a 0.2 s ramp down of laser power at the end (3 mice; Fig. 2e; Methods). Lick initiation was blocked during the silencing of a frontal cortical area: anterior-lateral motor cortex (ALM; anterior 2.5 mm, lateral 1.5 mm from Bregma). This is consistent with the established role of ALM as a premotor cortex for licking[12,63,64]. Notably, following the transient silencing of ALM, the distribution of lick times shifted significantly later, with the median lick time delayed by 0.79 (0.59 - 0.97) s (mean and 95% confidence interval). This suggests a causal role of ALM in controlling the lick time.

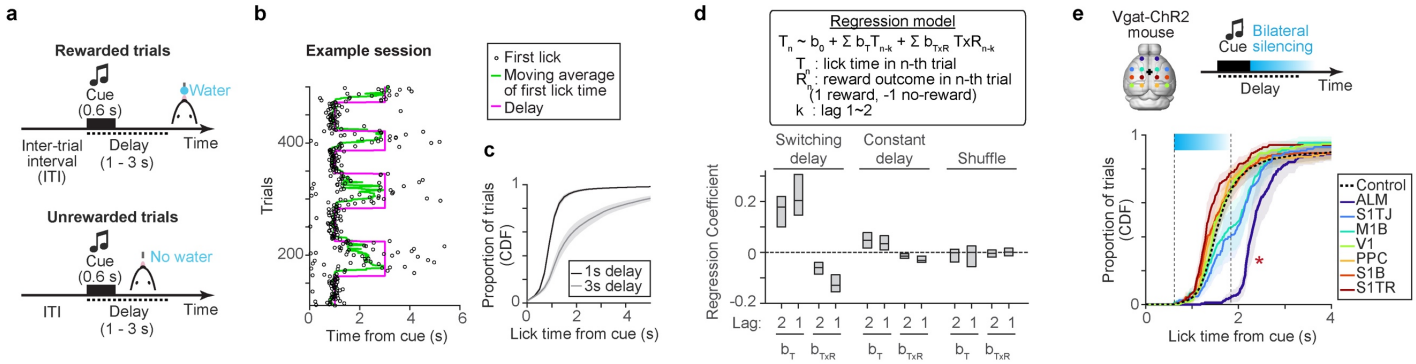

Figure 2. ALM is required for lick timing control

a. Flexible lick timing task. The delay epoch started at the trial onset signaled by the cue. The first lick after the delay was rewarded, and a premature lick during the delay aborted the trial.

b. Example session. Only part of an entire session is shown.

c. Cumulative distribution of lick time in 1 s and 3 s delay blocks (56 sessions, 10 mice). Shades, SEM (hierarchical bootstrap).

d. Regression coefficients of the regression model based on previous lick time (T) and interaction (T × R), with 2 lags. Switching delay (n = 30 mice). Constant delay (n = 13 mice).

e. Top, optogenetic loss of function screening of dorsal cortical areas during delay. Bottom, cumulative distribution of lick time in trials with different silenced areas (3879 control trials, 172 ± 5 silencing trials per brain region; mean ± SEM; 9 sessions, 3 mice). Cyan bar, silencing window (1.2 s). Shades, SEM (hierarchical bootstrap). *, $p < 0.001$ (hierarchical bootstrap with a null hypothesis that the lick times in control trials are later than or equal to the ones in silencing trials). Regions adjacent to ALM (M1B and S1TJ) exhibited weaker effects, attributed to the limited spatial resolution of the manipulation[62].

Similar scalable timing dynamics in ALM and striatum

To investigate neural activity associated with lick time control in ALM, we conducted high-density silicon probe recordings (4467 putative pyramidal neurons recorded in 172 sessions, 45 mice; Extended Data Table. 1). Many ALM neurons displayed ramping activity during the task, peaking around the onset of lick (Fig. 3a; Out of 4467 neurons, 1363 showed a significant increase and 1809 showed a significant decrease in spike rate before the lick compared to the baseline; Methods). The ramping speed of these neurons varied across trials and predicted lick timing (Fig. 3a). Notably, temporal warping[11,23,48,49] which normalizes the temporal axis between cue and lick in each trial, significantly reduced across-trial variability in spike rate in 68.1% of neurons (Fig. 3a bottom and Extended Data Fig. 4ab). This indicates that two-thirds of ALM neurons exhibited temporal scaling (stretching or shrinking) of activity patterns across trials, with the speed of their dynamics anticipating lick time.

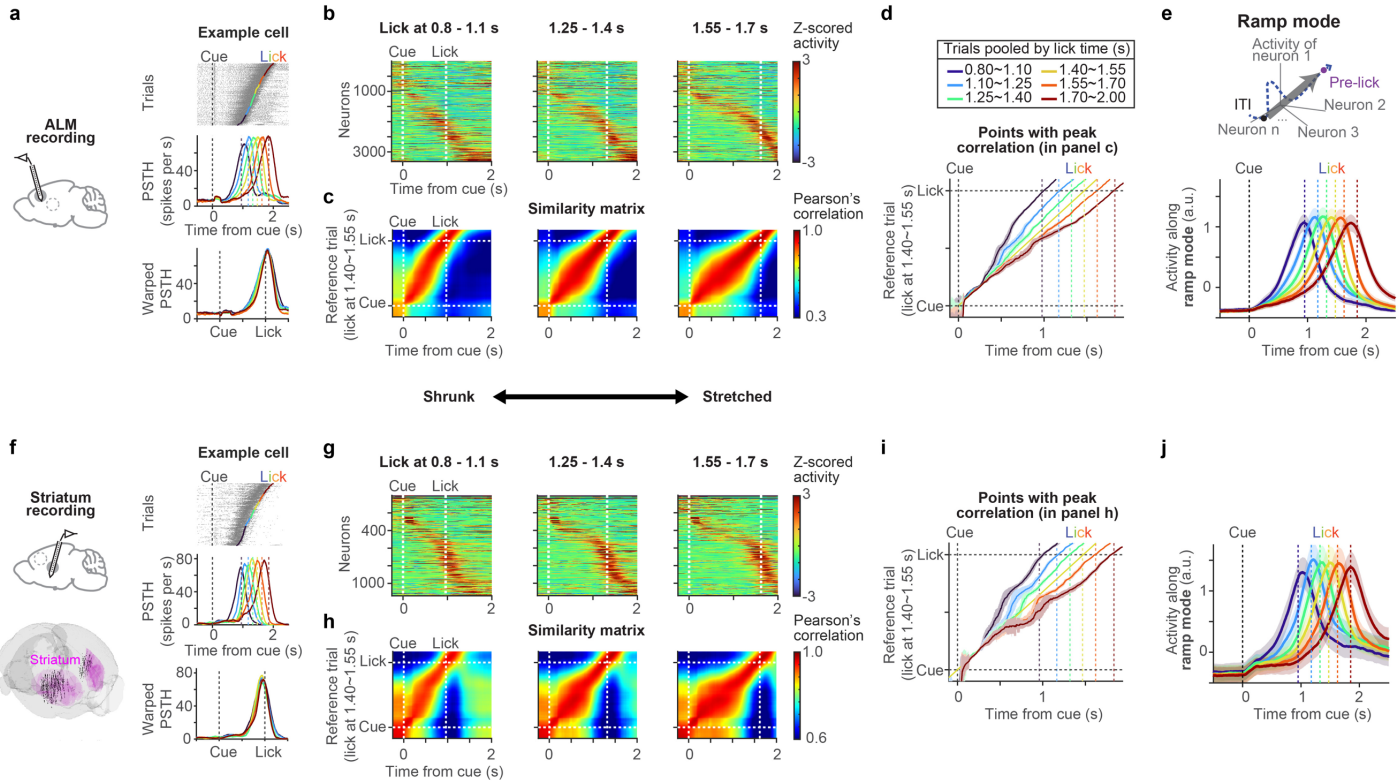

Figure 3. Similar scalable timing dynamics in ALM and striatum

a. An example ALM cell. Top, spike raster. Trials are sorted by lick time. Colored dots, first lick time, divided into six groups with equal trial numbers. Middle, trial averaged peri-stimulus time histogram (PSTH) of the six lick time groups. Vertical dotted lines in the same color indicate the corresponding lick time. Bottom, PSTH following temporal warping (Methods).

b. ALM activity (Z-scored) across trials with different lick time ranges. Neurons are sorted in the same order for all panels (by peak firing time in trials with lick between 1.40 ~1.55 s). Only neurons with more than 10 trials across all 6 lick time ranges were analyzed (n = 3261 neurons, 45 mice).

c. Similarity matrix (Pearson's correlation) comparing ALM population activity between trials with different lick time ranges (reference trial: lick between 1.40 and 1.55 s).

d. Time points with the peak correlation in the similarity matrix (**c**), across 6 different lick time ranges. Shades, SEM (hierarchical bootstrap).

e. Top, schema of ramp mode. Bottom, ALM population activity projected along the ramp mode, with trials grouped by lick time (same color scheme as **d**). n = 3261 neurons, 45 mice. Shades, SEM (hierarchical bootstrap).

f-j. The same as in **a-e** but for striatal recording (n = 1073 neurons, 16 mice). Bottom in **f**, the location of all recorded units registered to Allen common coordinate framework.

Consistent with individual cells, the population activity patterns in ALM appeared to temporally scale with lick time (Fig. 3b). To quantify this observation, we calculated similarity matrices using Pearson's correlations, comparing population activity between trials with differing lick times (Fig. 3c). We then identified the points with peak correlations in these similarity matrices (Fig. 3d). This analysis revealed that although similar population activity patterns emerged across trials, the speed at which these patterns unfolded varied as a function of lick time (Fig. 3d).

To further characterize the population activity patterns between the cue and the lick, we defined three modes (directions in population activity space) to differentiate the spike rate during specific time windows after the cue from the baseline activity during the inter-trial interval (ITI; 0 - 1 s before the cue): 0 - 300 ms after the cue (cue mode, CM), 500 - 800 ms before the lick (middle mode, MM), and 200 - 500 ms before the lick (ramp mode, RM). These modes were further orthogonalized to each other (Methods). Population activity along these three modes collectively spanned the time from cue to lick and explained most of the task-modulated activity (74 %; Extended Data Fig. 5e). Notably, population activity along the MM (MM activity) and RM (RM activity) displayed temporal scaling (Fig. 3e and Extended Data Fig. 5b-d). Thus, a large proportion of ALM activity between the cue and the lick, which we refer to as 'timing dynamics', demonstrated temporal scaling, similar to observations in other timing tasks and species[9,11,22,24,49,65].

Many excitatory neurons in ALM project to the striatum, spanning various sectors, including the ventrolateral striatum[66–68] (VLS; Extended Data Fig. 6a). Considering the involvement of the striatum in various timing behaviors[11,23,24,41–44,46], it likely cooperates with ALM to control lick timing. We performed electrophysiological recordings in the striatum using Neuropixels probes (Fig. 3f and Extended Data Fig. 6b and d). In total, we recorded 1972 neurons (97 sessions, 16 mice), with the majority classified as putative striatal projection neurons (SPN; 64%) or fast-spiking interneurons (FSI; 30%) based on spike features[69] (Extended Data Fig. 6e). Because both SPN and FSI showed similar activity patterns (consistent with previous study[69]) and comprised ~95% of the population together, we have pooled all cell types for analysis. Overall, striatal activity between the cue and lick was analogous to those in ALM[11,69,70]. Firstly, the activity of individual striatal neurons was well-explained by temporal scaling (in 56% of cells; Fig. 3fg and Extended Data Fig. 4c), and many neurons displayed ramping activity (out of 1972 neurons, 741 showed a significant increase and 400 showed a significant decrease in spike rate before the lick compared to the baseline). Secondly, striatal population activity (similarity matrix, MM activity, and RM activity) showed a temporal profile similar to that in ALM with temporal scaling (Fig. 3h-j and Extended Data Fig. 5f-j).

In some sessions, we simultaneously recorded from ALM and the striatum (16 sessions, 8 mice). At each time point, we used the population activity from either ALM or the striatum to predict how much time was left until lick using a k-nearest neighbor (kNN) decoder ('time to lick ($T_{\text{to lick}}$)'; Extended Data Fig. 6f-h; Methods). The decoded $T_{\text{to lick}}$ in these two brain areas was significantly correlated across individual trials (Extended Data Fig. 6i-k). Thus, ALM and the striatum demonstrate similar scalable timing dynamics coupled at the single-trial level.

ALM, but not the striatum, encodes trial history

What determines the speed of timing dynamics and guides lick time? Given that trial history influences lick time in this task (Fig. 2d), we hypothesized that certain neurons encode trial history. Such neural activity could establish the initial conditions of a network[10,60,71] and/or provide inputs to guide the speed of timing dynamics and action timing.

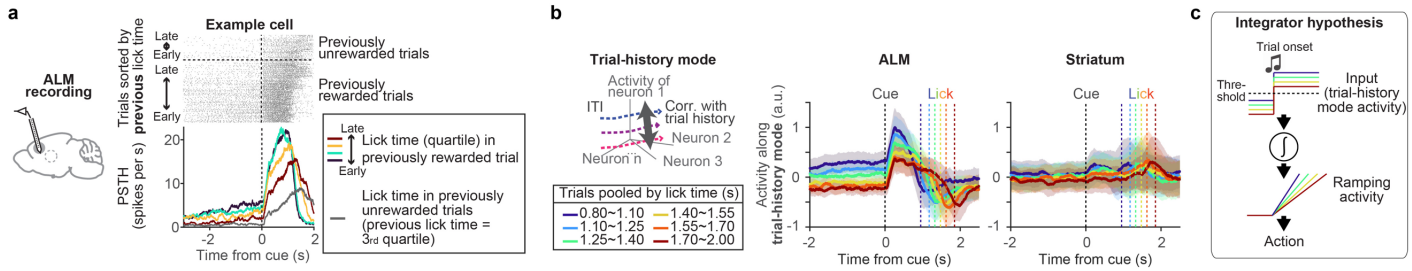

Figure 4. Sensory perturbation is integrated into ramping activity

a. An ALM example cell whose activity is modulated by the lick time and reward outcome in the previous trial. Top, spike raster, grouped by reward outcome in the previous trial and sorted by the lick time in the previous trial. In this example cell, the ITI activity is higher in trials after rewarded trials and with earlier licks. Bottom, mean PSTH. Lick times of the previous rewarded trials were divided into quartiles indicated by different colors. The gray trace, trials following previously unrewarded trials with previous trial's lick times within the 3rd quartile.

b. Left, schema of trial-history mode. Middle, ALM population activity projected along trial-history mode. Right, striatal population activity projected along trial-history mode. Shades, SEM (hierarchical bootstrap).

c. Schema of the integrator hypothesis.

In line with our hypothesis, some ALM neurons exhibited spiking activity correlated with the lick time and reward outcome of the previous trial (Fig. 4a and Extend Data Fig. 4ef). During the ITI, 11.0 % of ALM cells significantly encoded the previous lick time (proportion of neurons having $p < 0.05$ with Kruskal-Wallis tests, testing whether ITI activity is indistinguishable across trials with different previous lick times). ALM neurons encoding previous lick time also tended to encode previous reward outcome (Extended Data Fig. 4g). Because upcoming lick time and trial history are correlated (Fig. 2d), we calculated the partial correlation between the activity of ALM neurons and previous lick time while removing the effect of upcoming lick time (Methods). This partial correlation was significantly higher than trial shuffle and session permutation controls[72] both before and after the cue. This implies that ALM neurons encode trial history throughout the trial (Extended Data Fig. 4hi; Methods).

If trial-history information encoded in ALM influences subsequent lick times, such neural correlates may be absent in contexts where trial history is not used. Supporting this idea, ALM activity was not correlated with trial history above chance level under the constant delay condition, where mice did not rely on trial history to guide their lick times (Extended Data Fig. 4k and n).

To examine the evolution of ALM activity encoding trial history at the population level, we defined a "trial-history mode" defined by the rank correlation between the ITI activity and trial history (the predicted lick time based on previous trials using trial-history regression; Fig. 4b, left; Methods). ALM activity along the trial-history mode was modulated after reward outcome (Extended Data Fig. 4l3), exhibited a graded persistent activity during the ITI, and showed a step-like increase upon the cue presentation at the trial onset (Fig. 4b, middle). The amplitude of this activity anticipated upcoming lick times throughout the trial (Fig. 4b, middle). Thus, activity along this mode carries trial-history information and predicts upcoming lick time across trials. This activity is amplified following trial onset, potentially allowing it to influence other activity patterns. Notably, temporal integration of a step input (with a threshold to prevent integration before the step begins; Fig. 4c, dashed line) produces a ramp that initiates at the step onset. The speed of the ramp is determined by the amplitude of the step. Thus, the activity profile of the trial-history mode closely matches with hypothetical input to an integrator that generates scalable RM activity (Fig. 4c).

In contrast, the spiking activity of striatal neurons was not correlated with trial history above chance level (Extended Data Fig. 4j and m), and no clear activity patterns were observed along the trial-history mode (Fig. 4b, right). Thus, while timing dynamics are similar in ALM and the striatum, ALM, but not the striatum, encodes trial history information throughout trials. This ALM activity may function as an input to an integrator, generating RM activity through temporal integration and guiding proper lick timing.

ALM silencing pauses the timer

Neural correlates of temporal integration (i.e., trial-history mode and RM) in ALM alone are insufficient to conclude that ALM is the integrator (Extended Data Fig. 1). Likewise, while the shift in lick time following ALM silencing (Fig. 2e) suggests a causal role for ALM in controlling lick timing, behavioral effect alone is insufficient to conclude its function (Extended Data Fig. 1; even in models where ALM is not an integrator, ALM silencing can shift lick time). Therefore, to examine whether ALM functions as an integrator, we recorded ALM activity using silicon probes during calibrated silencing.

Strong cortical silencing can induce post-silencing rebound activity that triggers actions[73,74] complicating the interpretation of subsequent dynamics and behavior. To address this, we calibrated the silencing protocol to minimize rebound activity. Silicon probe recordings confirmed nearly complete ALM silencing during photostimulation using 1.5 mW 488 nm laser in Vgat-ChR2-EYFP mice (8 spots, bilateral). Limiting the silencing duration to 0.6 s (including a 0.3 s ramp down) minimized post-silencing rebound activity and licking (Extended Data Fig. 7a-d). However, rebound activity and licking increased over sessions due to an unknown adaptation mechanism[75,76] (Extended Data Fig. 7j-k). Therefore, we restricted our analysis to the initial two days of ALM silencing (similarly, restricted to the first day for striatal manipulation in Fig. 6; Extended Data Fig. 7).

Silencing ALM during the delay epoch (0.6 s after cue onset) with this protocol shifted the median lick time by 0.47 (0.37 - 0.56) s (mean and 95% confidence interval, n = 14 mice; Fig. 5ab and Extended Data Fig. 8a-d). Interestingly, this temporal shift was close to the silencing duration (0.6 s), as if the 'timer' was paused during silencing. In contrast, silencing ALM before the cue did not shift the lick time distribution, suggesting that this manipulation has no long-lasting effect on lick time (ALM activity rapidly recovered after ITI silencing, implying that trial-history information is robustly maintained across brain areas before the trial onset[77–79]; Extended Data Fig. 9q).

Silicon probe recordings of ALM during delay silencing (590 neurons, 14 mice) revealed that, despite near-complete silencing during photostimulation, population activity patterns resembling those just before the silencing suddenly reemerged after the silencing ceased, and ALM dynamics continued to unfold from there in parallel to the unperturbed condition (quantified by similarity matrix; Fig. 5fg). Consistently, at the end of ALM silencing, RM activity rapidly recovered (Fig. 5h and Extended Data Fig. 9a). However, RM activity after silencing was significantly lower than in the unperturbed condition (Extended Data Fig. 9a), indicating that it is not a recovery to the original trajectory as depicted in Fig. 1c right. Instead, RM activity recovered close to the pre-silencing level, after which activity began to slowly ramp up again (Fig. 5h and Extended Data Fig. 9a). To quantify this observation further and to also appreciate the non-RM activity, we used a kNN decoder to analyze the evolution of decoded $T_{to\ lick}$. We compared ALM silencing and unperturbed trials with matched decoded $T_{to\ lick}$ before the silencing (0.6 s after the cue). During the silencing, the decoded $T_{to\ lick}$ diverged between ALM silencing and unperturbed trials, and such shift persisted in parallel after the silencing (Extended Data Fig. 9d). Thus, after the rapid post-silencing recovery to the pre-silencing level, ALM dynamics unfolded in parallel to those observed under unperturbed conditions, explaining the shift in lick time close to the silencing duration.

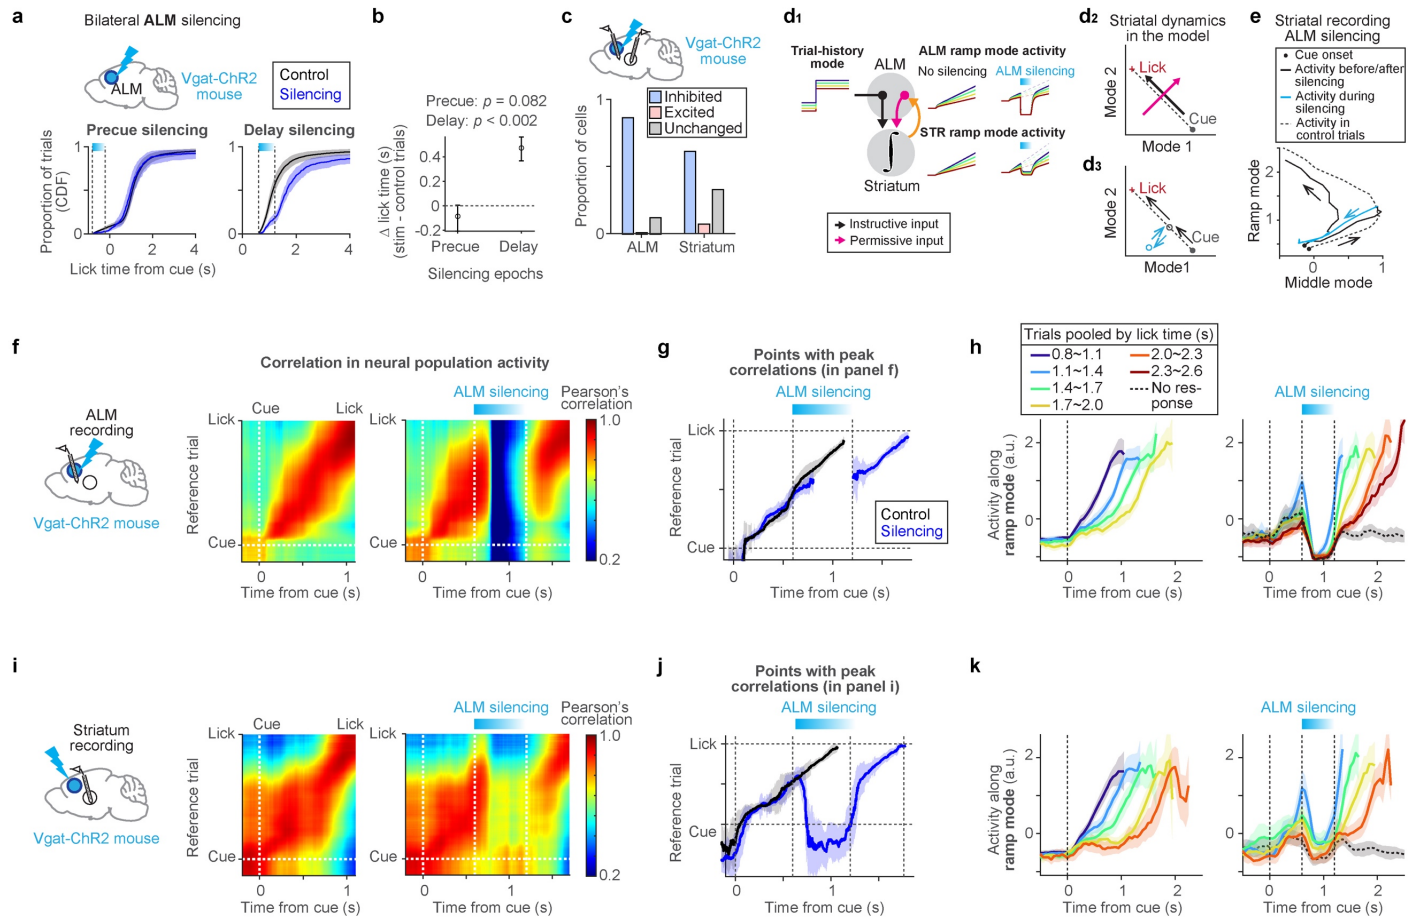

Figure 5. Transient optogenetic perturbation of ALM

a. Top, bilateral ALM silencing in Vgat-ChR2 mouse. Bottom, cumulative distribution of lick time in trials with precue (left) or delay silencing (right). Black, control trials. Blue, silencing trials. Shades, 95 % confidence interval (hierarchical bootstrap). n = 28 sessions, 14 mice.

b. Change in median lick time in **a**. Error bars, 95% confidence interval (hierarchical bootstrap). *P*-value, hierarchical bootstrap with a null hypothesis that there is no change in lick time.

c. The proportion of ALM pyramidal neurons (left) and striatal neurons (right) with spiking activity significantly inhibited or excited ($p < 0.05$; Unchanged, p >= 0.05; rank sum test) by ALM silencing. Neurons with a mean spike rate above 1 Hz in control trials during the silencing window were analyzed. This includes 391 neurons in ALM (n = 28 sessions, 14 mice) and 238 neurons in the striatum (n = 14 sessions, 7 mice).

d. Schema of a model explaining the data (**d1**). Schema of how instructive and permissive inputs from ALM drive activity in striatum (**d2**). Schema of striatal activity during ALM silencing (**d3**). Black arrows, activity before/after silencing. Cyan arrows, activity during silencing. Black circle, activity at the onset and end of silencing. Cyan circle, activity in the middle of silencing.

e. Striatal activity in two-dimensional space along the ramp and middle modes during ALM silencing. The mean trajectory of trials with silencing (lick between 1.4-1.7 s) is shown. Arrows indicate the direction the trajectory evolves.

f. Left, recording of ALM during ALM silencing. Middle and right, similarity matrix of ALM population activity between trial types. Middle, control vs. ref trial (lick between 1.4 - 1.7 s). Right, ALM silencing trials vs. the ref trial.

g. Points with the peak correlation in the similarity matrix in **f**. Shades, SEM (hierarchical bootstrap). Black, control. Blue, silencing.

h. ALM population activity along the ramp mode (n = 28 sessions, 14 mice), trials grouped by lick time. Lick time ranges that exist in at least two-thirds of the analyzed sessions were shown. Therefore, the plotted lick time ranges vary depending on the manipulation conditions. Shades, SEM (hierarchical bootstrap). Left, control. Right, ALM silencing.

i-k. Same as in **f-h** but for striatal recording. n = 14 sessions, 7 mice.

