Title: Cerebellar associative learning underlies skilled reach adaptation

Authors: Dylan J. Calame, Matthew I. Becker, Abigail L. Person

Abstract:

Cerebellar output has been shown to enhance movement precision by scaling the decelerative phase of reaching movements in mice. We hypothesized that during reach, initial kinematics cue late-phase adjustments through cerebellar associative learning. We identify a population-level response in mouse PCs that scales inversely with reach velocity, suggesting a candidate mechanism for anticipatory control to target limb endpoint. We next interrogate how such a response is generated by combining high-density neural recordings with closed-loop optogenetic stimulation of cerebellar mossy fiber afferents originating in the pontine nuclei during reach, using perturbation schedules reminiscent of classic adaptation paradigms. We found that reach kinematics and PC electrophysiology adapt to position-locked mossy fiber perturbations and exhibit aftereffects when stimulation is removed. Surprisingly, we observed partial adaptation to position-randomized stimulation schedules but no opposing aftereffect. A model that recapitulated these findings provided novel insight into how the cerebellum deciphers cause-and-effect relationships to adapt.

Introduction:

In humans and animals with altered cerebellar function, movement is disorganized, causing hallmark symptoms of endpoint dysmetria[1,2] and impaired abilities to adapt movements in the face of novel conditions[3–7]. These observations lend to the idea that the cerebellum acts as a feedforward controller that learns anticipatory control[8–10] making movements fast, smooth, and accurate[11,12].

Exploring these learned anticipatory signals, researchers have focused on two dominant paradigms: associative learning and motor adaptation. In associative delay eyeblink conditioning, a neutral cue predicts an aversive corneal air puff through repeated pairing. Over many trials, the subject learns to anticipate the air puff with a well-timed eyeblink. Importantly, the neutral cue (the conditioned stimulus) can be fully replaced by stimulation of cerebellar mossy fibers[13,14] and the air puff (the unconditioned stimulus) can be fully replaced by stimulation of climbing fibers[14] – the teaching signals that drive complex spike (Cspk) mediated plasticity at the Purkinje cell (PC) dendrites[15,16]. PCs learn to pause at the predicted time of the unconditioned stimulus causing bursts of activity in cerebellar nucleus neurons that drive eyeblink closure (the conditioned response). Related associative learning paradigms that pair neutral cues with cerebellar-driven limb movements extend these principles to multiple effectors[17–19]. Associative learning has been thematically linked to motor adaptation of simple movements, such as in VOR adaptation, where sensory signals such as vestibular information conveyed via mossy fibers may act as conditioned stimuli used as predictive cues for novel cerebellar output over learning[20–22]. Perturbations that result in movement error and associated Cspks lead to adaptation of PC simple spike tuning with correlated changes in behavior[23–25]. This iterative tuning of simple spikes is interpreted as a remapping of mossy fiber inputs conveying sensorimotor information onto novel cerebellar outputs[25].

Yet, with more complex motor skills such as reach, both cerebral cortex and cerebellum are proposed as sites of learning[26–28], with long-range loops likely functioning synergistically[29]. Thus, learning-related changes in cerebellum could either be inherited from cortical plasticity, generated by cerebellar plasticity, or both[30]. Notably, acute disruption of either cortical input to the cerebellum or output from the cerebellum impairs skilled reach kinematics[2,31,32], indicating that cerebellar learning may play an important role in regulating the execution of skilled movements. Indeed, in well-coordinated reaches, early kinematic features like peak reach velocity covary with subsequent limb deceleration, which is causally scaled by graded cerebellar nuclear activity[2,33]. We hypothesize that cerebellar output during reach is mechanistically akin to a conditioned response, such that decelerative output is cued by predictive mossy fiber activity that encodes within-movement kinematic features, thereby reducing error.

Here, we performed experiments designed to link the cerebellum's role in associative learning with its generation of anticipatory control of reaching movements. We superimposed a variant of associative learning onto a skilled reach task through repeated optogenetic manipulation of mossy fiber activity in

closed loop with reach, triggered at a consistent kinematic landmark. In addition to monitoring the kinematic effect of stimulation over trials, we also measured changes in PC response to stimulation with adaptation. To test reliance of temporal specificity on learning, we randomized the position of stimulation. A cerebellar model of timed adaptation within a movement recapitulated our key experimental findings and gives mechanistic insight into the circuit properties underlying cerebellar reach adaptation. Together, these experiments unify the frameworks of cerebellar associative learning and motor adaptation in skilled movements, an important step in understanding mechanisms of motor learning.

