

A cerebellar granule cell-climbing fiber computation to learn to track long time intervals

Highlights

- First simultaneous recordings of cerebellar granule cells and climbing fibers
- Mice learn reward timing during operant conditioning, with cerebellar contributions
- Sustained granule cell firing links action to future climbing fiber reward spiking
- Granule cell-climbing fiber activity is suited to learning to track long intervals

Authors

Martha G. Garcia-Garcia,
Akash Kapoor, Oluwatobi Akinwale, ...,
Mark J. Schnitzer, Liqun Luo,
Mark J. Wagner

Correspondence

mark.wagner@nih.gov

In brief

Via the first simultaneous recordings of cerebellar granule cell and climbing fiber activity, Garcia-Garcia et al. demonstrate a new type of cerebellar computation suited to learning to track long time intervals via sustained neural activity ramping. This strategy could enable cerebellar contributions to longer-timescale cognitive processes requiring temporal learning.



Article

A cerebellar granule cell-climbing fiber computation to learn to track long time intervals

Martha G. Garcia-Garcia,^{1,5} Akash Kapoor,^{1,5} Oluwatobi Akinwale,^{1,5} Lina Takemaru,^{1,5} Tony Hyun Kim,² Casey Paton,¹ Ashok Litwin-Kumar,³ Mark J. Schnitzer,^{2,4} Liqun Luo,² and Mark J. Wagner^{1,6,*}

¹National Institute of Neurological Disorders & Stroke, National Institutes of Health, Bethesda, MD 20894, USA

²Department of Biology and Howard Hughes Medical Institute, Stanford University, Stanford, CA 94305, USA

³Zuckerman Mind Brain Behavior Institute, Department of Neuroscience, Columbia University, New York, NY 10027, USA

⁴Department of Applied Physics, Stanford University, Stanford, CA 94305, USA

⁵These authors contributed equally

⁶Lead contact

*Correspondence: mark.wagner@nih.gov

<https://doi.org/10.1016/j.neuron.2024.05.019>

SUMMARY

In classical cerebellar learning, Purkinje cells (PkCs) associate climbing fiber (CF) error signals with predictive granule cells (GrCs) that were active just prior (~150 ms). The cerebellum also contributes to behaviors characterized by longer timescales. To investigate how GrC-CF-PkC circuits might learn seconds-long predictions, we imaged simultaneous GrC-CF activity over days of forelimb operant conditioning for delayed water reward. As mice learned reward timing, numerous GrCs developed anticipatory activity ramping at different rates until reward delivery, followed by widespread time-locked CF spiking. Relearning longer delays further lengthened GrC activations. We computed CF-dependent GrC → PkC plasticity rules, demonstrating that reward-evoked CF spikes sufficed to grade many GrC synapses by anticipatory timing. We predicted and confirmed that PkCs could thereby continuously ramp across seconds-long intervals from movement to reward. Learning thus leads to new GrC temporal bases linking predictors to remote CF reward signals—a strategy well suited for learning to track the long intervals common in cognitive domains.

INTRODUCTION

The cerebellum is widely viewed as a structure for learning predictions^{1–5} from inputs ranging from the body to neocortical cognition centers.^{6–9} These diverse inputs propagate through a uniform circuit (Figure 1A). Purkinje cells (PkCs) receive input from ~100,000 granule cells (GrCs)¹⁰ and just one climbing fiber (CF). PkC computation depends on integrating GrC inputs^{11,12} using synaptic strengths that can be modified by CF instructive signals.¹³ Specifically, when a PkC receives a CF spike, any GrC inputs active in the preceding ~150 ms are weakened via long-term depression (LTD).^{14–19} Similarly, reduced CF activity triggers long-term potentiation (LTP) on coincidentally active GrC inputs.^{20–22} Given the brief plasticity window to “sense” active GrCs, a core cerebellar function is learning short-latency associations between predictive events and CF “errors.”^{23–27} Thus, decades of theory have posited GrC representations of time, or “basis sets,” as key to cerebellar learning.^{28–36} Yet, limited available data suggest varied types of GrC representations.^{37–41} Instead of a single uniform GrC basis, an intriguing alternative is that GrCs assume differing profiles to suit specific computational needs. In this scenario, understanding cerebellar computations for different learned behaviors would

require characterizing the associated GrC bases and their relationship to CF teaching signals.

Cognitive behaviors, including those with cerebellar contributions,^{4,43,44} often require linking events separated by seconds or more. Cerebellar computations in such contexts remain obscure, and the associated GrC bases and their interplay with CF teaching signals are uncharacterized. Due to technical difficulty, GrC and CF inputs onto PkCs have yet to be simultaneously observed and related to learning.

Here, we recorded simultaneous GrC and CF activity using dual-color two-depth two-photon Ca²⁺ imaging over days of operant conditioning for delayed reward. Mice learned to lick in anticipation of an expected reward—a behavior with cerebellar contributions. Meanwhile, GrCs and CFs developed two key features. First, after learning, numerous GrCs ramped activity up and down at differing rates from forelimb movement until reward 1 s later; subsequent relearning of a 2-s delay further lengthened GrC profiles commensurately over days. Second, the reward itself triggered widespread time-locked CF spiking that persisted throughout learning. To understand the computations these GrC-CF patterns subserve, we predicted GrC → PkC synaptic changes based on canonical plasticity rules. Because many GrCs developed activity that spanned the delay, CF-driven plasticity near reward sufficed to grade numerous GrC → PkC



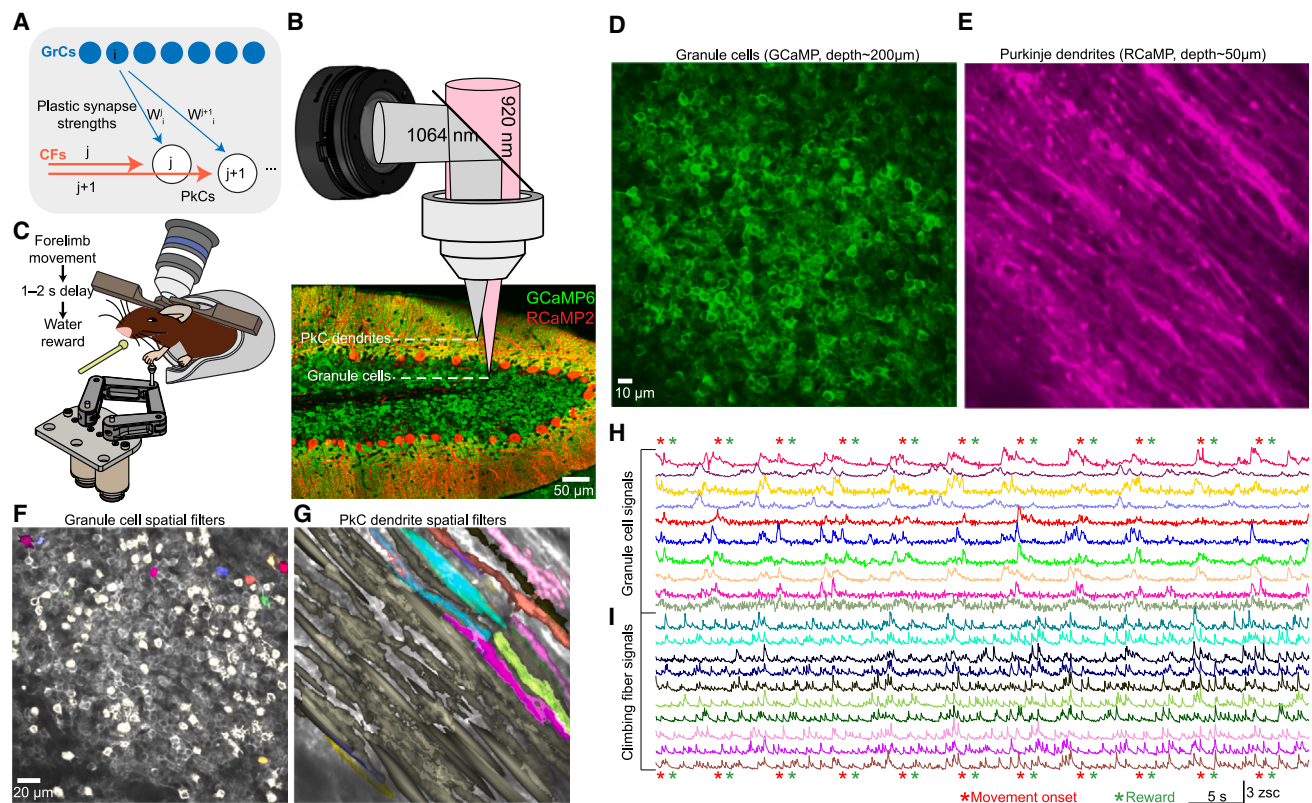


Figure 1. Two-color two-depth two-photon Ca^{2+} imaging of cerebellar GrCs and CFs

(A) Schematic of cerebellar microcircuit: $\sim 100,000$ GrCs and one CF innervate each PkC, with GrC \rightarrow PkC synapses adjusted by CF-dependent plasticity. GrC “ j ” synapses on PkCs “ j ” and “ $j+1$ ” with weights W_i^j and W_i^{j+1} .
 (B) Imaging schematic and histology. GrCs transgenically expressed GCaMP6f, whereas PkCs virally expressed R-CaMP2. PkC dendritic Ca^{2+} reports complex spikes⁴² and thus CF activity. Through one objective, a 920 nm laser excited GCaMP in GrC somas, while a remotely focused 1,064 nm laser excited PkC dendrites.
 (C) Mice grasped a robotic manipulandum and self-initiated 8-mm-maximum forward pushes of at least 6–7 mm for water reward following a delay (main GrC-CF data: 1.1 s; all other studies: 1 or 2 s). 2 s after reward time, the handle automatically returned over the following 2 s.
 (D and E) Example *in vivo* mean two-photon simultaneous images of GrCs (D) and PkC dendrites (E).
 (F and G) Example extracted GrCs (F) and PkC dendrites (G). Detected spatial filters of active cells are superposed in pale yellow or for 10 GrCs and PkC dendrites in colors corresponding to the traces in (H) and (I).
 (H and I) For 10 GrCs (H) and PkC dendrites (I), color matched to cells in (F) and (G), traces show time-varying fluorescence of each neuron. Stars show forelimb movements and rewards. PkC dendritic spiking hereafter referred to as CFs. zsc, Z scores.

weights by GrC activity timing throughout the delay. We therefore predicted that PkC spiking could adapt to track the interval from forelimb movement until reward up to 2-s later. PkC recordings confirmed such delay-spanning activity ramps. Thus, the emergence during learning of new GrC bases, which link predictive events to reward-evoked CF teaching signals seconds later, permits long-timescale associations. This outlines a cerebellar computation suited to learning neural ramps that track long intervals common in cognitive domains.

RESULTS

Simultaneous GrC-CF recordings during reward-driven learning

We sought to simultaneously record GrC and CF activity over learning, which has previously been impossible. Due to tight packing and small cell size, two-photon imaging is necessary

to track mammalian GrC populations.^{9,40,41} Two-photon imaging can also record CF activity via PkC dendritic Ca^{2+} , which reports complex spikes.⁴² Thus, in principle, GrCs and CFs could simultaneously be imaged in a single color. However, because $\sim 100,000$ GrC axons physically intersect with each PkC dendrite, it would be impossible to spatially disambiguate their fluorescence *in vivo*. Instead, we leveraged spectral separation via dual-color imaging.

Cerebellar areas related to cognition include lobules $\sim \text{VI-VIII}$.⁴ Lobule VI (LVI) communicates with multiple cognitive circuits,^{45–49} exhibits nonmotor signaling,^{9,41} and receives most inputs to its GrCs from pons,⁵⁰ which relays neocortical information. Although these features are common in the cerebellum, to deeply characterize GrC-CF computation in one region, we centered our imaging on LVI right vermis (ipsilateral to the reaching forelimb). We transgenically expressed GCaMP6f in GrCs (Math1-Cre/Ai93/zTA⁴¹) and virally expressed R-CaMP2⁵¹ in

PkCs (AAV-L7-6⁵²-R-CaMP2) (Figures 1B and S1A histology), whose dendritic Ca²⁺ reliably reports complex spikes and thus CF input.^{42,53} Next, to simultaneously visualize both, we imaged at two depths: a 920 nm laser illuminated GrC somas ~150–300 μ m below the brain surface, and a remotely focused 1,064 nm laser illuminated PkC dendrites at a depth of ~50 μ m (Figure 1B).

As animals learned an operant behavior (Figure 1C; details below), we simultaneously monitored GrCs (143 \pm 5 per session, 117 sessions in 20 mice) and PkC dendrites (hereafter referred to as CFs; 55 \pm 2 PkC dendrites per session), either at 30 or 22.5 Hz (example: Figures 1D and 1E; Video S1), over 7.5 \pm 0.4 days (Figures S1B–S1G). We extracted spatial locations (Figures 1F and 1G) and time-varying fluorescence (Figures 1H and 1I) of each neuron and further extracted spike times for CFs⁴² (STAR Methods). Importantly, because we imaged GrCs directly below the imaged PkC dendrites (field of view: ~235 \times 235 μ m²), axons of these GrCs likely physically intersected most PkC dendrites, and GrCs synapse onto ~50% of the PkC dendrites they intersect⁵⁴ (Figures S1H and S1I). Thus, we visualized both input streams—CF and some of the GrCs—to the same dozens of PkCs over learning.

During imaging, water-restricted mice grasped a robotic arm to self-initiate forward pushes (Figure 1C). Following successful movement and a delay (main GrC-CF imaging dataset: 1.1 s, all other studies: 1 or 2 s), animals received water reward. Delays were thus far longer than classical cerebellar association time-scales.⁵⁵ Mice acclimated to the task for ~1–3 sessions before we started imaging (STAR Methods). Thus, “day 1” hereafter refers to the first imaging session. After ~1–2 imaging sessions, we also withheld water on a random 20% of trials (“omitted reward”). Mice learned the task over ~1 week: (1) movement duration decreased (Figure S2A), (2) execution rate increased (Figures S2B–S2D), and (3) anticipatory licking increased (Figure 2A).^{41,56}

Mice learn the timing of delayed reward during operant conditioning

We first tested whether mice learned reward timing. To distinguish the amount of anticipatory licking from its timing, we examined the temporal distribution of licks prior to reward in mice trained on a 1.1-s delay (Figure 2B). Novice mice licked more just after forelimb movement than just before reward, whereas expert mice partly inverted this pattern (Figures 2C and S2E; late vs. early preference: -0.13 ± 0.02 novice vs. $+0.2 \pm 0.01$ expert; these and all subsequent quantifications mean \pm SEM across observations, see also Table S1). On omitted-reward trials, experts reduced licking after the expected reward time (Figures 2D and S2F; licking off-time, last time lick rate exceeded 70% of peak). Thus, mice learned to concentrate their licking closer to the expected reward.

To assay reward timing learning directly, we tested whether behavior adapted to delay duration. Mice trained on a 1-s delay, and then retrained on a 2-s delay for another week (Figure 2E). We examined licking on omitted-reward trials aligned to forelimb movement, such that differences after $t = 0$ s could not be attributed to exogenous stimuli (Figures 2F and S2G; Video S2). Whereas 1-s experts licked more shortly after forelimb move-

ment compared with after 2 s, 2-s-experts exhibited the opposite pattern (Figure 2G). Moreover, when 1-s-experts initially switched to the 2-s paradigm, their licking was diffuse across the delay, consistent with timing uncertainty (Figures S2H and S2I). Thus, mice adapted their peak licking to the respective reward times (Figure 2H), demonstrating reward timing learning over delays of 1 or 2 s.

Cerebellum contributes to the execution and learning of expectation-driven licking

To test cerebellar contributions to this behavior, we implanted windows over right LVI in PCP2-cre/Ai32 mice, which express ChR2 in all PkCs.^{57–59} Activating inhibitory PkCs suppresses the cerebellar nuclei (CN), thus reducing the circuit's influence on downstream brain regions. Mice trained for a week on the 1-s-delay operant task. In experts, on a random 20%–40% of trials, we activated PkCs at 40 Hz for 0.8 s, triggered during the delay period (200 ms after a 7-mm-reach-distance threshold; STAR Methods). Stimulation largely abolished anticipatory licking (Figures 2I and S2M; Video S3). After laser offset, reward delivery triggered rapid recovery of licking (Figure 2I). However, after laser offset on reward-omission trials, licking remained disrupted throughout the trial (Figure 2J). First, mice recovered less licking in the early post-omission period (Figures 2J and 2K). Second, long after the omission of the expected reward, mice exhibited aberrantly elevated licking (Figures 2J and 2K). Thus, transient cerebellar perturbation imparted long-lasting disturbance to the amount and timing of expectancy-driven licking but not reward-evoked licking.

We performed several further controls. First, to test whether mice simply needed more recovery time after stimulation, we repeated the experiment using briefer PkC stimulation early in the delay on half of reward-omission trials. Nevertheless, recovery of anticipatory licking remained weak and poorly timed (Figure S2J). We next repeated laser illumination in opsin-negative mice (Figure S2K), and finally, we repeated PCP2/Ai32 stimulation in a less relevant region (vermis lobule IX, Figure S2L)—both without effect.