The rank order of RM activity across trials predicted lick time under unperturbed conditions (Fig. 5h, left). Notably, this rank order collapsed during silencing but recovered afterward (Fig. 5h and Extended Data Fig. 9b). Similarly, ALM population activity before and after silencing both significantly predicted upcoming lick time on single trials using the kNN decoder, whereas this predictability was lost during silencing (Extended Data Fig. 9c; Methods). These findings suggest that following silencing, ALM activity rapidly reverted to a state closely resembling its pre-perturbation condition at the individual trial level.

The recovery of ALM activity to pre-silencing levels and the shift in lick time close to the silencing duration challenge multiple network models where ALM acts as the only integrator or is solely driven by external inputs (Extended Data Fig. 1). Instead, our data suggest that ALM silencing momentarily pauses the temporal integration because ALM provides input to an integrator (a 'timer') in another brain area (Fig. 1c, middle). Additionally, the rapid post-perturbation recovery of ALM ramping activity to pre-perturbation levels implies that ALM ramping activity follows the dynamics of the external integrator, which was paused during the ALM silencing. This indicates that ALM acts both as an input to and a follower of the external integrator (i.e., suggesting an implementation that combines Fig. 1c middle and right; Extended Data Fig. 1f).

The striatum maintains timing information during ALM silencing

During ALM silencing, other brain areas must retain the timing information to restore ALM dynamics afterward. Given the prominent timing dynamics observed in the striatum (Fig. 3), the striatum may serve such a role. To test this, we recorded striatal activity during ALM silencing using Neuropixels probes[80] (372 neurons, 7 mice). A significant portion of striatal neurons (60%) decreased their spike rates during ALM silencing, indicating that ALM provides a major excitatory drive to the striatum (Fig. 5c, right, and Extended Data Fig. 10e-h). Consistently, upon ALM silencing, RM activity in the striatum rapidly decayed, and the population activity patterns (similarity matrix) resembled those observed during ITI when the spike rate is low (Fig. 5i-k). However, striatal activity was not entirely abolished during ALM silencing (Extended Data Fig. 10f). The residual striatal activity during ALM silencing maintained the rank order of RM activity and predicted lick time (Fig. 5k and Extended Data Fig. 9fg). Thus, despite the reduced mean spiking activity, the striatum retained timing information during ALM silencing.

The observed multi-regional dynamics can be replicated by a model where the striatum functions as an integrator, and ALM acts as both an input to and a follower of the striatal integrator (Fig. 5d$_1$ and Extended Data Fig. 1f). In this model, to explain the pause of time representation and the large reduction in striatal activity during ALM silencing, ALM provides two types of input to the striatum:

1) Instructive input: This input is temporally integrated by the striatum to generate scalable timing dynamics (Fig. 5d$_{1\text{-}2}$, black arrows). ALM then follows the timing dynamics generated in the striatum, resulting in similar timing dynamics between ALM and the striatum (Fig. 5d$_1$, orange arrow). In the state space, the instructive input aligns with the direction in which timing dynamics evolve (i.e. along the direction of integration, Fig. 5d$_2$). During ALM silencing, the lack of instructive input pauses the evolution of activity along this direction, effectively pausing the representation of time.

2) Excitatory drive ('permissive input'): This input is not temporally integrated by the striatum (Fig. 5d$_{1\text{-}2}$, pink arrows). In the state space, the permissive input is orthogonal to the direction in which timing dynamics evolve, controlling the amplitude of striatal activity without affecting the representation of time (Fig. 5d$_2$). During ALM silencing, the lack of this permissive input results in a large reduction in striatal activity.

Altogether, due to the loss of these two inputs during ALM silencing, the striatum maintains timing information at a reduced activity level (cyan circle in Fig. 5d3). Once ALM silencing ends, the permissive input returns, and the striatal activity recovers to pre-silencing levels. Additionally, the recovery of the instructive input after ALM silencing allows the timing dynamics to evolve from the pre-silencing level along a normal trajectory. Altogether, this results in a parallel shift in timing dynamics after ALM silencing, matching the experimental data (Fig. 5d, and Extended Data Fig. 1f). Striatal activity in a two-dimensional space defined by RM and MM (Fig. 5e), which captures the majority of activity during the task, evolves consistently with the model (Fig. 5d3; in the data, dynamics develop more like an arc instead of line).

The key assumption of this model is that the striatum serves as a part of the integrator; otherwise, an alternative brain area would be necessary. To test this prediction, we conducted perturbation experiments in the striatum.

Striatum inhibition rewinds striatal timing dynamics

The striatum contains two major projection cell types: D1 receptor-expressing direct pathway SPN (D1-SPN) and D2 receptor-expressing indirect pathway SPN (D2-SPN)[81]. Consistent with the anti-kinetic function of D2-SPN[81–83], D2-SPN silencing and non-cell-type-specific striatal silencing using the soma-targeted light-dependent chloride channel, stGtACR1[84], triggered licking during photostimulation (Extended Data Fig. 8m-s). This response made it unsuitable for testing the function of the striatum in lick timing control. Consequently, we focused on inhibiting D1-SPNs to investigate the effects of transient striatal perturbation.

To inhibit D1-SPN, we crossed Drd1-cre FK150[85] line with cre-dependent stGtACR1 transgenic reporter mice[73], and bilaterally implanted tapered fiber optics[86] (1.0 mm taper, NA 0.37) to deliver light to a large volume in the striatum (Fig. 6b and Extended Data Fig. 6c). We confirmed the inhibition using optrode recordings (488 nm, 0.25-0.5 mW): 9 out of 25 SPNs (36%) significantly reduced spike rates without axonal excitation[73,87,88] or post-silencing rebound (we cannot distinguish SPN subtypes based on spike features[68,89], but since half of SPNs are D1 SPN[81,90], we estimated that we inhibited ~70% of D1 SPN around the fiber optics; Fig. 6e, left, Extended Data Fig. 7hi, and Extended Data Fig. 10i-l).

Transient bilateral inhibition of D1 SPN in the ventrolateral striatum (VLS) during the delay epoch (0.6 s duration, starting 0.6 s after the cue onset; 488 nm, 0.25-0.5 mW) significantly increased the no-response rate by 38 (5.7 - 41)% (mean, 95% confidence interval, n = 6 mice; Fig. 6cd). In trials where mice licked after the inhibition, the median lick time was shifted later by 1.0 (0.69 - 1.4) s (mean, 95% confidence interval), significantly longer than the duration of photostimulation, and the effect of ALM silencing (Fig. 5).

The shift in lick time caused by unilateral inhibition was approximately half that of bilateral inhibition, suggesting an additive effect of this manipulation (Fig. 6cd, light blue). Inhibiting VLS before the cue did not affect the subsequent lick time distribution, implying that the striatum is specifically involved in lick timing control after the cue and that this manipulation has no long-lasting effect (Fig. 6cd). The behavioral effect of VLS inhibition was stronger than dorsomedial striatum (DMS) inhibition (Extended Data Fig. 8i-l, n = 6 mice), consistent with the strong anatomical and functional connections between VLS and ALM[66–68].

To measure the impact of D1-SPN inhibition on striatal dynamics, we performed optrode recordings (103 neurons, 5 mice). During unilateral inhibition, population activity patterns comprising all striatal cell types in VLS, appeared

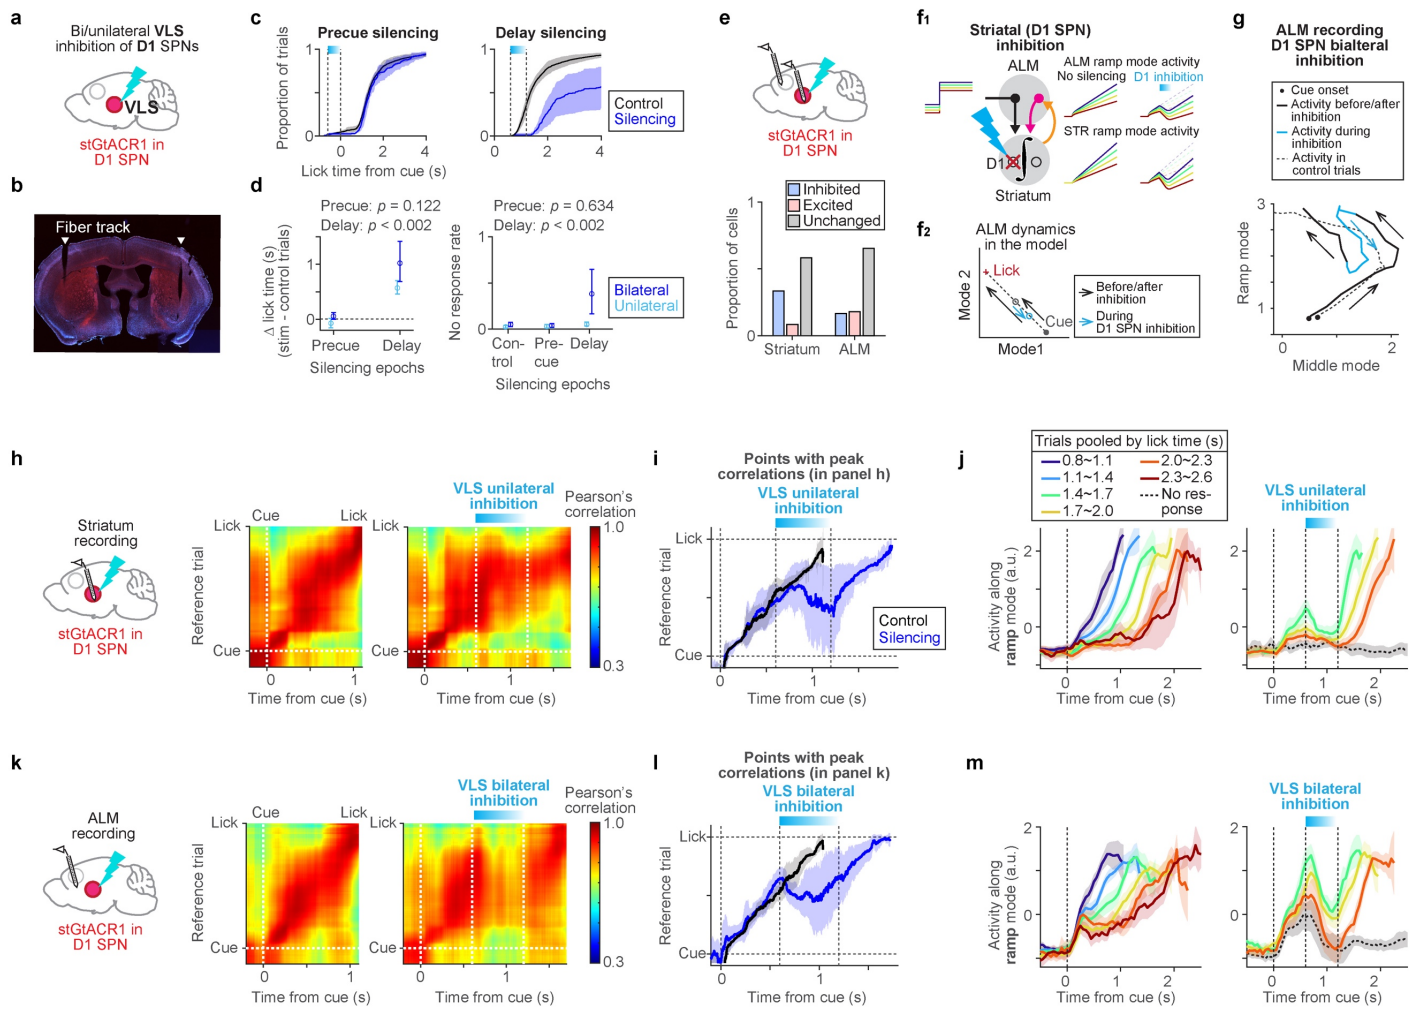

Figure 6. Transient optogenetic perturbation of D1 SPN in the striatum

a. Schema of inhibiting D1 SPN in ventral lateral striatum (VLS).

b. Coronal section of a Drd1-cre;cre-dependent-stGtACR1:fusion-red mouse with bilateral tapered optic fiber implanted in the VLS (white arrows). Red, fusion red. Blue, Nissl staining.

c. Cumulative distribution of lick time in trials with precue (left) or delay silencing (right). Black, control trials. Blue, trials with bilateral inhibition. Shades, 95 % confidence interval (hierarchical bootstrap).

d. Left, change in median lick time in **c**. Right, no-response rate in **c**. Error bars, 95% confidence interval (hierarchical bootstrap). *P*-value, hierarchical bootstrap with a null hypothesis of no change compared to control. Blue, bilateral inhibition. Light blue, unilateral inhibition.

e. The proportion of striatal neurons (left) and ALM pyramidal neurons (right) with spiking activity significantly inhibited or excited ($p < 0.05$; Unchanged, $p >= 0.05$; rank sum test) by inhibiting D1 SPN in VLS. Neurons with a mean spike rate above 1 Hz in control trials during the silencing window were analyzed. This includes 25 striatal projection neurons (n = 5 sessions, 5 mice) and 156 ALM neurons (n = 6 sessions, 6 mice).

f. Schema of the model explaining the data (**f1**). Schema of ALM dynamics during striatal inhibition in the model (**f2**).

g. ALM activity in a two-dimensional space along the ramp and middle mode during inhibition of D1 SPN in VLS. The mean trajectory of trials with silencing (lick between 1.4-1.7 s) is shown. Arrows indicate the direction the trajectory evolves.

h. Left, striatum recording during unilateral D1 SPN inhibition in VLS. Middle and right, similarity matrix of striatal population activity between trial types. Middle, control vs. ref trial (lick between 1.4 - 1.7 s). Right, D1 SPN inhibition trials vs. ref trial.

i. Points with the peak correlation in **h**. Shades, SEM (hierarchical bootstrap). Black, control. Blue, inhibition.

j. Striatal population activity along the ramp mode (n = 10 sessions, 5 mice), trials grouped by lick time. Shades, SEM (hierarchical bootstrap). Left, control. Right, D1 SPN inhibition trials.

k-m. Same as in **h-j** but for ALM recording during bilateral inhibition of D1 SPN in VLS. n = 6 sessions, 6 mice.

to effectively 'rewind' their progression: they stopped unfolding, with a slight recession in the points with peak correlation (quantified by similarity matrix, Fig. 6hi). After the inhibition, activity patterns developed from the post-inhibition state, parallel to the unperturbed condition (Fig. 6i).

The similarity matrix reflects the activity of non-task modulated cells as well. To quantify the effect of inhibition on striatal dynamics representing time, we analyzed the striatal RM activity and decoded $T_{to\ lick}$. Unlike ALM silencing, which rapidly decreased the striatal activity at the stimulation onset (Fig. 5k), both striatal RM activity and decoded $T_{to\ lick}$ gradually decayed during photoinhibition (Fig. 6j and Extended Data Fig. 9i). Because significantly photoinhibited cells (i.e., putative D1 SPN expressing stGtACR1) showed rapid silencing at light onset (within 50 ms; Extended Data Fig. 10l), the gradual decay in timing dynamics is unlikely due to slow photoinhibition but more likely driven by network effects. After inhibition, the ramp restarted in parallel to control conditions from the post-inhibition level (Fig. 6j and Extended Data Fig. 9l). Hence, striatal timing dynamics slowly rewind during unilateral VLS D1-SPN inhibition, as if the D1 inhibition is integrated into the timing dynamics. This is consistent with D1 SPN being a part of the integrator (Fig. 1c, left, and 6f).

Striatum is required for ALM timing dynamics

In our network model (Fig. 6f and Extended Data Fig. 1f), ALM timing dynamics follow those in the striatum. If so, VLS D1-SPN inhibition should rewind ALM timing dynamics as well. To test this hypothesis, we recorded ALM activity during bilateral VLS D1-SPN inhibition (255 neurons, 6 mice). VLS D1-SPN inhibition had relatively minor effects on ALM spiking activity both during ITI and delay epoch: 16.7% of ALM neurons were significantly inhibited and 18.0% were excited during D1-SPN inhibition (Fig. 6e, right). The mean spiking activity decreased by 0.17 ± 0.14 spikes per second (mean ± SEM) during D1-SPN inhibition in the ITI. Thus, VLS D1-SPN is not the major excitatory drive of ALM activity.

Notably, however, VLS D1-SPN inhibition during the delay significantly impacted the timing dynamics in the ALM. During VLS D1-SPN inhibition, ALM population activity patterns paused their progression with a slight recession in the points showing peak correlation (Fig. 6kl). ALM RM activity and decoded $T_{to\ lick}$ gradually decayed during VLS D1-SPN inhibition but resumed ramping in parallel to control trials after inhibition ended, without a rapid recovery phase (Fig. 6m and quantifications in Extended Data Fig. 9m and p). Consistent with 'rewinding', during VLS D1-SPN inhibition, ALM dynamics in the two-dimensional space defined by RM and MM evolved in nearly the opposite direction from the normal trajectory (Fig. 6g). This contrasts with ALM dynamics during ALM silencing, where activity moved toward zero point (See Extended Data Fig. 11 for quantification of rewinding vs. activity moving toward zero). Unilateral VLS inhibition induced a similar, albeit weaker, decay in RM and decoded $T_{to\ lick}$ in the ipsilateral ALM (Extended Data Fig. 12e-h), similar to how unilateral VLS inhibition affected striatal timing dynamics (Fig. 6j). Thus, while D1-SPN is not a major excitatory drive of ALM activity, it strongly influences timing dynamics in ALM. These results are consistent with a model where ALM timing dynamics follow those generated in the striatum (Extended Data Fig. 1f), rather than ALM maintaining an independent integrator (Extended Data Fig. 1d).

The impact of ALM and VLS inhibition on ALM timing dynamics differs qualitatively. Even a weak ALM inhibition (0.3 mW instead of the 1.5 mW used in Fig. 5) caused a weak yet rapid decay in RM activity at the onset of photostimulation, followed by a recovery of ramping during photostimulation (Extended Data Fig. 12a-d) and a mild behavioral effect (shifted the median lick time by 0.13 (0.028 - 0.26) s; mean, 95% confidence interval; n = 5 mice). Thus, the gradual decay in ALM timing dynamics during D1-SPN inhibition cannot be explained by its weak inhibitory effect.

Overall, VLS D1-SPN inhibition has a stronger impact on behavior than ALM silencing, despite its weaker effect on mean spike rates. This strong behavioral effect is due to its specific impact on timing dynamics. The absence of rapid recovery of timing dynamics after striatal inhibition (Extended Data Fig. 9m) suggests there is no other independent area functioning as a timer to restore the activity. These findings support a model where the striatum is a crucial component of the integrator/timer generating timing dynamics, with ALM timing dynamics reflecting those generated by the striatum (Fig. 6f).

Discussion

The frontal cortex and striatum often exhibit similar activity patterns and are essential for motor timing and many other behaviors[11,23,25,44,53,70,91–93] (Fig. 3). This poses a challenge in distinguishing their functional roles. To address this, we conducted a series of transient perturbations coupled with multi-regional electrophysiology. Across conditions, both ALM and striatum population activity predicted subsequent lick timing even after perturbations (Extended Data Fig. 9), indicating a tight causal link between the dynamics in these brain areas and motor timing. All transient manipulations temporally shifted subsequent timing dynamics in parallel to unperturbed conditions beyond the perturbation period and affected lick time. Furthermore, the extent of the shift depended on the strength of manipulations (unilateral vs. bilateral striatal inhibition, and ALM silencing with different laser powers; Extended Data Fig. 12). These findings support the hypothesis that an integrator mediates the generation of timing dynamics.

Importantly, depending on the manipulated brain areas, the multi-regional dynamics during and after perturbation exhibited striking differences: Silencing ALM paused the 'timer' without erasing the timing information in the striatum, whereas inhibiting the striatum effectively rewound the 'timer' in both areas. Additionally, ALM, but not the striatum, encoded the trial history that guides upcoming lick time. Our findings support a model where the striatum functions as an integrator generating timing dynamics in response to inputs provided via ALM (Fig. 5d and Extended Data Fig. 1f).

The neocortex providing inputs to control the striatal integrator may be a general mechanism of temporal integration within the forebrain. First, singing mice (*Scotinomys teguina*) adjust their song durations based on social context. Silencing the frontal cortex reduces this context-dependent modulation without affecting the execution of the song, supporting our model where the frontal cortex adjusts the timer based on contextual information[94]. Second, temperature manipulation of the striatum influences the perception of duration[48], implying that the striatum is essential for both controlling action timing and perceiving time. Third, ramping activity correlated with temporal integration of external signals, or evidence accumulation, is observed in the forebrain during decision-making tasks [15,18,22,35–38,95], and the striatum has been shown to be one of the key sites for evidence accumulation. Our network model may explain the mechanism of temporal integration across motor and cognitive behaviors.

Neural correlates and the causality of manipulation on behavior are insufficient to distinguish multi-regional dynamics models (Extended Data Fig. 1). Specifically, prolonged manipulations that encompass the entire trial (such as muscimol infusion and chemogenetics) have limited ability to differentiate between models (Extended Data Fig. 1, bottom). Therefore, combining transient perturbation with large-scale electrophysiology is critical. This strategy can be employed to dissect multi-regional dynamics across behaviors. Importantly, stringent consideration of behavioral adaptation to perturbations[75,76], and calibration of photostimulation conditions to prevent rebound are critical for interpretable and reproducible perturbation experiments (Extended Data Fig. 7j-q).

In a memory-guided licking task with a sensory cue signaling proper lick time, transient ALM silencing was followed by a rapid recovery of ramping activity to the original trajectory (akin to Fig. 1c right; or switching to another ramping trajectory for a different action), without any temporal shift in ramping activity, regardless of the strength of inhibition[58,59]. In contrast, identical ALM manipulation during the timing task led to a temporal shift in ramping activity, with the extent of the shift depending on the inhibition strength. Thus, depending on the task, the same brain area (ALM) and similar neural activity (ramping activity) are governed by different dynamical systems, presumably to optimize computation for the task at hand.

Similarly, ALM encoded trial history across trials only when mice used trial history to adapt lick timing. We expected that silencing ALM during the ITI would alter the subsequent lick time by affecting this trial history information. Unexpectedly, ALM activity recovered rapidly following ITI silencing and did not affect lick time (Fig. 5a and Extended Data Fig. 9q). This implies that trial history information is robustly maintained across brain areas during the ITI (consistent with the observation of such activity patterns across areas[77–79]), yet after the cue, it is signaled to the striatum via ALM to control timing dynamics.

In a dynamical system, the evolution of activity states is shaped by both the initial conditions and external inputs of a network. There is substantial evidence that initial conditions critically influence the speed of dynamics[10,60,96]. Although we treated trial history information as an input to the integrator in our network models, it is probable that both initial conditions and external inputs synergistically determine action timing. In addition, temporal integration can be implemented by tuned feedback, lateral inhibition, and/or a feedforward network[29,32–34,97]. Our results do not distinguish between these models, which require further investigation. The striatum predominantly comprises inhibitory neurons with relatively sparse lateral connections[81]. It indirectly modulates thalamic activity through other basal ganglia nuclei, while intralaminar thalamic nuclei provide direct excitatory input to the striatum[67,81,82,98,99]. Future perturbation experiments across areas within this long-range subcortical loop may help elucidate how the integrator is fully implemented in the forebrain.

Acknowledgments: We thank N. Li, D. Fitzpatrick, M. Sarvestani, L. Colgan, B. Mensh, R. Heldman, K. Citrin, T. Yiu, X. Yu, and Inagaki lab members for comments on the manuscript, P. Scarpinato and N. Spiller for DeepLabCut pipeline, K. Shirley for imaging, and H. Shearin and other MPFI ARC members for animal care. This work was funded by Max Planck Florida Institute for Neuroscience (H.K.I.), Max Planck Free Floater Program (H.K.I.), NIH New Innovator Award (NINDS and OD; 1DP2NS132108; H.K.I.), Searle Scholars Program (H.K.I.), Klingenstein-Simons Fellowship (H.K.I.), and McKnight Scholar Award (H.K.I.).

Author contributions: Conceptualization: HKI; Investigation: ZY, MI, CG, LF, HKI; Funding acquisition: HKI; Supervision: HKI; Writing – original draft: ZY, HKI; Writing – review & editing: ZY, HKI.

Declaration of interests: Authors declare that they have no competing interests.

Extended Data Fig. 1

Extended Data Fig 1. Two-regional network models replicating the ramping activity in ALM and striatum
Multi-regional network models of ALM and striatum: both regions contain two neurons and exhibit ramping activity with temporal scaling to replicate the data in Fig. 3. Additionally, step input mimicking the trial-history mode activity is provided to ALM based on Fig. 4 (except for **a**). The network configuration, i.e., the connectivity matrix, varies across models (Methods), leading to different location(s) of integrator(s) and various responses to transient perturbations. ALM silencing was implemented by silencing both neurons in ALM, and striatal inhibition was implemented by silencing one of the striatal neurons (mimicking the silencing of D1 SPNs, which constitute roughly 50% of all striatal neurons). Note that prolonged silencing does not distinguish between models (bottom row), underscoring that transient perturbation with concurrent multi-regional recording is essential to differentiate models. The redundant model (**c**) can be modified to include another neuron in ALM that signals input to the striatum along the integrator mode. This would allow for the recovery of activity to pre-perturbation levels after ALM silencing, as seen in the data (Fig. 5). However, in this case, ALM would continue to show ramping activity during striatal inhibition, which would be inconsistent with the data (Fig. 6). Therefore, only **f** replicates the data.

- **a.** Externally driven model. 1st row, schema of the model, and the connectivity matrix used for the simulation. 2nd row, input (I) into the network. 3rd and 4th row, ALM and striatal ramp mode activity without perturbation. 5th and 6th row, ALM and striatal ramp mode activity during transient ALM silencing (the ALM activity along the ramp mode was cropped for visualization purposes. In all cases, ALM activity during ALM silencing decreased to 0). 7th and 8th row, ALM and striatal ramp mode activity during transient striatum inhibition. 9th and 10th row, ALM and striatal ramp mode activity during prolonged ALM silencing throughout the trial. Dashed lines, control conditions overlaid.
- **b.** Same as in **a** but for integration via interareal connections (distributed) model.
- **c.** Same as in **a** but for coupled local integrators (redundant) model.
- **d.** Same as in **a** but for ALM being the integrator and striatum being the follower (specialized) model.
- **e.** Same as in **a** but for striatum being a leaky integrator and ALM being the follower (specialized) model.
- **f.** Same as in **a** but for striatum being the integrator and ALM being the follower (specialized) model that replicates the data.