Results:

A pause-based PC population code tuned to reach velocity

Neurons in the anterior interposed nucleus fire proportional to reach velocity and causally scale limb deceleration, such that the limb lands on target despite initial kinematic variability[2,33], consistent with the cerebellum implementing anticipatory control. To determine whether upstream PCs may drive these decelerative bursts in the cerebellar nuclei, we combined kinematic and electrophysiological recordings in mice engaged in a skilled head-fixed reach task. After mice were proficient at the task, we recorded reach kinematics with high-speed cameras via an IR-reflective marker affixed to the mouse's hand (Fig. 1a, Fig. S1). Acute recordings in cerebellar cortex were made simultaneously, using either single electrodes or Neuropixel probes (Fig. 1a, Video S1). Recordings were targeted to a cerebellar cortical site situated between Lob 4/5 and Simplex known to influence forelimb movements in mice[17]. We found that activity in PCs was highly modulated around the time of the reach across cells and sessions (Fig 1c,f, Fig. S3b), often preceding reach onset and following reach termination by 100s of ms.

To test the prediction that decelerative signals in the cerebellar nuclei derive from Purkinje neuron activity patterns during reach, we first sought to understand what individual PCs encode. We used least absolute shrinkage and selection operator (LASSO) regression to model PC simple spike firing rate using limb kinematics on a trial-by-trial basis, with a ten-fold cross validation step to avoid over-fitting[34] (see Methods). On average, kinematics of the limb alone could explain 15.6 ± 1.7 % (mean ± SEM) of the variance in simple spike firing rate on individual trials, although trial-averaged data was a much closer fit (59.1 ± 3.6 %, Fig. 1c, d) consistent with other studies of PC simple spike tuning to limb movements in primates[35–39]. Kinematic encoding was not a result of generic movement-related modulation, but was specific to the kinematics of individual reaches as demonstrated by a reach shuffled control that reassigned reaches with PC firing recorded during separate reaches, and a spike shuffled control where simple spike times on each trial were time shuffled and regressed against kinematics. In both cases, regression performance on the empirical data was significantly higher than the shuffled controls indicating that simple spike firing rates encode reach kinematics on a reach-by-reach basis (Fig.

1d; n = 46 cells; empirical vs. reach shuffle: $p = 4.21 \times 10^{-11}$; empirical vs spike shuffle: $p = 2.80 \times 10^{-14}$, Wilcoxon signed rank test). The regression model performance did not change across the spatial trajectory of the movement, suggesting kinematic encoding is continuous in individual cells (Fig. 1e). To assess which kinematic variables in the regression model were the most important in modeling simple spike firing rate, we repeated the regression with each variable independently time shuffled and measured the change in variance explained relative to the complete model[40] (Fig. S3d). Positional terms -- outward, upward, and lateral – accounted for approximately 10% of the explained variance of the complete model, with each of the remaining 20 variables accounting for < 5%, although there was a wide variety in the relative importance of different kinematic variables across cells.

As has been noted previously, PCs tend to show complex patterns that predominate as positive or negative spike rate modulations[41,42]. Of 46 recorded cells, 22 displayed increases in activity during the reach epoch (bursters) and 24 showed a decrease (pausers) (Fig. 1f). When grouped as bursters and pauses, both populations showed a deeper modulation during reaches that extended farther in the outward direction than closer reaches (Fig. 1g). Binning the endpoint position of reaches across sessions relative to the median for each session showed that the modulation in firing rate during outreach linearly increased with the distance of reaches for bursters and pausers (Fig. 1h; Bursters: n = 22 cells, 1051 reaches; $R^2 = 0.40$, slope = 25.2 with 95% CI [6.7, 43.6], p = 0.011; Pausers: n = 24 cells, 1255 reaches; $R^2 = 0.75$, slope = -34.5 with 95% CI [-46.4, -22.6], $p = 2.9 \times 10^{-5}$). The population activity in bursters and pausers was nearly equivalent, aside from a notch in the burster population, raising the question of how these diverging patterns may contribute to a coordinated control policy – i.e. a mapping of input to output that achieves a goal.