These data suggested that the cerebellum contributes to expectation-driven licking, leading us next to test its involvement in learning reward-timed licking. We retrained 1-s-expert PCP2/Ai32 mice on a 2-s delay. 90% of trials were rewarded but also laser-ON ([1.6, 2.2] s from movement); 10% of trials were laser-OFF reward-omission “probe” trials (Figure 2L). Over 1-week of training, licking on probe laser-OFF reward-omission trials never adapted to the 2-s reward time (Figures 2M and 2N). After 3–7 subsequent days of 100% laser-off training, mice shifted their licking closer to 2 s. These results suggest that the posterior cerebellar cortex contributes to the learning and execution of licking driven by the expectation of upcoming reward but do not exclude similar effects for other regions untested in the present study.

GrCs ramp from movement until reward, which triggers time-locked CF spiking

To examine GrC-CF task representations in 1.1-s-delay experts, we computed average activity on rewarded trials (day 7+) (Figures 3A and 3B). We noted numerous GrCs with sustained

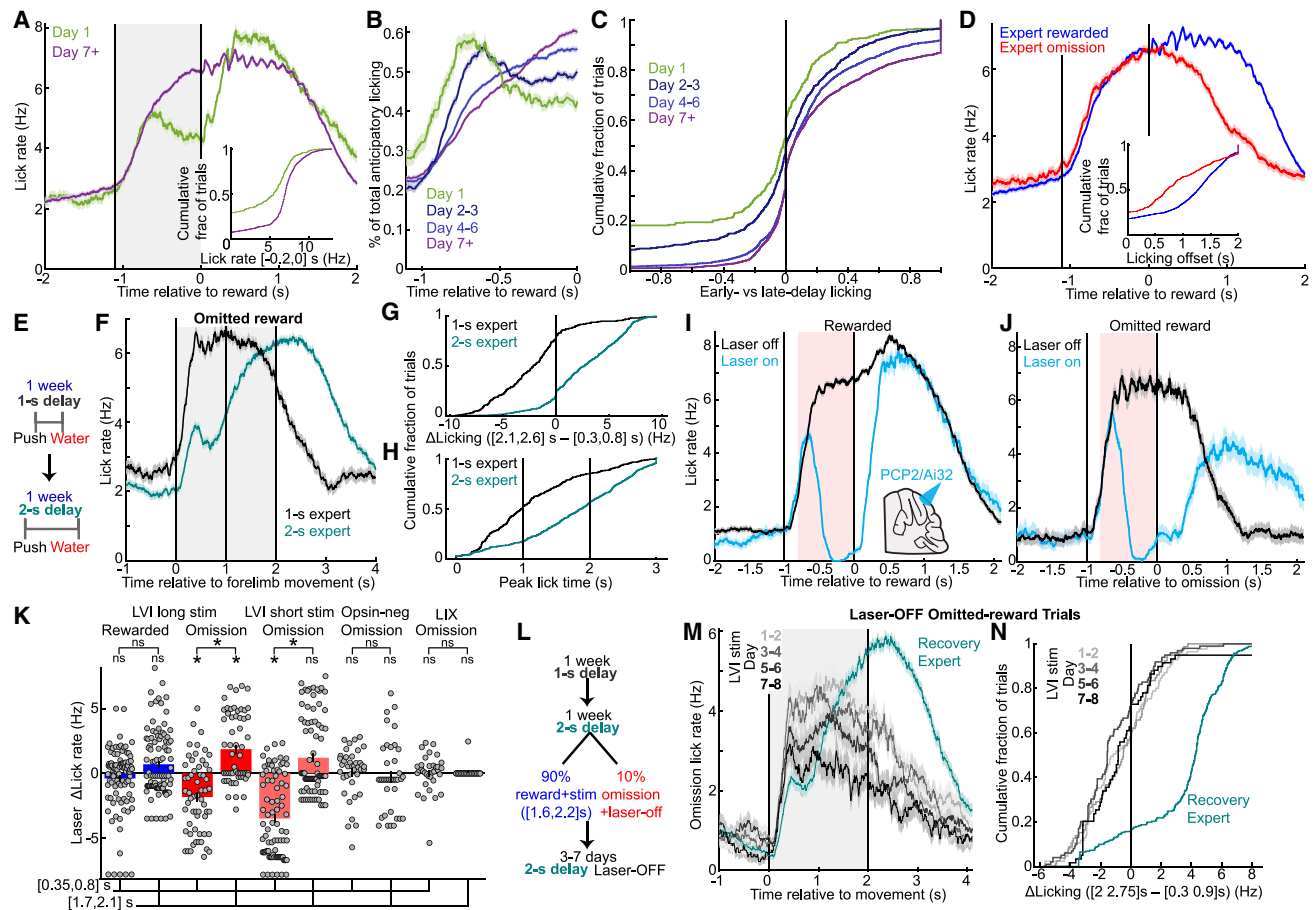


Figure 2. Mice learn to elevate licking near the time of expected reward, with cerebellar contributions

(A) Lick rate before reward was higher on day 7+ than day 1 (334 day 1 trials from 10 mice and 1,570 day 7+ trials from 20 sessions in 12 mice; this and subsequent gray regions denote delay period, and vertical lines denote movement and reward). Inset, $p < 10^{-6}$.

(B) Distribution of licks across delay (400/9, 919/16, 1,820/21, and 2,065/22 trials/sessions).

(C) Novice mice licked earlier than late in the delay, whereas expert mice inverted this pattern ([late – early]/[late + early]); early: [-0.8, -0.6] s, late: [-0.2, 0] s. $p < 10^{-6}$. These and all subsequent centers denote means, and shaded regions and error bars denote SEM across observations (see also Table S1).

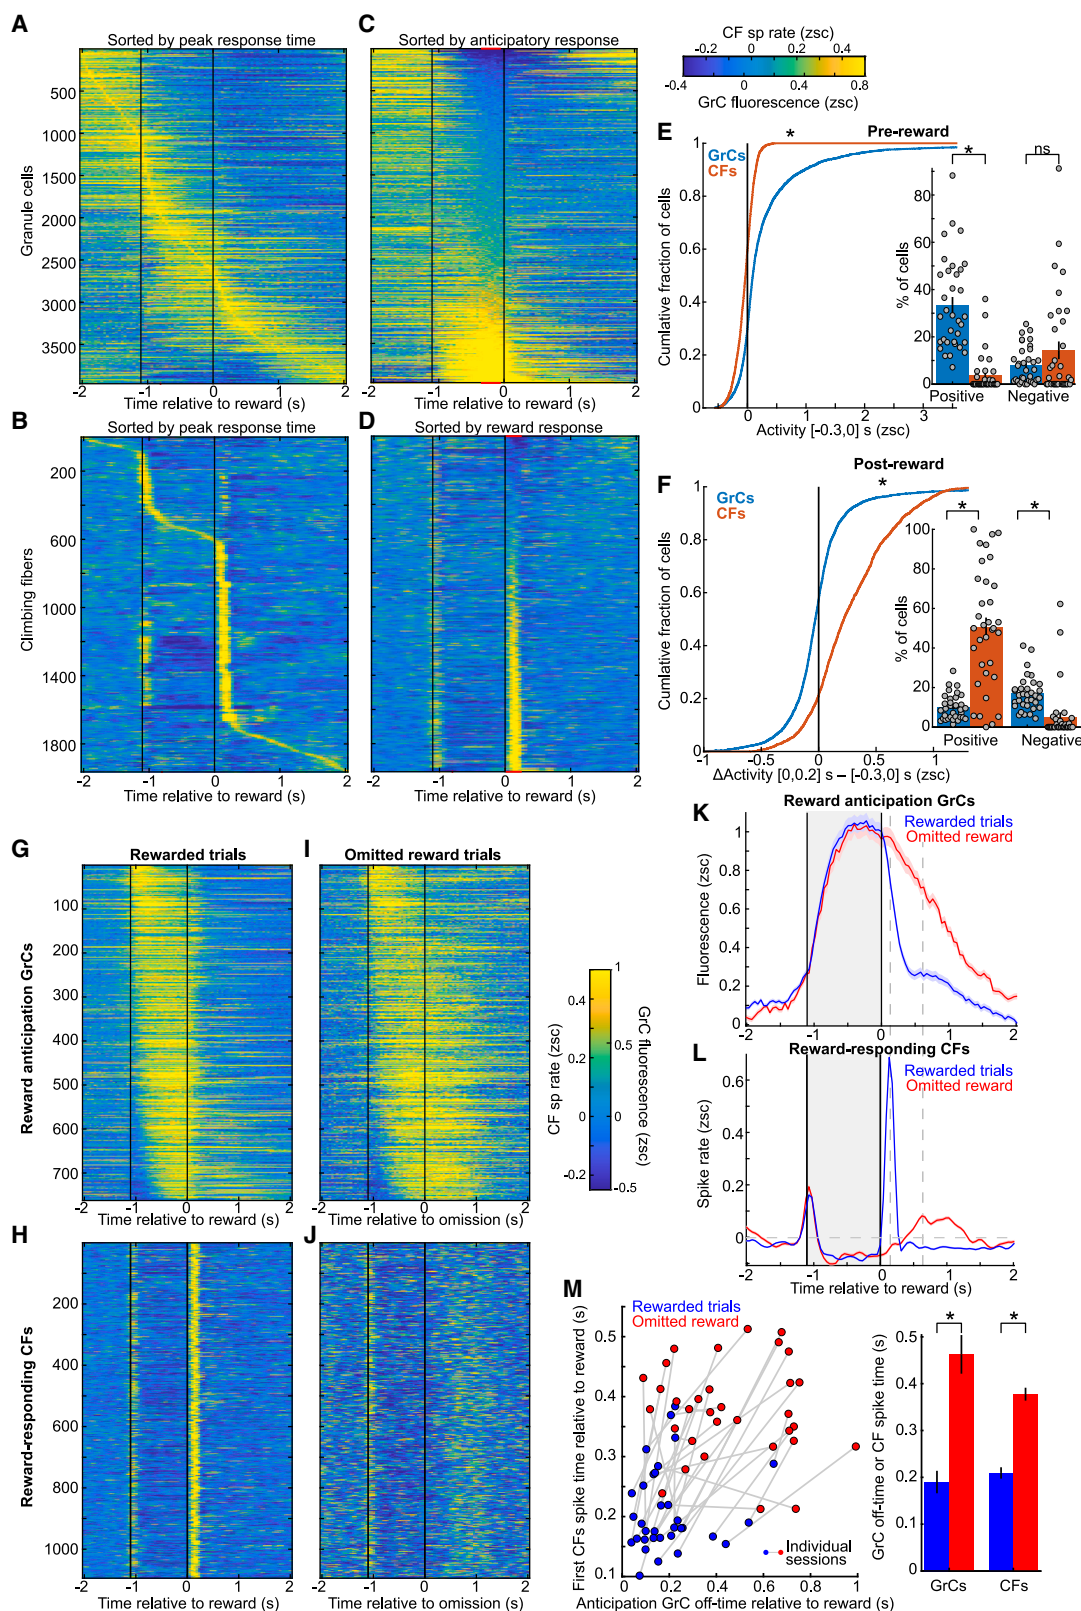
(D) Expert lick rate during rewarded and omitted-reward trials (1,570 and 495 trials from 20 sessions in 12 mice). Inset, last time (max 2 s) at which lick rate exceeded 70% of the prior peak ([-0.25, 0.5] s; $p < 10^{-6}$; 0 if licking never fell below 50% of peak).

(E–H) (E) Some mice trained with a 1-s delay followed by a 2-s delay. (F) Licking on omitted-reward trials for 1- vs. 2-s experts (283 1- and 519 2-s trials from 14 mice, 16 sessions each). (G) 2-s expert omission licking was higher after 2 s than early in the delay, whereas 1-s experts showed the opposite pattern ($p < 10^{-6}$). (H) Omission licking peaked near respective expected reward times ($p < 10^{-6}$, peak over [0, 3] s).

(I–N) PCP2Cre/Ai32 PkC stimulation studies. (I and J) PCP2Cre/Ai32 mice with windows over right cerebellum LVI trained on the 1-s-delay task. Starting 0.2 s after mice pushed >7 mm, we activated ChR2 for 0.8 s on 20%–40% of trials (I, rewarded: 478 laser-off, 99 laser-on; J, omitted reward: 68 laser-off, 59 laser-on; 7 mice). Stimulation abolished anticipatory licking. Reward triggered recovery of normal licking, but on omission trials licking remained weaker and less well timed (controls, Figures S2J–S2L). Pink: mean laser period.

(K) Quantification of Figures 2I, 2J, and S2J–S2L. “ Δ Lick rate”: each laser-on trial’s licking minus the mean laser-off licking, averaged from [0.35, 0.8] or [1.7, 2.1] s from reward. Rewarded licking (Figure 2I) changed little in either window ($p = 0.8$ and 0.06 ; difference $p = 0.1$). Omission licking was reduced just after reward time ($p = 3 \times 10^{-5}$ and $< 10^{-6}$ for long/Figure 2J and short/Figure S2J stim paradigms), but aberrantly elevated later (long stim: $p = 3 \times 10^{-5}$; brief-stim: $p = 0.06$; both paradigms late vs. early: $p < 10^{-6}$). Opsin-negative and LIX controls were not significant (bars $p = 0.2, 0.4, 0.3, 1$; differences $p = 0.1, 0.7$; laser-on omitted-reward trial counts: 99, 59, 85, 39, and 34; mice/session counts: 7/9, 7/7, 3/3, and 3/3).

(L–N) PCP2-Ai32 mice that were experts on the 1-s delay retrained with 2-s delay but with PkC stimulation on every rewarded trial from [1.6, 2.2] s from movement (90% of all trials, L). 10% of trials were laser-off reward-omission probe trials. Over 1 week of perturbed training—but evaluated on probe laser-OFF trials—mice never learned to lick near the 2-s reward delivery time (M, grayscale curves). During subsequent laser-OFF training, mice learned to lick near 2-s (M, green; 85, 79, 99, 59, and 261 reward-omission trials per condition from 7 mice). (N) Δ lick rate during expected reward time minus early delay (recovery > laser-on: $p < 10^{-6}$; laser-on across learning: $p = 0.07$).



(legend on next page)

activity during the delay and many CFs that were active just after the reward (Figures 3C and 3D). Quantitatively, cells active during the delay comprised $34\% \pm 3\%$ of GrCs but only $4\% \pm 1\%$ of CFs (Figure 3E). Conversely, cells that activated at reward comprised $50\% \pm 5\%$ of CFs but only $10\% \pm 1\%$ of GrCs (Figure 3F). A smaller $22\% \pm 5\%$ of CFs activated at movement (Figure S2P). In addition to CF spiking after reward, we also observed modest but variable suppression of CFs prior to reward ($14\% \pm 4\%$ of CFs, Figure 3F). Experts thus featured widespread sustained GrC delay activity and CF reward spiking.

Many GrCs and CFs appeared linked with anticipating and receiving reward. To characterize these features, we identified “reward anticipation” GrCs with elevated delay activity (Figures 3G, S2N, and S2O, 19% of GrCs, from 33/34 sessions in 19/20 mice). Similarly, we identified CFs with reward spiking (Figure 3H, 55% of CFs, from all 34 sessions/20 mice). This revealed that, whereas GrC anticipation terminated after reward, it was prolonged after unexpected reward omission (Figure 3I). CF reward spiking was absent on omission trials (Figure 3J). Instead, CF spiking rose later after reward omission, grossly coinciding with the delayed off-time of GrC anticipation (Figures 3K and 3L). To quantify, for each trial, we computed the times after reward (or omission) when (1) CF spiking rose and (2) GrC anticipation terminated; both increased significantly on omitted-reward trials (Figure 3M). Thus, these GrC-CF phenomena likely relate first to anticipation and then either reward delivery or recognition of its omission.

Instead of signaling reward anticipation, could GrCs simply signal licking? When animals awaited reward, both anticipatory GrC signals and anticipatory licking ramped together. However, GrC activity otherwise diverged from licking: (1) after reward delivery, GrC anticipation terminated (Figure S3A; decay time 0.19 s post-reward)—yet, licking further increased to its highest values (Figure S3A; decay time 1.1 s post-reward); (2) compared with rewarded trials, following reward omission, GrC anticipation was substantially more prolonged (Figure S3B; signals extended by 130%)—yet, licking was substantially weaker and more brief (Figure S3B; 37% less licking). Finally, anticipatory GrCs were not modulated by individual lick onsets (Figures S3C and S3D). We also did not observe selective delay period body, eye, or whisker movements consistent with anticipatory GrC temporal profiles (Figures S3J–S3Q). Thus, anticipatory GrC signals were most consistent with a temporally building expectation of reward.

Overall, we found that reward anticipation (1) followed forelimb movement and modest CF spiking, (2) rose with sustained GrC activation, and (3) terminated at reward (or omission) when GrC anticipation decayed, and many CFs spiked. This hinted that many expert GrCs spanned the delay from movement until reward-evoked CF spiking. Thus, we next examined whether this was linked to learning reward timing.

Learning increases GrC reward anticipation but leaves reward-evoked CF spiking unchanged

To test whether GrC reward anticipation and CF reward spiking changed with learning, we imaged repeatedly as novice mice learned the 1.1-s delay (15 mice, 7.5 ± 0.4 sessions per mouse, Figures S1D and S1E). We examined reward anticipation GrCs on days 1 and 7+ (Figure 4A, spread across all day-1 sessions/mice and 33/34 day-7+ sessions from 19/20 mice). The magnitude of GrC anticipation grew, and its prevalence increased moderately (Figure 4C). Thus, learning enhanced sustained GrC delay activity.