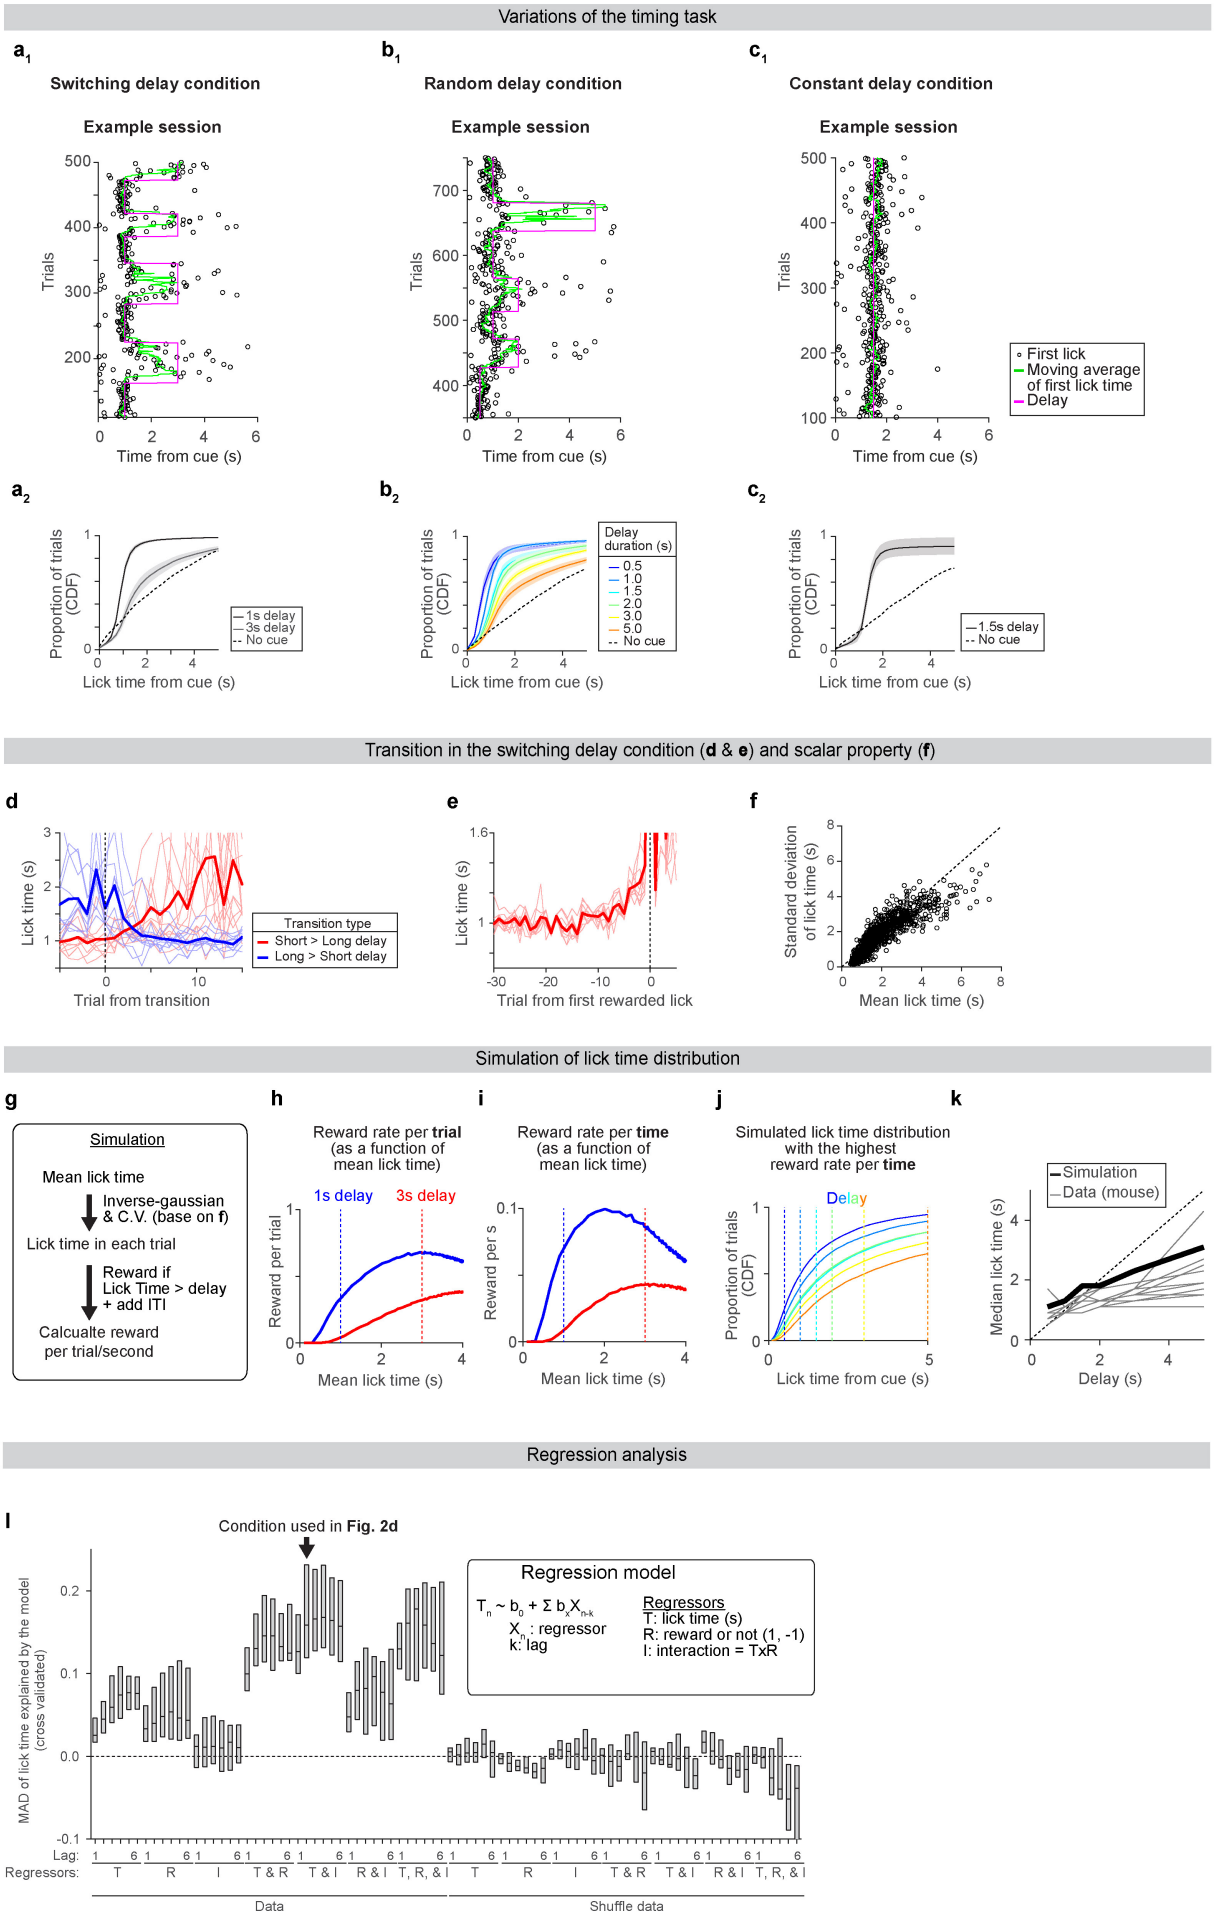

Extended Data Fig. 2

Extended Data Fig 2. Characterization of the lick time distribution in the lick-timing task
We either changed the delay duration in blocks of trials (switching delay condition, when switching between two delays, **a**; random delay condition, when switching across multiple delays, **b**) or kept the delay identical across sessions/trials (constant delay condition, **c**). When delays switched, mice adjusted their lick times within 10 trials (**d, e**). The lick time distribution exhibited scalar properties, similar to many other timing tasks across species (**f**). Mice often licked earlier than the delay duration, not maximizing reward per trial. However, this lick time distribution (especially in the best-performing mice) is close to the simulated distribution that maximizes the reward amount per time, given the short inter-trial interval and the inverse normal distribution of lick times[100–104] (**g-k**). 42 regression models were screened to identify the best model explaining the lick time (**l**).

- **a.** Lick time distribution under switching delay condition. Example session (**a1**). Cumulative distribution of lick time in 1 s and 3 s delay blocks (**a2**). Duplicated from Figure 2bc for comparison.
- **b.** Same as in **a** but for the random delay condition. n = 276 sessions, 17 mice.
- **c.** Same as in **a** but for the constant delay condition. n = 71 sessions, 13 mice.
- **d.** Change in lick time after transitioning between 1 and 3 s delay blocks. 0, last trial before delay transition. Thick lines, the mean across mice. Thin lines, individual mice (n =10 mice).
- **e.** Change in lick time before the first rewarded lick following transitions from a short to a long delay. Thick line, the mean across mice. Thin lines, individual mice (n = 10 mice).
- **f.** Relationship between the mean and the standard deviation of lick time. Circles, individual sessions (n = 153 sessions, 30 mice). The mean and standard deviation of lick time are correlated, consistent with scalar properties reported across species[105].
- **g.** Simulated lick time distribution following an inverse-Gaussian distribution and coefficient of variation (CV) in **f**. Subsequently, based on the task structure, we calculated the reward rate (Methods).
- **h.** Simulated reward rate per **trial** as a function of mean lick time in 1 s delay block (blue) and 3 s delay block (red).
- **i.** Simulated reward rate per **time** as a function of mean lick time. Note that the peak reward rate is attained with a shorter mean lick time compared to that in **h**.
- **j.** Following the procedure described in **g** and **i**, we estimated the optimal mean lick time for each delay duration (delay duration is indicated by colored vertical dotted lines). Then, we plotted the distribution of simulated lick times in these conditions, which yielded the highest reward rate per unit of time. Note that in a large proportion of trials, the licks occurred before the end of the delay period, replicating what was observed in the data.
- **k.** Thick line, the median lick time of simulated optimal lick time distribution in **j**. Thin lines, experimental data (individual mice). The optimal lick time appears to align with the upper limit of the experimental data.
- **l.** Median absolute deviation (MAD) of lick time explained by different trial-history regression models under random delay condition (n = 276 sessions, 17 mice; results were consistent in the switching delay condition). Regressors and lags in each model are indicated at the bottom. Arrow, the condition that best explained the data (Fig. 2d).

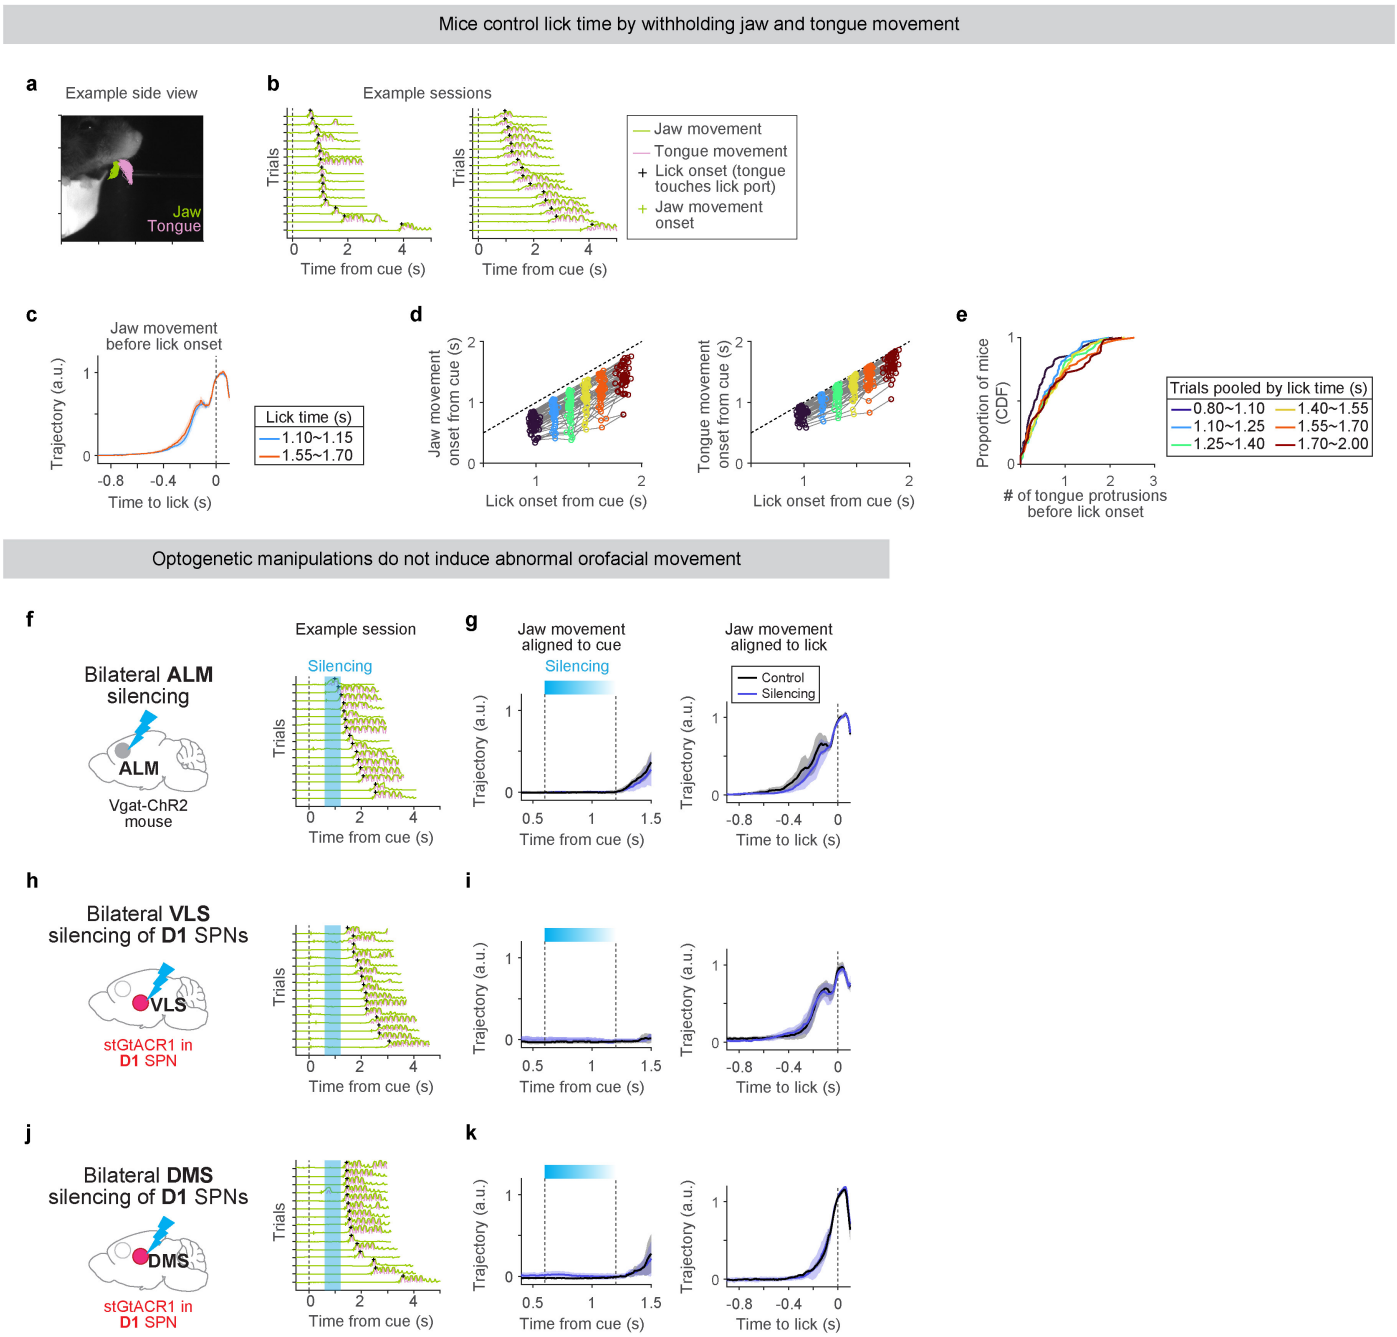

Extended Data Fig. 3

Extended Data Fig 3. Characterization of orofacial movements in the lick-timing task
Jaw and tongue movements in most recording sessions were tracked using high-speed videography and DeepLabCut[106] (**a;** Methods). This revealed that mice adjusted lick times (detected by contact of the tongue with the lick port) by controlling the onset rather than the speed of tongue/jaw movements (**b-e**). Additionally, we did not detect any abnormal movements during ALM silencing (**f, g**), VLS inhibition (**h, i**), and DMS inhibition experiments (**j, k**).

a. An example side view clip of a mouse. Movement of the jaw (green) and tongue (purple) were tracked. Trajectories of individual trials are overlaid.

b. Vertical jaw (green) and tongue (purple) movements in two example sessions. The left session is from the animal shown in **a**. 15 randomly selected trials sorted by lick time are shown.

c. Average trajectories of vertical jaw movement aligned to the lick onset. The kinematics of jaw movement remains consistent regardless of the lick timing (indicated in different colors). Shades, SEM (bootstrap). n = 58 sessions, 34 mice.

d. Relationship between lick onset (the timing when the tongue contacted the lick port) and jaw movement onset (left) or tongue movement onset (right). Trials were grouped into six ranges. The onset of movement was tightly correlated with lick onset. Circles, individual mice. Dotted line, the unity line. n = 58 sessions, 34 mice.

e. The cumulative distribution of the number of subthreshold tongue protrusions not reaching the lick port before lick onset. The distribution was not different across lick times (shown in different colors). **c-e** conclude that mice did not change their kinematics but instead the onset of movement when they lick at different timings. n = 58 sessions, 34 mice.

f. Tongue and jaw movement trajectories in an example session with ALM silencing. Same format as in **b**. Cyan bar, silencing.

g. Jaw movement aligned to the cue (left) or the lick onset (right). Trials with lick after the silencing were analyzed. Shades, SEM (bootstrap). n = 26 sessions, 9 mice. No abnormal movement was detected during silencing, and the animals followed normal kinematics to lick even in the silencing trials.

h-i. Same as in **f-g**, but with D1 VLS silencing. n = 6 sessions, 6 mice.

j-k. Same as in **f-g**, but with D1 DMS silencing. n = 6 sessions, 6 mice.

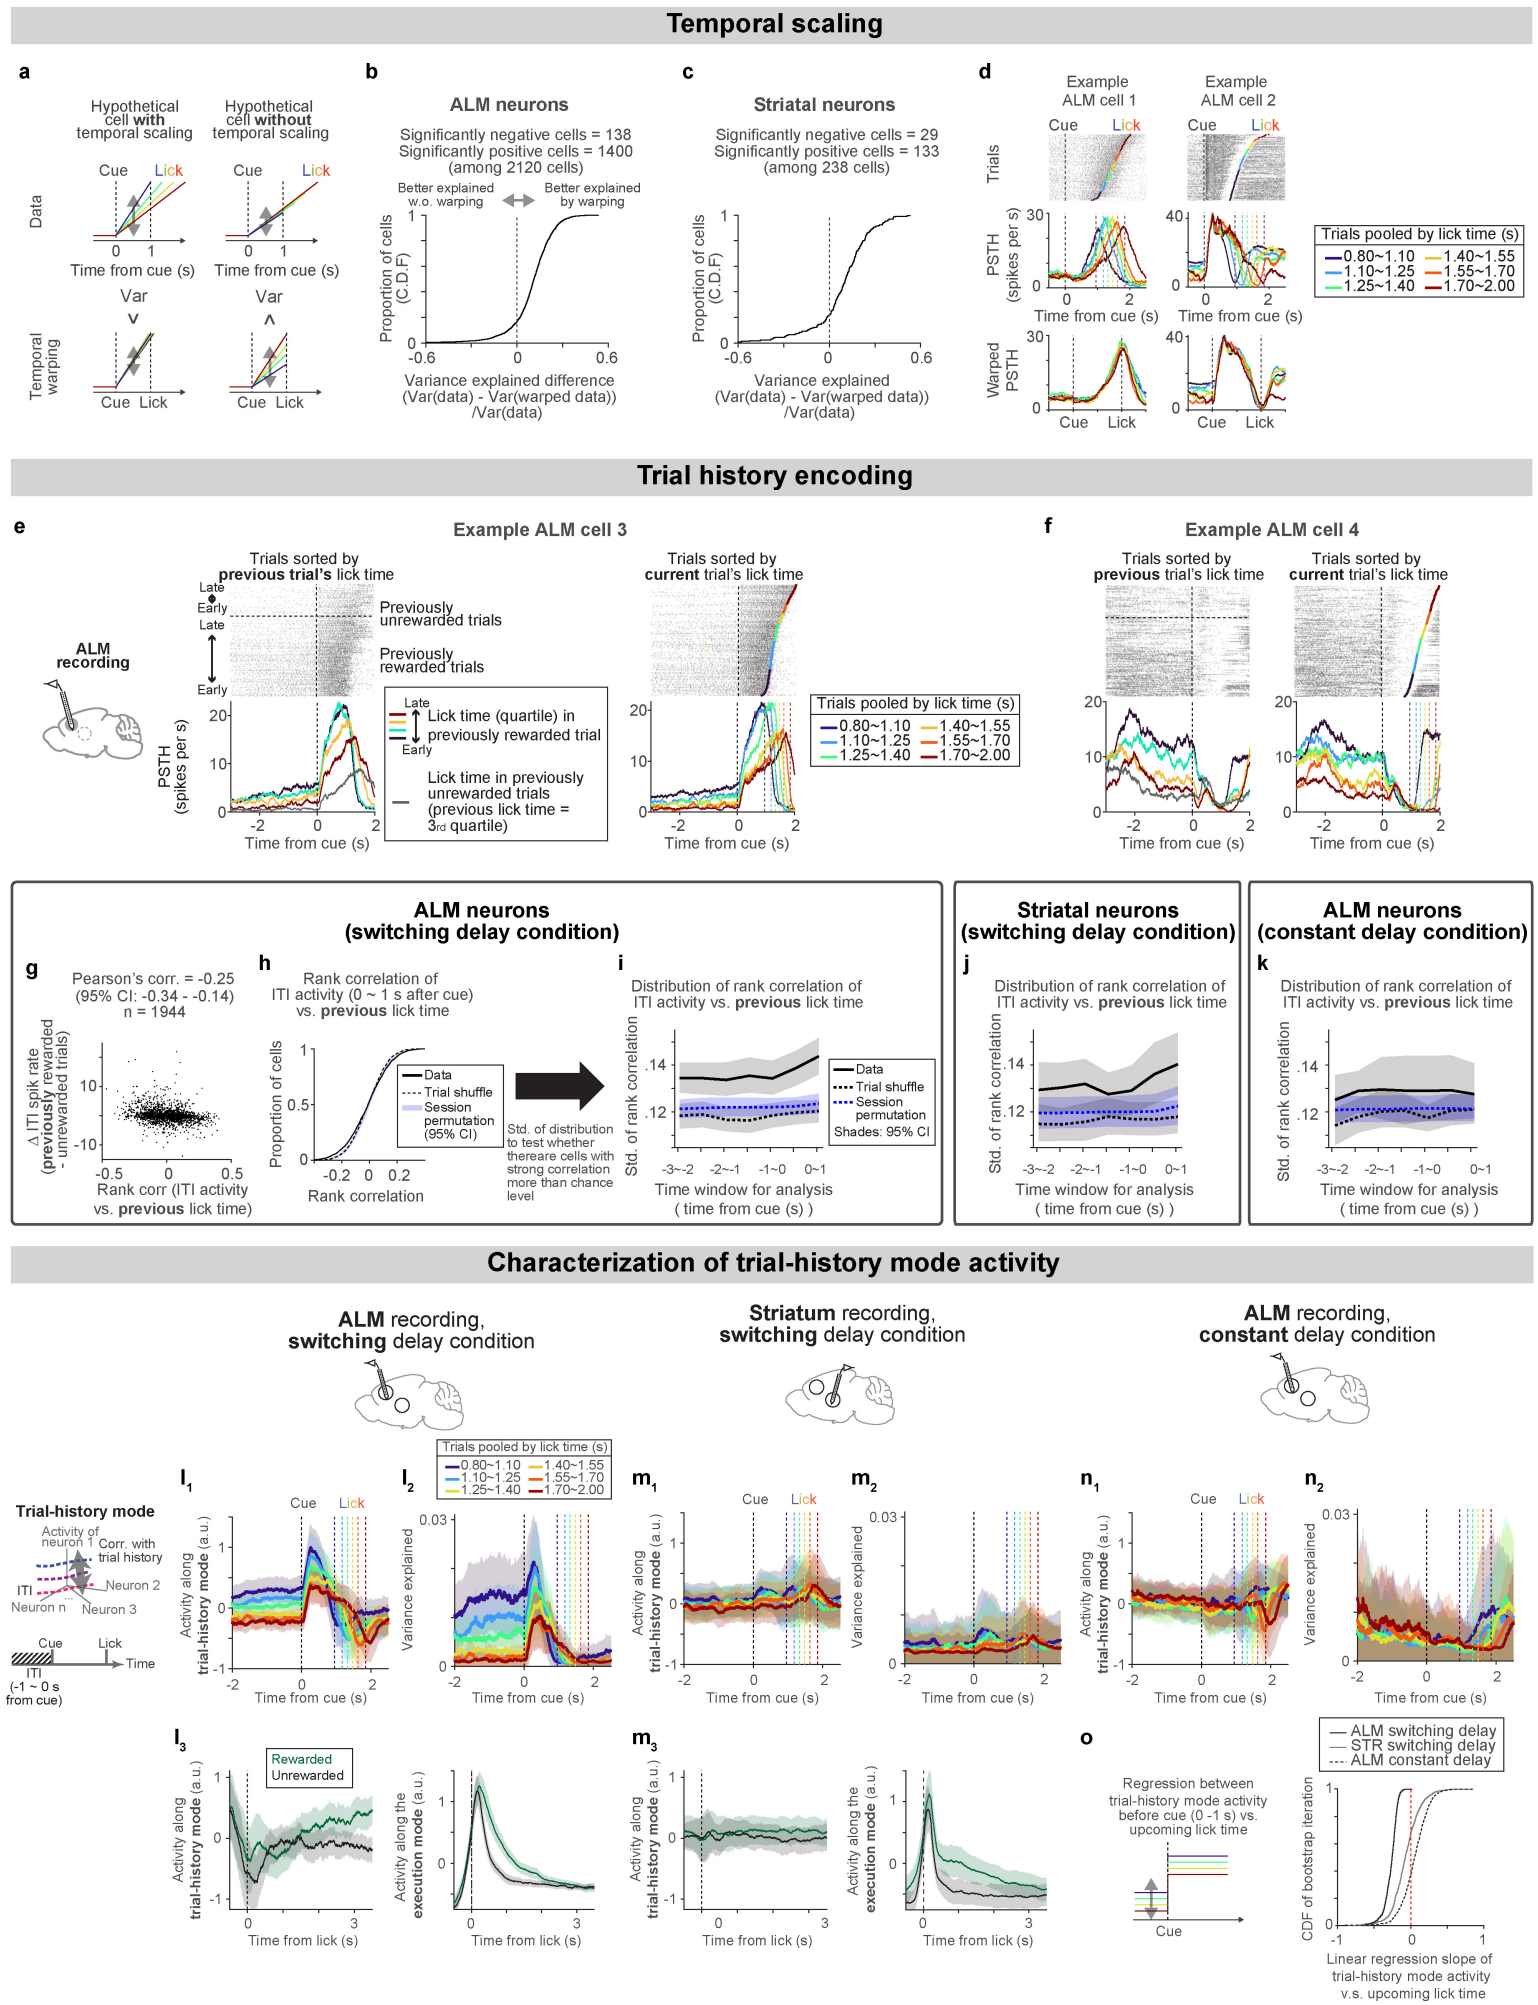

Extended Data Fig. 4

Extended Data Fig 4. Temporal scaling and encoding of trial history in ALM and striatum

Temporal warping of spike times revealed that the activity in a large proportion of ALM and striatal neurons can be better explained by temporal scaling rather than a null model without scaling (**a-d**). The encoding of trial history (previous trial's lick time and reward outcome) by the ITI activity in ALM is strong under the switching delay condition (**e-i**) but not above chance level under the constant delay condition (**k**) or in the striatum (**j**). Consistently, activity along the trial-history mode in the ALM under the switching delay condition predicts upcoming lick time (**l** and **o**). However, activity along trial-history mode in the striatum under the switching delay condition (**m**) or in ALM under the constant delay condition (**n**) does not predict upcoming lick time.