Populations of ~40 PCs converge onto single nuclear cells[43]. In the oculomotor vermis – where bursting and pausing profiles of PCs strongly resemble the patterns we saw during reach – grouping PCs into populations across bursting and pausing classes revealed much stronger kinematic relationships with saccades[41,44]. Speculating that independent PCs firing patterns might have combined population activity inverse of decelerative nuclear bursts, we next grouped all PCs across all animals and looked at average activity for reaches binned by outward velocity. Strikingly, the broad onset of activity seen in individual Purkinje cells vanished at the population level. Rather, there were sharp drops in net activity during the reach epoch that scaled with the velocity of outreach (Fig. 1i). Quantifying the simple spike firing rate during this pause (see Methods) showed a strong negative relationship with outreach velocity (Fig. 1j; n = 46 cells, 2077 reaches; $R^2 = 0.83$, slope = -2.19 with 95% CI [-2.67, -1.71], $p = 1.0 \times 10^{-8}$). The timing of this pause coincided with the time of peak outward acceleration during outreach, just before the transition to the decelerative phase of reach (Fig. 1k).

Summarizing, we found that individual PCs are privately tuned to specific kinematic features of reach but weakly related to previously observed patterns of firing in the cerebellar nuclei. Yet, at the population level, PC activity shows scaled drops in activity shortly before deceleration, consistent with a disinhibitory mechanism driving decelerative bursts in nuclear cells. We hypothesize that these pauses may be mechanistically akin to conditioned responses seen in delay eyeblink conditioning – learned pauses that produce anticipatory movements in response to predictive cues. Both the precise timing and scaling of the pausing population activity observed here are consistent with learned cerebellar responses linked to motor and sensory contingencies to control movement. As such, this behavior offers a unique opportunity to test theories relating motor adaptation to associative learning in service of skilled movement[24,45–47].

Cspks signal movement onset and reach outcome

To probe mechanisms that might shape cerebellar cortex scaling of output as a function of kinematics, we first identified cerebellar recordings with Cspks, the drivers of learning in PCs. In 14 of 46 recorded cells, Cspks could be sorted across the experiment (see Methods). Within these cells, Cspk probability increased dramatically shortly before movement onset, consistent with reports of synchronized Cspk activity occurring at movement initiation[48–50], then remained elevated during and after reach (Fig. 2a). We analyzed the 13 cells that displayed Cspks during or shortly after reach endpoint (reach onset to 250 ms after limb reversal) and found that in 5 of 13 cells, trials with Cspks had significantly different endpoint positions in the outward and lateral directions (Fig. 2b; $p < 0.05$, paired t-test). The remaining 8 cells showed no discernable tuning of Cspks to position. Additionally, trials with Cspks had significantly elevated simple spike rates during outreach compared to non-Cspk trials (Fig. 2c; $p = 0.0199$, $n = 13$ cells; paired t-test). Thus, Cspks are responsive to movement features and suggest they could act as a homeostatic check on movements where simple spike rate is elevated[23,51–54]. These results imply that PC simple spike rate patterns during reach are likely modified by behavior-specific Cspk activity.

Repeated closed-loop optogenetic perturbation of cerebellar inputs during reach causes hallmark characteristics of sensorimotor adaptation

Next, we sought to probe how cerebellar inputs shape movement kinematics and how erroneous PC simple spike activity alters tuning to these inputs. Previous work has shown that stimulation of pontine afferents to the cerebellum perturbs reaching movements in mice[32]. This effect is interpretable as corrupted cortical information entering the cerebellum which initially drives an erroneous cerebellar control policy and acute kinematic effects. If cerebellar associative learning underlies the formation of an anticipatory control policy, a number of predictions emerge: pontocerebellar mossy fiber stimulation

that drives errant reaching will, when repeated over many reaches, lead to adaptation of PC responsivity. Removing the perturbation should lead to aftereffects due to accumulated learning of new contingencies. Finally, adaptation and aftereffects will be dependent upon the temporal context of the perturbation within the movement, where learning only accumulates when perturbations are temporally locked to the execution of the movement.