We next considered CF reward signals. In leading theories, CFs signal errors that drive GrC → PkC plasticity, which reduces future errors.⁶⁰ Both behavioral and GrC signatures of reward anticipation improved with learning (Figures 2A–2C, 4A, and 4C). Thus, if reward-evoked CF spikes signaled prediction errors, such spiking should decrease after learning. To quantify this, we identified CFs activated by reward (Figure 4B; spread across all days 1 and 7+ sessions). We found that after learning, reward-evoked CF spiking marginally increased in magnitude and remained equally prevalent (Figure 4D), contrary to an error.⁶¹ CF anticipatory suppression magnitudes and prevalence also increased (Figure S2Q). As an alternative to errors, CF reward spiking could mark the end of the expectation period spanned by anticipatory GrC activity.

Learning lengthens GrC profiles to span the delay until reward

To directly test whether operant learning lengthened GrC profiles to span the delay, a cohort of mice was trained with a 1-s delay for a week and then retrained with a 2-s delay (Figure 4E). We examined GrC-CF profiles at three points: expert 1-s delay performance (Figure 4F), just after switching to the 2-s delay (Figure 4G), and expert 2-s delay performance (Figure 4H). In 1-s experts, most anticipatory GrCs were active mainly during [0, 1] s after forelimb movement, which was similar in 2-s-delay-novices (Figures 4I–4L black/brown). After a week of 2-s training,

Figure 3. Reward-anticipating GrC activity followed by reward-evoked CF spiking in expert mice

(A–D) Each row shows the fluorescence (GrCs, A and C) or spike rate (CFs, B and D) of a single neuron, aligned to reward delivery and averaged across expert rewarded trials (3,965 GrCs, 1,964 CFs, 34 sessions/20 mice). Cells sorted by time of peak activity (A and B) or magnitude of anticipatory (C) or reward-evoked activity (D). Red lines denote sorting quantification windows.

(E and F) Activity quantifications either during delay (E) or post-reward (F), shown as histograms and (insets) binary statistical categories ($p < 0.05$ and magnitude exceeding ± 0.2 Z scores [zsc]). $^*p < 10^{-6}$. ns, $p = 0.54$.

(G–J) Rasters of anticipatory GrCs (G and I) and reward-activated CFs (H and J) on rewarded (G and H) or omitted-reward trials (I and J 20% omissions). 762 GrCs spread across 33/34 expert sessions in 19/20 mice; >0.1 zsc comparing $[-0.3, -0.03]$ s with both $[-1.3, -1]$ and $[+0.3, +0.5]$ s; 1,094 CFs spread across all 34 expert sessions/20 mice; spike rate during $[0, 0.2]$ s > 0.1 zsc, and > 0.1 zsc higher than pre-reward $[-0.3, -0.03]$ s. See also Figures S2N and S2O.

(K and L) For neurons in (G)–(J), averages across trials and cells. Gray lines show times of first CF peaks following reward delivery and omission, respectively.

(M) Dots show mean across trials of GrC anticipatory off-time (x axis, when fluorescence fell to $<50\%$ of peak over $[-0.5, 0]$ s), vs. time of first population CF spiking (y axis, when average rose above 20th percentile of the reward response $[0, 0.2]$ s). From sessions contributing to (G)–(J), 33/34 sessions, 19/20 mice, on trials with elevated CF reward spiking (117 ± 6 trials per session). Inset, average across sessions ($p = 1.6 \times 10^{-6}$ for both, 33 sessions).




Figure 4. With behavioral learning, GrCs increasingly spanned the delay, whereas CFs persistently signaled reward

(A and B) Average across GrCs with elevated delay-period activity (A, $[-0.3, -0.03] > [-1.5, -1.3]$ and $> [0.3, 0.5]$ s), on day 1 vs. day 7+ sessions. (B) Average across CFs with elevated reward spiking (from $[0, 0.2]$ s). Cell/session/mouse counts: (A) day 1: 752/15/15, day 7+: 1,135/33/19; (B) day 1: 657/15/15, day 7+: 1,462/34/20.

(C and D) With learning, GrC anticipation increased in magnitude (C, $p < 10^{-6}$, 752/1,135 day 1/7+ GrCs) and in prevalence ($p = 0.02$, 33/15 day 1/7+ sessions). CF reward responses increased in magnitude modestly (D, $p = 4 \times 10^{-4}$, 657/1,462 day 1/7+ CFs) and remained equally prevalent ($p = 0.56$, 15/34 day 1/7+ sessions). Significance proportions: thresholded magnitude differences from (A) and (B) at 0.2 zsc and single-cell $p < 0.05$; total #cells: 839/1,964 day 1/7+ CFs; 2,334/3,965 day 1/7+ GrCs. Dots show sessions.

(E) A cohort of mice trained for a week on the 1-s delay task, before switching to a 2-s delay for another week.

(F–H) 1-to-2-s retraining GrC responses in five mice. GrCs in each image are sorted by center of time points with elevated activity during $[0, 2]$ s, relative to pre-movement levels $[-0.8, -0.3]$ s. Cell/sessions counts: 257/5, 263/5, and 523/6. 2-s-novice data averaged over first 50 trials to highlight earliest exposures; 1-s-expert and 2-s-expert averages over random 50-trial subsets. Color-bar applies to all panels.

(I and J) Duration of GrC elevated activity between $[0, 2]$ s, relative to pre-movement levels $[-1, 0]$ s was significantly higher in 2-s-experts ($p < 10^{-6}$; 257, 263, and 523 GrCs, respectively).

(K) The proportion of GrCs with elevated activity for >1.6 s during $[0, 2]$ s rose with learning ($p = 0.02$, dots show 5/5/6 sessions, respectively).

(L) Average activity of reward anticipation GrCs (criteria: activity higher in final 0.3 s before reward compared with early in the delay $[0.1, 0.4]$ s; $n = 89/257$, 105/263, and 247/523 GrCs per learning phase). Dashed pink lines denote 1- or 2-s reward.

(M) Average activity of reward-responding CFs (rate $[0, 0.2]$ s >0.1 zsc higher than $[-0.3, -0.03]$ ms. 17/61, 34/71, and 35/143 CFs per condition).

elevated GrC activity largely spanned [0, 2] s following forelimb movement (Figures 4I–4L turquoise; elevated activity sustained for ~50% longer, from 0.8 ± 0.04 s to 1.2 ± 0.03 s). Thus, learning lengthened the duration of many individual GrCs' activity ramps. In addition, GrC differences between 1-s and 2-s delays were again inconsistent with licking motor signals per se: during [1, 2] s after movement, when 1-s experts licked robustly to consume reward, 2-s-experts licked less, as they delayed anticipatory licking—opposite to the differences in GrC activation (Figures S3E–S3I). Consistent with 1s-delay data (Figure 4B), 2-s novices and experts both exhibited strong CF reward responses (Figure 4M). Thus, over days of training, GrC anticipatory profiles lengthened to span the delay between reaching and reward, which evoked persistent time-locked CF spiking. We hypothesized that lengthening GrC profiles to span the delay might help CF-guided GrC readout in PkCs.

Learning increases reward timing information in GrC populations

After learning, many GrCs developed activity that spanned the delay, but for what computational purpose? Because animals successfully learned reward timing (Figure 2), we tested the quality of GrC reward time information via single-trial population linear time decoding (Figures 5A–5C; 10-fold cross-validated predictions). Learning substantially increased GrC delay-time decoding (by 133%; Figures 5D and 5E). Did time decoding generically improve, e.g., due to more consistent GrC activity? To test this, we decoded time after reward consumption and found persistently low accuracy (Figure 5F); by contrast, time decoding during reward consumption was already maximally accurate in novices (Figure S4A). Similarly, when comparing 1- and 2-s paradigms, we found that 2-s expert GrCs were best able to decode 2-s of time passage, whereas 1-s decoding was more similar across conditions (Figures S4B–S4F). Thus, learning specifically enhanced GrC timing accuracy during the learned delay.

Learning generated new GrC temporal bases that enhanced population reward timing information. Prior data demonstrated that PkCs^{62,63} and CN cells⁶⁴ ramp in ways that may track delay periods. How might PkCs appropriately integrate numerous GrC inputs to readout their population timing signals?

LTD simulated on anticipatory GrC bases predicts a gradient of GrC-PkC synaptic changes

Learning yielded GrCs with increasing reward timing information—could this information be extracted by PkCs? Among many other factors,^{65–68} a major contributor to PkC output is the integration of GrC inputs, partly directed by CF-driven synaptic plasticity. Classical cerebellar LTD weakens GrC → PkC synapses for GrCs that were active in the ~150 ms preceding a CF spike (for vermis¹⁸) (Figure 5G). We thus aimed to use the first simultaneous GrC-CF recordings to predict GrC → PkC LTD.

We recorded both the CF input and up to several hundred GrC inputs for several dozen PkCs. Therefore, for each PkC, we predicted LTD for each GrC input (simplifying by assuming that every GrC innervated every PkC). For each reward-evoked CF spike, we tabulated each GrC's mean activity in the preceding plasticity window as a prediction of LTD (Figure 5H). We rectified

these quantities to include only activity above baseline. We then computed the mean LTD for each GrC onto each PkC over all trials. To maintain overall PkC input strength, as likely achieved biologically by homeostatic⁶⁹ and opposing^{20,70} processes, such as LTP, we normalized the range of this vector (STAR Methods). Finally, we also computed the mean LTD across PkCs for each GrC, yielding a per session prediction of GrC → PkC synaptic weights.

To visually compare GrC profiles with their predicted LTD, we sorted GrCs by LTD onto the average PkC. Surprisingly, this ordered GrCs by delay activity timing—even when peak activity long preceded the LTD window (Figures 5I and S4G–S4I). To quantify this effect, we described each GrC by the center of its delay activity. We scattered this “anticipatory center time” against predicted LTD (Figures 5J and S4J–S4L). The two quantities were correlated and most strongly in experts—even for “center times” long preceding the LTD window (Figures 5K and S4O). As a control, this effect was absent when computing LTD using randomly reordered GrC traces (Figure 5K, gray). Thus, reward LTD is sensitive to GrC timing across the preceding delay.

This result reflects specific GrC properties in our task: activity levels in the 150 ms LTD window near reward provided a powerful snapshot of GrC activity timing up to 2 s prior. To visualize, we examined activity of GrCs grouped by predicted LTD magnitude (Figure 5L, normalized in magnitude to highlight differences in timing). Indeed, progressively lower-LTD GrC groups exhibited progressively earlier delay activity ramps—even as these ramps shifted hundreds of ms before the LTD window. However, these differences collapsed after reward.

Although this strategy is mathematically effective for the 1.1-s delay, the basis in Figure 5L would be ineffective for 2-s delays. We thus tested generalization to a longer delay by computing LTD on our 2-s-delay expert data. This revealed a GrC basis analogous to that in 1.1-s-experts but “stretched” in time (Figures 5M, S4M, and S4N). This demonstrates that learning a longer delay yielded a new GrC basis more suited to tracking the longer interval when guided by a similar but more temporally remote reward-evoked CF teaching signal.

Thus, we find that because learning yielded numerous GrCs that ramped activity up and down at different rates from forelimb movement until reward (Figure 3G), GrCs' activity at reward provided a snapshot of their prior temporal profiles—even for GrCs that peaked up to 2-s earlier (Figures 5M and S4O). Thus, CF-driven GrC integration might enable PkCs to track delay passage.

LTD from CF spiking at reward computationally suffices to readout GrC timing information

If LTD is sensitive to GrC anticipatory timing, what computation could this facilitate in PkCs? We produced a minimal readout by summing GrCs weighted by the LTD weight vectors (Figure 6A), thereby predicting a possible component of PkC output (Figures 6B and S5A–S5C). The resulting weighted sums of GrCs correlated substantially with time through the delay and increasingly so as learning progressed (Figures 6D, S5I, and S5J). In experts, LTD-weighted GrC sums resembled optimal

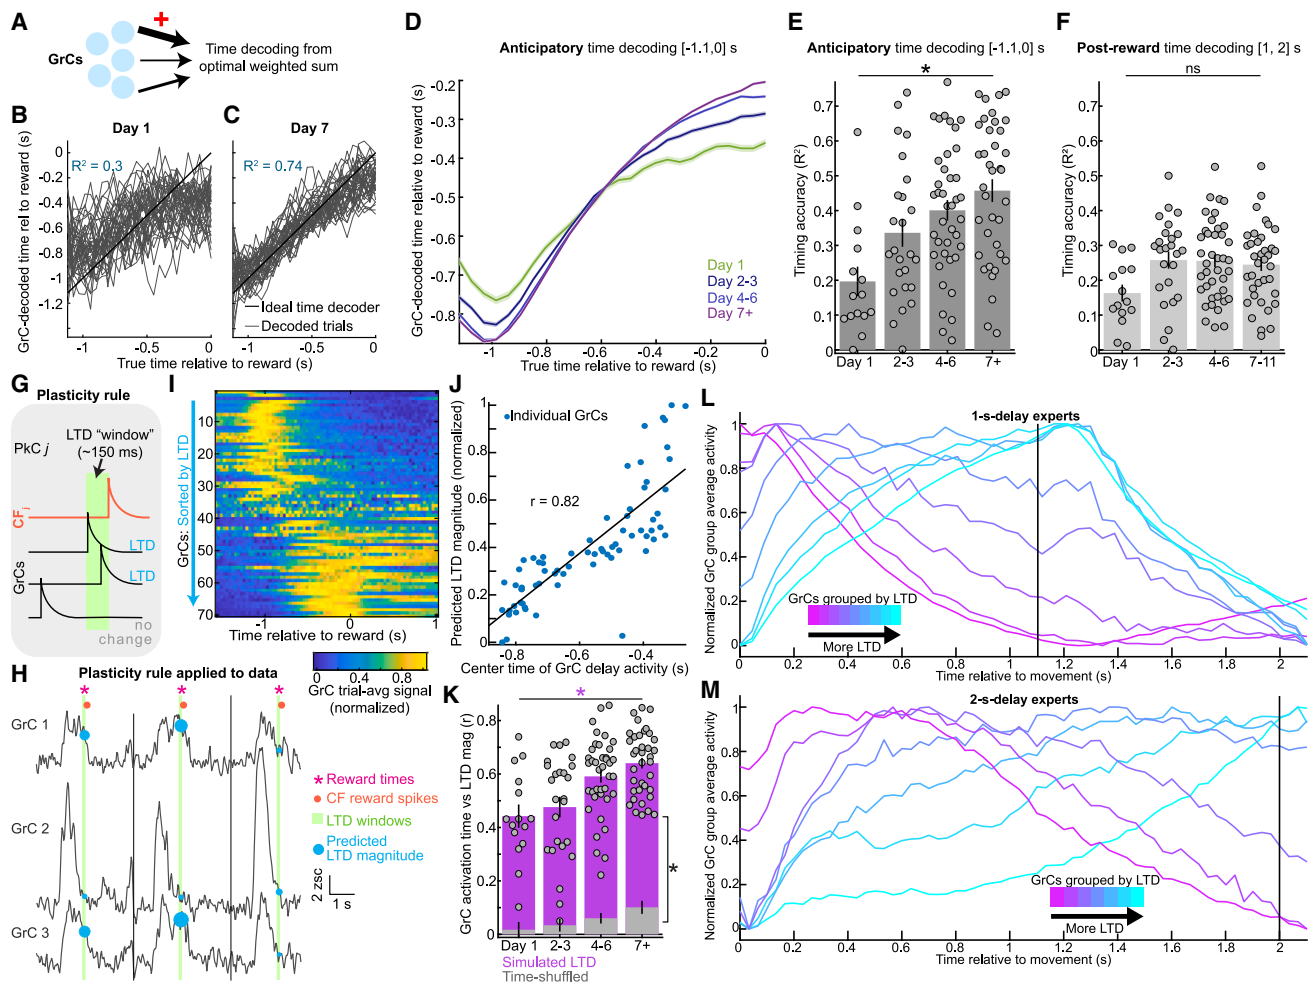


Figure 5. Learning yields GrC reward timing information that is computationally accessible to LTD

(A–C) Linear regression (A, least squares, 10-fold cross-validated) to decode time to reward via weighted sum of GrCs ([–1.1, 0] s). Example decoding performance day 1 (B) and day 7 (C) (41 trials, 127 and 111 cells).

(D) GrC time decoding output averaged across all trials (1,047, 2,242, 4,355, and 4,708 trials from 20 mice).

(E and F) GrC delay-time decoding accuracy (E) averaged across sessions ($p = 1.4 \times 10^{-5}$). Dots show 15, 25, 40, and 37 sessions). Accuracy post-reward (F, [1, 2] s) was persistently low ($p = 0.1$) and substantially poorer than anticipatory decoding in experts ($p < 10^{-6}$). See also [Figures S4A–S4D](#).

(G) CF-dependent GrC → PkC plasticity rule. When a PkC (j) receives a CF spike, GrC inputs active in the prior ~150 ms are weakened (LTD, top two GrCs), but other GrC inputs are not (bottom GrC).

(H) Simulating LTD on CF–GrC data. Example CF and GrCs in three trials centered on reward and concatenated (black lines denote trial breaks). Orange dots denote CF spikes within [0, 250] ms of reward. Each GrC's activity in green LTD window ([–150, 25] ms from CF spike) was tabulated as predicted LTD between that GrC and the CF-recipient PkC. A logistic function bounded each LTD event between [0, 1] ($\frac{1}{1+e^{-F/s}}$; F = avg GrC signal in LTD window; s = 95th percentile fluorescence per cell, [STAR Methods](#)).

(I) Example session: GrC profiles sorted top to bottom by predicted LTD averaged over trials roughly ordered GrCs by time of peak activity during the delay (70 GrCs, 95 trials). Additional examples, [Figures S4G–S4I](#).