- **a.** Schema illustrating two hypothetical cells that encode time differently through ramping activity[49]. Left, a cell with ramping activity encoding relative time, where the ramping speed changes as the lick time varies (i.e., temporal scaling). In this scenario, the across-trial variance (double-headed arrows) decreases following temporal warping (bottom). Right, a cell with ramping activity encoding absolute time, where the spike rate increases as time progresses. In this case, the across-trial variance increases following temporal warping (right). Vertical dotted lines, time window used for variance calculation.
- **b.** The cumulative distribution of the difference of variance explained between data and warped data across ALM neurons (higher value represents data explained better by temporal scaling; see **a**). n = 2139 neurons, 31 mice. Neurons with more than 200 trials were analyzed. Significant cells, $p < 0.05$ with bootstrap of trials.
- **c.** Same as in **b** but for striatal cells, n = 595 neurons,10 mice.
- **d.** Two example ALM cells showing temporal scaling. Same format as **Fig. 3a and f**: trials are sorted by the current trial's lick time and grouped into 6 ranges. Colored dots on the spike raster indicate the first lick time.
- **e.** An ALM example cell whose ITI activity is modulated by the lick time and reward outcome in the previous trial. The same cell as in Fig. 4a for comparison. Top, spike raster, grouped by reward outcome in the previous trial and sorted by the lick time in the previous trial. In this example cell, the ITI activity is higher in trials after rewarded trials and with earlier licks. Bottom, mean PSTH. Lick times of the previous rewarded trials were divided into quartiles indicated by different colors. The gray trace, trials following previously unrewarded trials with previous trial's lick times within the 3rd quartile. Right, the same cell but trials sorted by lick time in current trials. The spike rate during ITI predicts the upcoming lick time.
- **f.** Another ALM example cell showing trial history modulation during ITI, same format as **e**.
- **g.** Relationship between the encoding of previous lick time (partial rank correlation between spike rate vs the lick time in previous unrewarded trials, removing the effect of upcoming lick time; Methods) and whether the animal received a reward or not in the previous trial (based on activity during ITI: 0 - 1s before the cue). Dots, individual neurons. Neurons with more than 20 previous unrewarded, 20 previous rewarded trials, and 75 current trials were analyzed.
- **h.** The cumulative distribution of partial rank correlation between the inter-trial-interval (ITI) activity (0 - 1 s after the cue) and the lick time in previous unrewarded trials (removing the effect of upcoming lick time; Methods) across ALM neurons. Black solid line, data. Black dashed line, trial shuffle control. Blue shades, 95% confidence interval from session permutation test[72] (Methods). Neurons with more than 75 previous unrewarded trials were analyzed (n = 1944 neurons, 40 mice).
- **i.** The standard deviation of partial rank correlation across cells (shown in **g**) was calculated across different time windows. A higher standard deviation than the trial shuffle and session permutation test implies a significant encoding of the previous lick time by the ITI activity.
- **j.** Same as in **h-i** but for striatal neurons. n = 1457 neurons, 16 mice.
- **k.** Same as in **h** but for ALM neurons under constant delay condition. n = 1588 neurons, 11 mice.
- **l.** ALM population activity along the trial-history mode, under switching delay condition (**l1**). Duplicated from Fig. 4b for comparison. Variance of spiking activity explained by trial-history mode (**l2**). Colors, different lick

times. Shades, SEM (hierarchical bootstrap). n = 3261 neurons. ALM activity along the trial-history mode gradually diverged after rewarded vs. unrewarded licks (left) under the switching delay condition (**l3**). This trend is not significant at 3.5 s from lick $p = 0.058$ (hierarchical bootstrap with a null hypothesis that the trial-history mode activity following rewarded trials is smaller than or equal to the one following unrewarded trials). In contrast, ALM activity along the execution mode, which captures the activity during the lick (Methods), showed a transient change in activity after the lick (up to ~2 s). This transient change along the execution mode most likely reflects differences in lick patterns between these trial types. Since the trial-history mode activity started diverging after the execution modes converged, the divergence of trial-history mode activity is likely not due to movement. Rewarded (green) and unrewarded (black) trials with similar lick times (lick between 1.4 and 1.8 s after the cue) were analyzed.

m. Same as in **l** for striatal recording during the switching delay condition. $p = 0.401$ in **m3** (left).

n. Same as in **a** for ALM recording during the constant delay condition.

o. The relationship between trial-history mode activity and upcoming lick time. We performed a linear regression analysis between trial-history mode activity during 0-1 s before the cue and the upcoming lick time for each iteration of the hierarchical bootstrap (left, schema). We then plotted the distribution of the linear regression coefficient across these iterations as a cumulative distribution function (CDF, right). A negative value indicates higher trial-history mode activity precedes earlier lick. $p < 0.001$, = 0.3274, 0.6310 for ALM under switching delay, striatum under switching delay, ALM under constant delay, respectively (with a null hypothesis that the slope of the linear regression is larger than or equal to 0). Thus, only under the switching delay condition in ALM, the trial-history mode activity predicts the upcoming lick time.

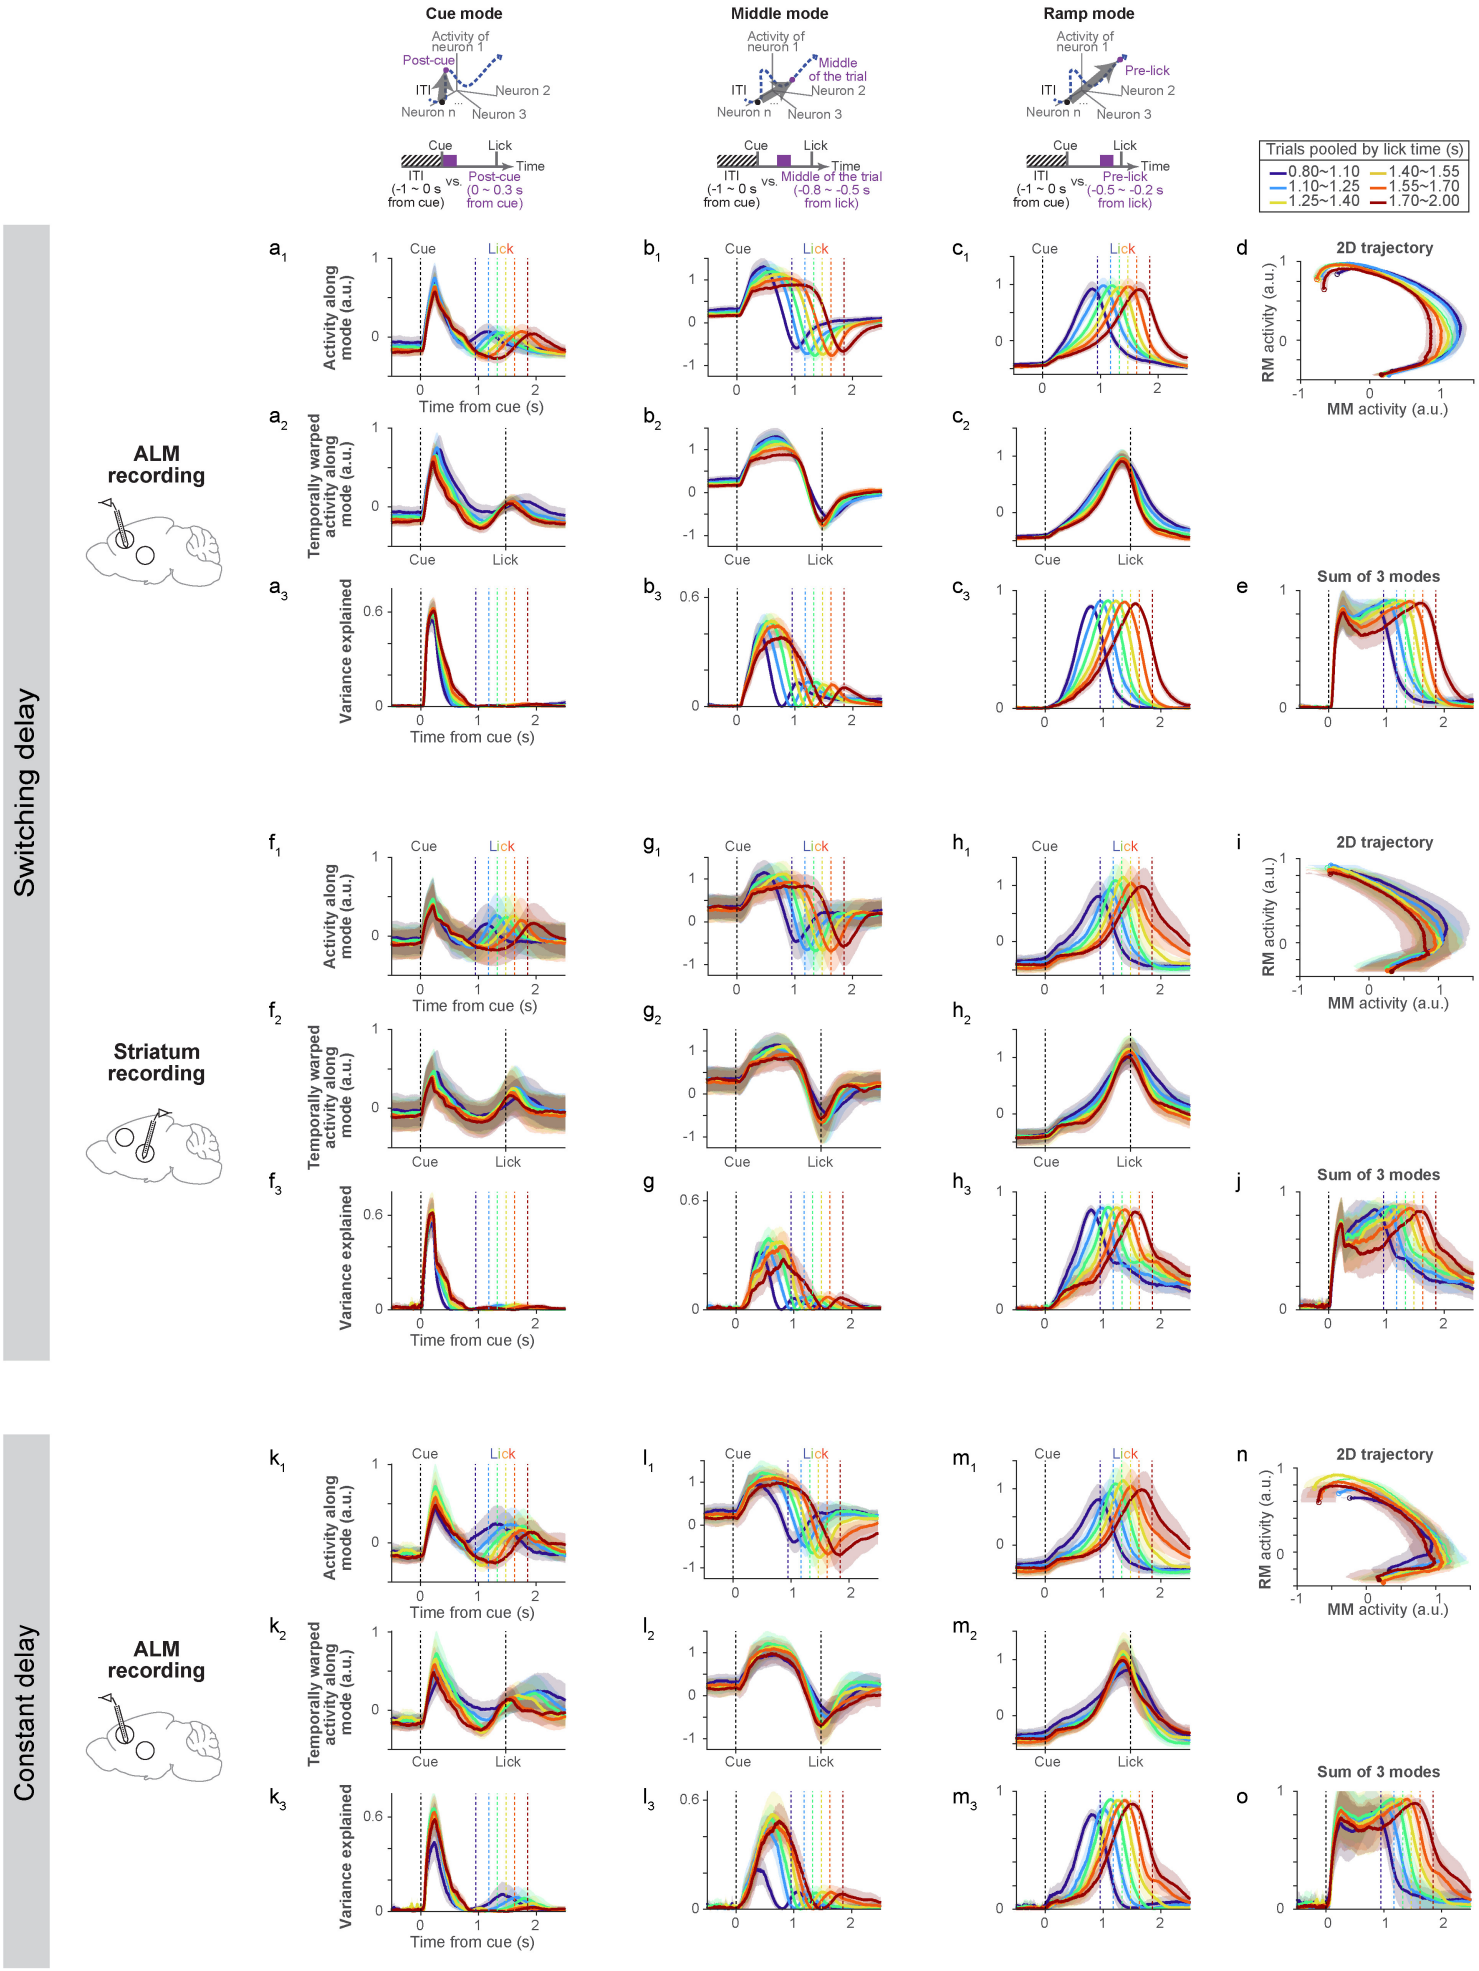

Extended Data Fig. 5

Extended Data Fig 5. Comparison of activity modes across brain regions/conditions

In both ALM and striatum, we observed three modes of population activities (cue mode, middle mode, and ramp mode) that together tiled the time from trial onset to lick, explaining around 80 % of the variance in task-modulated spiking activity (Methods). The cue-triggered peak of cue mode activity did not show temporal scaling (differences in activity across trials with different lick time ranges became larger after scaling). In contrast, the second peak of cue mode activity around lick, as well as middle and ramp mode activities showed temporal scaling across conditions. Also, note that activity patterns along these modes and the variance explained are qualitatively similar between ALM and striatum.

a. ALM population activity along the cue mode, under switching delay condition (**a1**). Cue mode activity temporally warped between cue and lick (**a2**, Methods). Variance of spiking activity explained by cue mode (**a3**, Methods). Colors, different lick times. Shades, SEM (hierarchical bootstrap). n = 3261 neurons.

b. Same as in **a** but for ALM population activity along the middle mode.

c. Same as in **a** but for ALM population activity along the ramp mode. Duplicated from Fig. 3e for comparison.

d. Two-dimensional space defined by the population activity along the ramp mode (RM) and middle mode (MM). Regardless of the lick time, the activity follows a similar trajectory, but the speed varies across different lick times.

e. The combined variance explained by the three modes.

f-j. Same as in **a-e** but for striatal neurons under switching delay condition. n = 1070 cells.

k-o. Same as in **a-e** but for ALM neurons under constant delay condition. n = 1040 cells.

Neuropixels recording in striatum

a

AAV-GFP injection site (ALM) | Bregma, AP -1.0 ~ -0.5 mm | Bregma, AP -0.5 ~ 0.0 mm | Bregma, AP 0.0 ~ 0.5 mm | Bregma, AP 0.5 ~ 1.0 mm

Projection of ALM neurons

GFP
NeuN

b

Location of units and their ramping activity

Ramp amplitude (spikes per s)
Pre lick activity - baseline

c

Location of fiber optics

DMS
VLS

d

Striatum recording

Recording tracks

e

Proportion of inter spike interval > 2 s

Post spike suppression (ms)

spike width (ms)

Cell numbers
Fast spiking interneuron (FSI): 584
Striatal projection neuron (SPN): 1217
Tonically active neuron (TAN): 127
Other Interneurons: 44

Decoding analysis

f

Estimate lick timing using kNN decoder

A data point in a test trial

Activity of neuron 1

1) Find k nearest neighbor (kNN) data points in ref trials

Neuron 2

Neuron 3

Neuron n

2) Average "time to lick" of the kNN ref data points

Decoded time to lick of the test data point

g

p (all vs. IM & RM) = 0.133
p (IM vs. all) = 3.93 e-16
p (RM vs. all) = 1.21 e-07

Proportion of sessions

All neurons
ITI mode (IM)
Ramp mode (RM)
Cue mode (CM)
IM & RM
RM & CD
All, shuffle

Pearson's correlation between decoded vs. actual lick time (decoded at cue onset)

h

Pearson's correlation between decoded vs. actual lick time (decoded at cue onset)

Correlation

ALM recording
STR recording

Numebr of neurons per session

i

ALM & striatum

ALM
Striatum

Example trials

Decoded time to lick (s)

Time from cue (s)

j

Example session

Corr = 0.79
Shuffle:0.51-0.62

Decoded time to lick (s) based on ALM activity

Decoded time to lick (s) based on STR activity

k

Corr. = 0.69 ± 0.10
(mean ± std)
p < 0.05 in 16 / 16 sessions

Data

Session

Shuffle

Extended Data Fig. 6

Extended Data Fig 6. Striatal anatomy and recording

ALM projects to a large portion of the striatum (**a**), and we observed task-modulated activity across these sectors (**b**). See **d** for example recording tracks and **e** for the classification of cell types based on spike features. We developed a kNN-based decoder to estimate the time to lick using simultaneously recorded units at single time points (**f-h**). When ALM and striatum were simultaneously recorded, the decoded time to lick in both regions (decoded independently) was correlated at the single trial level, implying that the representations of time are highly synchronized between these areas (**i-k**).

- **a.** ALM neurons project across sectors in the striatum. AAV-GFP was injected into ALM.
- **b.** The spatial distribution of recorded striatal neurons in the Allen CCF. Colors, ramping increases in spiking activity from the baseline (0-1 s before the cue). Black dots, neurons that do not ramp up.
- **c.** Locations of the tips of the tapered fiber optics implanted for bilateral striatal silencing (only showing one hemisphere as the two fibers were implanted symmetrically). Related to Fig.6.
- **d.** An example brain image of recording tracks acquired by a light sheet microscopy (Methods). Striatal recording tracks are labeled with CM-DiI (magenta). Coronal view, maximal intensity projection of 415 µm optical section.
- **e.** Striatal cell types were classified based on three spike features[69] (Methods).
- **f.** Schema depicting a k-nearest neighbor (kNN) method to decode the time to lick using population neural dynamics at each time point (Methods).
- **g.** Comparison of the decoding accuracy among kNN decoders based on activity along different dimensions. The kNN decoder based on activity along two dimensions (trial-history mode and ramp mode, magenta) performed nearly as well as the kNN decoder based on the top 90% principal components (all neurons, red; this is the condition we used for analyses). This suggests that these two modes contain most of the timing information. Decoding based on trial-history mode or ramp mode alone performed worse. Decoding accuracy was quantified by Pearson's correlation between actual lick time vs. lick time decoded at cue onset (Methods).
- **h.** The performance of the kNN decoder as a function of the number of simultaneously recorded neurons. The performance increased with more recorded neurons.
- **i.** Decoded $T_{\text{to lick}}$ was highly correlated between ALM and striatum at a single trial level. Four example trials from an example session are shown. Traces end at the time of lick.
- **j.** The relationship between decoded lick times estimated from ALM and striatal neurons in an example session. Dots, all time points (50 ms bin; from cue to lick) in the example session. Pearson's correlation across all time points (0.79) was significantly higher than that of the trial shuffle control (0.51-0.62, 95% confidence interval).
- **k.** Pearson's correlation of decoded time between ALM and striatum was significantly higher than trial shuffle controls in all simultaneously recorded sessions (16 sessions). Thus, ALM and striatal timing dynamics are synchronized at a single trial level.

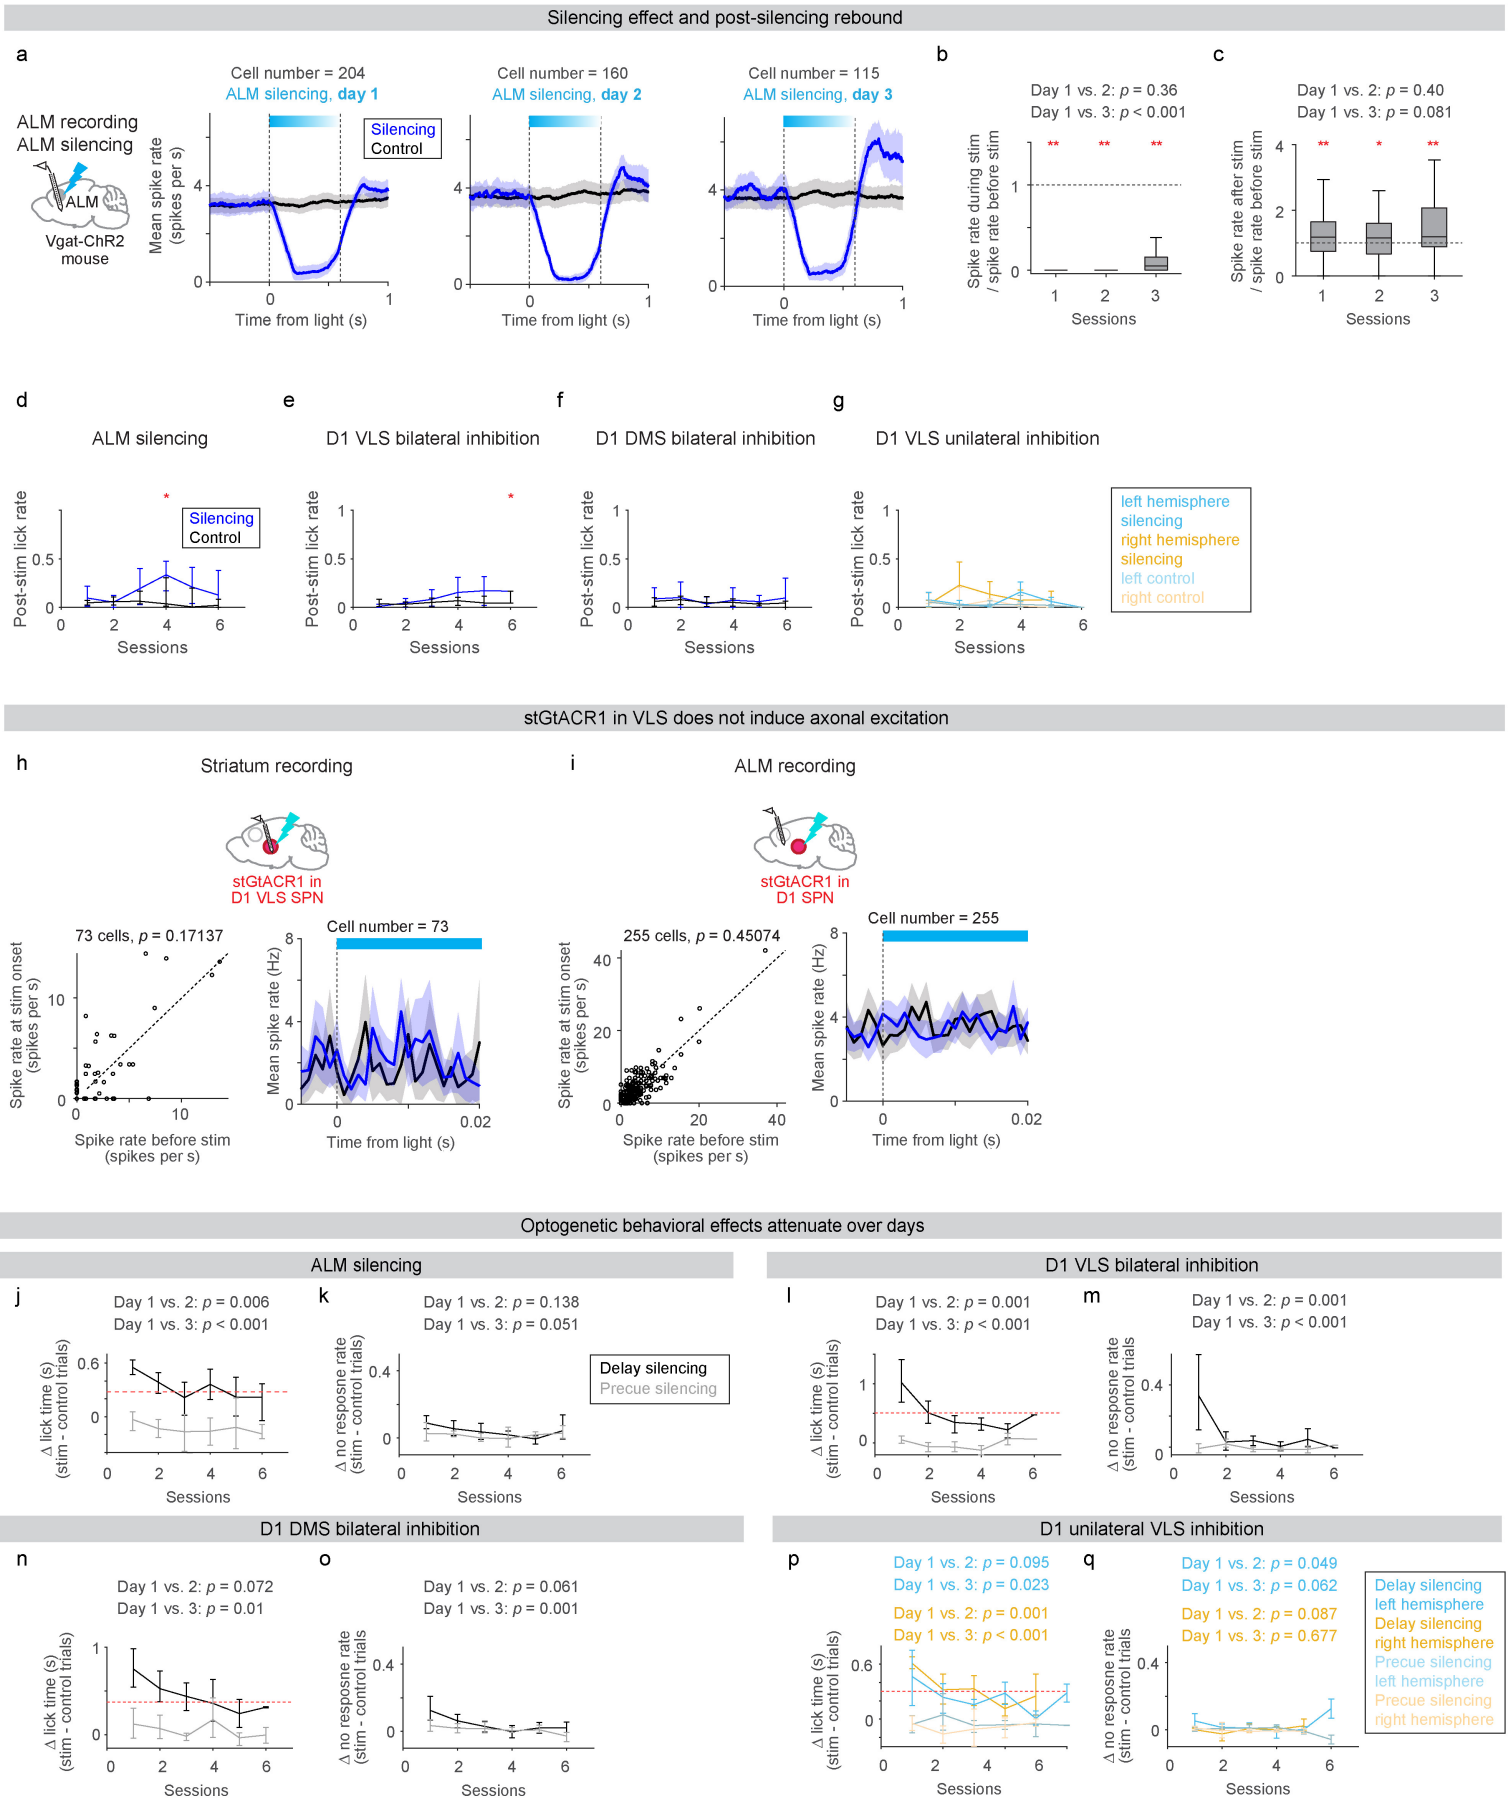

Extended Data Fig. 7

Extended Data Fig 7. Characterization of optogenetic effect across sessions

Here, we characterized possible caveats of optogenetic manipulations. First, we measured the post-silencing rebound activity and lick, over days (**a-g**). We noticed that for ALM silencing, post-silencing rebound activity increased over days (**a** and **c**), accompanied by the increase of post-silencing lick (**d**), which explains the decrease in the shift of lick timing over sessions (**j-k**). Therefore, we restricted our analysis to data from the first two days. Second, we confirmed the lack of axonal excitation in both striatum and ALM when we used stGtACRl in D1-SPN[73,87,88] (**h-i**). Third, we observed that the behavioral effect of D1-SPN inhibition showed a drastic decline over sessions regardless of manipulated sectors, so we focused our analysis on the data from the first manipulation day. These observations are consistent with the short-lived effect of manipulations observed in other species and manipulations[75,76].