To drive erroneous activity in PCs during reaching movements, we injected AAV-expressing hSyn-ChR2 into the pontine nuclei in mice, a major hub relaying motor commands from motor cortex to the cerebellum [32,55–59] (Fig. 3, Fig. S4a). Recordings of PCs showed that optogenetic stimulation of mossy fiber afferents in the cerebellar cortex drove both increases and decreases in simple spike firing rates (Fig. S4b; 15/34 cells, 9 increase, 6 decrease; $p<0.05$, paired t-test). This diverging stimulation effect is likely due to network properties in the cerebellar input layers leading to either net excitatory or inhibitory drive onto PCs[32,60]. Interestingly, cells with sorted Cspks (see Methods) showed an increase in Cspk rate in response to mossy fiber stimulation (Fig. S4c; n = 19 cells; $p = 0.026$, paired t-test), consistent with previous findings during electrical mossy fiber stimulation[61]. Cspks time-locked to mossy fiber stimulation suggest that optogenetically-driven simple spikes may engage plasticity mechanisms to respond to perturbation.

To assess whether repeated closed-loop stimulation could engage cerebellar learning mechanisms to produce sensorimotor adaptation, optical fibers were implanted in cerebellar cortex at the interface between Lobule Simplex and Lobules 4/5 (Fig. S5). Experiments were structured in a block format where animals reached unperturbed in a baseline block, followed by a stimulation block where closed-loop stimulation of pontocerebellar axons (50-ms train) was delivered on every reach when the hand passed a 1-cm threshold in the outward direction, and finally a washout block where stimulation was removed to assess any aftereffects of learning. Each block was roughly 15-30 reaches long determined by each individual animal's endurance in the task (Fig. 3a; baseline: 23.1 ± 6.24 reaches; stimulation: 22.4 ± 5.77 reaches; washout: 20.56 ± 6.65 reaches; mean ± SD; n = 5 animals, 104 sessions). Early in the stimulation block, we found that stimulation caused acute changes in reach kinematics: in 4/5 animals it caused hypermetric reaches in outward position and in 1 animal it caused hypometric reaches (Fig. 3b-c, Fig. S6a examples 1 and 2). To assess the relative change in hand position over the stimulation block, we measured the magnitude of the stimulation effect over the block, defining the initial direction of the stimulation effect on hand position as positive and the opposing direction as negative. We found that the magnitude of the stimulation effect decreased over the stimulation block. When the stimulation was removed, early in the washout block reaches deviated in the direction opposite the initial stimulation direction, before eventually correcting back to baseline at the end of the washout block (Fig 3d,e; n = 5 animals, 104 sessions; end baseline to early stimulation: $p = 0.012$, early stimulation to middle

stimulation: $p = 1.27 \times 10^{-3}$, end baseline to early washout: $p = 0.014$, late stimulation to early washout: $p = 0.043$, early washout to late washout: $p = 6.25 \times 10^{-3}$; paired t-test). The magnitude of the initial stimulation effect predicted the magnitude of the initial washout aftereffect across animals (Fig. 3f, $R^2 = 0.820$, slope = -0.608 with 95% CI [-1.13, -0.0853], $p = 0.0342$); however, hypometric effects were generally larger than hypermetric effects (both during stimulation and washout), possibly due to biomechanical constraints of the limb and reaching apparatus imposing a ceiling effect on hypermetric movements. Interestingly, the aftereffect did not manifest until the time that stimulation would have been delivered during outreach (Fig. 3c, Fig. S6d). In control experiments using red light (635 nm) we observed no kinematic deviations or adaptation profiles as seen with blue-light stimulation (Fig. S6e). Further, blue-light stimulation at rest produced negligible movements (Fig. S6f; n = 4 animals, 21 sessions; maximum outward velocity during stimulation: 0.26 cm/s).

To summarize, we have shown that animals adapt to a precisely timed internal perturbation of pontocerebellar mossy fibers and this learning is reflected in opposing aftereffects when the perturbation is removed. Adaptation was temporally precise, with changes in limb kinematics early in the washout block timed to the predicted point of perturbation.

PC recordings show electrophysiological correlates of adaptation at the time of perturbation