(J) For session in (I), correlation between each GrC's center of delay activity (x axis, over [–1.1, 0] s) vs. its predicted LTD magnitude (y axis; 70 GrCs, $r = 0.82$, $p < 10^{-6}$, diagonal line, linear fit). Additional examples: [Figures S4J–S4L](#).

(K) Correlation from (J) across sessions grew with learning (magenta; $p < 10^{-6}$). Dots show 116 sessions; days 7+ positive r at $p < 10^{-6}$, 37 expert sessions). LTD computed on time-shuffled GrC data had smaller correlations ($p < 10^{-6}$, 37 expert sessions). See also [Figure S4O](#).

(L) GrCs grouped by predicted LTD magnitude (percentiles computed for each session; percentile bin edges: [0, 20, 40, 50, 60, 70, 80, 90, 95, and 100]). Traces: average activity per GrC group (normalized to [0, 1] to highlight differences in timing). Cell counts from bottom to top: 2,077, 2,068, 1,025, 1,057, 1,032, 1,036, 1,038, 517, and 522; 76 day 4+ sessions. See [Figures S4M and S4N](#).

(M) Same as (L), for 2-s experts (bin edges: [0, 30, 50, 65, 80, 90, 95, and 100]; counts: 157, 103, 80, 78, 53, 26, and 26; 6 sessions).

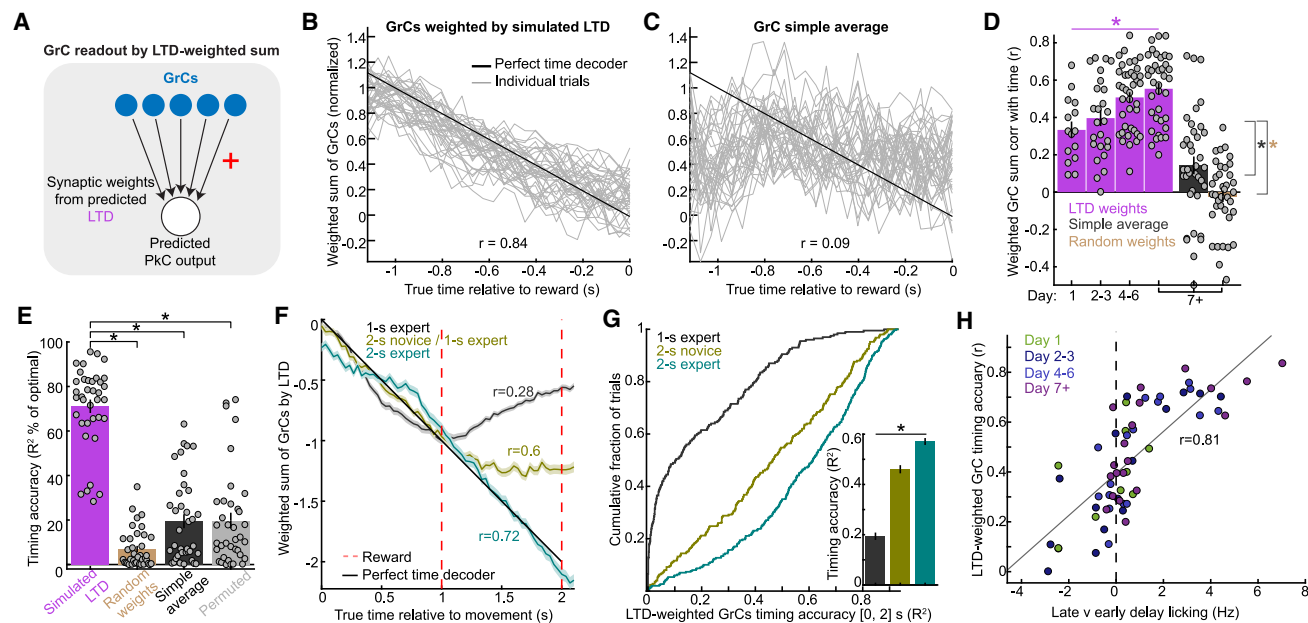


Figure 6. Simulated LTD-weighted GrC averages track time to reward for up to 2 s

(A–C) GrC readout using LTD predictions of GrC → PkC weights (A). Example day-7 single-trial LTD-weighted GrC sums (B) or simple GrC average (C). Correlation: time vs. weighted sums (40 trials shown of 68 total, 111 GrCs). Normalization: range of trial-average scaled to [1.1, 0]. Additional examples, Figures S5A–S5C.

(D) Correlation between time and LTD-weighted GrC sums rose with learning ($p = 1.6 \times 10^{-5}$, dots show 116 sessions from 20 mice). Simple GrC average or averages using randomly reordered LTD weights (random weights) poorly correlated with time (weaker than LTD weighted, $p < 10^{-6}$).

(E) In experts, GrC sums using LTD weights were far closer to optimal time decoding than those using random weights, simple averages, or weights from LTD computed on time-shuffled GrCs ($p < 10^{-6}$, dots show 37 sessions).

(F) For 1-to-2-s retraining data, LTD-weighted GrC sums using CF reward spikes (i.e., either 1- or 2-s-post-movement). Trial counts: 250, 250, and 299; mice/sessions: 5/5, 5/5, and 5/6.

(G) LTD-weighted GrC sum timing accuracy (R^2) over [0, 2] s for data in (F) ($p < 10^{-6}$, 799 trials).

(H) Session-by-session behavioral performance (late [−0.2, 0] s minus early [−0.8, −0.6] s delay licking) vs. LTD-weighted GrC sum timing accuracy. Spearman $r = 0.81$, $p < 10^{-6}$, 62 sessions. Diagonal line, linear fit.

GrC time readouts in both accuracy and weights (Figures S5A–S5G and S5I–S5L). Thus, learning led to GrC activity patterns that were increasingly well suited to LTD-based readout of delay passage until reward.

Given the prevalence and prominence of GrC anticipation signals in experts (Figures 3 and 4), might far-simpler GrC readouts suffice to extract timing information? We considered three trivial readouts: (1) simple average of all GrCs (Figure 6C), (2) GrCs weighted randomly (by reordering the above LTD weights), or (3) LTD simulated on time-shuffled GrCs. In each case, resulting GrC readouts only weakly correlated with time in experts (Figures 6D and 6E). Thus, although learning enhanced GrC timing signals, this information could not be extracted trivially. Therefore, PkCs cannot automatically inherit timing signals, but setting synaptic strengths via LTD (with likely symmetric contributions from LTP, Figure S5K) sufficed to integrate GrCs into interval-tracking signals.

Are LTD-weighted GrC sums specific to the delay duration? We computed LTD on data from mice trained on a 1-s delay followed by retraining on a 2-s delay (Figure 6F). This yielded three results: (1) in 1-s-delay experts, GrC sums tracked time from [0, 1] s, but the readout decayed from [1, 2] s; (2) in 2-s-delay novices, GrC sums tracked time from [0, 1] s but then saturated from

[1, 2] s; and (3) only in 2-s experts did GrC sums track time throughout [0, 2] s. Accuracy of 2-s timing thus improved substantially (Figure 6G). LTD triggered near reward could therefore extract at least 2 s of GrC timing signals—but only after learning appropriately lengthens GrC activity patterns to span the delay. Finally, to test whether readout quality related to behavior, we quantified lick timing specificity, which covaried with timing accuracy of LTD-weighted GrC sums (Figures 6H and S5H). Together, these data demonstrate that learning endows GrCs with anticipatory timing signals that could be integrated into a delay-tracking output via classical plasticity driven by reward-evoked CF spiking.

PkC delay spiking ramps bear out predictions of LTD readout of GrC timing signals

To test the prediction that a component of PkC simple spiking (SS) could track time to reward, we used Neuropixels (imec) to record PkCs in 1-s-delay experts (Figures S6A and S6B). We identified confirmed PkC SS units by their SS pauses after complex spikes (64 cells, Figures S6C–S6G) and putative PkCs based on their physical proximity ($\pm 100 \mu\text{m}$) to confirmed PkCs⁷¹ and similar spiking metrics (98 cells, Figures S6H–S6J and S7E). We computed mean SS rates aligned to reward in

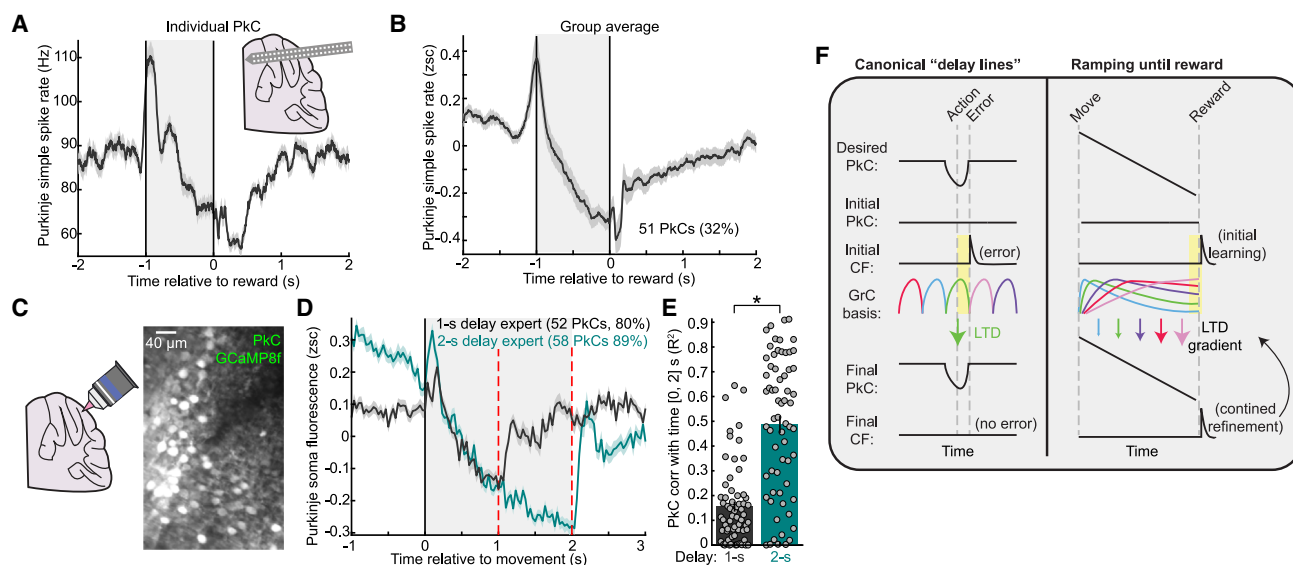


Figure 7. PkC simple spike ramps track interval from movement to expected reward for up to 2 s

(A and B) Neuropixels PkC recordings. (A) Example expert PkC simple spike (SS) rate ramped downward during the delay (124 trials). (B) Average Z-scored SS rate of PkCs whose delay SS rate decreased below baseline (negative slope $[-1, 0]$ s and negative zsc $[-150, -25]$ ms; 32%, 51/162 PkCs in 4 mice/16 day 7+ sessions; positive ramping cells in Figure S7).

(C–E) Alternatively, we imaged PkC two-photon somatic Ca^{2+} in the region of our GrC-CF recordings, (C, mean *in vivo* image). (D) Trial-averaged fluorescence for PkCs with negative slope from $[0, 1]$ s for 1-s-expert data (65/4/2 total PkCs/sessions/mice) or $[0, 2]$ s for 2-s-expert data (65/3/2 total PkCs/sessions/mice). Timing accuracy $[0, 2]$ s was higher in 2-s-expert PkCs (E, $p < 10^{-6}$).

(F) Schematic. Left, a canonical strategy using a “delay line” GrC basis: erroneous actions trigger CF spiking. GrCs sequentially activate at distinct times for short durations. Learning adjusts synapses of GrCs temporally coincident with CF signal, eliminating both future error and CF error signal.

Right, strategy to produce PkC delay-tracking ramps from action to reward. During learning, GrC profiles lengthen to densely span the delay with varying kinetics. Pre-reward LTD window provides snapshot of GrC ramping kinetics (similar to Figures 5L and 5M). Using classical LTD, CF reward spiking grades many GrC \rightarrow PkC synapses by GrC anticipatory timing. LTD-weighted GrCs yield PkC spiking ramps from predictor to reward. Thus, new GrC basis sets that emerge with learning enable new types of PkC computation.

experts (Figure S7A). In both individual PkCs (Figure 7A) and among 32% of all PkCs (Figure 7B), SS rates gradually decreased with delay passage, as predicted (Figure 6B). Interestingly, an additional 25% of PkCs exhibited positive ramping SS, which tracked the delay with polarity opposite to our prediction (Figure S7B, discussion). SS rates of many PkCs thus tracked the delay as predicted by the readout of the GrC temporal basis.

Due to cerebellar and Neuropixels probe geometry, our electrophysiological recordings spanned multiple lobules and were more ventral than our dorsal-surface imaging sites (Figure S6A), such that few, if any, PkC recordings likely derived from the small region where we performed GrC-CF imaging. As an alternative, we used two-photon imaging to densely sample PkCs in dorsal-surface right vermis LVI. By monitoring Ca^{2+} in PkC somas with jRCaMP8f,⁷² we obtained correlates of PkC SS rates⁷³ (Figure 7C). Most PkCs exhibited downward delay ramping of somatic fluorescence (Figures 7D and S7F). Quantified by fluorescence correlation with time or its magnitude near reward, $\sim 80\%$ of PkCs negatively ramped with delay passage (Figures S7G and S7H), as predicted by GrC-CF signals in the same region.

To test the prediction that learning generates PkC ramps to span the learned delay duration (Figures 6F and 6G), we retrained PkC somatic imaging mice on a 2-s delay. As predicted and unlike 1-s experts, 2-s-expert PkC fluorescence progres-

sively decreased through the 2-s delay and returned to baseline after reward (Figure 7D, green). Thus, only 2-s-expert PkCs conveyed 2 s of timing information (Figures 7E and S7I). The differences in PkC activity between 1-s vs. 2-s experts were unlikely to be explained by licking because PkCs exhibited large differences between 1-s and 2-s delays during $[1, 2]$ s after movement that qualitatively conflicted with corresponding differences in licking (Figures S3F and S3G). Overall, PkC recordings bore out multiple predictions of LTD-based readout of delay-ramping GrCs. Because many PkCs tracked the delay until reward (Figures 7A–7E) and because trivial GrC readouts cannot produce such signals (Figures 6D and 6E), it is thus likely that GrC integration via LTD and/or computationally similar mechanisms, such as LTP (Figure S5K), contribute to delay-tracking PkC SS ramps.

DISCUSSION

Using simultaneous GrC-CF imaging during operant conditioning for delayed reward, we found that learning yielded numerous GrCs that ramped activity for up to 2 s from forelimb movement until reward; reward then triggered widespread time-locked CF spiking (Figures 3 and 4). Via simulation, we found that canonical LTD applied to this GrC temporal basis could grade GrC \rightarrow PkC synaptic weights by GrC anticipatory timing (Figure 5). We thus

predicted and confirmed that many PkCs could generate interval-tracking SS ramps from movement until reward (Figures 6 and 7). Newly learned GrC bases were also specific to behavioral timescale: GrCs in 1-s-experts were ill suited to 2-s interval tracking, but 2-s-delay retraining stretched GrC profiles to span the longer delay (Figures 4, 5, and S4). This process mirrored behavioral learning (Figures 2 and S2) and accounted for PkC ramp differences between 1-s and 2-s delays (Figures 7 and S7). Given the behavioral importance of long-interval-tracking, ubiquitous neural ramping in the forebrain and its importance for perceived time passage,^{74,75} and the recent findings of cerebellar PkC and CN neural ramps,^{62–64,76,77} these GrC temporal bases and their interplay with CF reward signals may be common elements of cerebellar computation.

Learning yields GrCs that link predictors to temporally remote CF reward signals

Decades of theory posited GrC representations of time,^{28–36} such as the “delay line” GrC basis²⁶ (Figure 7F, left). Contrastingly, in our expert data, up to 40% of GrCs simultaneously activated for up to 2 s until reward (Figure 7F, right). Because different anticipatory GrCs activated with different kinetics, each GrC’s activity level just before reward provided a snapshot of its prior temporal profile (Figures 5L and 5M). Rather than redundancy, dense GrC activity more finely sampled the delay. By sustaining activity from movement until reward, GrCs that activated early were linked to reward-evoked CF spiking much later. Thus, GrCs in our data are well suited to generate interval-spanning ramps via a remote CF instructive signal, whereas a delay line basis will generally not (Figures 7F, left, and S5M–S5P). Because these GrC bases scaled temporally across delay durations (Figure 4), it is possible that this basis might also “compress” to shorter durations and thus help drive anticipatory PkC suppression in classical cerebellar short-latency associative learning.^{78,79}

Our central finding is that learning yielded new GrC activity profiles that rose and fell over seconds to link a predictor to temporally remote reward-evoked CF teaching signals. This claim is agnostic to the mechanism for GrC alterations. Nevertheless, sustained activity ramps often reflect recurrent excitatory circuits,^{80,81} as found in the cortico-cerebellar loop and the nucleocortical loop, and the latter exhibits structural plasticity that could help amplify specific mossy fiber input patterns to GrCs.⁸²

Persistent CF reward signals

Canonically, CFs were thought to signal errors, such that learning eliminates both error and CF error signal. Often, CF representations are more complex.^{61,83–86} Here, reward-evoked CF spiking persisted throughout learning, contrary to an error (Figure 4). We posit two possible explanations for this result. First, expert CF spiking could help transition or initiate subsequent behavioral states, in this case, from reward anticipation to consumption. Intriguingly, other studies have also reported behaviorally time-locked CF spiking that either emerged, solidified, or persisted during learning, with stronger expert CF spiking sometimes associated with better performance.^{42,60,79,83} Second, expert CF reward spiking could help to continually refine

GrC integration. In contrast to a sparse GrC basis where persistent CFs might excessively weaken target GrCs relative to those active just outside the LTD window (Figure 7F, left), we predict that CF reward spiking grades many GrC→PkC synapses at once (Figure 7F, right). If paired with homeostatic or opposing mechanisms^{20,69} to maintain overall PkC input strength, this strategy could yield stable synaptic weights and learning.