- **a.** Mean spike rate of ALM putative pyramidal neurons during ALM silencing (during no cue trials) over 3 consecutive days (spike rate was smoothed by 200 ms boxcar causal filter). Shades, SEM (hierarchical bootstrap). Note the increase in the post-silencing rebound activity over days.
- **b.** Quantification of ALM spiking activity during ALM silencing: spike rate during photostimulation (50 - 250 ms from the photostimulation onset) divided by spike rate 0 - 500 ms before the photostimulation. The central line in the box plot, median. Top and bottom edges, 75% and 25% points. Whiskers, the lowest/highest datum within the 1.5 interquartile range of the lower/upper quartile. $^{**}p < 0.001$, $^{*}p < 0.01$ (signed rank test with *Bonferroni* correction for multiple comparisons). *P*-value for across-day comparisons, rank-sum test comparing data between days.
- **c.** Quantification of post-silencing rebound activity during ALM silencing: spike rate 500 ms after photostimulation divided by spike rate 500 ms before the photostimulation onset. The box plot follows the same format as in **b**. $^{**}p < 0.002$, $^{*}p < 0.02$ (signed rank test with *Bonferroni* correction for multiple comparisons;). *P*-value for across-day comparisons, rank-sum test comparing data between days.
- **d.** Post-stim lick rate (Proportion of trials with lick within 600 ms following the photostimulation offset time) across days for ALM silencing. $^{*}p < 0.001$ (hierarchical bootstrap with *Bonferroni* correction for multiple comparisons; null hypothesis is that the post-stim lick rate in photostimulation trials is lower or equal to that in control trials).
- **e.** Same as in **d** but for D1 VLS bilateral inhibition.
- **f.** Same as in **d** but for D1 DMS bilateral inhibition.
- **g.** Same as in **d** but for D1 VLS unilateral inhibition.
- **h.** When stGtACR1 exhibits leaky expression at the axon[73,87,88], it can induce short-latency axonal excitation at the onset of photostimulation. Thus, we examined the spiking activity at the onset of D1-SPN inhibition (during no cue/precue silencing trials) in VLS, which did not show signs of axonal excitation. Left, spike rate during baseline (20 ms before the photostimulation) vs. spike rate within 20 ms of the photostimulation onset for individual striatal neurons. *P*-value, signed rank test. Right, mean spike rate of SPNs (no smoothing of spike rate). n = 73 neurons.
- **i.** Same as in **g** but for ALM neurons during bilateral D1 VLS silencing. n = 255 neurons. Neither the striatum (**f**) nor ALM (**h**) showed signs of axonal excitation caused by GtACR1.
- **j.** The shift in lick time caused by ALM silencing became weaker over consecutive days with ALM silencing. Error bars, 95% confidence interval (hierarchical bootstrap). The red dashed line, half of the effect observed in session 1. *P*-value, bootstrap with a null hypothesis that the effect of delay silencing on session 1 is smaller than or equal to the subsequent compared day. n = 14 mice.
- **k.** The change in no response over consecutive days with ALM silencing. Error bar, 95% confidence interval (hierarchical bootstrap). *P*-value, bootstrap with a null hypothesis that the effect of delay silencing on session 1 is smaller than or equal to the subsequent compared day.

l-m. Same as in **i-j** but for bilateral D1 SPN inhibition in VLS. n = 6 mice.

n-o. Same as in **i-j** but for bilateral D1 SPN inhibition in DMS. n = 6 mice

p-q. Same as in **i-j** but for unilateral D1 SPN inhibition in VLS. n = 5 mice. Inhibition of the left hemisphere (blue) over 5 days was followed by Inhibition of the right hemisphere (yellow). Note that the effect became weaker over days but recovered after switching the hemispheres. Thus, the adaptation to optogenetic manipulation is local to each hemisphere.

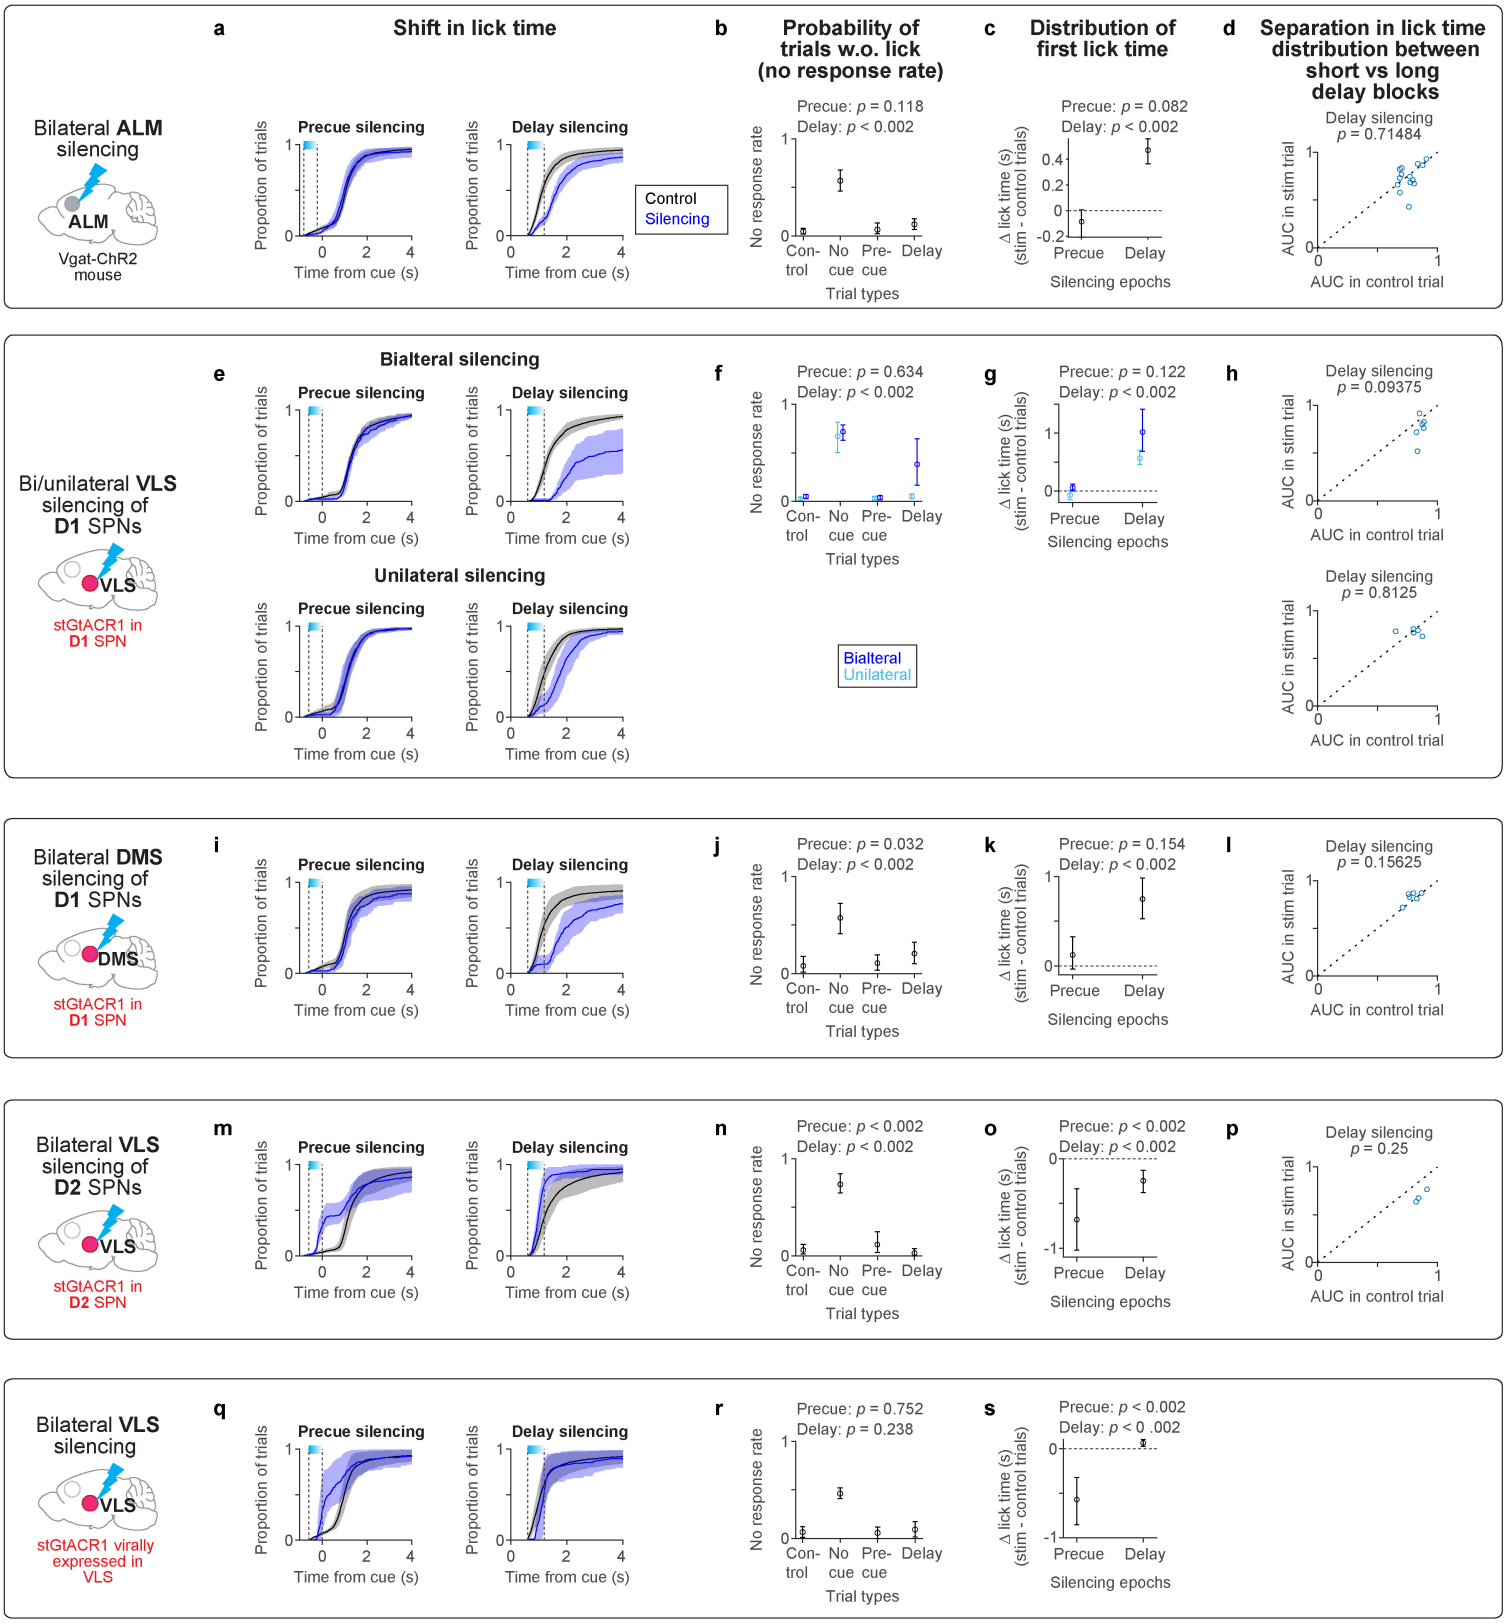

Extended Data Fig. 8

Extended Data Fig 8. Summary of the behavioral effects observed across optogenetic manipulations

a. Lick time distribution during precue or delay ALM silencing, duplicated from Figure 5a for comparison. n = 14 mice.

b. No-response rate in control, no cue, precue silencing, and delay silencing trials (Methods). No cue trial refers to randomly interleaved trials without a cue, serving to monitor the spontaneous lick rate not triggered by a cue. This provides the upper bound of the no-response rate. *P*-value, hierarchical bootstrap with a null hypothesis that no-response rate in control trials is the same as in silencing trials. Error bars, 95% confidence interval.

c. Shift in median lick time caused by precue or delay silencing. Duplicated from Figure 5b for comparison.

d. ALM silencing does not affect the separation in lick time distribution between delay blocks (short vs. long delay blocks in the switching delay condition). First, we calculated the area under the curve (AUC) of the ROC analysis comparing lick time distributions in short vs. long delay blocks. Then, we compared the AUC values (indicating how well the lick time distributions between delay blocks are separated) between control and photostimulation trials. Circles, individual animals. *P* value, signed rank test. The absence of change in the separation of lick time distributions suggests that ALM silencing does not erase the information of intended lick time, consistent with the recovery of dynamics after the silencing (Fig. 5).

e-h. Same as in **a-d** but for D1 SPN inhibition in VLS. Top, bilateral D1 SPN inhibition in VLS (n = 6 mice). Duplicated from Figure 6c for comparison. Bottom, unilateral inhibition (n = 5 mice).

i-l. Same as in**a-d** for D1 SPN inhibition in DMS. n = 6 mice.

m-p. Same as in **a-d** for D2 SPN inhibition in VLS. n = 3 mice.

q-s. Same as in **a-c** for cell-type-nonspecific striatal silencing in VLS. n = 2 mice.

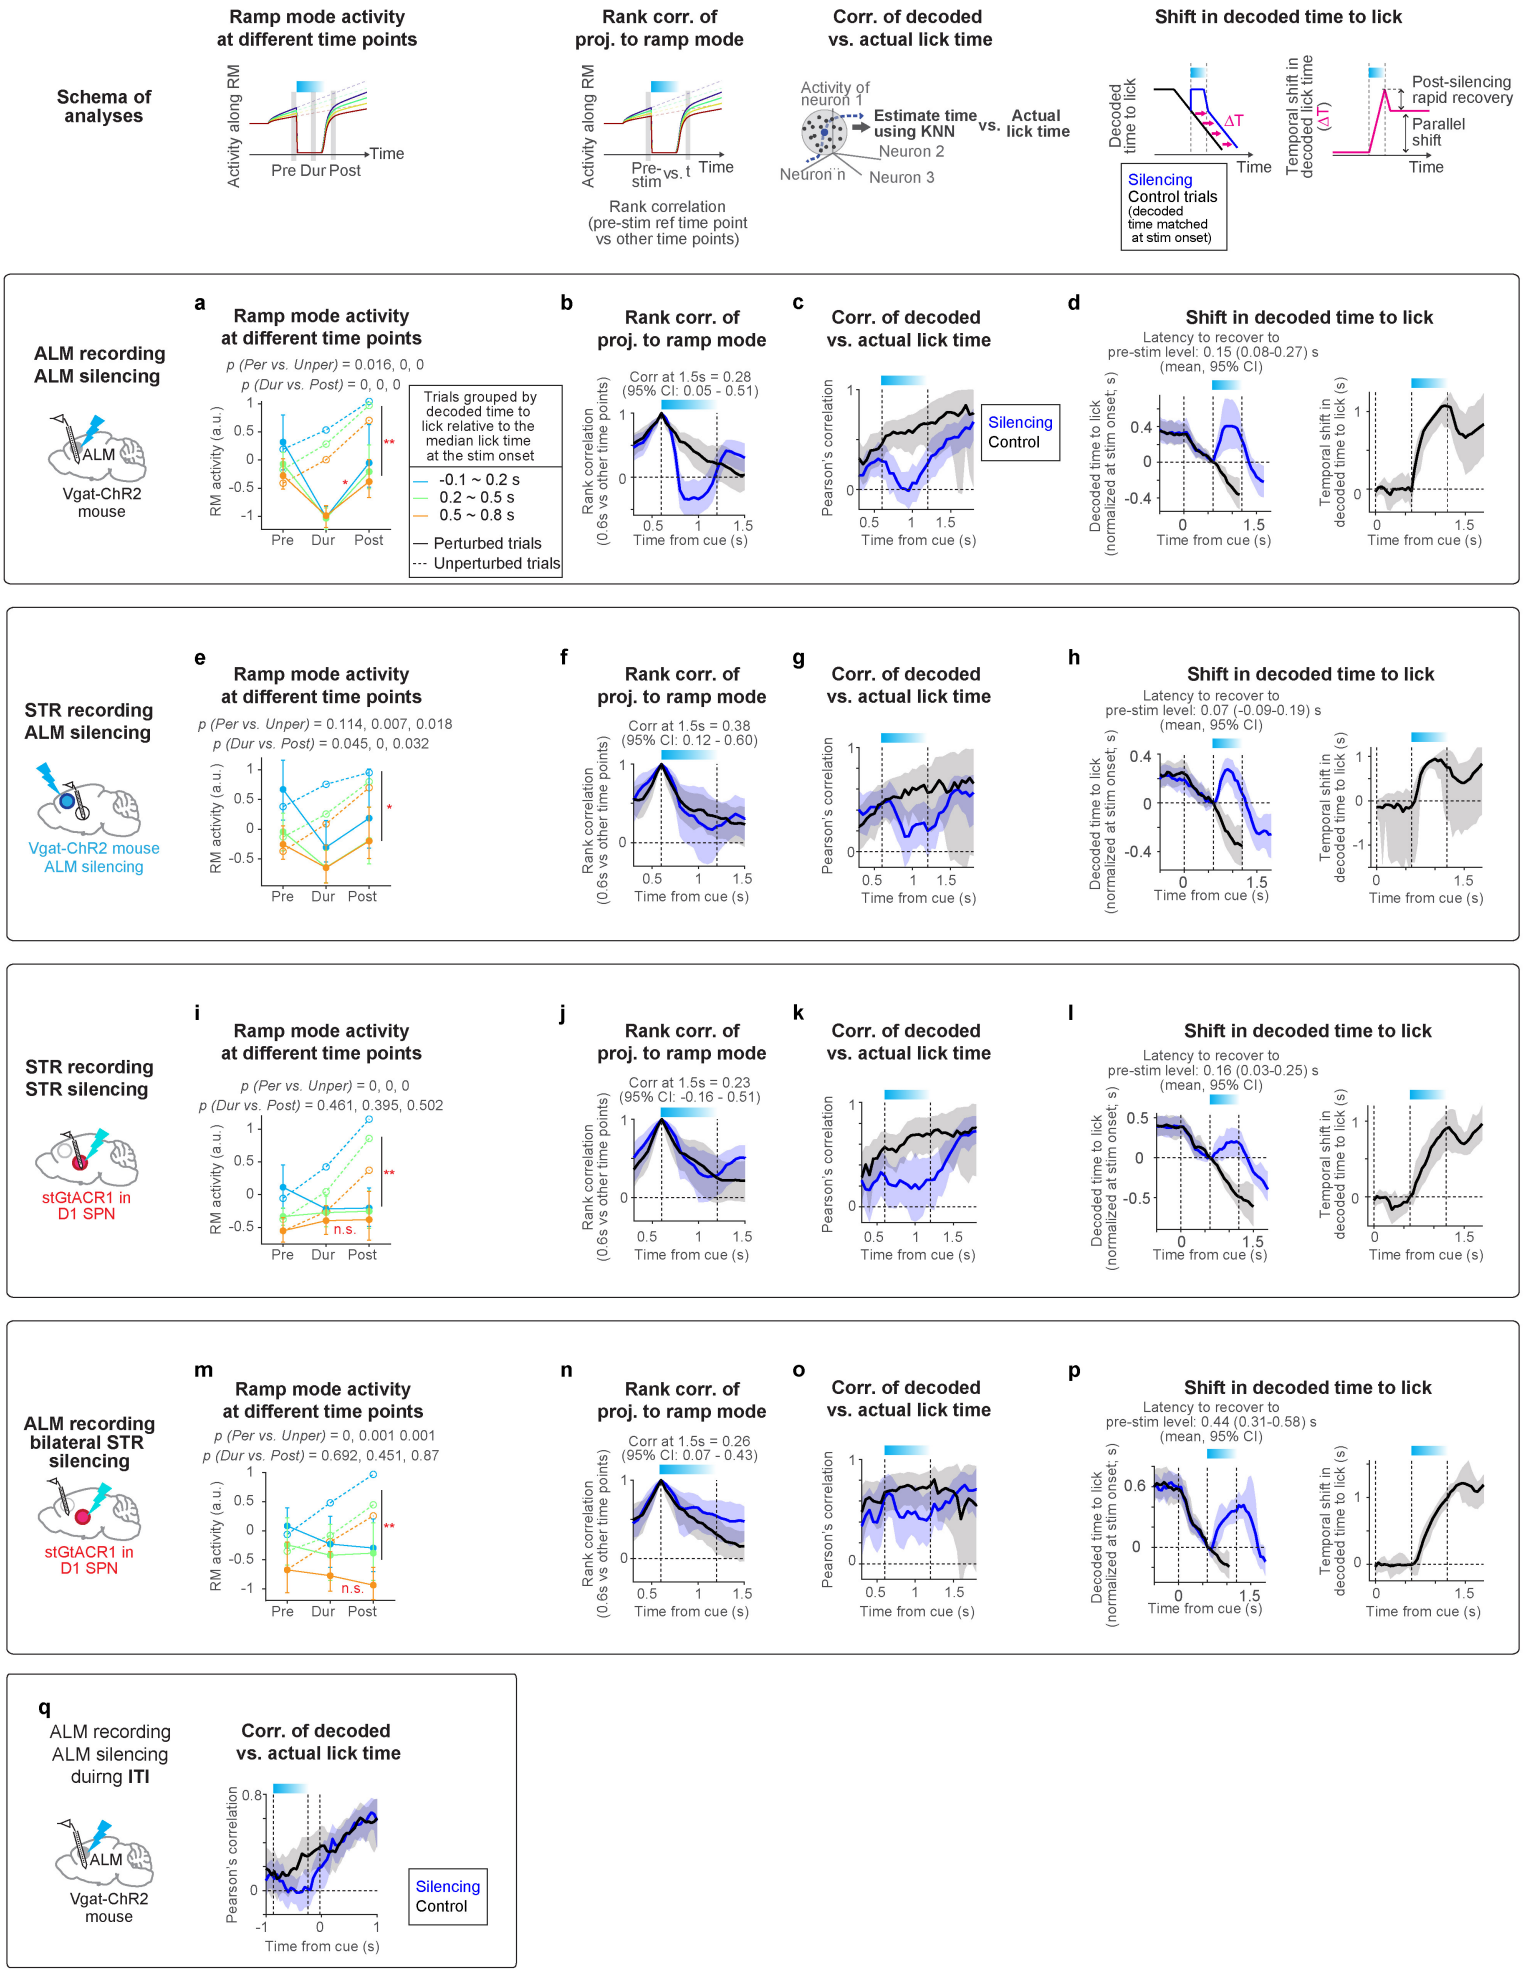

Extended Data Fig. 9

Extended Data Fig 9. Quantification of timing dynamics across perturbation experiments

In addition to the ramp mode (RM) activity shown in Fig. 5 and 6, we present additional quantifications of the timing dynamics across perturbation experiments. Firstly, we quantified the change in RM activity during and after transient perturbation. Upon ALM perturbation, the RM activity in both areas showed a V-shaped profile, reflecting rapid silencing and rapid recovery of the activity (**a** and **e**). This contrasts with the linear decay profile observed in both areas upon D1-SPN inhibition (**i** and **m**). Secondly, the rank correlation of projection along the ramp mode across trials collapsed in ALM during ALM silencing (**b**), but not in other conditions (**f, j, and n**). Thirdly, the correlation of predicted upcoming lick time (based on kNN decoder) versus actual lick time collapsed in ALM during ALM silencing (**c**), but not in other conditions (**g, k, and o**). Fourthly, the decoded time to lick was shifted during perturbation and the shift persisted after the perturbation across conditions (**d, h, l, and p**). In contrast to silencing during the delay epoch, silencing before the cue was followed by a recovery of decoded lick time (**r**).

a. Quantification of the change in RM activity during and after ALM silencing. RM activity is shown before (0.6 s after the cue, 'Pre'), during (0.9 s after the cue, 'Dur'), and after (1.25 s after the cue, 'Post') silencing. Perturbed and unperturbed trials with matched $T_{to\ lick}$ at the silencing onset (0.6 s after the cue) are shown. Different colors, trials with different decoded $T_{to\ lick}$ relative to the median lick time. Dotted lines, unperturbed trials. Solid lines, perturbed trials. *P*-values, hierarchical bootstrap with the null hypotheses as follows: *p (Per vs. Unper)*, tests the null hypothesis that activity in perturbed trials is higher than in matched unperturbed trials at post-silencing; *p (Dur vs. Post)*, tests the null hypothesis that in perturbed trials, activity at post-silencing is lower than during silencing. This tests whether there is a recovery of activity from the middle of the silencing to right after (0.05 s after) the end of silencing. P-values for each trial type are shown (-0.1~0.2, 0.2~0.5, 0.5~0.8, respectively). **, $p < 0.005$ for all trial types; *, $p < 0.05$ for all trial types; n.s., non-significant for all trial types. Error bars, 95% confidence interval (hierarchical bootstrap).

b. Rank correlation of ALM activity along the ramp mode across trials. The rank order at the pre-silencing condition (0.6 s after the cue) is compared with that at other time points. Shades, 95% confidence interval (hierarchical bootstrap). The rank correlation in silenced trials (blue) collapsed during the silencing but recovered to the control (black) level afterward.

c. Pearson's correlation of lick time decoded by kNN decoder (based on ALM population activity) vs. the actual lick time. Shades, 95% confidence interval (hierarchical bootstrap). ALM activity predicted upcoming lick time before and after the silencing, but such correlation disappeared during the silencing. **b and c** imply that after the silencing the time information recovered.

d. Decoded $T_{to\ lick}$ from each time point based on kNN decoding analysis of population activity (left). Shades, 95% confidence interval (hierarchical bootstrap). The decoded $T_{to\ lick}$ was normalized by subtracting the decoded $T_{to\ lick}$ at stimulus onset to account for different lick times across trials and sessions. Across conditions, the decoded $T_{to\ lick}$ shifted during the perturbation, and the shift persisted after the perturbation. The temporal shift in decoded $T_{to\ lick}$ between control and silencing conditions (right) shows that immediately after silencing, the temporal shift rapidly decayed, reflecting the rapid recovery of activity (schema on top). After this rapid phase, the temporal shift remains constant, implying that the decoded $T_{to\ lick}$ is shifted in parallel between control and silencing trials.

e-h. Same as in **a-d** but for striatum recording during ALM silencing. **e** is duplicated from Figure 5k for comparison.

i-l. Same as in **a-d** but for striatum recording during unilateral D1 SPN inhibition in VLS. **i** is duplicated from Figure 6j for comparison.

m-p. Same as in **a-d** but for ALM recording during D1 SPN inhibition in VLS. **i** is duplicated from Figure 6m for comparison.

q. Same as in **c** but for ALM recording during ALM precue silencing. Note that the correlation decays during the silencing, but recovered after the silencing.