To investigate cellular correlates of learning in PCs during behavioral adaptation to this circuit-level perturbation, we performed stimulation experiments while recording near the optical fiber with a Neuropixel probe. To assure that any firing rate changes were not attributable to unstable cell isolation across the experiment, we assessed the stability of every PC using two metrics: a correlation of spike template waveforms and the displacement of units along the electrode in the baseline and washout blocks (Fig. S7, see Methods). 126 of 231 recorded cells were stable across the experiment, 85 of which were modulated with reach. To assess optogenetic stimulus responsivity in these neurons, we compared simple spike firing rates between baseline and stimulated reaches within the 50-ms stimulation epoch. Consistent with mossy fiber stimulation at rest, we observed a diverging effect pattern with stimulation during reach: 19 cells showed significant increases in simple spike firing and 10 cells showed decreases (Fig. 4a,b, $p < 0.05$, paired t-test). Grouping cells by the initial stimulation effect (increase or decrease) and analyzing simple spike firing rate differences from the last 5 baseline reaches revealed patterns of perturbation and adaptation consistent with kinematic data. The depth of modulation to stimulation was greatest during the first stimulated trial, then progressively restored to baseline levels across the stimulation block, a pattern which held for both stimulation-increase and -decrease cells (Fig. 4c,d). We performed statistical analysis with a sliding window and observed significant differences appearing at the timepoint of stimulation early in the block. By contrast, in later trials, no significant changes in firing

rate were evident relative to baseline (Fig. 4c,d; Stimulation-increase cells: n = 19 cells; first stim: $p = 3.81 \times 10^{-6}$, first 5 stim: $p = 9.45 \times 10^{-3}$, middle 5 stim: p = 0.08, last 5 stim: p = 0.21; Stimulation-decrease cells: n = 10 cells; first stim: $p = 1.95 \times 10^{-3}$, first 5 stim: $p = 9.77 \times 10^{-3}$, middle 5 stim: p = 0.32, last 5 stim: p = 0.49; one sample Wilcoxon test).

To determine the electrophysiological basis for aftereffects in the behavior, we analyzed cells that had significant differences between baseline and the washout block in the stimulation window (Fig. S8a,b). While cells with a stimulation effect and adaptation profile had mixed responses during washout, we identified a population that decreased during the stimulation block, then increased back to baseline during washout (Fig. 4e). Like stimulation responsive cells, significant differences were timed to the stimulation window, and faded over the course of the washout block (n = 25 cells; first wash: $p = 5.96 \times 10^{-8}$, first 5 wash: p = 0.011, middle 5 wash: p = 0.43, last 5 wash: p = 0.73; one sample Wilcoxon test). Interestingly, some of these aftereffect cells showed no acute response to the optogenetic stimulus, indicating that adaptation may be distributed across more PCs than receive erroneous inputs (Fig. 4e), explored in a model below.

Dissociation of adaptation and aftereffects with a randomized perturbation schedule

In both behavioral and electrophysiological data, we have shown that adaptation is temporally specific. We hypothesized that the temporal specificity of perturbation within the reach produced a fixed association between active inputs and error, facilitating adaptation. We therefore predicted that by presenting spatially inconsistent stimuli trial to trial, mice would exhibit little adaptation to stimulation. To test this, we repeated block-stimulation experiments, but rather than stimulating when the hand passed the 1-cm outward plane, we stimulated at a pseudorandomized position in the outward direction uniformly distributed between 0.3 and 1.8 cm (Fig 5a,b). To assess the effect of stimulation at different points in the reach, we aligned reaches to the time of stimulation and measured the difference in position compared to aligned baseline block reaches. Baseline subtracted reach profiles showed a characteristic change in outward position aligned to the time of stimulation, similar to results in fixed-position stimulation experiments. Interestingly, even though perturbation positions were distributed across the stimulation block, we found that animals still exhibited adaptation to the stimulation early in the stimulation block, but this adaptation plateaued to intermediate levels between middle and late block epochs (Fig. 5c,d; n = 5 animals, 60 sessions; Early stim to mid stim: p = 0.037, paired t-test). To assess the presence of aftereffects, we analyzed the positional differences between baseline and washout reaches near the mean of the distribution of stimulus thresholds (50-100 ms after crossing the 1-cm outward plane). Despite evidence for adaption to the randomized stimulation, there were no consistent aftereffects; instead, reaches tended to have a greater distribution of positional differences that averaged

to roughly zero (Fig. 5e; early washout difference from end baseline: -2.58 x 10^{-3} cm; p = 0.957, one-sample t-test).