Impact of delay-spanning GrCs on PkCs

Recent studies found extended ramping in PkCs^{62,63,76} and CN^{64,77} neurons, raising the question whether ramps were “inherited” from upstream or could result from cerebellar computation. PkCs in our task also ramped over seconds-long delays (Figures 7 and S7). But our GrC-CF recordings permitted two key discoveries: (1) interval-tracking PkCs cannot be generated by trivial GrC readout (Figures 6C–6E), despite widespread sustained GrC activity, but (2) LTD from CF reward spiking computationally sufficed to integrate diverse GrCs into seconds-long-interval-tracking PkC ramps (Figure 6F). We thus posit that learning involves both progressive changes to GrC activity and changes in its readout in PkCs. An open question is whether these changes occur together or sequentially (Figure S5H).

A limitation of our technique is that we sampled GrC inputs physically closest to CF-recipient PkCs, but GrCs contribute from distances up to ~1 mm.⁸⁷ Although GrC profiles in our task varied minimally at this spatial scale,^{9,41} we cannot exclude differential roles for distant GrC inputs. Furthermore, our simulation assumed the classical plasticity rule in which CF spikes lead to GrC→PkC LTD, but our data do not provide direct evidence that CF reward spikes cause LTD. Nonetheless, although LTD is modulated by factors unobserved in this study, it is generally thought that most CF spikes lead to some LTD.⁸⁸ In addition, although we predicted GrC→PkC connection strengths, no technique can yet quantify connection strengths among hundreds of cells registered to their activity during learning. Instead, we tested and confirmed predictions about PkC ramping. Although many PkCs bore out predicted delay ramps, in our electrophysiological recordings, some also ramped upward—unexplainable by LTD alone but consistent with prior findings of bidirectional PkC modulation.²² Such PkC features could reflect numerous synergistic mechanisms: the balance of LTD vs. LTP (triggered by reduced CF spiking^{20–22}; Figure S5K); interplay of GrCs and interneurons, which could contribute to PkC ramping by driving SS below baseline (Figure 7); lasting rises in PkC Ca²⁺⁸⁹; and changes in PkC excitability.⁶⁹ On the other hand, upward-ramping PkCs were rare when imaging near the site of our GrC-CF recordings (Figures 7D, 7E, and S7F–S7I). Functionally, all ramps may help track intervals between events.

Function of ramping GrC bases and delay-tracking PkCs

The cerebellum has long been thought to synchronize actions to a short latency prior to an event⁹⁰ (Figure 7F, left). Increasing data suggest that the cerebellum may also contribute to continuous timing functions where ramping signals predominate.^{62–64,76,77} In our data, cerebellar delay ramping resembled a temporally building expectation of reward (Figure S3). Nevertheless, our data cannot adjudicate whether the cerebellum

generates such ramps (1) solely when needed to drive temporally building licking when awaiting reward or (2) also to help maintain internal estimates of time passage generally. We favor (2) because PkCs ramp in differing behavioral contexts and cerebellar regions^{64,77,91}; cerebellar ramps are tightly linked to those in neocortex,^{9,64,91} which have generalized temporal expectancy functions^{92–94}; and the diversity of CN targets^{95,96} implies broader utility than driving a fixed set of motor effectors.

How might such signals influence behavior? Activating PkCs during the delay transiently terminated licking, potentially by inhibiting CN cells projecting to areas driving motor output. The transient manipulation also disrupted subsequent lick timing relative to expected reward (Figures 2J, and S2J), possibly by disturbing internal estimates of time passage. Finally, repeatedly perturbing PkC activity prevented the relearning of new reward timing even when probed on laser-off trials (Figure 2M), possibly by disrupting signals used for temporal learning. However, it is important to note that this perturbation disrupted spontaneous PkC spiking in addition to GrC-driven spiking, leaving the possibility that direct GrC perturbation might yield a different constellation of effects.

Long-interval time tracking via ramping signals is supported by decades of forebrain studies, indicating its importance for both estimation and perception of time in multiple species.^{74,75} For the cerebellum, we demonstrate that learning leads to new sustained GrC bases that link a predictor to distant reward-evoked CF teaching signals—a strategy well suited to learning to track long time intervals. Cerebellar learning of interval-tracking ramps may characterize its interactions with the forebrain and contributions in cognitive domains generally.

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- KEY RESOURCES TABLE
- RESOURCE AVAILABILITY
 - Lead contact
 - Materials availability
 - Data and code availability
- EXPERIMENTAL MODEL DETAILS
 - Mice
- METHOD DETAILS
 - Virus
 - Viral injection
 - Histology
 - Window and headplate implantation
 - Behavioral data collection
 - Behavioral training
 - Two-photon microscopy
 - PkC somatic Ca²⁺ imaging and analysis using jRCaMP1f
 - Optogenetic studies
 - Neuropixels recordings
- QUANTIFICATION AND STATISTICAL ANALYSIS
 - Image processing
 - Cell identification and signal preprocessing
 - PkC somatic Ca²⁺ imaging analysis
 - Neural response analysis
 - Neural learning analysis
 - Temporal decoding analysis
 - Neuropixels analysis

- Statistics
- Plasticity simulations

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at <https://doi.org/10.1016/j.neuron.2024.05.019>.

ACKNOWLEDGMENTS

We thank H. Hirai for the L7-6 plasmid, H. Bito for the R-CaMP2 plasmid, C.K. Kim for plasmid cloning, and NIMH SI for fabrication. We thank J.S. Diamond, Y. Gu, A.J. Levine, E.S. Giniger, Z.M. Khaliq, R.J. Krauzlis, M.R. Carey, J.F. Medina, and S.G. Lisberger for comments on the manuscript. This research was supported in part by the Intramural Research Program of the NIH. Additional funding: to M.G.G.-G, an NINDS NCF; to M.J.W., a BWF CASI award and BBRF YIA; and to L.L. and M.J.S., NIH and HHMI funding.

AUTHOR CONTRIBUTIONS

M.J.W. conceived of and supervised the study. M.G.G.-G, A.K., O.A., and M.J.W. performed experiments. L.T., A.K., O.A., C.P., and M.J.W. processed data. L.T. and M.J.W. analyzed data. L.T., M.J.W., and A.L.-K. performed simulations. T.H.K. designed the optical approach. L.L. and M.J.S. advised and contributed reagents and equipment.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: January 1, 2024

Revised: March 31, 2024

Accepted: May 16, 2024

Published: June 12, 2024

REFERENCES

1. Raymond, J.L., and Medina, J.F. (2018). Computational Principles of Supervised Learning in the Cerebellum. *Annu. Rev. Neurosci.* 41, 233–253. <https://doi.org/10.1146/annurev-neuro-080317-061948>.
2. Hull, C. (2020). Prediction signals in the cerebellum: beyond supervised motor learning. *eLife* 9, e54073. <https://doi.org/10.7554/eLife.54073>.
3. Doya, K. (2000). Complementary roles of basal ganglia and cerebellum in learning and motor control. *Curr. Opin. Neurobiol.* 10, 732–739. [https://doi.org/10.1016/S0959-4388\(00\)00153-7](https://doi.org/10.1016/S0959-4388(00)00153-7).
4. Sokolov, A.A., Miall, R.C., and Ivry, R.B. (2017). The Cerebellum: Adaptive Prediction for Movement and Cognition. *Trends Cogn. Sci.* 21, 313–332. <https://doi.org/10.1016/j.tics.2017.02.005>.
5. Ivry, R.B., and Keele, S.W. (1989). Timing functions of the cerebellum. *J. Cogn. Neurosci.* 1, 136–152. <https://doi.org/10.1162/jocn.1989.1.2.136>.
6. Ishikawa, T., Shimuta, M., and Häusser, M. (2015). Multimodal sensory integration in single cerebellar granule cells in vivo. *eLife* 4, e12916. <https://doi.org/10.7554/eLife.12916>.
7. Huang, C.C., Sugino, K., Shima, Y., Guo, C., Bai, S., Mensh, B.D., Nelson, S.B., and Hantman, A.W. (2013). Convergence of pontine and proprioceptive streams onto multimodal cerebellar granule cells. *eLife* 2, e00400. <https://doi.org/10.7554/eLife.00400>.
8. Kelly, R.M., and Strick, P.L. (2003). Cerebellar loops with motor cortex and prefrontal cortex of a nonhuman primate. *J. Neurosci.* 23, 8432–8444. <https://doi.org/10.1523/JNEUROSCI.23-23-08432.2003>.
9. Wagner, M.J., Kim, T.H., Kadmon, J., Nguyen, N.D., Ganguli, S., Schnitzer, M.J., and Luo, L. (2019). Shared Cortex-Cerebellum Dynamics in the Execution and Learning of a Motor Task. *Cell* 177, 669–682.e24. <https://doi.org/10.1016/j.cell.2019.02.019>.

10. Hirano, T. (2018). Purkinje Neurons: Development, Morphology, and Function. *Cerebellum* 17, 699–700. <https://doi.org/10.1007/s12311-018-0985-7>.
11. Fujita, M. (1982). Adaptive filter model of the cerebellum. *Biol. Cybern.* 45, 195–206. <https://doi.org/10.1007/BF00336192>.
12. Dean, P., Porrill, J., Ekerot, C.-F., and Jörntell, H. (2010). The cerebellar microcircuit as an adaptive filter: experimental and computational evidence. *Nat. Rev. Neurosci.* 11, 30–43. <https://doi.org/10.1038/nrn2756>.
13. Sejnowski, T.J. (1977). Storing covariance with nonlinearly interacting neurons. *J. Math. Biol.* 4, 303–321. <https://doi.org/10.1007/BF00275079>.
14. Ito, M., Sakurai, M., and Tongroach, P. (1982). Climbing fibre induced depression of both mossy fibre responsiveness and glutamate sensitivity of cerebellar Purkinje cells. *J. Physiol.* 324, 113–134. <https://doi.org/10.1113/jphysiol.1982.sp014103>.
15. Coesmans, M., Weber, J.T., De Zeeuw, C.I., and Hansel, C. (2004). Bidirectional Parallel Fiber Plasticity in the Cerebellum under Climbing Fiber Control. *Neuron* 44, 691–700. <https://doi.org/10.1016/j.neuron.2004.10.031>.
16. Jörntell, H., and Ekerot, C.-F. (2002). Reciprocal Bidirectional Plasticity of Parallel Fiber Receptive Fields in Cerebellar Purkinje Cells and Their Afferent Interneurons. *Neuron* 34, 797–806. [https://doi.org/10.1016/S0896-6273\(02\)00713-4](https://doi.org/10.1016/S0896-6273(02)00713-4).
17. Medina, J.F., and Lisberger, S.G. (2008). Links from complex spikes to local plasticity and motor learning in the cerebellum of awake-behaving monkeys. *Nat. Neurosci.* 11, 1185–1192. <https://doi.org/10.1038/nn.2197>.
18. Suvrathan, A., Payne, H.L., and Raymond, J.L. (2016). Timing Rules for Synaptic Plasticity Matched to Behavioral Function. *Neuron* 92, 959–967. <https://doi.org/10.1016/j.neuron.2016.10.022>.
19. Rowan, M.J.M., Bonnan, A., Zhang, K., Amat, S.B., Kikuchi, C., Taniguchi, H., Augustine, G.J., and Christie, J.M. (2018). Graded Control of Climbing-Fiber-Mediated Plasticity and Learning by Inhibition in the Cerebellum. *Neuron* 99, 999–1015.e6. <https://doi.org/10.1016/j.neuron.2018.07.024>.
20. Gutierrez-Castellanos, N., Da Silva-Matos, C.M., Zhou, K., Canto, C.B., Renner, M.C., Koene, L.M.C., Ozyildirim, O., Sprengel, R., Kessels, H.W., and De Zeeuw, C.I. (2017). Motor Learning Requires Purkinje Cell Synaptic Potentiation through Activation of AMPA-Receptor Subunit GluA3. *Neuron* 93, 409–424. <https://doi.org/10.1016/j.neuron.2016.11.046>.
21. Gao, Z., van Beugen, B.J., and De Zeeuw, C.I. (2012). Distributed synergistic plasticity and cerebellar learning. *Nat. Rev. Neurosci.* 13, 619–635. <https://doi.org/10.1038/nrn3312>.
22. De Zeeuw, C.I. (2021). Bidirectional learning in upbound and downbound microzones of the cerebellum. *Nat. Rev. Neurosci.* 22, 92–110. <https://doi.org/10.1038/s41583-020-00392-x>.
23. Lisberger, S.G. (1998). Cerebellar LTD: A Molecular Mechanism of Behavioral Learning? *Cell* 92, 701–704. [https://doi.org/10.1016/S0092-8674\(00\)81397-5](https://doi.org/10.1016/S0092-8674(00)81397-5).
24. Koekoek, S.K.E., Hulscher, H.C., Dortland, B.R., Hensbroek, R.A., Elgersma, Y., Ruigrok, T.J.H., and De Zeeuw, C.I. (2003). Cerebellar LTD and Learning-Dependent Timing of Conditioned Eyelid Responses. *Science* 301, 1736–1739. <https://doi.org/10.1126/science.1088383>.
25. Kalmbach, B.E., Voicu, H., Ohshima, T., and Mauk, M.D. (2011). A Subtraction Mechanism of Temporal Coding in Cerebellar Cortex. *J. Neurosci.* 31, 2025–2034. <https://doi.org/10.1523/JNEUROSCI.4212-10.2011>.
26. Yamazaki, T., and Tanaka, S. (2009). Computational models of timing mechanisms in the cerebellar granular layer. *Cerebellum* 8, 423–432. <https://doi.org/10.1007/s12311-009-0115-7>.
27. Albergaria, C., Silva, N.T., Pritchett, D.L., and Carey, M.R. (2018). Locomotor activity modulates associative learning in mouse cerebellum. *Nat. Neurosci.* 21, 725–735. <https://doi.org/10.1038/s41593-018-0129-x>.
28. Barri, A., Wiechert, M.T., Jazayeri, M., and DiGregorio, D.A. (2022). Synaptic basis of a sub-second representation of time in a neural circuit model. *Nat. Commun.* 13, 7902. <https://doi.org/10.1038/s41467-022-35395-y>.
29. Medina, J.F., and Mauk, M.D. (2000). Computer simulation of cerebellar information processing. *Nat. Neurosci.* 3 (Suppl), 1205–1211. <https://doi.org/10.1038/81486>.
30. Bullock, D., Fiala, J.C., and Grossberg, S. (1994). A neural model of timed response learning in the cerebellum. *Neural Netw.* 7, 1101–1114. [https://doi.org/10.1016/S0893-6080\(05\)80161-3](https://doi.org/10.1016/S0893-6080(05)80161-3).
31. Marr, D. (1969). A theory of cerebellar cortex. *J. Physiol.* 202, 437–470. <https://doi.org/10.1113/jphysiol.1969.sp008820>.
32. Albus, J.S. (1971). A theory of cerebellar function. *Math. Biosci.* 10, 25–61. [https://doi.org/10.1016/0025-5564\(71\)90051-4](https://doi.org/10.1016/0025-5564(71)90051-4).
33. Babadi, B., and Sompolsky, H. (2014). Sparseness and expansion in sensory representations. *Neuron* 83, 1213–1226. <https://doi.org/10.1016/j.neuron.2014.07.035>.
34. Billings, G., Piasini, E., Lörincz, A., Nusser, Z., and Silver, R.A. (2014). Network structure within the cerebellar input layer enables lossless sparse encoding. *Neuron* 83, 960–974. <https://doi.org/10.1016/j.neuron.2014.07.020>.
35. Cayco-Gajic, N.A., Clopath, C., and Silver, R.A. (2017). Sparse synaptic connectivity is required for decorrelation and pattern separation in feed-forward networks. *Nat. Commun.* 8, 1116. <https://doi.org/10.1038/s41467-017-01109-y>.
36. Lanore, F., Cayco-Gajic, N.A., Gurnani, H., Coyle, D., and Silver, R.A. (2021). Cerebellar granule cell axons support high-dimensional representations. *Nat. Neurosci.* 24, 1142–1150. <https://doi.org/10.1038/s41593-021-00873-x>.
37. Kennedy, A., Wayne, G., Kaifosh, P., Alviña, K., Abbott, L.F., and Sawtell, N.B. (2014). A temporal basis for predicting the sensory consequences of motor commands in an electric fish. *Nat. Neurosci.* 17, 416–422. <https://doi.org/10.1038/nn.3650>.
38. Knogler, L.D., Markov, D.A., Dragomir, E.I., Štih, V., and Portugues, R. (2017). Sensorimotor Representations in Cerebellar Granule Cells in Larval Zebrafish Are Dense, Spatially Organized, and Non-temporally Patterned. *Curr. Biol.* 27, 1288–1302. <https://doi.org/10.1016/j.cub.2017.03.029>.
39. Lin, Q., Manley, J., Helmreich, M., Schlumm, F., Li, J.M., Robson, D.N., Engert, F., Schier, A., Nöbauer, T., and Vaziri, A. (2020). Cerebellar Neurodynamics Predict Decision Timing and Outcome on the Single-Trial Level. *Cell* 180, 536–551.e17. <https://doi.org/10.1016/j.cell.2019.12.018>.
40. Giovannucci, A., Badura, A., Deverett, B., Najafi, F., Pereira, T.D., Gao, Z., Ozden, I., Kloth, A.D., Pnevmatikakis, E., Paninski, L., et al. (2017). Cerebellar granule cells acquire a widespread predictive feedback signal during motor learning. *Nat. Neurosci.* 20, 727–734. <https://doi.org/10.1038/nn.4531>.
41. Wagner, M.J., Kim, T.H., Savall, J., Schnitzer, M.J., and Luo, L. (2017). Cerebellar granule cells encode the expectation of reward. *Nature* 544, 96–100. <https://doi.org/10.1038/nature21726>.
42. Wagner, M.J., Savall, J., Hernandez, O., Mel, G., Inan, H., Rumyantsev, O., Lecoq, J., Kim, T.H., Li, J.Z., Ramakrishnan, C., et al. (2021). A neural circuit state change underlying skilled movements. *Cell* 184, 3731–3747.e21. <https://doi.org/10.1016/j.cell.2021.06.001>.
43. Moberget, T., and Ivry, R.B. (2016). Cerebellar contributions to motor control and language comprehension: searching for common computational principles. *Ann. N. Y. Acad. Sci.* 1369, 154–171. <https://doi.org/10.1111/nyas.13094>.