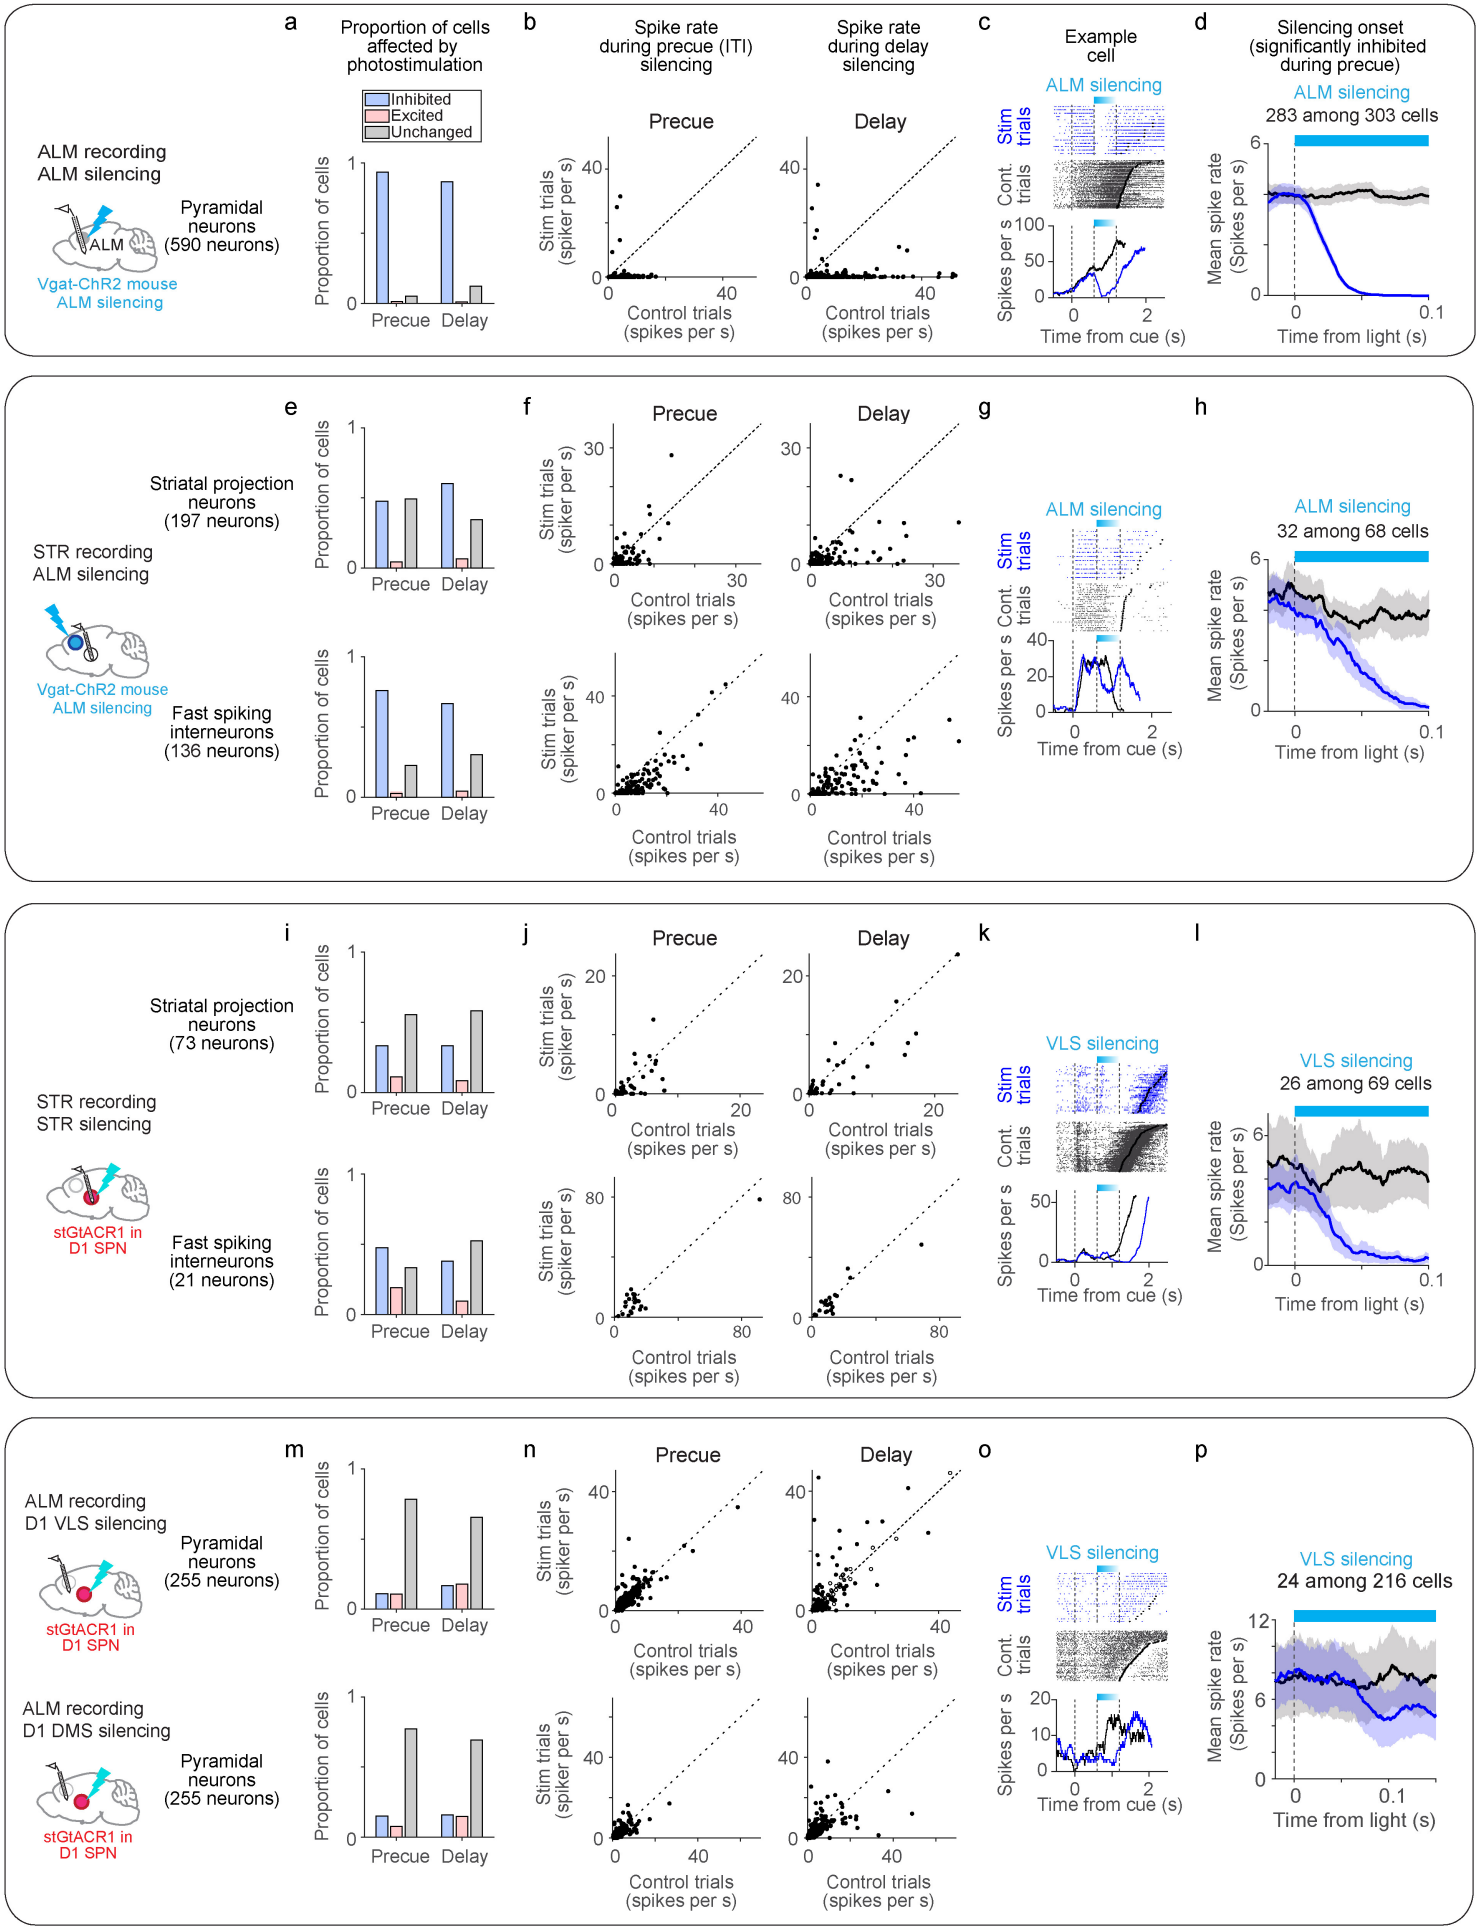

Extended Data Fig. 10

Extended Data Fig 10. Summary of optogenetic effect on spiking activity

We characterized the effects of different optogenetic manipulations on spiking activity, including the proportion of inhibited/excited/unaffected cells (**a, e, i,** and **m**), changes in spike rate (**b, f, j,** and **n**), and the onset of the optogenetic effect in significantly inhibited neurons (analyzed during the precue epoch to avoid confounds caused by behavior-related activity; **d, h, l,** and **p**).

a. The proportion of ALM pyramidal neurons affected by bilateral ALM silencing was assessed. Neurons were categorized as inhibited or excited ($p < 0.05$) or unchanged ($p > 0.05$) using a rank sum test. Cells with a mean spike rate higher than 1 Hz in control trials during the photostimulation window were considered. Duplicated from Figure 5c for comparison. Units from the first two manipulation sessions were included.

b. Spike rate of ALM neurons during precue (left) or delay silencing (right) compared to control unperturbed trials. Filled circles indicate significantly affected cells ($p < 0.05$, rank-sum test). The spike rate was calculated between 50-250 ms from the photostimulation onset.

c. An example ALM neuron with a reduced spike rate during delay ALM silencing. Top, spike raster. Blue, silencing trials. Black, control trials. Bottom, PSTH up to the median lick time for each trial type. The spike rate was smoothed using a 200 ms boxcar causal filter.

d. The onset of the optogenetic effect in ALM neurons with a significant reduction in spike rate during ALM silencing ($p < 0.05$, rank sum test). We analyzed silencing during the precue epoch to avoid confounds caused by behavior-related activity. Neurons with more than 10 trials and a spike rate higher than 1 Hz in the control condition were considered. The spike rate was smoothed using a 30 ms boxcar causal filter and aligned to the photostimulation onset. Blue, silencing trials. Black, control trials. Shading, SEM (hierarchical bootstrap).

e-h. Same as in **a-d** but for striatal recording during ALM silencing. Top, striatal projection neurons. Bottom, striatal fast-spiking interneurons. Units from the first two manipulation sessions were included.

i-l. Same as in **a-d** but for striatal recording during D1 SPN unilateral inhibition. Top, striatal projection neurons. Bottom, striatal fast-spiking interneurons. Data from the first session was included for **i-k**. However, for **l**, data was pooled across 3 sessions, as the onset of silencing did not appear to change across days, despite significant behavioral changes likely caused by rebound or other activity changes after the silencing.

m-p. Same as in **a-d** but for ALM recording during D1 silencing in VLS (top), or in DMS (bottom). Units from the first manipulation session were included.

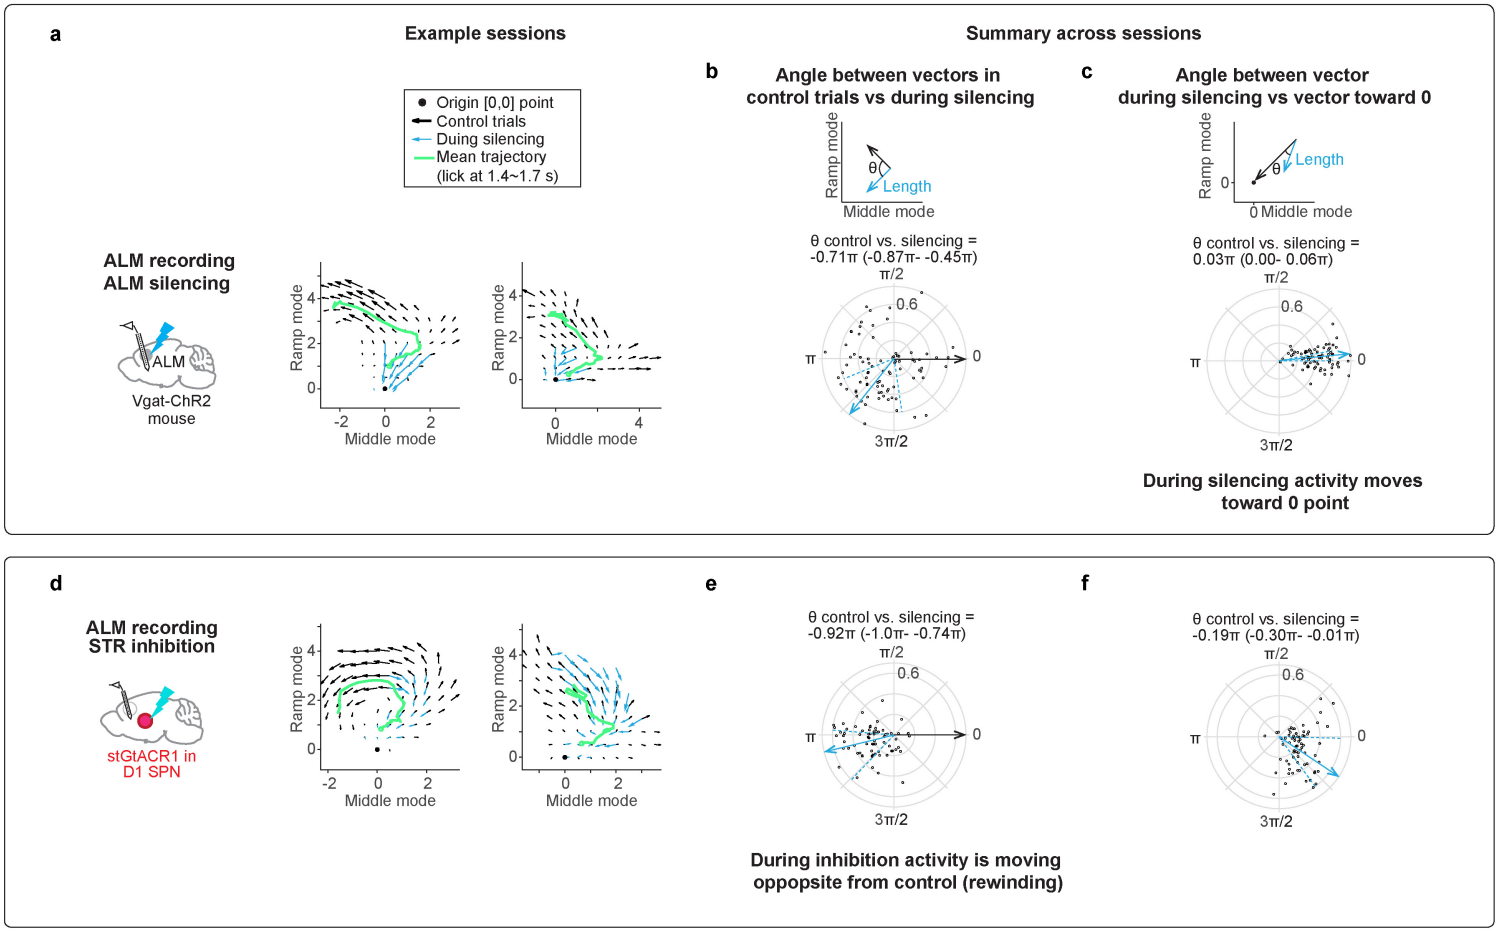

Extended Data Fig. 11

Extended Data Fig 11. Rewinding of ALM dynamics during D1-SPN silencing in VLS

We analyzed how activity evolves in the two-dimensional space defined by RM and MM, which captures a large fraction of task-modulated activity (Extended Data Fig. 5). For each session, we calculated the direction in which the activity developed from individual states (arrows in the two-dimensional space). Specifically, we pooled all time points (50 ms bin) between cue and lick in control trials, calculated the direction of activity evolution (in the next 50 ms bin), and averaged these to acquire the vector field per session (black arrows in **a** and **d**; for states with more than 30 data points). Similarly, we acquired the vector field during inhibition by pooling all time points during inhibition in silenced trials. We then calculated the angle between vectors in control versus inhibition across all states and sessions (circles in **b** and **e**) to test whether activity evolved in the opposite direction from the normal trajectory during inhibition. Additionally, we calculated the angle between vectors during silencing and the vector toward the zero point to test whether activity evolved toward the zero point. During ALM silencing, the trajectory is better explained by activity moving toward zero (compare **b** vs. **c**). In contrast, during D1-SPN silencing, the trajectory is better explained by activity moving in the opposite direction from the normal trajectory (compare **e** vs. **f**). In **b, c, e,** and **f**, the angle indicates the mean and 95 % confidence interval.

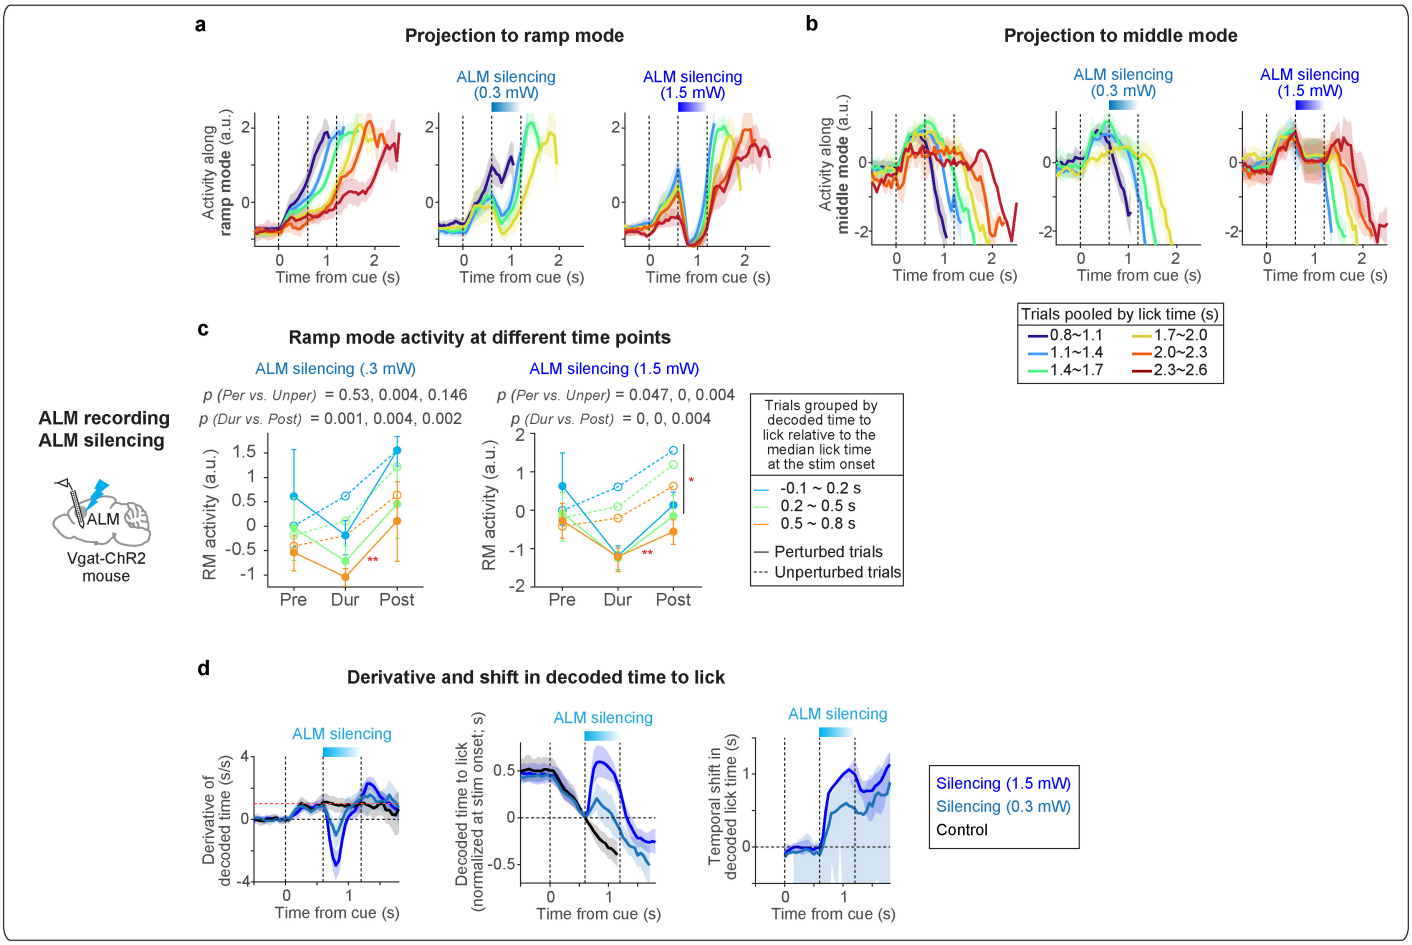

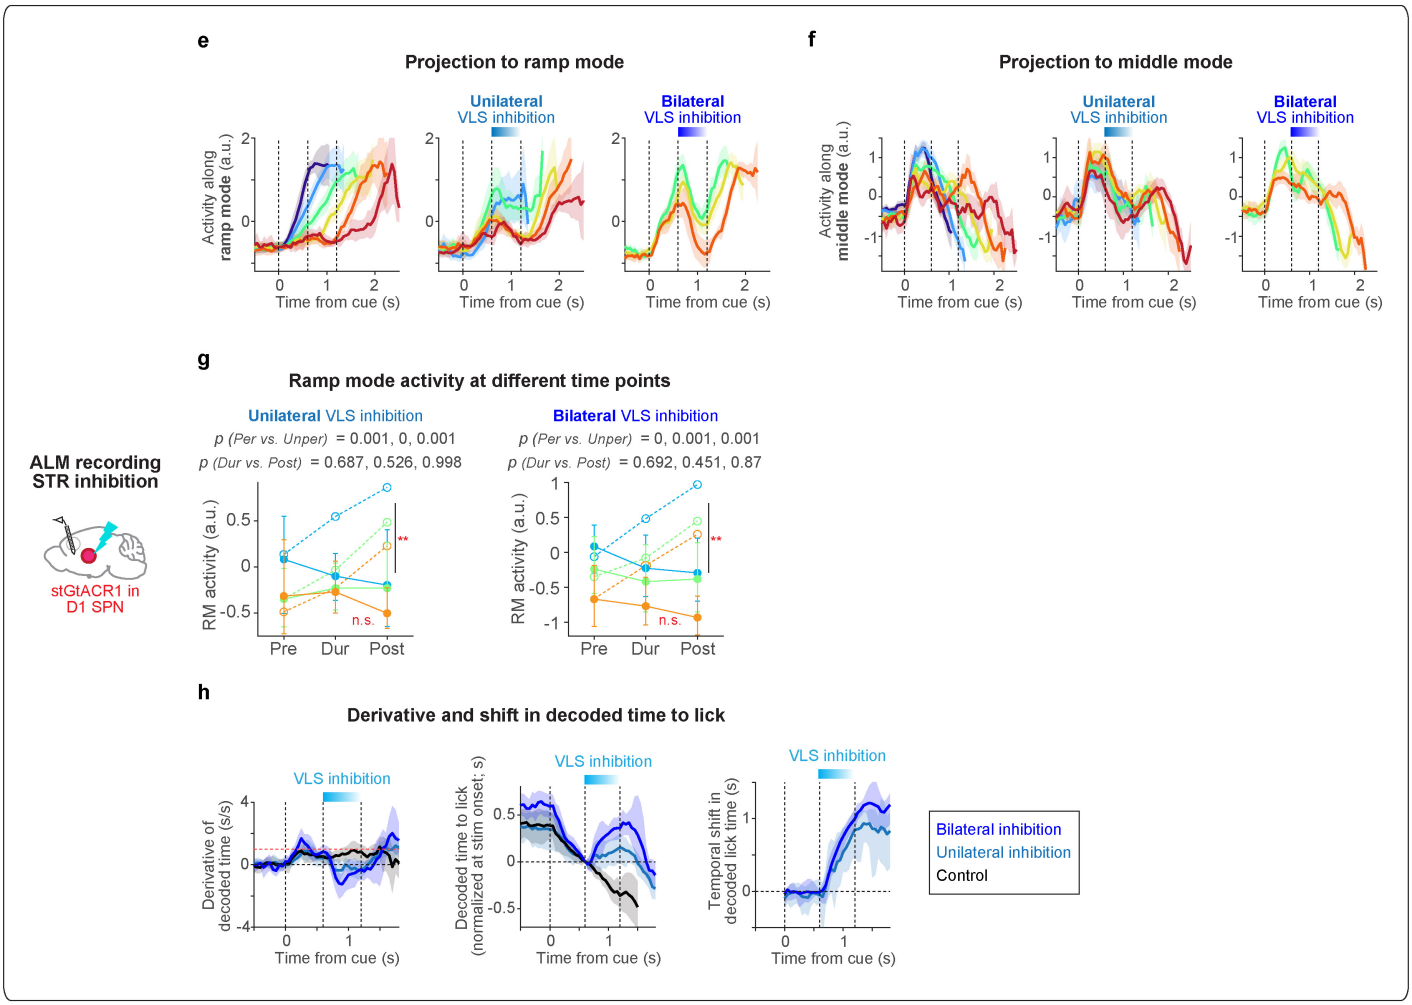

Extended Data Fig. 12

Extended Data Fig 12. Graded and persistent effect of perturbations

We recorded ALM activity while perturbing it with different laser powers (0.3 vs 1.5 mW; **a-d**) or while unilaterally (ipsilateral) or bilaterally inhibiting D1-SPN in the VLS (**e-h**). In both cases, the perturbation with different intensities resulted in graded and long-lasting changes in ramping activity and decoded time beyond the duration of perturbations, consistent with the notion that the perturbation was integrated into the timing dynamics.