A cerebellar model accounts for experimental adaptation and aftereffect dissociation

To better understand the non-intuitive adaptation profile of position-randomized stimulation, we modified a simple model of PC firing based on a previously published study[62]. As an input, the model takes 2 populations of 500 parallel fibers, each modulated briefly for 30 ms, that as a population tiled a 400-ms epoch, modeling a hypothetical movement (Fig. 6a). At equilibrium, the two populations of parallel fibers are perfectly balanced during the movement and cause no deviation in the PC firing rate from trial to trial. The model employed a learning rule such that any elevation of the PC rate from this equilibrium would lead to depressing the weights of parallel fibers active at the time of deviation through a Cspk-like error signal, as in cerebellar LTD. Conversely, parallel fibers with depressed weights relax back to baseline levels in the absence of Cspks. We titrated the learning rate to match that observed in fixed-position stimulation experiments (see Methods).

First, we modelled fixed-position optogenetic-perturbation experiments by artificially increasing parallel fiber activity in a random subset of 100 parallel fibers for 50 ms in the middle of the hypothetical movement (Fig. 6a). Initially, this caused a large deviation in the PC firing rate in the stimulated window, resulting in an error and synaptic depression of the concomitantly active parallel fibers (Fig. 6b). Over several repeated perturbation trials this reweighting minimized the effect of the perturbation, correcting PC firing rate back to baseline. After 20 trials, we removed the perturbation. The model output then exhibited opposing aftereffects in PC firing rate at the previous time of perturbation, before eventually relaxing back to baseline. The adaptation profile was quantitatively similar to the empirically observed behavior. Importantly, we note that the aftereffect seen in the PC firing profile is a consequence of depressed weights in both perturbed parallel fibers and other unperturbed parallel fibers that were coincidentally active at the time of the perturbation (Fig. 6c,d). Thus, the model was unable to distinguish the difference between parallel fibers that caused or did not cause a deviation from the target PC activity within the perturbation epoch.

Next, we modeled the position-randomized mossy fiber stimulation paradigm (Fig 6e-g). As with the empirical results, we saw a reduction in the magnitude of the perturbation effect, consistent with high probabilities of Cspks around the time of a perturbation – that is, the perturbed inputs are subject to learning because they are always aligned to the error that follows (Fig. 6e). While the magnitude of adaptation was smaller than observed in the fixed position model, we found that the model learning plateaued late in the perturbation block, similar to empirical observations (Fig. 6g). When the perturbation was removed, there were minimal aftereffects, also consistent with experimental data.

Model weights at the end of perturbation show that this absence of aftereffects is explained by the lack of accumulated learning in coincidentally active parallel fibers; *i.e.*, when perturbations are distributed across the movement, coincidently active parallel fibers are different from trial to trial, and therefore subjected to only transient plasticity (Fig. 6f). Thus, in randomized stimulation, the presence of adaptation illustrated a mechanism by which the cerebellum distinguishes cause-and-effect using time: adaptation is explained by the conserved causal relationship between stimulated PC inputs and error, while the absence of an aftereffect is the result of unaccumulated trial-over-trial learning in coincidentally active non-stimulated inputs. By contrast, aftereffects in the fixed position paradigm are a consequence of the system generalizing attribution of error to fibers that were merely coincidently active relative to perturbation but did not necessarily drive error.

Discussion

Here we discovered a naturally occurring PC population pause during mouse reaching movements that scaled with the velocity of outreach and occurred shortly before the transition to the decelerative phase of movement, reminiscent of emergent PC population kinematic coding in oculomotor vermis during saccades[41]. We speculate that this pause is a type of conditioned response: sensorimotor information relayed through mossy fibers act as learned cues for PCs to scale the decelerative phase of movement via disinhibition of the anterior interposed nucleus. We further demonstrate kinematic effects of mossy fiber stimulation that decrease over trials, akin to sensorimotor adaptation, with concordant changes in PC activity that imply cerebellar associative learning. We observed a surprising dissociation of the adaptation and aftereffect profile seen with block perturbation experiments when randomizing the position of stimulation during reach, designed to test the reliance of adaptation on perturbation context. A model demonstrated that aftereffects are a consequence of misattribution of error to consistently coactive parallel fibers. Conversely, the dissociation of adaptation and aftereffects reflects a lack of accumulated plasticity at a single point during the movement.