44. Schmahmann, J.D., Guell, X., Stoodley, C.J., and Halko, M.A. (2019). The Theory and Neuroscience of Cerebellar Cognition. *Annu. Rev. Neurosci.* 42, 337–364. <https://doi.org/10.1146/annurev-neuro-070918-050258>.
45. Guell, X., Gabrieli, J.D.E., and Schmahmann, J.D. (2018). Triple representation of language, working memory, social and emotion processing in the cerebellum: convergent evidence from task and seed-based resting-state fMRI analyses in a single large cohort. *NeuroImage* 172, 437–449. <https://doi.org/10.1016/j.neuroimage.2018.01.082>.
46. Stoodley, C.J., MacMore, J.P., Makris, N., Sherman, J.C., and Schmahmann, J.D. (2016). Location of lesion determines motor vs. cognitive consequences in patients with cerebellar stroke. *NeuroImage Clin.* 12, 765–775. <https://doi.org/10.1016/j.nicl.2016.10.013>.
47. Stoodley, C.J., and Schmahmann, J.D. (2009). Functional topography in the human cerebellum: a meta-analysis of neuroimaging studies. *NeuroImage* 44, 489–501. <https://doi.org/10.1016/j.neuroimage.2008.08.039>.
48. Stoodley, C.J., and Schmahmann, J.D. (2010). Evidence for topographic organization in the cerebellum of motor control versus cognitive and affective processing. *Cortex* 46, 831–844. <https://doi.org/10.1016/j.cortex.2009.11.008>.
49. King, M., Hernandez-Castillo, C.R., Poldrack, R.A., Ivry, R.B., and Diedrichsen, J. (2019). Functional boundaries in the human cerebellum revealed by a multi-domain task battery. *Nat. Neurosci.* 22, 1371–1378. <https://doi.org/10.1038/s41593-019-0436-x>.
50. Shuster, S.A., Wagner, M.J., Pan-Doh, N., Ren, J., Grutzner, S.M., Beier, K.T., Kim, T.H., Schnitzer, M.J., and Luo, L. (2021). The relationship between birth timing, circuit wiring, and physiological response properties of cerebellar granule cells. *Proc. Natl. Acad. Sci. USA* 118, e2101826118. <https://doi.org/10.1073/pnas.2101826118>.
51. Inoue, M., Takeuchi, A., Horigane, S., Ohkura, M., Gengyo-Ando, K., Fujii, H., Kamijo, S., Takemoto-Kimura, S., Kano, M., Nakai, J., et al. (2015). Rational design of a high-affinity, fast, red calcium indicator R-CaMP2. *Nat. Methods* 12, 64–70. <https://doi.org/10.1038/nmeth.3185>.
52. Nitta, K., Matsuzaki, Y., Konno, A., and Hirai, H. (2017). Minimal Purkinje Cell-Specific PCP2/L7 Promoter Virally Available for Rodents and Non-human Primates. *Mol. Ther. Methods Clin. Dev.* 6, 159–170. <https://doi.org/10.1016/j.omtm.2017.07.006>.
53. Mukamel, E.A., Nimmerjahn, A., and Schnitzer, M.J. (2009). Automated analysis of cellular signals from large-scale calcium imaging data. *Neuron* 63, 747–760. <https://doi.org/10.1016/j.neuron.2009.08.009>.
54. Isope, P., and Barbour, B. (2002). Properties of unitary granule cell→Purkinje cell synapses in adult rat cerebellar slices. *J. Neurosci.* 22, 9668–9678. <https://doi.org/10.1523/JNEUROSCI.22-22-09668.2002>.
55. Medina, J.F., Nares, W.L., Ohyama, T., and Mauk, M.D. (2000). Mechanisms of cerebellar learning suggested by eyelid conditioning. *Curr. Opin. Neurobiol.* 10, 717–724. [https://doi.org/10.1016/s0959-4388\(00\)00154-9](https://doi.org/10.1016/s0959-4388(00)00154-9).
56. Wagner, M.J., Savall, J., Kim, T.H., Schnitzer, M.J., and Luo, L. (2020). Skilled reaching tasks for head-fixed mice using a robotic manipulandum. *Nat. Protoc.* 15, 1237–1254. <https://doi.org/10.1038/s41596-019-0286-8>.
57. Yamada, T., Yang, Y., Valnegri, P., Juric, I., Abnoui, A., Markwalter, K.H., Guthrie, A.N., Godec, A., Oldenborg, A., Hu, M., et al. (2019). Sensory experience remodels genome architecture in neural circuit to drive motor learning. *Nature* 569, 708–713. <https://doi.org/10.1038/s41586-019-1190-7>.
58. de Groot, A., van den Boom, B.J.G., van Genderen, R.M., Coppens, J., van Veldhuijzen, J., Bos, J., Hoedemaker, H., Negrello, M., Willuhn, I., De Zeeuw, C.I., and Hoogland, T.M. (2020). NiNscope, a versatile miniscope for multi-region circuit investigations. *eLife* 9, e49987. <https://doi.org/10.7554/eLife.49987>.
59. Jackman, S.L., Beneduce, B.M., Drew, I.R., and Regehr, W.G. (2014). Achieving High-Frequency Optical Control of Synaptic Transmission. *J. Neurosci.* 34, 7704–7714. <https://doi.org/10.1523/JNEUROSCI.4694-13.2014>.
60. Ohmae, S., and Medina, J.F. (2015). Climbing fibers encode a temporal-difference prediction error during cerebellar learning in mice. *Nat. Neurosci.* 18, 1798–1803. <https://doi.org/10.1038/nn.4167>.
61. Heffley, W., and Hull, C. (2019). Classical conditioning drives learned reward prediction signals in climbing fibers across the lateral cerebellum. *eLife* 8, e46764. <https://doi.org/10.7554/eLife.46764>.
62. Avila, E., Flierman, N.A., Holland, P.J., Roelfsema, P.R., Frens, M.A., Badura, A., and De Zeeuw, C.I. (2022). Purkinje Cell Activity in the Medial and Lateral Cerebellum During Suppression of Voluntary Eye Movements in Rhesus Macaques. *Front. Cell. Neurosci.* 16, 863181. <https://doi.org/10.3389/fncel.2022.863181>.
63. Gaffield, M.A., Sauerbrei, B.A., and Christie, J.M. (2022). Cerebellum encodes and influences the initiation, performance, and termination of discontinuous movements in mice. *eLife* 11, e71464. <https://doi.org/10.7554/eLife.71464>.
64. Chabrol, F.P., Blot, A., and Mrcic-Flogel, T.D. (2019). Cerebellar Contribution to Preparatory Activity in Motor Neocortex. *Neuron* 103, 506–519.e4. <https://doi.org/10.1016/j.neuron.2019.05.022>.
65. Tsutsumi, S., Hidaka, N., Isomura, Y., Matsuzaki, M., Sakimura, K., Kano, M., and Kitamura, K. (2019). Modular organization of cerebellar climbing fiber inputs during goal-directed behavior. *eLife* 8, e47021. <https://doi.org/10.7554/eLife.47021>.
66. Tsutsumi, S., Yamazaki, M., Miyazaki, T., Watanabe, M., Sakimura, K., Kano, M., and Kitamura, K. (2015). Structure-function relationships between aldolase C/zebrin II expression and complex spike synchrony in the cerebellum. *J. Neurosci.* 35, 843–852. <https://doi.org/10.1523/JNEUROSCI.2170-14.2015>.
67. Dorgans, K., Demais, V., Bailly, Y., Poulain, B., Isope, P., and Doussau, F. (2019). Short-term plasticity at cerebellar granule cell to molecular layer interneuron synapses expands information processing. *eLife* 8, e41586. <https://doi.org/10.7554/eLife.41586>.
68. Brown, A.M., Arancillo, M., Lin, T., Catt, D.R., Zhou, J., Lackey, E.P., Stay, T.L., Zuo, Z., White, J.J., and Sillitoe, R.V. (2019). Molecular layer interneurons shape the spike activity of cerebellar Purkinje cells. *Sci. Rep.* 9, 1742. <https://doi.org/10.1038/s41598-018-38264-1>.
69. Yang, Z., and Santamaria, F. (2016). Purkinje cell intrinsic excitability increases after synaptic long term depression. *J. Neurophysiol.* 116, 1208–1217. <https://doi.org/10.1152/jn.00369.2016>.
70. Salin, P.A., Malenka, R.C., and Nicoll, R.A. (1996). Cyclic AMP mediates a presynaptic form of LTP at cerebellar parallel fiber synapses. *Neuron* 16, 797–803. [https://doi.org/10.1016/s0896-6273\(00\)80099-9](https://doi.org/10.1016/s0896-6273(00)80099-9).
71. Beau, M., Herzfeld, D.J., Naveros, F., Hemelt, M.E., D'Agostino, F., Oostland, M., Sánchez-López, A., Chung, Y.Y., Maibach, M., Stabb, H.N., et al. (2024). A deep-learning strategy to identify cell types across species from high-density extracellular recordings. Preprint at bioRxiv. <https://doi.org/10.1101/2024.01.30.577845>.
72. Zhang, Y., Rózsa, M., Liang, Y., Bushey, D., Wei, Z., Zheng, J., Reep, D., Broussard, G.J., Tsang, A., Tsegaye, G., et al. (2023). Fast and sensitive GCaMP calcium indicators for imaging neural populations. *Nature* 615, 884–891. <https://doi.org/10.1038/s41586-023-05828-9>.
73. Ramirez, J.E., and Stell, B.M. (2016). Calcium Imaging Reveals Coordinated Simple Spike Pauses in Populations of Cerebellar Purkinje Cells. *Cell Rep.* 17, 3125–3132. <https://doi.org/10.1016/j.celrep.2016.11.075>.
74. Wittmann, M. (2013). The inner sense of time: how the brain creates a representation of duration. *Nat. Rev. Neurosci.* 14, 217–223. <https://doi.org/10.1038/nrn3452>.
75. Curtis, C.E., and D'Esposito, M. (2003). Persistent activity in the prefrontal cortex during working memory. *Trends Cogn. Sci.* 7, 415–423. [https://doi.org/10.1016/S1364-6613\(03\)00197-9](https://doi.org/10.1016/S1364-6613(03)00197-9).