a. ALM population activity along the ramp mode in control trials (left) and ALM silencing trials with different powers (middle, right). Shades, 95% confidence interval (hierarchical bootstrap). n = 130 cells, 4 mice. Trials with different laser powers were randomly interleaved in the same sessions.

b. Same as in **a** but for middle mode.

c. Quantification of the change in RM activity during and after ALM silencing. Same format as in Extended Data Fig. 9a, e, i, and m. Left, weak ALM silencing. Right, strong ALM silencing. Note that the plot has a V-shaped profile regardless of laser power, whereas in D1-SPN inhibition, it showed a linear decay profile regardless of whether the inhibition was unilateral or bilateral.

d. Analysis of decoded time based on ALM population activity (using kNN decoder). Left: derivative of decoded time (computed over a 200 ms window). In control trials (block), the derivative changes from 0 to 1 (red horizontal dashed line) after the cue. In stimulated trials, the derivative becomes negative at light onset, reflecting rapid decay in activity, and increases after the light, reflecting rapid recovery. Middle and right: same format as in Extended Data Fig 9d. Shades, 95% confidence interval.

e-h. Same as in **a-d** for unilateral vs bilateral inhibition of D1-SPNs in VLS. Data from different mice were combined (bilateral data is duplicated from Extended Data Fig. 9m; unilateral data, n = 112 cells, 4 mice). Unlike ALM silencing, the derivative of decoded time does not show a large change even during inhibition, reflecting a gradual change in timing dynamics.

METHOD DETAILS

EXPERIMENTAL MODEL AND SUBJECT DETAILS

Mice

This study is based on both adult male and female mice (age > P60). We used five mouse lines: C57Bl/6J (JAX# 000664), VGAT-ChR2-EYFP[107] (JAX #14548), Drd1-cre FK150[85], Adora2-cre KG126[85], R26-LNL-GtACR1-Fred-Kv2.1[108] (JAX #33089). See Extended Data Table 1 for mice used in each experiment.

All procedures were in accordance with protocols approved by the MPFI IACUC committee. We followed the published water restriction protocol[109]. Mice were housed in a 12:12 reverse light: dark cycle and behaviorally tested during the dark phase. A typical behavioral session lasts between 1 and 2 hours. Mice obtained all of their water in the behavior apparatus (approximately 0.6 ml per day). Mice were implanted with a titanium headpost for head fixation[109] and single-housed. For cortical photoinhibition, mice were implanted with a clear skull cap[62]. For D1/D2 SPN silencing, tapered fiber optics (1.0 mm taper, NA 0.37, core diameter 200 µm, Doric lenses) were bilaterally implanted during the headpost surgery. For bilateral striatal manipulations, tapered fiber optics were implanted around the following target coordinates (relative to Bregma and dura, in mm; Extended Data Fig. 6c): AP -0.3, ML ± 3, DV 3.5 for the ventral lateral striatum (VLS); AP 0.6, ML ± 1.5, DV 3 for the dorsal medial striatum (DMS). Craniotomies for recording were made after behavioral training.

Viral injection

To virally express stGtACR1 in the striatum (Extended Data Fig. 8q), we followed published protocols (dx.doi.org/10.17504/protocols.io.bctxiwpn) for virus injection. AAV2/5 CamKII-stGtACR1-FusionRed (titer: 9.5×10^12) was injected into AP -0.3 mm (posterior to Bregma), ML 3 mm, DV 2.75 and 3.5 mm, 100nl each depth. The same tapered fiber optics described above were bilaterally implanted at DV 3.5 mm.

Behavior

At the beginning of each trial, an auditory cue was presented, which consisted of three repeats of pure tones (3 kHz, 150 ms duration with 100 ms inter-tone intervals, 74 dB). A delay epoch started from the onset of the cue presentation. Licking during the delay epoch aborted the trial without water reward, followed by a 1.5 s timeout epoch. Licking during the 10 s answer epoch following the delay was considered a 'correct lick', and a water reward (approximately 2 µL/drop) was delivered immediately, followed by a 1.5 s consumption epoch. If mice did not lick during the 10 s answer period, the trial would end without a reward. Trials were separated by an inter-trial interval (ITI) randomly sampled from an exponential distribution with a mean of 3 s with 1 s offset (with a maximum ITI of 7 s). This prevented mice from predicting the trial onset without cue. Animals had to withhold licking during the full ITI epoch for the next trial to begin (otherwise the ITI epoch will be repeated). In approximately 10% of randomly interleaved trials, the auditory cue was omitted to assess spontaneous lick rate ('no cue' trials). No water reward was delivered in no cue trials.

We followed the protocol described in Majumder et al, 2023[110] for training. In brief, the delay duration increased from 0.1 s to 1.8 s gradually based on the animal's performance[110]. Once mice reached 1.8 s delay, we started either the switching delay (Extended Data Fig. 2a), the random delay (Extended Data Fig. 2b), or the constant delay conditions (Extended Data Fig. 2c). In the switching delay condition, we switched delay between 1 vs. 3 s or 1 vs. 1.8 s every 30 ~ 70 trials (the number of trials was randomly selected and not contingent upon behavior). Similarly, in the random delay condition, we randomly switched delay among 0.5, 1.0, 1.5, 2.0, 3.0, or 5.0 s every 30 ~ 70 trials. For the constant delay condition, mice were trained with a constant delay of 1.5 s across sessions for at least two weeks. Otherwise, the task design and reward contingency remained the same. To avoid human bias, the behavior was automatically controlled by Bpod (Sanworks) and custom MATLAB codes.

Optogenetics

Photostimulation was deployed on < 25% in randomly selected trials. To prevent mice from distinguishing photostimulation trials from control trials using visual cues, a 'masking flash' (1 ms pulses at 10 Hz) was delivered using 470 nm LEDs (Luxeon Star) throughout the trial. For both ChR2 and stGtACR1, we used a 488 nm laser (OBIS 488 - 150C, Coherent).

The ChR2-assisted photoinhibition of the dorsal cortices was performed through clear-skull cap[62] (Fig. 2e) or craniotomy (in case of simultaneous recording; Fig. 5). We scanned the 488 nm laser light using Galvo mirrors. We stimulated GABAergic interneurons in Vgat-ChR2-EYFP mice starting at 0.6 s after the cue, lasting for 1.2 s (including 0.2 s ramping down; Fig. 2e) or 0.6 s duration (including 0.3 s ramping down; Fig. 5). Time-averaged laser power was 1.5 (or 0.3) mW per spot (8 spots in total: 4 spots in each hemisphere centered around the target coordinates with 1 mm intervals; We photoinhibited each spot sequentially at the rate of 5 ms per step). For Fig. 2e, the targeted brain area was randomly selected for each photostimulation trial.

To silence D1 or D2 expressing SPNs using stGtACR1 (Fig. 6), we delivered photostimuli (0.25 or 0.5mW, 488 nm) bilaterally (Fig. 6k-m) or unilaterally (in case of optrode; Fig. 6h-j) in the striatum starting 0.6 s after the cue and lasting for 0.6 s (including 0.3 s ramping down). The light was delivered through implanted fiber optics and intensity was measured at the fiber tip.

Extracellular electrophysiology

A small craniotomy (diameter, 0.5 - 1 mm) was made over the recording sites one day before the first recording session. Extracellular spikes were recorded using 64-channel two-shank silicon probes (H-2, Cambridge Neurotech) for ALM, and Neuropixels probe 1.0[80] for striatum. Voltage signals were multiplexed, recorded on a PCI6133 board (National instrument), and digitized at 400 kHz (14-bit). The signals were demultiplexed into 64 voltage traces sampled at 25 kHz and stored for offline analysis. All recordings were made with the open-source software SpikeGLX (http://billkarsh.github.io/SpikeGLX/). During recordings, the craniotomy was immersed in a cortex buffer (125 mM NaCl, 5 mM KCl, 10 mM glucose, 10 mM HEPES, 2 mM $MgSO_4$, 2 mM $CaCl_2$; adjust pH to 7.4). Brain tissue was allowed to settle for at least five minutes before recordings.

For the optrode recordings (Fig. 6h-j), we used 64-channel two-shank silicon optrodes with a 1.0 mm taper fiber optic attached adjacently, NA 0.22, core diameter 200 µm (Cambridge Neurotech). Optrode was acutely inserted each session and the light delivery protocol was identical to that used for behavioral experiments described in ***Optogenetics***. Neuropixels probe and optrode tracks labeled with CM-DiI were used to determine recording locations[111].

Histology

Mice were perfused transcardially with PBS, followed by 4% PFA / 0.1 M PBS. Brains were post-fixed overnight and transferred to 20 or 30 % sucrose PB before sectioning on a freezing microtome. To reconstruct recording tracks, we either generated coronal sections followed by conventional imaging (protocol described in Inagaki et al, 2022[112]), or cleared the brain followed by light-sheet microscopy. To clear the brain, we used the EZ Clear method[113]. We followed the previous protocol to map the recording tracks to Allen Common Coordinate Framework (CCF)[111]. Extended Data Fig. 6a is based on data collected in Guo et al, 2017[114].

QUANTIFICATION AND STATISTICAL ANALYSIS

Behavioral analysis

We analyzed the time of the first lick after the cue onset in each trial. Lick time was measured by detecting the tongue's contact with the lick port using an electrical lick detector. For optogenetic experiments (Fig. 5ab, 6cd, and Extended Data Fig. 8), we analyzed trials with the first lick time occurring after the onset of photostimulation in both control and photostimulated trials to compare the effect of photostimulation on behavior. The no-response rate (Fig. 6d, right, and Extended Data Fig. 8) was calculated as the probability of mice not responding within 5 s after the cue. The shift in lick time (Δ lick time; Fig. 5b, 6d, and Extended Data Fig. 8) was based on the median lick time. Post-stim lick rate (Extended Data Fig. 7d-g) was calculated as the probability of mice licking within 0.6 s after the photostimulation offset. To analyze behavior while the mice were engaged in the task, we analyzed all trials between the first occurrence of five consecutive cue trials with licks and 20 trials before the last occurrence of three consecutive no-response trials.

To analyze whether optogenetic manipulation affects the separation of lick time distributions between delay blocks in the switching delay condition (short, 1 s, vs. long delay, 1.8 or 3 s, blocks; Extended Data Fig. 8d, h, l, and p), we performed a receiver operating characteristic (ROC) analysis. First, we conducted ROC analysis to distinguish lick time distributions between the two delay blocks (for control and photostimulation trials separately). We then analyzed the area under the curve (AUC) as a quantification of separation and compared the AUC values between control and photostimulation trials.

Due to the attenuation of behavioral effects of optogenetic manipulation (Extended Data Fig. 7j-q), we restricted analyses of both behavioral and physiological data to the first (for striatal manipulation) or the first two (for ALM manipulation) manipulation sessions per mouse. All analyses, including the calculation of confidence intervals and p-values, were performed using a hierarchical bootstrap, unless stated otherwise. First, we randomly selected animals with replacements. Second, we randomly selected sessions for each animal with replacement. For each selected session, we calculated the behavioral metrics described above. This procedure was repeated 1000 times to estimate the mean, confidence intervals, and statistics.

Behavioral simulation

We simulated the optimal lick time in the timing task (Extended Data Fig. 2g-k). Based on the mean and standard deviation (std) of the lick time distribution in the data (std lick time = -0.2 + 1.1 × mean lick time; Extended Data Fig. 2f), and assuming an inverse Gaussian distribution of lick time[101–103], we randomly sampled lick times in 1000 trials for each mean lick time. Mimicking our task structure (ITI = 2.3 s, consumption time of 1.5 s, and delay duration as indicated in the plot), and assuming mice remember whether they received a reward or not within a 30 s time window, we calculated the estimated reward amount per trial or time for each mean lick time over the 30 s window.

Trial-history regression analysis

For the linear regression analysis in Fig. 2d and Extended Data Fig. 2l, we tested 42 combinations of regressors with 1 - 6 lags. The median absolute deviation (MAD) of lick time explained by different regression models was calculated as 1 - R1 / R2, where R1 is the median of the absolute value of the model residuals, and R2 is the median of the absolute value of the null model residuals. To calculate MAD-explained, we performed 5-fold cross-validation.

Videography analysis

High-speed (300 Hz) videography of orofacial movement (side view) was acquired using a CMOS camera (Chameleon3 CM3-U3-13Y3M-CS, FLIR) with IR illumination (940nm LED). We used DeepLabCut[106] to track the movement of the tongue and jaw. Movements along the dorsoventral direction were analyzed and plotted in Extended Data Fig. 3. Trajectories were normalized by subtracting the mean position before the cue from each

trajectory and then dividing by the minimum value across trials (thus, the downward movement of the tongue and jaw looks upward in the plot). The onset of jaw movement in each trial is the first time point after the cue when the normalized movement trajectory exceeds 10% of the max value. The onset of tongue movement is when DeepLabCut first detects the tongue after the cue. In two out of 34 mice, we occasionally observed jaw movement within 200 ms of the cue. These trials were excluded from the analysis for the average jaw and tongue onset of individual mice (in Extended Data Fig. 3d), as these rare early subthreshold movements are likely startled responses to the cue.

Extracellular recording analysis

Spike sorting and cell type classification

JRClust[115] (https://github.com/JaneliaSciComp/JRCLUST) with manual curations was used for spike sorting. We used a combination of quality metrics (described in Majumder et al, 2023) to select single units for analysis[110]. Units with a total trial number of less than 75 were excluded from analyses.

For ALM recording, units with a mean spike rate above 0.5 Hz were analyzed. This includes 4467 putative pyramidal neurons (spike width >= 0.5 ms[62]) out of a total of 5093 neurons. For striatal recording, units within the striatum (regions annotated as "striatum", "caudoputamen", and "fundus of striatum" after registration to the Allen CCF) with a mean spike rate above 0.1 Hz were analyzed. This contains 1217 striatal projection neurons (spike width >= 0.4 ms and with post-spike suppression duration <= 40 ms), 584 fast-spiking interneurons (spike width < 0.4 ms and with less than 10% chance of having a long interspike interval), 127 tonically active neurons (TAN), and 44 other unidentified interneurons[69]. For the single session analyses (decoding and projection to modes), only putative pyramidal neurons were analyzed for the ALM recording, while all neurons were included for the striatal recording data. See Extended Data Table 1 for the number of recorded neurons in each experiment.

Similarity matrices

To plot the similarity matrix of population activity of recorded neurons in unperturbed trials (Fig. 3c and h), we calculated the mean spike activity of individual neurons across different lick time ranges to yield a population activity matrix, with the number of rows equal to the number of neurons and the number of columns equal to the number of time points (200 ms bin). Then, we calculated pairwise Pearson's correlation between these population activity matrices in trials with lick times between 1.40 - 1.55 s (reference trials) and the test trials with other lick time ranges. For Fig. 5 and 6, we compared the pairwise Pearson's correlation between unperturbed trials with lick times between 1.4 - 1.7 s (reference trials) and the trials with photostimulation in which lick time was after the end of photostimulation (Fig. 5f and i, and Fig. 6h and k). As a control, we selected unperturbed trials with lick times after the end of photostimulation, and randomly subselected the same number of trials as in the photostimulation condition. For each similarity matrix, we identified the points along the Y-axis with the maximum correlation (above 0.8) for each time point along the X-axis, and repeated this procedure with the hierarchical bootstrap (Fig. 3d and i, Fig. 5g and j, and Fig. 6i and l).

Single-cell analyses

For the peri-stimulus time histogram (PSTH) of example cells (Fig. 3a and f, Fig. 4a, and Extended Data Fig. 4d-f), we grouped trials based on lick time. PSTHs were calculated based on 1 ms time bin. The PSTHs were smoothed with a 200 ms causal boxcar filter unless specified otherwise. To temporally warp the PSTH for individual cells, we linearly scaled the spike timing after the cue, based on the time from cue to lick. Specifically, Spike $\text{time}_{\text{warped}}$ = Spike $\text{time}_{\text{original}}$ / ($\text{LT}_{\text{trial to be warped}}$ / $\text{LT}_{\text{target warp time}}$), where LT denotes the first lick time in each trial, and $\text{LT}_{\text{target warp time}}$ = 1 s.

Across-trial variance (Extended Data Fig. 4a-c) was calculated as the variance of spiking activity across trials during the first 1s after the cue for the original or temporally warped data (the across-trial variance was calculated for five 200 ms time windows within the 1 s and then averaged). To temporally warp the population activity along certain modes, we linearly scaled the activity in the time range of $T_{cue} > 0.1$ by $(LT_{target} - 0.1)/(LT_{trial\ to\ be\ warped} - 0.1)$, where T_{cue} denotes time after the cue[110].

To calculate the proportion of cells affected/unaffected by photostimulation (Fig. 5c, 6e, and Extended Data Fig. 10), we analyzed the spikes within the time window of 50 - 250 ms from the photostimulation onset. To quantify the effect of photostimulation, trials with licks before the photostimulation onset were excluded from the analysis. For individual cells, the spike rate in control and photostimulation trials was compared using the rank sum test. Cells with a mean spike rate above 1 Hz during this window and more than 10 trials per condition were analyzed.

In Extended Data Fig. 4, we analyzed the partial rank correlation between the spike rate (in specific time windows) and the lick time in previous trials (without reward), removing the effect of upcoming lick time, for each cell (Extended Data Fig. 4g-k; we only analyzed previous unrewarded trials to avoid confound caused by the representation of rewards; analysis of previous rewarded trials yielded similar results). Specifically, we calculated the rank correlation between spike rate (R) vs. previous lick time (P) (ρRP), rank correlation between R vs. upcoming lick time (U) (ρRU), and rank correlation between P and U (ρPU). Then the partial correlation between spike rate vs. lick time in the previous trial removing the effect of upcoming lick time is as follows:

$$\rho RP \cdot U = \frac{\rho RP - \rho RU \cdot \rho PU}{\sqrt{1-\rho^2 RU}\sqrt{1-\rho^2 PU}}$$

The distribution of rank correlation was quantified by the standard deviation (the higher the value, the broader the distribution of rank correlation, indicating more cells encoding trial history). As controls, we performed a trial shuffle test, which shuffles the trial order and destroys trial history, and a session permutation test to avoid the confound of nonsensical correlations[72]. Specifically, we used behavioral data from randomly selected sessions to analyze the ITI activity in a given session and repeated this procedure 1000 times.

Dimensionality reduction

We characterized population activity patterns between the cue and the lick by defining modes that differentiate the baseline activity during the inter-trial interval (ITI; 0 - 1 s before the cue) from the activity during specific 300 ms time windows after the cue: 0 - 0.3 s after the cue (cue mode, **CM**), 0.5 - 0.8 s before the lick (middle mode, **MM**), 0.2 - 0.5 s before the lick (ramp mode, **RM**), and 0 - 0.3 s after the lick (execution mode, **EM**).

Specifically, to calculate **RM** for a population of n recorded neurons, we looked for an $n \times 1$ unit vector that maximally distinguished the mean activity before the trial onset (0 - 1 s before cue; $\boldsymbol{r}_{before\ cue}$) and the mean activity before the first lick (0.2 - 0.5 s before the first lick; $\boldsymbol{r}_{before\ lick}$) in the *n-dimensional* activity space. We defined a population ramping vector: $\boldsymbol{w} = \boldsymbol{r}_{before\ lick} - \boldsymbol{r}_{before\ cue}$. **RM** is $\boldsymbol{w}$ normalized by its norm. Similarly, we defined **CM, MM**, **EM** using different time windows, and **MM** was orthogonalized to **RM**, and **CM** was orthogonalized to both **MM** and **RM** using the Gram-Schmidt process. Thus, the upper limit of the sum of variance explained (**CM**+**MM**+**RM**) in Extended Data Fig. 5 is 1. **EM** was orthogonalized to **RM** (Extended Data Fig. 4l3m3).

To calculate the trial-history mode, we first calculated the rank correlation between the ITI activity (0 - 1 s before the cue) and the predicted lick time across trials for each neuron. To predict the lick time for each trial, we used the linear regression model described in Fig. 2d following 5-fold cross-validation. We obtained an n × 1 unit vector representing the rank correlation of each neuron and normalized it by its norm to calculate trial-history mode.

In Fig. 3, 4, and Extended Data Fig. 4l-o and 5, we have pooled cells recorded across sessions (i.e., pseudo-sessions). For each cell, we randomly selected 50 unperturbed control trials to define the mode. These unperturbed trials met the following criteria: the first lick occurred within 1 to 3 s after the cue, and there were no licks 3 s before the cue onset. Then, we selected a different set of trials to project the activity along these modes. Only neurons with more than 10 trials within all six lick time ranges were included. The six lick time ranges are: 0.80 - 1.10, 1.10 - 1.25, 1.25 - 1.40, 1.40 - 1.55, 1.55 - 1.70, and 1.70 - 2.00 s.

To calculate the variance of spiking activity explained (variance explained) by individual modes (Extended Data Fig. 5), we calculated the squared sum of the activity along individual modes after subtracting the baseline activity (0 - 0.2 s before the cue), and then divided that by the squared sum of the spike rate across neurons after subtracting the baseline activity. To calculate the variance explained by the sum of **CM**, **MM**, and **RM** reported in the main text, we calculated the variance explained between 0.2s from cue (around when task modulation started) to lick for individual lick time ranges and then averaged across them. We calculated the variance explained by trial-history mode activity similarly but without subtraction of the baseline activity (Extended Data Fig. 4l-n). In Extended Data Fig. 4o, We performed a linear regression analysis between trial-history mode activity during 0 - 1 s before the cue and the upcoming lick time for each iteration of the hierarchical bootstrap. We then plotted the distribution of the linear regression coefficient across these iterations as a cumulative distribution function.

Single session analyses

For single-session analyses (Fig. 5e, h, and k, 6g, j, and m, Extended Data Fig. 9, 11, and 12), sessions with more than 300 trials and five neurons were analyzed. Spiking activity was binned per 50 ms time window. Activity between 1 s prior to the cue and the first lick in each trial was analyzed (i.e., post-lick activity was excluded as we focused on timing dynamics prior to the first lick). Dimensionality reduction was performed in the same manner as in the pseudo-session analysis, but modes were defined individually for each session.

To decode the time to lick ($T_{\text{to lick}}$) from simultaneously recorded neural population activity, we conducted a k-nearest neighbor (kNN) regression analysis. Within each experimental session, trials were partitioned into two sets: a test set comprising randomly selected 100 unperturbed trials and all perturbed trials, and a training set consisting of the remaining trials. For each moment in a test trial (50 ms window), we searched the training set to identify k data points with the most similar population activity patterns (Mahalanobis distance based on the top principal components explaining 90% of variance). We averaged the lick times in trials, from which these k data points came, to estimate the lick time of the test set. This analysis allowed us to predict lick time at the trial onset significantly better than trial shuffle controls (Extended Data Fig. 6g). We tested "k" values between 20-50 (which are close to the square root of the number of data points in the training dataset) and found they yielded similar results and did not change conclusions (data not shown). In the paper, we report the results with k = 30. Minor sessions showed low decodability due to a small number of recorded neurons, trials, and/or lack of task-modulated cells, these sessions were excluded from analyses (Extended Data Fig. 6g). Specifically, for Fig. 5h and k, 6j and m, and Extended Data Figs. 9, 11, and 12, we analyzed sessions in which the kNN decodability (Pearson's correlation between decoded lick time at the stim onset, i.e., 0.6 s after the cue vs. actual lick time) was higher than 0.25.

To analyze the effect of perturbations systematically, we compared unperturbed vs. perturbed trials (Extended Data Fig. 9a, c-d, e, g-h, i, k-l, m and o-q, and 12cd and gh) after matching the number of trials and decoded time at the onset of perturbation between conditions. Specifically, we randomly resampled animals, sessions, and trials hierarchically (hierarchical bootstrap; 1000 iterations). For each perturbed trial in each bootstrap iteration, we identified an unperturbed trial within the same session with the closest decoded time at the onset of perturbation (cue onset for cue manipulation and 0.6 s after the cue for optogenetics). Then, we pooled these trials. This procedure allowed us to examine how decoded time (and projection along each mode) changed after the perturbation

in conditions where their activity patterns were similar before the perturbation. The confidence intervals of projections in Fig. 5 and 6, and rank correlation along the ramp mode (Extended Data Fig. 9b, f, j, and n) were also based on hierarchical bootstrap (but without matching decoded time).