By demonstrating remapping of inputs to outputs of the cerebellar cortex, we link concepts developed in delay eyeblink conditioning to adaptation of a skilled volitional movement. Specifically, the mossy fiber stimulation used here to drive reach perturbations is analogous to mossy fiber stimulation used as a conditioned stimulus in eyeblink conditioning. We speculate that motor plan or early kinematic information acts endogenously as such a conditioned stimulus associated with reach outcome that, when erroneous, drives cerebellar learning[47]. We note some nuanced differences between paradigms, however. For instance, adaptation to pontocerebellar stimulation occurs within tens of trials, many fewer than conditioned eyeblink responses, which require hundreds of pairings [63]. However, non-human primates and cats exhibit rapid adaptation consistent with our results in other sensorimotor adaptation

paradigms[30,64]. One possible explanation for these different learning speeds is the richness of a temporal basis set that may emerge in the cerebellar granule cell layer in response to inflow of efferent commands and sensory feedback during movement compared to a relatively impoverished basis set from an invariant, unimodal conditioned stimulus. Indeed, locomotion concurrent with eyeblink training expedites learning[65], consistent with this view. Further, eyeblink responses are generated *de novo,* where PCs must develop a novel response to an unfamiliar stimulus. Conversely, pontocerebellar stimulation during reach alters the execution of a movement for which a mossy fiber to cerebellar output mapping may already exist. Thus, PCs must simply adjust already existing responses, potentially speeding the rate of learning.

Another conspicuous departure from learning seen in eyeblink conditioning is that mossy fiber stimulation during reach drives an error. Thus, the unconditioned stimulus is not externally imposed but is rather the erroneous behavior that results from the perturbed mossy fiber activity. In this sense, the mossy fiber activity that interferes with cerebellar control acts as both a conditioned stimulus and generates a movement error that acts as an unconditioned stimulus to drive learning. In addition, eyeblink conditioning involves the cerebellum associating two stimuli that are not causally linked (*i.e.*, the tone does not cause an air puff) while reach adaptation associates sensorimotor information that is causal to reach error (*i.e.*, the erroneous motor commands generated by the cerebellum cause a reach error). Because this conditioned stimulus cannot be decoupled from the movement error, adaptation should always occur with repeated stimulation even under randomized stimulus conditions. In the case of external perturbations of limb movements, randomizing the direction of perturbations on reaching movements manifests as reach adaptation on the subsequent trial[66], but adaptation does not accumulate because the cause of errors cannot be predicted. Our randomization experiments have a key difference: the internally perturbed mossy fibers are a consistent source of error, allowing the system to drive adaptation to these inputs. Importantly, because these perturbed inputs have no temporal correlation with the movement, no aftereffects are produced in the native population of mossy fibers active in the absence of stimulation.

Isolating a locus of skilled reach adaptation to the cerebellum poses an important conceptual hurdle. Cerebral cortex is a major input to the pontine nuclei – the focus of perturbation in this study – thus learning in motor cortex must be accounted for in cerebellar contributions to movement. Likewise, cerebellar outputs relay information back to motor cortex indirectly via thalamus[67–70]. Previous work has demonstrated that reach-associated pontocerebellar stimulation drives activity in motor cortex[32], meaning each learner in this loop stays apprised of the activity in the other. Could plasticity sites outside the cerebellum account for our observations? Our data argue for a locus of learning in the cerebellum in two major ways: First, we observe a reduced efficacy of mossy fiber drive onto Purkinje cells over

many repeated trials. A parsimonious explanation is that highly-plastic parallel fiber synaptic weights are changing during adaptation rather than cortical commands overriding these proximal perturbations. Second, if PC firing rate changes were caused by modulated afferents to the cerebellum, it would be difficult to reconcile such a mechanism with adaptation to randomized stimulation because these compensatory cerebellar inputs could not predict the time of stimulation.

How might multiple connected brain regions, all of which are implicated in learning, accomplish learning a task in parallel? In our study, mice were expertly trained when we introduced optogenetic perturbation of inputs. Thus, stimulating pontocerebellar fibers, we corrupted the relationship of action directed by motor cortex and the established cerebellar response tuned to that action. Through adaptation, the cerebellum learned to assist movements with these newly modified inputs as evidenced by the diminishing kinematic effect on the limb; when stimulation was removed, the novel mismatch of cortical and adapted cerebellar contribution to the movement again manifests as movement errors. While our data are in line with associative learning paradigms as an explanation for motor adaptation (*i.e.*, direct policy learning[71]), they do not preclude the possibility that associative mechanisms are being used to tune a forward internal model that, when erroneous, changes to adapt the efficacy of the inputs and thus the accuracy of the model.