76. Popa, L.S., Streng, M.L., and Ebner, T.J. (2017). Long-Term Predictive and Feedback Encoding of Motor Signals in the Simple Spike Discharge of Purkinje Cells. *eNeuro* 4, ENEURO.0036-17.2017. <https://doi.org/10.1523/ENEURO.0036-17.2017>.
77. Gao, Z., Davis, C., Thomas, A.M., Economo, M.N., Abrego, A.M., Svoboda, K., De Zeeuw, C.I., and Li, N. (2018). A cortico-cerebellar loop for motor planning. *Nature* 563, 113–116. <https://doi.org/10.1038/s41586-018-0633-x>.
78. Jirenhed, D.-A., Bengtsson, F., and Hesslow, G. (2007). Acquisition, Extinction, and Reacquisition of a Cerebellar Cortical Memory Trace. *J. Neurosci.* 27, 2493–2502. <https://doi.org/10.1523/JNEUROSCI.4202-06.2007>.
79. Ten Brinke, M.M., Boele, H.-J., Spanke, J.K., Potters, J.-W., Kornysheva, K., Wulff, P., Ijpelaar, A.C.H.G., Koekkoek, S.K.E., and De Zeeuw, C.I. (2015). Evolving models of pavlovian conditioning: cerebellar cortical dynamics in awake behaving mice. *Cell Rep.* 13, 1977–1988. <https://doi.org/10.1016/j.celrep.2015.10.057>.
80. Wong, K.-F., and Wang, X.-J. (2006). A recurrent network mechanism of time integration in perceptual decisions. *J. Neurosci.* 26, 1314–1328. <https://doi.org/10.1523/JNEUROSCI.3733-05.2006>.
81. Wang, X.-J. (2001). Synaptic reverberation underlying mnemonic persistent activity. *Trends Neurosci.* 24, 455–463. [https://doi.org/10.1016/S0166-2236\(00\)01868-3](https://doi.org/10.1016/S0166-2236(00)01868-3).
82. Gao, Z., Proietti-Onori, M., Lin, Z., Ten Brinke, M.M., Boele, H.J., Potters, J.W., Ruigrok, T.J.H., Hoebeek, F.E., and De Zeeuw, C.I. (2016). Excitatory Cerebellar Nucleocortical Circuit Provides Internal Amplification during Associative Conditioning. *Neuron* 89, 645–657. <https://doi.org/10.1016/j.neuron.2016.01.008>.
83. Bina, L., Romano, V., Hoogland, T.M., Bosman, L.W.J., and De Zeeuw, C.I. (2021). Purkinje cells translate subjective salience into readiness to act and choice performance. *Cell Rep.* 37, 110116. <https://doi.org/10.1016/j.celrep.2021.110116>.
84. Kostadinov, D., Beau, M., Blanco-Pozo, M., and Häusser, M. (2019). Predictive and reactive reward signals conveyed by climbing fiber inputs to cerebellar Purkinje cells. *Nat. Neurosci.* 22, 950–962. <https://doi.org/10.1038/s41593-019-0381-8>.
85. Larry, N., Yarkoni, M., Lixenberg, A., and Joshua, M. (2019). Cerebellar climbing fibers encode expected reward size. *eLife* 8, 533653. <https://doi.org/10.7554/eLife.46870>.
86. Heffley, W., Song, E.Y., Xu, Z., Taylor, B.N., Hughes, M.A., McKinney, A., Joshua, M., and Hull, C. (2018). Coordinated cerebellar climbing fiber activity signals learned sensorimotor predictions. *Nat. Neurosci.* 21, 1431–1441. <https://doi.org/10.1038/s41593-018-0228-8>.
87. Huang, C.-m., Miyamoto, H., and Huang, R.H. (2006). The mouse cerebellum from 1 to 34 months: Parallel fibers. *Neurobiol. Aging* 27, 1715–1718. <https://doi.org/10.1016/j.neurobiolaging.2005.09.025>.
88. Zang, Y., and De Schutter, E. (2019). Climbing Fibers Provide Graded Error Signals in Cerebellar Learning. *Front. Syst. Neurosci.* 13, 46. <https://doi.org/10.3389/fnsys.2019.00046>.
89. Tanaka, K., Khiroug, L., Santamaria, F., Doi, T., Ogasawara, H., Ellis-Davies, G.C.R., Kawato, M., and Augustine, G.J. (2007). Ca²⁺ requirements for cerebellar long-term synaptic depression: role for a postsynaptic leaky integrator. *Neuron* 54, 787–800. <https://doi.org/10.1016/j.neuron.2007.05.014>.
90. Johansson, F., Hesslow, G., and Medina, J.F. (2016). Mechanisms for motor timing in the cerebellar cortex. *Curr. Opin. Behav. Sci.* 8, 53–59. <https://doi.org/10.1016/j.cobeha.2016.01.013>.
91. Zhu, J., Hasanbegović, H., Liu, L.D., Gao, Z., and Li, N. (2023). Activity map of a cortico-cerebellar loop underlying motor planning. *Nat. Neurosci.* 26, 1916–1928. <https://doi.org/10.1038/s41593-023-01453-x>.
92. Asaad, W.F., Rainer, G., and Miller, E.K. (2000). Task-Specific Neural Activity in the Primate Prefrontal Cortex. *J. Neurophysiol.* 84, 451–459. <https://doi.org/10.1152/jn.2000.84.1.451>.
93. Inagaki, H.K., Fontolan, L., Romani, S., and Svoboda, K. (2019). Discrete attractor dynamics underlies persistent activity in the frontal cortex. *Nature* 566, 212–217. <https://doi.org/10.1038/s41586-019-0919-7>.
94. Jocham, G., Hunt, L.T., Near, J., and Behrens, T.E.J. (2012). A mechanism for value-guided choice based on the excitation-inhibition balance in prefrontal cortex. *Nat. Neurosci.* 15, 960–961. <https://doi.org/10.1038/nn.3140>.
95. Kebschull, J.M., Richman, E.B., Ringach, N., Friedmann, D., Albarran, E., Kolluru, S.S., Jones, R.C., Allen, W.E., Wang, Y., Cho, S.W., et al. (2020). Cerebellar nuclei evolved by repeatedly duplicating a conserved cell-type set. *Science* 370, eabd5059. <https://doi.org/10.1126/science.abd5059>.
96. Fujita, H., Kodama, T., and du Lac, S. (2020). Modular output circuits of the fastigial nucleus for diverse motor and nonmotor functions of the cerebellar vermis. *eLife* 9, e58613. <https://doi.org/10.7554/eLife.58613>.
97. Pachitariu, M., Steinmetz, N., Kadir, S., Carandini, M., and Kennedy D, H. (2016). Kilosort: realtime spike-sorting for extracellular electrophysiology with hundreds of channels. Preprint at bioRxiv, 061481. <https://doi.org/10.1101/061481>.
98. Rossant, C., Kadir, S.N., Goodman, D.F.M., Schulman, J., Hunter, M.L.D., Saleem, A.B., Grosmark, A., Belluscio, M., Denfield, G.H., Ecker, A.S., et al. (2016). Spike sorting for large, dense electrode arrays. *Nat. Neurosci.* 19, 634–641. <https://doi.org/10.1038/nn.4268>.
99. Siegle, J.H., Jia, X., Durand, S., Gale, S., Bennett, C., Graddis, N., Heller, G., Ramirez, T.K., Choi, H., Luviano, J.A., et al. (2021). Survey of spiking in the mouse visual system reveals functional hierarchy. *Nature* 592, 86–92. <https://doi.org/10.1038/s41586-020-03171-x>.
100. Mathis, A., Mamidanna, P., Cury, K.M., Abe, T., Murthy, V.N., Mathis, M.W., and Bethge, M. (2018). DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. *Nat. Neurosci.* 21, 1281–1289. <https://doi.org/10.1038/s41593-018-0209-y>.
101. Barski, J.J., Dethleffsen, K., and Meyer, M. (2000). Cre recombinase expression in cerebellar Purkinje cells. *Genesis* 28, 93–98. [https://doi.org/10.1002/1526-968X\(200011/12\)28:3/4<93::AID-GENE10>3.0.CO;2-W](https://doi.org/10.1002/1526-968X(200011/12)28:3/4<93::AID-GENE10>3.0.CO;2-W).
102. Madisen, L., Mao, T., Koch, H., Zhuo, J.M., Berenyi, A., Fujisawa, S., Hsu, Y.W.A., Garcia, A.J., 3rd, Gu, X., Zanella, S., et al. (2012). A toolbox of Cre-dependent optogenetic transgenic mice for light-induced activation and silencing. *Nat. Neurosci.* 15, 793–802. <https://doi.org/10.1038/nn.3078>.
103. Pnevmatikakis, E.A., and Giovannucci, A. (2017). NoRMCorre: an online algorithm for piecewise rigid motion correction of calcium imaging data. *J. Neurosci. Methods* 291, 83–94. <https://doi.org/10.1016/j.jneumeth.2017.07.031>.
104. Pnevmatikakis, E.A., Soudry, D., Gao, Y., Machado, T.A., Merel, J., Pfau, D., Reardon, T., Mu, Y., Lacefield, C., Yang, W., et al. (2016). Simultaneous Denoising, Deconvolution, and Demixing of Calcium Imaging Data. *Neuron* 89, 285–299. <https://doi.org/10.1016/j.neuron.2015.11.037>.
105. Nguyen, T.M., Thomas, L.A., Rhoades, J.L., Ricchi, I., Yuan, X.C., Sheridan, A., Hildebrand, D.G.C., Funke, J., Regehr, W.G., and Lee, W.-C.A. (2023). Structured cerebellar connectivity supports resilient pattern separation. *Nature* 613, 543–549. <https://doi.org/10.1038/s41586-022-05471-w>.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Antibodies		
GFP	Aves Labs	GFP-1010; RRID: AB_2307313
RFP	Abcam	62341; RRID: AB_945213
Alexa 488 Goat anti-Chicken IgY (H+L)	Thermo	A-11039; RRID: AB_2534096
Alexa 594 Donkey anti-Rabbit IgG (H+L)	Thermo	A-32754; RRID: AB_2762827
Bacterial and virus strains		
AAV-L7-6-R-CaMP2	Inoue et al. ⁵¹ and Nitta et al. ⁵²	N/A
AAV-ef1a-jGCaMP8f	Addgene	Cat# 176756; RRID: Addgene_176756
Chemicals, peptides, and recombinant proteins		
Isoflurane	Henry Schein	CAS 26675-46-7
DAPI mounting medium	Southern Biotech	0100-20
Tribromoethanol	N/A	N/A
C&B Metabond Quick Adhesive	Parkell	UN1247
Deposited data		
Source data	This paper	Dryad: https://doi.org/10.5061/dryad.bk3j9kdm6
Experimental models: Organisms/strains		
Mouse: Math1-Cre	Jackson Labs	Stock# 011104; RRID: IMSR_JAX:011104
Mouse: Ai93 (TITL-GCaMP6f)-D	Jackson Labs	Stock# 024103; RRID: IMSR_JAX:024103
Mouse: ztTA	Jackson Labs	Stock# 012266; RRID: IMSR_JAX:012266
Mouse: PCP2-Cre	Jackson Labs	Stock# 004146; RRID: IMSR_JAX:004146
Mouse: Ai32	Jackson Labs	Stock# 024109; RRID: IMSR_JAX:024109
Mouse: Wild-type CD1	Charles River	Stock# 022; RRID: IMSR_CRL:022
Software and algorithms		
MATLAB	Mathworks	https://www.mathworks.com
NoRMCorre	Simons Foundation/Flatiron	https://github.com/flatironinstitute/NoRMCorre
ScanImage	MBF Biosciences	https://www.mbfbioscience.com/products/scanimage/
LabVIEW	National Instruments	http://www.ni.com/en-us/shop/labview.html
PCA/ICA	Mukamel et al. ⁵³	https://github.com/mukamel-lab/CellSort
Kilosort 2.5	Pachitariu et al. ⁹⁷	https://github.com/MouseLand/Kilosort
phy	Rossant et al. ⁹⁸	https://github.com/cortex-lab/phy
ecephys	Siegle et al. ⁹⁹	https://github.com/AllenInstitute/ecephys_spike_sorting
DeepLabCut	Mathis et al. ¹⁰⁰	https://github.com/DeepLabCut/DeepLabCut
cNMF	Simons Foundation/Flatiron	https://github.com/dylkot/cNMF
Source code for all figures	This paper	https://github.com/wagnerlabnih/garcia-garcia-neuron-2024
Other		
Laser scanning confocal microscope	Zeiss	LSM 510
Stereotaxic Instrument	Kopf	Model 940 Small Animal
Micro syringe pump injector	World Precision Instruments	UMP3T-1
Two-axis robotic manipulandum	Wagner et al. ⁵⁶	https://github.com/mjwagner/haptic-for-mice
488 nm laser	Coherent	OBIS LX
920 nm + 1064 nm laser	Spark	Alcor 920/1064 2W Dual

(Continued on next page)

Continued

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Resonant-linear galvo	Avantor	6215H, CRS 8k or 12k
Ti:Sapph tunable laser	Spectra Physics	MaiTai
Focus tunable lens	Optotune	EL-10-30-CI-NIR-LD-MV, EL-16-40-CI-NIR-LD-MV
Fiberport	Thorlabs	PAF2-7A
Steel multimode patch cable	Thorlabs	MR83L01

RESOURCE AVAILABILITY

Lead contact

Further information and requests for resources and reagents should be directed to and will be fulfilled by the lead contact, Mark Wagner (mark.wagner@nih.gov).

Materials availability

This study did not generate new unique reagents.

Data and code availability

- Data will be deposited at Dryad at the time of publication (<https://doi.org/10.5061/dryad.bk3j9kdm6>).
- Source code will be deposited at GitHub at the time of publication (<https://github.com/wagnerlabnih/garcia-garcia-neuron-2024>).
- Any additional information required to reanalyze the data reported in this paper is available from the [lead contact](#) upon request.

EXPERIMENTAL MODEL DETAILS

Mice

All GrC-CF imaging experiments used Math1-Cre / Ai93 (TIGRE-LSL-TRE-GCaMP6f) / ztTA (R26-CAG-LSL-tTA) mice, aged 6 – 16 weeks, which, in the cerebellum, express GCaMP6f selectively in GrCs.⁴¹ Optogenetics experiments used double transgenic crosses of PCP2-Cre¹⁰¹ mice to Ai32 mice¹⁰² (Jackson labs), while electrophysiology and PkC somatic imaging used WT mice. All procedures were approved by both the NIH and Stanford animal care and use committees.

METHOD DETAILS

Virus

An abridged PkC-specific promoter (L7-6)⁵² drove expression of the red Ca²⁺ indicator R-CaMP2,⁵¹ packaged into AAV-L7-6-R-CaMP2. PkC imaging used AAV-ef1A-jGCaMP8f virus. Virus injected at ~10¹² genomes/mL.

Viral injection

Mice were anesthetized via isoflurane (~2% in ~1L/min of O₂). We cleaned the scalp and removed hair. We made a sagittal incision in the skin to expose the underlying skull, which we then lightly scraped free of soft tissue. We drilled a ~300 μm hole over the vermis of Lobule VI (~0 to 0.3 mm right of midline, on the post-lambda suture). We inserted a ~25 μm diameter-tip glass capillary ~300 μm below the pial surface and injected 500 nL of AAV-L7-6-R-CaMP2 (or jGCaMP8f for PkC surgeries) into the tissue and waited 5 minutes after injection before withdrawing the tip. Virus was typically allowed to express for 10 to 14 days prior to imaging.

Histology

Mice were deeply anesthetized via IP injection of 5% tribromoethanol in PBS, and then transcardially perfused with PBS followed by 10% formalin, after which the brain was removed and left overnight in 10% formalin. We cut 60 μm sagittal sections from the cerebellum using a vibratome (Leica VT 1000S). The sections were then washed in PBS.

For GrC-CF histology, we performed immunostaining first by 1h of blocking in 10% FBS in PBST. We dual-stained the sections using chicken anti-GFP (Aves Labs, GFP-1010) and rabbit anti-RFP antibodies (AB-62341) for ~48h, followed by a 0.1% PBST wash, followed by Alexa 488 anti-chicken (A-11039) and Alexa 594 anti-rabbit (A-32754) secondary antibody stainings for ~3 h. After mounting the sections, we imaged the resulting slides using a confocal microscope with 488 nm and 561 nm lasers (Zeiss LSM 510).

For PkC Neuropixels histology, we directly mounted the sagittal sections in DAPI mounting medium and imaged using the 405nm and 531nm lasers.

Window and headplate implantation

Implantation followed viral injection. We removed a patch of skin with a mediolateral span from ear to ear, and a rostrocaudal span from the back of the ears to the eyes. We scraped the entire exposed skull surface free of soft tissue, and then used VetBond (3M) adhesive to seal the edge of the skin incision to the skull. We drilled a ~3.5 to 4 mm diameter cranial window centered over the right vermis of lobule VI. We then affixed a #0 or #1 3 mm diameter glass cover slip onto a 3 mm outer diameter, 1 mm height, 2.7 mm inner diameter steel ring, using UV curing optical adhesive (Thorlabs NOA 81). We stereotactically inserted the glass-ring combination into the cranial window at a 20°–25° angle counter-clockwise from the sagittal axis and a 40°–45° degree azimuthal angle from the vertical axis. We depressed the glass to a depth below the average depth of the underside of the skull at the window perimeter. We then sealed the outer face of the steel ring to the skull using Metabond (Parkell).

We then implanted a custom headplate with a 5 mm central opening and two extensions with two holes for screws to affix the headplate to fixation bars. We stereotactically lowered the headplate, surface parallel to the glass window, onto the skull with the central opening centered on the window, and then covered the entire exposed skull with Metabond up to the top surface of the headplate.

Behavioral data collection

As described previously,⁵⁶ our operant device was a two-axis robotic manipulandum consisting of two DC motors, two high-resolution optical rotary encoders, and a 4-linkage two-degree-of-freedom manipulandum configuration. The manipulandum was controlled via a custom programmed nested series of feedback loops across a field-programmable gate array (FPGA), PC with real-time Linux OS, and Windows PC (together forming the National Instruments cRIO platform). In addition to handling real-time control tasks, the apparatus recorded the x and y position of the handle of the manipulandum at 200 Hz. We also synchronously acquired at 200 Hz: a readout of the solenoid trigger that released water reward; the frame counter from the two-photon microscope; in some cases, a capacitive lick sensor reading; and in some cases, the optogenetic laser trigger pulses. Reward delivery was triggered either via windows PC (the main GrC-CF 1.1-s-delay dataset), or directly in the Real-Time OS (all other datasets: 1-s-to-2-s learning data; PCP2-Cre/Ai32 ChR2 data; Neuropixels recordings; PkC somatic imaging). The latter eliminated roughly ~100 ms of delay variability; the resulting data was qualitatively similar between methodologies.

In addition, in some cases, we collected behavioral video data at 30 – 120 Hz (Imaging Source). This was synchronized to the microscopy data either by hardware frame trigger or by automatically identifying the frames on which the IR laser shutters opened and closed. Finally, we used DeepLabCut¹⁰⁰ to semi-automatically annotate 2D locations of the mouse's two forepaws, nose, tongue, eye pupil, or whiskers.

For some mice, we obtained capacitive sensor lick readings (26 mice total; main dataset: 12 mice; PCP2-Cre/Ai32 ChR2 dataset: 7 mice; 1-s-vs-2-s imaging dataset: 5 mice; PkC somatic Ca²⁺ imaging dataset: 2 mice). Capacitance measurements were sensitive to gross changes in positioning or degree of contact, such that the sensor sometimes became 'stuck' in the 'on' state. We automatically excluded from licking analysis any trials in which the sensor was "high" for more than 90% of any continuous 1.5-s block of time. For the remaining trials, we computed lick rate by identifying lick onset events in the raw binary sensor contact traces. Finally, we filtered the binary traces to produce smoothly-varying rate signals.

Behavioral training

During active behavioral training, mice were water restricted to 1 mL per day, weighed daily for excess weight loss, and monitored for signs of lethargy, coat deterioration, hunching, and general distress. Mice obtained water until satiety during behavioral training and received remaining water up to the 1 mL dosage in the home cage.

Reaching tasks employed our custom two axis robotic manipulandum.⁵⁶ The linear reaching task had the following structure. Mice self-initiated trials by pushing the handle of the robotic arm. Trials were terminated when either of two conditions were met: (1) animals reached the 8 mm virtual "wall" where the robot immediately terminated movement; or (2) animals ceased pushing the handle at any distance >3 mm for more than 100 ms. Following a delay after the end of a successful (>6-7 mm) movement, the computer dispensed (via a gravity-fed reservoir through a solenoid valve) a water droplet from a gavage needle in front of the animal's mouth. The main GrC-CF imaging dataset had a 1.1-s-delay; for all other datasets including optogenetics, 1-s-to-2-s behavioral and imaging studies, Neuropixels recordings, and PkC somatic imaging, delays were either 1-s or 2-s as indicated. Following an additional ~2 s delay, the robotic arm automatically returned to the animal to self-initiate the subsequent trial.

Mice underwent typically 1 to 3 days of pretraining during which we did not perform imaging, until animals reached ~50 trials in 30 minutes. The function of pretraining was to minimize sessions in which animal brains were exposed to laser light without producing any behavioral imaging data. After this, we imaged during all subsequent training sessions for ~1 week, returning to the same field of view each day. After the week of chronic imaging, we sometimes imaged distinct imaging fields for 1 to 3 additional days.

Two-photon microscopy

We used a custom two-photon microscope with a resonant-galvo x-y scanhead (Cambridge Technology. Galvo, 6215H; resonant, either CRS 8k or CRS 12k). The brain was illuminated by two lasers, one at 920 nm for GCaMP6f (typical power: ~60 mW), and one at 1064 nm for R-CaMP2 (typical power: ~40 mW). In some cases, 920 nm was provided by a fixed wavelength 920 nm laser (Spark

Alcor 920 2W), while in other cases it was provided by a tunable Ti:Sapph (Spectra Physics MaiTai). 1064 nm illumination was provided by a fixed wavelength laser (Spark Alcor 1064 2W). The 1064 nm laser path passed through an electrically tunable lens (Optotune, either EL-10-30-CI-NIR-LD-MV or EL-16-40-CI-NIR-LD-MV), in some cases with an additional offset lens (75 mm, 125 mm, or 150 mm focal length). The emission path split red and green onto two PMTs.

The main dataset was acquired using a 40x magnification 0.8 numerical aperture (NA) objective (Olympus XLUMPlan). The image size was 512 x 512 pixels over a field of view of 234 x 234 μm . In some cases (those using the CRS 12kHz scanner), the 1064 nm (red) and 920 nm (green) data were acquired via alternating frames by alternately shuttering the two lasers, resulting in an effective frame rate of 22.5 Hz. In other cases (those using the CRS 8 kHz), the two channels were acquired simultaneously at 30 Hz.

To select a field of view, we began by locating the region of cerebellum with viral expression of R-CaMP2 in PkCs. From this restricted region, we then optimized for the region with the best joint SNR in both GrCs and PkCs. We then optimized GrC imaging depth using the objective z-piezo, and finally optimized the PkC dendrite imaging depth using the Optotune remote focusing adjustment.