For the two-dimensional plots and vector field analysis (Extended Data Fig. 11), we analyzed how activity evolves in the two-dimensional space defined by RM and MM. Spiking activity was binned in 50 ms time windows, and activity between the cue and the first lick in each trial was analyzed. For each session, we projected the activity of ALM neurons along RM and MM. The projection was divided by the standard deviation of activity among control trials but was not subtracted by the mean so that 0 represents 0 spike activity. For individual data points (x) in control trials, we calculated the vector $r_x^{control}$representing the direction activity evolves in the next time point (50 ms time bin) in the two-dimensional state. Then, we calculated the mean vector for individual states in the two-dimensional space by averaging all vectors within a spatial bin of 0.5 along both the MM- and RM- axes (if the spatial bin contains more than 30 data points): $r_{XY}^{control}$, where X and Y denotes the location of the state along MM and RM axes, respectively Similarly, we acquired the vector field during inhibition by pooling all time points during inhibition (100 ms - 400 ms from the inhibition onset) in photostimulation trials to acquire r_{XY}^{stim}. Then, we calculate the direction between $r_{XY}^{control}$ and r_{XY}^{stim} for all states where both control and stim vectors exist. We have excluded points where $r_{XY}^{control}$ is within $\pi/6$ from tanh(Y/X) because if the activity is evolving against the zero point under control conditions, we cannot distinguish between whether during inhibition the activity is rewinding or moving toward the zero point.

Network models

Using a dynamical systems approach, we consider four variables representing the average membrane currents (*h*) and spike rates (*r = f(h)*, where *f(h)* is the neural activation function) of neuronal populations in two regions of the brain. Specifically, we modeled two neurons in each brain area. In this model, the membrane potential of neuron i, h_i(t), was governed by the following nonlinear differential equation:

$$\tau\frac{dh_i(t)}{dt} = -h_i(t) + \sum W_{ij} h_j(t) + I_i^{base}(t) + I_i^{ext}(t) + I_i^{stim}(t)$$

$$r_i(t) = f(h_i(t))$$

Where τ is the membrane time constant, *Wij* is the element of the connectivity matrix between presynaptic neuron j and postsynaptic neuron i, $I_i^{base}(t)$ is the baseline input current, $I_i^{ext}(t)$ is the external input current, and $I_i^{stim}(t)$ is the negative current mediated by optogenetics to neuron i. The membrane current h_i(t) was converted to the spike rate by applying a threshold-linear activation function *f(h) = max(h,0)*.

The baseline input currents were chosen so that the system displays a stable fixed point at a low spike rate (lower attractor) with a spike rate of 5 spikes per second, consistent with the baseline firing rate observed in the experimental data). The connectivity matrix *W* and the external input $I^{stim}(t)$ are shown in Extended Data Fig. 1. For each area, we defined the ramp mode using the same criteria as in the experimental data, and we then plotted spike rate activity along the ramp mode (Extended Data Fig. 1).

We tested models with different connectivity matrices reflecting distinct computational roles of ALM and striatum (Extended Data Fig. 1). In the *externally driven* model (Extended Data Fig. 1a), both ALM and striatum received a ramping input that scaled with the desired lick times, progressively shifting the location of the fixed point in time. In the *distributed* model (Extended Data Fig. 1b), integration was achieved via interareal connections between ALM

and striatum; in the absence of these long-range connections, neither ALM nor striatum displayed slow temporal dynamics. Conversely, in the *redundant* model, ALM and striatum implemented two identical integrators (Extended Data Fig. 1c). Although weakly connected, their behavior was independent of each other's input. In the *specialized ALM integrator* model (Extended Data Fig. 1d), ALM served as the integrator while the striatum followed ALM dynamics. The roles of ALM and striatum were inverted in the *specialized striatum leaky integrator* model (Extended Data Fig. 1e), where the striatum integrated the stimuli but faced substantial leaking, while ALM's role was limited to providing excitatory input to the striatum. The model that best matched the neural dynamics observed in our data featured the striatum as a perfect integrator and the ALM as a crucial input region (*specialized striatum integrator*, Extended Data Fig. 1f).

To simulate ALM silencing, both ALM neurons (A1 and A2 in Extended Data Fig. 1) received a negative current $I^{ext}(t)$. To stimulate D1-SPN inhibition, we injected a negative current into one of the striatum neurons (S1). To mimic transient perturbations in experiments, the negative current was introduced at 0.6 s after the cue and lasted for 0.6 s including a 300 ms ramp down. For prolonged perturbations, the negative current was applied throughout the trial without a ramp-down. The negative current for ALM silencing in all models was set to $I^{ext}(t)$ = -10. To maintain similar perturbation effects on striatal ramp mode activity across different conditions, the negative current for D1-SPN inhibition was varied across models as follows: -0.3 for the externally driven model (Extended Data Fig. 1a), -0.1 for the distributed model (Extended Data Fig. 1b), -0.1 for the redundant model (Extended Data Fig. 1c), -0.2 for the specialized ALM integrator model (Extended Data Fig. 1d), -0.02 for the specialized striatum leaky integrator model (Extended Data Fig. 1e), and -0.03 for the specialized striatum integrator model (Extended Data Fig. 1f).

Statistics

The sample sizes are similar to the sample sizes used in the field. No statistical methods were used to determine the sample size. During spike sorting, experimenters could not tell the trial type and therefore were blind to conditions. All *signed rank* and *rank sum* tests were two-sided. All bootstrap was done over 1,000 iterations.

Reagent and data availability

The recording data in NWB format, along with codes for data analyses and simulations, will be shared on DANDI at the time of publication.

REFERENCES

1. Caro, T. M. *Antipredator Defenses in Birds and Mammals*. (University of Chicago Press, 2005).
2. Henderson, J., Hurly, T. A., Bateson, M. & Healy, S. D. Timing in free-living rufous hummingbirds, Selasphorus rufus. *Curr. Biol. CB* **16**, 512–515 (2006).
3. Bateson, M. Interval timing and optimal foraging. in *Functional and neural mechanisms of interval timing* 113–141 (CRC Press/Routledge/Taylor & Francis Group, Boca Raton, FL, US, 2003). doi:10.1201/9780203009574.ch5.
4. Buhusi, C. V. & Meck, W. H. What makes us tick? Functional and neural mechanisms of interval timing. *Nat. Rev. Neurosci.* **6**, 755–765 (2005).
5. Mauk, M. D. & Buonomano, D. V. The Neural Basis of Temporal Processing. *Annu. Rev. Neurosci.* **27**, 307–340 (2004).
6. Paton, J. J. & Buonomano, D. V. The Neural Basis of Timing: Distributed Mechanisms for Diverse Functions. *Neuron* **98**, 687–705 (2018).
7. Merchant, H., Harrington, D. L. & Meck, W. H. Neural Basis of the Perception and Estimation of Time. *Annu. Rev. Neurosci.* **36**, 313–336 (2013).
8. Tsao, A., Yousefzadeh, S. A., Meck, W. H., Moser, M.-B. & Moser, E. I. The neural bases for timing of durations. *Nat. Rev. Neurosci.* **23**, 646–665 (2022).
9. Jazayeri, M. & Shadlen, M. N. A Neural Mechanism for Sensing and Reproducing a Time Interval. *Curr. Biol.* **25**, 2599–2609 (2015).
10. Remington, E. D., Egger, S. W., Narain, D., Wang, J. & Jazayeri, M. A Dynamical Systems Perspective on Flexible Motor Timing. *Trends Cogn. Sci.* **22**, 938–952 (2018).
11. Wang, J., Narain, D., Hosseini, E. A. & Jazayeri, M. Flexible timing by temporal scaling of cortical responses. *Nat. Neurosci.* **21**, 102–110 (2018).
12. Inagaki, H. K. *et al.* Neural Algorithms and Circuits for Motor Planning. *Annu. Rev. Neurosci.* **45**, 249–271 (2022).
13. Romo, R. & Schultz, W. Role of primate basal ganglia and frontal cortex in the internal generation of movements. *Exp. Brain Res.* **91**, 396–407 (1992).
14. 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).
15. Gold, J. I. & Shadlen, M. N. The Neural Basis of Decision Making. *Annu. Rev. Neurosci.* **30**, 535–574 (2007).
16. Desmurget, M. & Turner, R. S. Motor Sequences and the Basal Ganglia: Kinematics, Not Habits. *J. Neurosci.* **30**, 7685–7690 (2010).
17. Kunimatsu, J., Suzuki, T. W., Ohmae, S. & Tanaka, M. Different contributions of preparatory activity in the basal ganglia and cerebellum for self-timing. *eLife* **7**, e35676 (2018).
18. Thura, D. & Cisek, P. The Basal Ganglia Do Not Select Reach Targets but Control the Urgency of Commitment. *Neuron* **95**, 1160-1170.e5 (2017).
19. Catanese, J. & Jaeger, D. Premotor Ramping of Thalamic Neuronal Activity Is Modulated by Nigral Inputs and Contributes to Control the Timing of Action Release. *J. Neurosci. Off. J. Soc. Neurosci.* **41**, 1878–1891 (2021).
20. Balcı, F. & Simen, P. A decision model of timing. *Curr. Opin. Behav. Sci.* **8**, 94–101 (2016).
21. Narayanan, N. S. Ramping activity is a cortical mechanism of temporal control of action. *Curr. Opin. Behav. Sci.* **8**, 226–230 (2016).
22. Murakami, M., Vicente, M. I., Costa, G. M. & Mainen, Z. F. Neural antecedents of self-initiated actions in secondary motor cortex. *Nat. Neurosci.* **17**, 1574–1582 (2014).
23. Bakhurin, K. I. *et al.* Differential Encoding of Time by Prefrontal and Striatal Network Dynamics. *J. Neurosci.* **37**, 854–870 (2017).
24. Mello, G. B. M., Soares, S. & Paton, J. J. A Scalable Population Code for Time in the Striatum. *Curr. Biol.* **25**, 1113–1122 (2015).
25. Emmons, E. *et al.* Temporal Learning Among Prefrontal and Striatal Ensembles. *Cereb. Cortex Commun.* **1**, tgaa058 (2020).
26. Vyas, S., Golub, M. D., Sussillo, D. & Shenoy, K. V. Computation Through Neural Population

Dynamics. *Annu. Rev. Neurosci.* **43**, 249–275 (2020).
27. Cunningham, J. P. & Yu, B. M. Dimensionality reduction for large-scale neural recordings. *Nat. Neurosci.* **17**, 1500–1509 (2014).
28. Khona, M. & Fiete, I. R. Attractor and integrator networks in the brain. *Nat. Rev. Neurosci.* **23**, 744–766 (2022).
29. Aksay, E. *et al.* Functional dissection of circuitry in a neural integrator. *Nat. Neurosci.* **10**, 494–504 (2007).
30. Simen, P., Balci, F., deSouza, L., Cohen, J. D. & Holmes, P. A Model of Interval Timing by Neural Integration. *J. Neurosci.* **31**, 9238–9253 (2011).
31. Miller, P., Brody, C. D., Romo, R. & Wang, X.-J. A Recurrent Network Model of Somatosensory Parametric Working Memory in the Prefrontal Cortex. *Cereb. Cortex* **13**, 1208–1218 (2003).
32. Seung, H. S. How the brain keeps the eyes still. *Proc. Natl. Acad. Sci.* **93**, 13339–13344 (1996).
33. Lim, S. & Goldman, M. S. Balanced cortical microcircuitry for maintaining information in working memory. *Nat. Neurosci.* **16**, 1306–1314 (2013).
34. Cannon, S. C., Robinson, D. A. & Shamma, S. A proposed neural network for the integrator of the oculomotor system. *Biol. Cybern.* **49**, 127–136 (1983).
35. Egger, S. W., Remington, E. D., Chang, C.-J. & Jazayeri, M. Internal models of sensorimotor integration regulate cortical dynamics. *Nat. Neurosci.* **22**, 1871–1882 (2019).
36. Brody, C. D. & Hanks, T. D. Neural underpinnings of the evidence accumulator. *Curr. Opin. Neurobiol.* **37**, 149–157 (2016).
37. Yartsev, M. M., Hanks, T. D., Yoon, A. M. & Brody, C. D. Causal contribution and dynamical encoding in the striatum during evidence accumulation. *eLife* **7**, e34929 (2018).
38. Pinto, L., Tank, D. W. & Brody, C. D. Multiple timescales of sensory-evidence accumulation across the dorsal cortex. *eLife* **11**, e70263 (2022).
39. Murakami, M., Shteingart, H., Loewenstein, Y. & Mainen, Z. F. Distinct Sources of Deterministic and Stochastic Components of Action Timing Decisions in Rodent Frontal Cortex. *Neuron* **94**, 908-919.e7 (2017).
40. Xie, T., Huang, C., Zhang, Y., Liu, J. & Yao, H. Influence of Recent Trial History on Interval Timing. *Neurosci. Bull.* **39**, 559–575 (2023).
41. Bakhurin, K. I. *et al.* Opponent regulation of action performance and timing by striatonigral and striatopallidal pathways. *eLife* **9**, e54831 (2020).
42. Gouvêa, T. S. *et al.* Striatal dynamics explain duration judgments. *eLife* **4**, e11386 (2015).
43. De Corte, B. J., Wagner, L. M., Matell, M. S. & Narayanan, N. S. Striatal dopamine and the temporal control of behavior. *Behav. Brain Res.* **356**, 375–379 (2019).
44. Emmons, E. B. *et al.* Rodent Medial Frontal Control of Temporal Processing in the Dorsomedial Striatum. *J. Neurosci.* **37**, 8718–8733 (2017).
45. Buhusi, C. V., Reyes, M. B., Gathers, C.-A., Oprisan, S. A. & Buhusi, M. Inactivation of the Medial-Prefrontal Cortex Impairs Interval Timing Precision, but Not Timing Accuracy or Scalar Timing in a Peak-Interval Procedure in Rats. *Front. Integr. Neurosci.* **12**, (2018).
46. Kunimatsu, J. & Tanaka, M. Striatal dopamine modulates timing of self-initiated saccades. *Neuroscience* **337**, 131–142 (2016).
47. Drew, M. R., Fairhurst, S., Malapani, C., Horvitz, J. C. & Balsam, P. D. Effects of dopamine antagonists on the timing of two intervals. *Pharmacol. Biochem. Behav.* **75**, 9–15 (2003).
48. Monteiro, T. *et al.* Using temperature to analyze the neural basis of a time-based decision. *Nat. Neurosci.* **26**, 1407–1416 (2023).
49. Banerjee, A., Chen, F., Druckmann, S. & Long, M. A. Neural dynamics in the rodent motor cortex enables flexible control of vocal timing. *BioRxiv Prepr. Serv. Biol.* 2023.01.23.525252 (2023) doi:10.1101/2023.01.23.525252.
50. Kim, J., Jung, A. H., Byun, J., Jo, S. & Jung, M. W. Inactivation of medial prefrontal cortex impairs time interval discrimination in rats. *Front. Behav. Neurosci.* **3**, 38 (2009).
51. Matell, M. S. & Meck, W. H. Cortico-striatal circuits and interval timing: coincidence detection of oscillatory processes. *Cogn. Brain Res.* **21**, 139–170 (2004).
52. Wolff, S. B. E., Ko, R. & Ölveczky, B. P. Distinct roles for motor cortical and thalamic inputs to

striatum during motor skill learning and execution. *Sci. Adv.* **8**, eabk0231 (2022).
53. Tunes, G. C. *et al.* Time encoding migrates from prefrontal cortex to dorsal striatum during learning of a self-timed response duration task. *eLife* **11**, e65495.
54. Bruce, R. A. *et al.* Complementary cognitive roles for D2-MSNs and D1-MSNs in interval timing. *eLife* **13**, (2024).
55. Wolff, S. B. & Ölveczky, B. P. The promise and perils of causal circuit manipulations. *Curr. Opin. Neurobiol.* **49**, 84–94 (2018).
56. Jazayeri, M. & Afraz, A. Navigating the Neural Space in Search of the Neural Code. *Neuron* **93**, 1003–1014 (2017).
57. O'Shea, D. J. *et al.* Direct neural perturbations reveal a dynamical mechanism for robust computation. 2022.12.16.520768 Preprint at https://doi.org/10.1101/2022.12.16.520768 (2022).
58. Li, N., Daie, K., Svoboda, K. & Druckmann, S. Robust neuronal dynamics in premotor cortex during motor planning. *Nature* **532**, 459–464 (2016).
59. Inagaki, H. K., Fontolan, L., Romani, S. & Svoboda, K. Discrete attractor dynamics underlies persistent activity in the frontal cortex. *Nature* **566**, 212–217 (2019).
60. Boucher, P. O. *et al.* Initial conditions combine with sensory evidence to induce decision-related dynamics in premotor cortex. *Nat. Commun.* **14**, 6510 (2023).
61. Zhao, S. *et al.* Cell-type Specific Optogenetic Mice for Dissecting Neural Circuitry Function. *Nat. Methods* **8**, 745–752 (2011).
62. Guo, Z. V. *et al.* Flow of Cortical Activity Underlying a Tactile Decision in Mice. *Neuron* **81**, 179–194 (2014).
63. Xu, D. *et al.* Cortical processing of flexible and context-dependent sensorimotor sequences. *Nature* **603**, 464–469 (2022).
64. Bollu, T. *et al.* Cortex-dependent corrections as the tongue reaches for and misses targets. *Nature* **594**, 82–87 (2021).
65. Tanaka, M. Cognitive Signals in the Primate Motor Thalamus Predict Saccade Timing. *J. Neurosci.* **27**, 12109–12118 (2007).
66. Hintiryan, H. *et al.* The mouse cortico-striatal projectome. *Nat. Neurosci.* **19**, 1100–1114 (2016).
67. Hunnicutt, B. J. *et al.* A comprehensive excitatory input map of the striatum reveals novel functional organization. *eLife* **5**, e19103.
68. Tang, Y. *et al. Opposing Regulation of Short-Term Memory by Basal Ganglia Direct and Indirect Pathways That Are Coactive during Behavior*. http://biorxiv.org/lookup/doi/10.1101/2021.12.15.472735 (2021) doi:10.1101/2021.12.15.472735.
69. Peters, A. J., Fabre, J. M. J., Steinmetz, N. A., Harris, K. D. & Carandini, M. Striatal activity topographically reflects cortical activity. *Nature* **591**, 420–425 (2021).
70. Zhou, S., Masmanidis, S. C. & Buonomano, D. V. Neural Sequences as an Optimal Dynamical Regime for the Readout of Time. *Neuron* **108**, 651-658.e5 (2020).
71. Kao, T.-C., Sadabadi, M. S. & Hennequin, G. Optimal anticipatory control as a theory of motor preparation: A thalamo-cortical circuit model. *Neuron* **109**, 1567-1581.e12 (2021).
72. Harris, K. D. Nonsense correlations in neuroscience. 2020.11.29.402719 Preprint at https://doi.org/10.1101/2020.11.29.402719 (2021).
73. Li, N. *et al.* Spatiotemporal constraints on optogenetic inactivation in cortical circuits. *eLife* **8**, e48622 (2019).
74. Guo, J.-Z. *et al.* Cortex commands the performance of skilled movement. *eLife* **4**, e10774 (2015).
75. Jeurissen, D., Shushruth, S., El-Shamayleh, Y., Horwitz, G. D. & Shadlen, M. N. Deficits in decision-making induced by parietal cortex inactivation are compensated at two timescales. *Neuron* **110**, 1924-1931.e5 (2022).
76. Bao, C. *et al.* The rat frontal orienting field dynamically encodes value for economic decisions under risk. *Nat. Neurosci.* **26**, 1942–1952 (2023).
77. Findling, C. *et al.* Brain-wide representations of prior information in mouse decision-making. 2023.07.04.547684 Preprint at https://doi.org/10.1101/2023.07.04.547684 (2023).
78. Hattori, R. & Komiyama, T. Context-dependent persistency as a coding mechanism for robust and widely distributed value coding. *Neuron* **110**, 502-515.e11 (2022).

79. Akrami, A., Kopec, C. D., Diamond, M. E. & Brody, C. D. Posterior parietal cortex represents sensory history and mediates its effects on behaviour. *Nature* **554**, 368–372 (2018).
80. Jun, J. J. *et al.* Fully integrated silicon probes for high-density recording of neural activity. *Nature* **551**, 232–236 (2017).
81. Gerfen, C. R. & Surmeier, D. J. Modulation of Striatal Projection Systems by Dopamine. *Annu. Rev. Neurosci.* **34**, 441–466 (2011).
82. DeLong, M. R. Primate models of movement disorders of basal ganglia origin. *Trends Neurosci.* **13**, 281–285 (1990).
83. Kravitz, A. V. *et al.* Regulation of parkinsonian motor behaviours by optogenetic control of basal ganglia circuitry. *Nature* **466**, 622–626 (2010).
84. Govorunova, E. G., Sineshchekov, O. A., Janz, R., Liu, X. & Spudich, J. L. NEUROSCIENCE. Natural light-gated anion channels: A family of microbial rhodopsins for advanced optogenetics. *Science* **349**, 647–650 (2015).
85. Gerfen, C. R., Paletzki, R. & Heintz, N. GENSAT BAC Cre-recombinase driver lines to study the functional organization of cerebral cortical and basal ganglia circuits. *Neuron* **80**, 10.1016/j.neuron.2013.10.016 (2013).
86. Pisanello, F. *et al.* Dynamic illumination of spatially restricted or large brain volumes via a single tapered optical fiber. *Nat. Neurosci.* **20**, 1180–1188 (2017).
87. Mahn, M. *et al.* High-efficiency optogenetic silencing with soma-targeted anion-conducting channelrhodopsins. *Nat. Commun.* **9**, 4125 (2018).
88. Messier, J. E., Chen, H., Cai, Z.-L. & Xue, M. Targeting light-gated chloride channels to neuronal somatodendritic domain reduces their excitatory effect in the axon. *eLife* **7**, e38506 (2018).
89. Shin, J. H., Kim, D. & Jung, M. W. Differential coding of reward and movement information in the dorsomedial striatal direct and indirect pathways. *Nat. Commun.* **9**, 404 (2018).
90. Gagnon, D. *et al.* Striatal Neurons Expressing D1 and D2 Receptors are Morphologically Distinct and Differently Affected by Dopamine Denervation in Mice. *Sci. Rep.* **7**, 41432 (2017).
91. Laboratory, I. B. *et al.* A Brain-Wide Map of Neural Activity during Complex Behaviour. 2023.07.04.547681 Preprint at https://doi.org/10.1101/2023.07.04.547681 (2023).
92. Chen, S. *et al.* Brain-wide neural activity underlying memory-guided movement. *Cell* **187**, 676-691.e16 (2024).
93. Kim, J., Kim, D. & Jung, M. W. Distinct Dynamics of Striatal and Prefrontal Neural Activity During Temporal Discrimination. *Front. Integr. Neurosci.* **12**, (2018).
94. Okobi, D. E., Banerjee, A., Matheson, A. M. M., Phelps, S. M. & Long, M. A. Motor cortical control of vocal interaction in neotropical singing mice. *Science* **363**, 983–988 (2019).
95. Erlich, J. C., Brunton, B. W., Duan, C. A., Hanks, T. D. & Brody, C. D. Distinct effects of prefrontal and parietal cortex inactivations on an accumulation of evidence task in the rat. *eLife* **4**, e05457 (2015).
96. Remington, E. D., Narain, D., Hosseini, E. A. & Jazayeri, M. Flexible Sensorimotor Computations through Rapid Reconfiguration of Cortical Dynamics. *Neuron* **98**, 1005-1019.e5 (2018).
97. Goldman, M. S. Memory without Feedback in a Neural Network. *Neuron* **61**, 621–634 (2009).
98. Mandelbaum, G. *et al.* Distinct Cortical-Thalamic-Striatal Circuits through the Parafascicular Nucleus. *Neuron* **102**, 636-652.e7 (2019).
99. Lee, K. *et al.* Gain Modulation by Corticostriatal and Thalamostriatal Input Signals during Reward-Conditioned Behavior. *Cell Rep.* **29**, 2438-2449.e4 (2019).
100. Kheifets, A. & Gallistel, C. R. Mice take calculated risks. *Proc. Natl. Acad. Sci.* **109**, 8776–8779 (2012).
101. Freestone, D. M., Balcı, F., Simen, P. & Church, R. M. Optimal response rates in humans and rats. *J. Exp. Psychol. Anim. Learn. Cogn.* **41**, 39–51 (2015).
102. Çavdaroğlu, B., Zeki, M. & Balcı, F. Time-based reward maximization. *Philos. Trans. R. Soc. B Biol. Sci.* **369**, 20120461 (2014).
103. Balcı, F. *et al.* Optimal Temporal Risk Assessment. *Front. Integr. Neurosci.* **5**, (2011).
104. Mitchell, S. H. Assessing delay discounting in mice. *Curr. Protoc. Neurosci. Editor. Board Jacqueline N Crawley Al* **66**, 8.30.1-8.30.12 (2014).

105. Gibbon, J. Scalar expectancy theory and Weber's law in animal timing. *Psychol. Rev.* **84**, 279–325 (1977).
106. Mathis, A. *et al.* DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. *Nat. Neurosci.* **21**, 1281–1289 (2018).
107. Zhao, S. *et al.* Cell type–specific channelrhodopsin-2 transgenic mice for optogenetic dissection of neural circuitry function. *Nat. Methods* **8**, 745–752 (2011).
108. Li, N. *et al.* Spatiotemporal constraints on optogenetic inactivation in cortical circuits. *eLife* **8**, e48622 (2019).
109. Guo, Z. V. *et al.* Procedures for Behavioral Experiments in Head-Fixed Mice. *PLOS ONE* **9**, e88678 (2014).
110. Majumder, S. *et al.* Cell-type-specific plasticity shapes neocortical dynamics for motor learning. 2023.08.09.552699 Preprint at https://doi.org/10.1101/2023.08.09.552699 (2023).
111. Liu, L. D., Chen, S., Economo, M. N., Li, N. & Svoboda, K. Accurate localization of linear probe electrodes across multiple brains. Preprint at https://doi.org/10.1101/2020.02.25.965210 (2020).
112. Inagaki, H. K. *et al.* A midbrain-thalamus-cortex circuit reorganizes cortical dynamics to initiate movement. *Cell* **185**, 1065-1081.e23 (2022).
113. Hsu, C.-W. *et al.* EZ Clear for simple, rapid, and robust mouse whole organ clearing. *eLife* **11**, e77419 (2022).
114. Guo, Z. V. *et al.* Maintenance of persistent activity in a frontal thalamocortical loop. *Nature* **545**, 181–186 (2017).
115. Jun, J. J. *et al.* Real-time spike sorting platform for high-density extracellular probes with ground-truth validation and drift correction. Preprint at https://doi.org/10.1101/101030 (2017).