In cases of chronic imaging, following the first day of two-photon imaging, we used the wide field camera image to move to the nearest major vasculature landmark, collected an image of the landmark, and recorded the distance from our chosen imaging field to the landmark via the X-Y-Z sample translation stage distances. To return to the same imaging field on subsequent days, we inverted this procedure by finding the landmark and moving the recorded distances to the imaging field. We also recorded the final Optotune remote focusing offset for the 1064 nm laser. We then fine-tuned the GrC imaging field by comparing the image to the mean image from the first day of chronic imaging, ensuring that landmark GrCs from the initial session were present in the daily field at several locations in the imaging frame.

PkC somatic Ca^{2+} imaging and analysis using jRCaMP8f

We followed the surgical preparation described for GrC-CF imaging, with the only difference being the injection of AAV1-EF1A-jRCaMP8f. After 10 days, we examined the mice under the two-photon using a Nikon 16x 0.8 NA objective, and localized the PkC layer, picking imaging fields with the greatest number of infected PkC bodies.

Optogenetic studies

We crossed PCP2-Cre¹⁰¹ mice to Ai32 mice¹⁰² to generate PCP2/Ai32 double transgenics and implanted cerebellar windows as described above. After a week of training, we acquired perturbation data in expert mice. We positioned a ferrule-terminated multi-mode optical fiber (200 μm core, 0.39 NA, Thorlabs MR83L01) ~1 mm above the glass centered on Lobule VI. A 488 nm laser (Coherent OBIS LX) delivered 5–15 mW (measured in CW at fiber tip) via a Fiberport (Thorlabs PAF2-7A). On interleaved laser-on trials, the robotic controller waited for the mouse's voluntary movement initiation and, when the handle crossed a 6–7 mm distance threshold, the FPGA waited a fixed period before delivering a fixed number of 5-ms TTL pulses with 15 ms interval to the OBIS modulation input. The following table summarizes parameters for each experiment:

Experiment	Laser-on (mean %)	Laser-on, mean % of rewarded	Laser-on, % of omitted	Laser-on time (s from reward)	Laser-off time (s from reward)	Number of pulses
Figures 2I and 2J	22	17	50	-0.83 \pm 0.006	-0.05 \pm 0.006	36 \pm 0.3
Figure S2J	10	0	50	-0.94 \pm 0.004	-0.58 \pm 0.003	15 \pm 0.06
Figure S2K	34	31	50	-1.1 \pm 0.01	0.88 \pm 0.02	61 \pm 0.4
Figure S2L	17	8	50	-0.8 \pm 0.002	-0.02 \pm 0.002	32
Figures 2M and 2N	90	100	0	1.66 \pm 0.003 (from reach)	2.25 \pm 0.003 (from reach)	25

As we did not mask the blue light, for blue light controls, we used opsin-negative transgenic imaging mice with cerebellar windows and performed the same procedure. For LIX controls, we exposed LIX beneath the soft tissue in the neck and positioned the fiber behind the cerebellum.

Neuropixels recordings

Under anesthesia as above, mice were implanted with a headplate. Mice were water restricted and trained for 1 to 2 days before initial recordings. The day before recording, we opened a <1 mm craniotomy in the same location used for imaging, and then covered the hole with Kwik-Cast (WPI). The following day, immediately before recording, we pierced a small hole in the dura with a needle. We then inserted the Neuropixels probe stereotactically and under visual control, directly through lobule VI from behind the animal, using an azimuthal angular range of 79–84 degrees, and a mediolateral angular range of 0–15 degrees. We varied the angle across recording days to avoid piercing the same tissue multiple times. We inserted roughly 2 mm, to reach the anterior end of the cerebellum, and thus passing primarily through Lobule VI and V. After confirming likely passage through

multiple PkC layers as viewed in real-time in spikeGLX, we allowed the tissue to settle for 5 to 10 minutes, before initiating a recording session during behavioral training for 18 minutes. We programmed our NI cRIO behavioral apparatus to output synchronization pulses to the Neuropixels NI DAQ. For the final two recording sessions for each animal, we coated the electrode in DiD for histological track visualization.

QUANTIFICATION AND STATISTICAL ANALYSIS

Image processing

Green and red channels were independently motion-corrected using a sequence of rigid followed by nonrigid NORMCORRE registration,¹⁰³ as we empirically found that brain motion at the molecular layer differed from that at the GrC layer. To correct for slow, full-frame changes in fluorescence that did not correspond to cellular activity, we fit a double exponential to the frame-averaged fluorescence across the entire movie. For each frame, we then divided every pixel by the exponential fit value for that frame, and then manually confirmed that the resulting frame-averaged fluorescence trace was roughly flat across the recording.

Cell identification and signal preprocessing

To identify the spatial locations of GrCs, we used a pipeline based on cNMF.¹⁰⁴ We initially used cNMF to identify candidate GrCs. We automatically discarded candidates whose central regions were too small ($<\sim 0.0005\text{mm}^2$) or too large ($>\sim 0.003\text{mm}^2$). We then manually discarded candidates in cases where either (1) the spatial filter clearly did not correspond to GrC somas, or (2) where no clear elevation of fluorescence was observed in the video at the spatial location of the cell filter during high points in the cNMF signal trace. We then manually annotated the movie for “missed” cells and “seeded” a second round of cNMF with these additional candidates and then repeated the above steps.

To identify the spatial locations of PkC dendrites, we used a pipeline based on PCA/ICA.⁵³ We again automatically discarded candidates based on size as above.

In some cases, we then applied a convolutional neural network classifier to automatically discard additional candidates whose spatial filters clearly did not correspond to Purkinje dendrites. We built our neural network with 15 layers: one image input layer, two repeats of a convolution-batch normalization-ReLU-average pooling layers sequence, a convolution-batch normalization-ReLU layers sequence, and a fully connected-soft max-classification layers sequence. We trained the classifier over eight epochs and shuffled the training data before each. We used an L2 regularization factor of 0.01, an initial learning rate of 0.001 with a piecewise learning rate schedule, and a mini batch size of 32. We trained this neural network on 7,400 previously sorted PkC dendrite spatial filters from five mice, across 74 sessions. The training set consisted of 4,213 (56.93%) images that visibly contained PkC dendrite spatial filters (class 1) and 3,187 (43.07%) images that did not (class 0). The final neural network had an accuracy of about 90%. We then discarded the candidates that were not classified as PkC dendrite spatial filters and manually inspected those remaining. In addition, we manually split spatial filters that clearly contained more than one PkC dendrite. Finally, we automatically searched for possible duplicate dendrites or candidate pieces of split dendrites to be merged, and manually removed identified duplicates and manually merged identified split dendrites.

For both GrCs and PkC dendrites, we then back applied the final set of spatial filters to the preprocessed movies, to yield the initial cell activity traces. For each cell, we next removed signal drifts on timescales slower than neural activity by subtracting off a 10th percentile filtered version of the signal (GrCs: 10 s moving window; PkC dendrites: 5 s).

Finally, we uniformly scaled the signal magnitudes across cells by normalizing the amplitude of their baseline (noise) fluctuations (initial signal magnitudes were ambiguous due to: lack of meaningful physical units; differences between cells in SNR & indicator expression levels; changes in SNR or brightness within a session; ambiguous impact of spatial filter pixel weighting). For each neuron, we first estimated the center of the noise distribution, by computing a (slow) moving median (GrCs: 2-minute window; PkC dendrites: 1-minute window), but we excluded any timepoints above the cell's 99th percentile value, which were Ca^{2+} transients. For each cell, we subtracted off this estimate to center the noise distributions on zero. Next, we determined the noise fluctuation magnitude, by computing a slow moving windowed standard deviation on all fluorescence values below zero (excluding the top half of the distribution removes neural activation transients). This corresponded to the lower half-normal distribution of fluorescence noise fluctuations. We divided each cell's trace by the estimated noise standard deviation $\sigma\sqrt{1 - \frac{2}{\pi}}$. Overall, this produced signals where zero was defined as the center of the noise distribution, and the noise standard deviation was normalized across cells to unity magnitude, i.e., z-scored.

Since PkC dendritic calcium transients reflect complex spikes⁴² and thus individual CF spikes, we also estimated the time of individual CF spikes. We deconvolved out the indicator kinetics ($\tau = 150$ ms) from the activity trace and thresholded the resulting signal (1.9 s.d.) to identify individual transient times. When further producing single-trial spike rates, we filtered the binary event traces with a 200 ms kernel.

PkC somatic Ca^{2+} imaging analysis

We again used cNMF to identify active PkC bodies. During manual curation, we accepted candidates that, in addition to having the expected size and shape, clearly exhibited elevated fluorescence in the z-scored movie during their nominally “brightest” frames.

Finally, since it was typically impossible to restrict the entire imaging field to the exact depth of the PkC layer, we often observed some contamination from PkC dendritic signals (complex spikes). We therefore manually excluded any detected PkC somas that exhibited many fast transients with the visual appearance of dendritic complex spikes (i.e., a single transient that decays in ~ 200 ms), as these were excessively contaminated by dendritic fluorescence. The resulting remaining signals were passed through the same postprocessing pipeline described in the previous section for GrCs.

Neural response analysis

For trial-aligned analyses, we began by identifying all reaching movements from the manipulandum position data. To align trials, we sequentially identified: the time of reward delivery; reach “end” and “midpoint” as the times when the y position first extended to within 0.5 mm of its maximal value or beyond 4 mm respectively; reach “start” as the time when a (100 ms-window-smoothed) velocity estimate fell to below 15 mm / s, prior to the midpoint. To restrict analysis to stereotyped trials, we identified reaches >7 mm in length. When analysis of the main dataset required simultaneous alignment to reward and reaching, we concatenated the reach-aligned and reward-aligned neural data such that movement midpoint uniformly occurred at -1.1 s relative to reward.

Neural learning analysis

Our main dataset consisted of 20 mice. Of these, in 15 we imaged across multiple days of learning, aiming for 7 days of repeated imaging of the same neurons, followed by an additional 1 to 3 days of imaging distinct fields of cells to sample more neurons. On average we obtained 7.2 ± 0.1 repeated imaging days in these 15 mice. Across all 20 mice, we obtained expert imaging data from, on average, 1.7 ± 0.1 distinct fields of view.

For analyses in which the relevant N values were unique cells (e.g., [Figure 3](#)), all included sessions were taken from distinct fields of view (no repeats).

For analyses in which the relevant N values were unique sessions (e.g., [Figures 5D–5F](#)), we included all sessions.

Temporal decoding analysis

We used a 10-fold cross-validated linear regression (ordinary least squares) analysis to determine the accuracy of GrC temporal decoding. For each session, we first produced a matrix of size $N_{\text{GrC}} \times T \times N_T$, where N_{GrC} was the number of GrCs, T was the number of imaging frames corresponding to 1.1 second, and N_T was the number of rewarded trials. We concatenated all trials along the time axis to produce a matrix of size $(N_T \times T) \times N_{\text{GrC}}$. We then produced a “target” temporal axis vector of length T , whose entries ran from -1.1 s to 0 s (to span the delay), and then repeated this vector N_T times to produce a target vector of size $(T \times N_T) \times 1$. From the cross-validated linear regression results, we retained the cross-validated prediction of the temporal axis for each trial and computed the resulting accuracy (R^2) and mean absolute error in ms. For post-reward temporal decoding, we repeated the identical procedure on 1 s of post-reward data and a time axis spanning 0 to 1 s.

Neuropixels analysis

We used Kilosort 2.5 to perform spike sorting. For all recordings (27 sessions across 4 mice), we applied the common average reference filter to the raw data prior to running Kilosort⁹⁷ and used the following parameters: threshold = [10, 3], high pass filter = 100 Hz, lowpass filter = 5000 Hz. To visualize and perform quality control on all of the clusters, we used the Phy⁹⁸ template GUI and the ecephys⁹⁹ quality metrics module.

Initially, we eliminated clusters with less than 100 spikes. We calculated the firing rate for each cluster in 30 second windows and discarded those that fell outside the appropriate ranges (40 Hz–200 Hz for simple spikes, 0.4 Hz–3 Hz for complex spikes).

We confirmed PkC identities based on cross-correlation analysis to putative complex spike clusters, with a characteristic SS “pause” of ~ 5 –20 ms after a CS ([Figure S6](#)). We then used confirmed PkCs as estimates of the physical location of PkC layers along the Neuropixels shank (± 10 channels/ $100 \mu\text{m}$ ⁷¹). Clusters within each PkC layer were considered potential simple spike clusters based on minimum spike rate (40 Hz), inter-spike interval violations ratio ($<2\%$), waveforms, and oscillatory autocorrelograms.

To handle cells that were not present for the entire recording session, we computed the spike rate per trial, and then the 75th percentile value across all trials with a mean rate >5 Hz, and defined a presence threshold as half this value. We then found the first and last trial that exceeded this threshold. Across 162 putative PkCs, by this metric the average cell was present for 13 ± 0.3 min, or $81 \pm 2\%$ of the total recording length.

Statistics

A detailed panel-by-panel summary of all statistical tests and observation counts is provided in [Table S1](#). As general principles, for all comparisons of two groups of unpaired samples, we used the Wilcoxon Rank-Sum (Mann-Whitney U) test. For all comparisons of two groups of paired samples, or to compare one group’s median to zero, we used the Wilcoxon Signed-Rank test. For all comparisons of two distributions of samples, we used the Kolmogorov-Smirnov two-sample test. For all analyses of groups of samples over stages of learning, we used an ANOVA evaluated on an ordinary least squares regression of the samples to the learning stage.

Plasticity simulations

Each imaged PkC dendrite reported activity of that PkC's CF input. Imaged PkCs also likely received input from most imaged GrCs, given our small imaging field. Classical GrC → PkC LTD depends only on relative activity of the CF and GrC input, to determine whether that GrC input synapse is weakened at each point in time. We reasoned that, even without access to all GrC inputs to a PkC, we could compute predicted changes in synaptic weight for all imaged GrCs based on simultaneous CF spiking. For simplicity, we assumed that every PkC in our small field of view received input from all imaged GrCs, although in reality the proportion would be lower.^{54,105}

The central question of our study was how persistent reward-evoked CF spikes can set meaningful GrC → PkC synaptic input strengths. Thus, we computed predicted LTD caused by CF spikes following reward within [0, 0.25] s. It remains unknown to what degree spontaneous 1 Hz background CF spikes—which far outnumber evoked CF spikes—play a similar role in driving synaptic changes, and we considered this broader question beyond the scope of the current study. We thus restricted our analysis to CF reward spikes, except in Figure S5J when we considered alternative inclusion windows. Similarly, we focused on CFs with detectable reward-evoked spiking, whose trial-averaged activity was higher after than prior to reward delivery (70% of all 6,439 CFs, spread across all 117 sessions in 20 mice).

To predict GrC → PkC LTD from reward-evoked CF spikes, for each session, each PkC, and each GrC input, we computed a “running tally” of that GrC input's LTD events. For every CF reward spike, we tabulated the GrC's mean activity during the preceding plasticity window [-150, -25] ms. This number—GrC activity just prior to CF reward spike—was taken to be the LTD event magnitude. To bound LTD magnitudes in a defined range [0, 1], we applied a logistic function to each LTD magnitude ($\frac{1}{1+e^{-F/s}}$; F = avg GrC signal in LTD window; s = 95th percentile fluorescence value per cell).

To summarize the LTD for each GrC, we averaged LTD events across trials and across timepoints. This yielded a single GrC weight vector for each PkC, size $N_{\text{GrC}} \times 1$. In reality, opposing and homeostatic mechanisms likely maintain each PkC's overall synaptic drive in a physiological range. To account for without modeling these unobserved factors, we simply normalized each PkC's GrC weight vector to unit sum.

Finally, we subtracted the mean of this vector. This was necessary: taking a strictly positive weighted sum of cells will generally produce an output that looks approximately like the average of the population. While GrC → PkC synapses are, in reality, strictly excitatory, an effect similar to mean-subtraction is likely achieved biologically by interneuron networks at the PkC output layer. Finally, given CF uniformity in our small imaging fields, we averaged GrC weight vectors across PkCs. This resulted in a single GrC weight vector per session, which we negated to account for the sign of LTD. These vectors gave rise to analyses in Figures 5 and S5.

To predict how GrC weights might affect PkC output, we used a minimal readout: a weighted sum of GrCs, with weights defined above. The resulting weighted sums produced analyses in Figure 6 and S5.

Finally, we computed three control weighted sums. First, we computed a simple average across GrCs (uniform weights, e.g., Figure 6C). Alternatively, (brown in Figure 6), we simply randomly reordered the LTD weight vector above. Each GrC thus had a random weight from the true distribution of LTD weights. In the final control, we repeated the entire above LTD procedure but instead using randomly reordered GrC data. Specifically, for every trial and for every GrC, we randomly permuted the timepoints in the trial. Thus, every GrC maintained the same distribution of fluorescence values, but its timing was random with respect to the CF reward spikes. In each case we then repeated the same analyses applied to the true data.

Parameters include: CF reward spike window ([0, 0.25] s); cutoff for included PkCs (average activity [0, 0.25] s > [-0.3, -0.025] s); LTD window ([-0.15, -0.025] s); logistic function bounding LTD event magnitudes.

All simulation panels operated based only on LTD described above, except Figure S5K which explicitly modeled LTP. To simulate LTP, for each CF we defined a threshold as the bottom 0.5% of all timepoints by firing rate, across trials. This threshold was empirically chosen to yield a total number of LTP events comparable to the total number of LTD events. Each LTP event used the same LTD “eligibility window” of [-150, -25] ms to measure preceding GrC activity. As for LTD, we then tabulated each GrC's activity in each LTP window to be the corresponding LTP event magnitude. The net plasticity effect was the sum of LTD and LTP event magnitudes.