Methods:

Animals

All procedures followed National Institutes of Health Guidelines and were approved by the Institutional Animal Care and Use Committee at the University of Colorado Anschutz Medical Campus. Animals were housed in an environmentally controlled room, kept on a 12 h light/dark cycle and had ad libitum access to food and water except during behavioral training and testing as described below. Adult C57BL/6 (Charles River Laboratories) mice of either sex (11 females, 6 males) were used in all experiments.

Surgical procedures

All surgical procedures were conducted under Ketamine/Xylazine anesthesia. After induction of anesthesia, the surgical site was cleaned and subcutaneously injected with bupivacaine (2.5 mg/mL). Pressure injections of approximately 150 nL of AAV2-hSyn-ChR2-mCherry were stereotaxically targeted to the left pontine nuclei (-4.0 mm anterior-posterior, -0.5 mm medial-lateral, -5.4 mm dorsal-ventral, measured from bregma) and animals were allowed to recover for a minimum of 8 weeks to ensure expression in mossy fiber terminals in the cerebellar cortex. Custom made aluminum head plates were affixed to the skull centered on bregma with luting (3M) and dental acrylic (Teet's cold cure). Optical fibers (105 mm core diameter, Thor Labs) attached to a ceramic ferrule (1.25 mm, Thor Labs) were implanted into the primary fissure, between Lob 4/5 and Simplex (-6.25 mm anterior-posterior, 1.9 mm medial-lateral, measured from bregma) at a depth of 1.2 mm[72]. For recording experiments, a craniotomy was made medial to the fiber placement and a recording chamber was secured with dental acrylic as previously described[73].

Behavioral task

Animals were allowed a minimum of 2 days of recovery after head fixation surgery, then were food restricted to 80-90% of their baseline weight for reach training. Mice were habituated to the headfixed apparatus by presenting food pellets (20 mg, BioServ #F0163) that could be retrieved with their tongue, then pellets were progressively moved further from the mouth until animals began reaching for food. Pellets were positioned to the right of the animal to encourage reaching with the right forelimb and moved to a consistent position specific to each mouse ~1.2 – 2.5 cm from reach start. Sessions lasted until animals successfully retrieved 20 pellets or until 30 minutes had elapsed, whichever came first. Mice were trained for a minimum of 15 days and were considered fully trained once they could successfully retrieve 50% of pellets 3 days in a row.

Kinematic tracking and closed-loop optogenetic stimulation

Hand position was tracked in real time using an infrared-based machine-vision motion-capture system (6 Optitrack Slim3U Cameras mounted with LED ring arrays, Motive software) at 120 frames-per-second as previously described[2]. Cameras were positioned in front and to the right of the animal and focused on the approximately 8 cm^3 spatial volume that covered the reach area of the right forelimb. 1-mm diameter retroreflective markers were used for camera calibration and affixed to the mouse hand for kinematic tracking. A custom-built calibration wand and ground plane were used to set position and orientation of the cameras in Optitrack Motive software. Camera calibration was refined monthly to account for any drift of the cameras over time. Calibrations that reported a mean triangulation error <0.05 mm were considered passes. The spatial origin was set to be at the center of the bar where mice placed their hand during rest. Spatial blocking and camera detection thresholds were adjusted to prevent erroneous tracking of minimally infrared-reflective objects.

Real-time hand positions were streamed into MATLAB (2018a) with a latency under 1 ms. Custom-written MATLAB code was used to detect when the hand passed a positional threshold 1-cm outward from the bar where the mice rested their hand then send a 'go' signal to an Arduino microcontroller (Uno) which triggered a laser (1.5-5 mW) with TTL pulses. To ensure low-latency closed-loop stimulation we used an open-source C++ dynamic link library[74] edited to reflect the parameters of laser stimulation (50-ms stimulation, 100 Hz, 2-ms train). This system has a closed-loop latency of 9.5 ms from the time of threshold crossing (120 fps camera frame rate, 0.5 ± 0.1 ms (mean ± SD) MATLAB-Arduino communication). Hand positions and stimulation times were streamed into MATLAB and saved for post-processing.

Kinematic analysis

All kinematic analysis was performed using custom-written MATLAB code. First, erroneously tracked objects were removed using a nearest neighbor analysis, which assessed the closest markers in subsequent frames and removed others, to produce a single positional trajectory of the hand marker over time. Any dropped frames where the marker was not detected were interpolated over, then data were filtered using a 2^{nd}-order low-pass Butterworth filter (10 Hz)[75] using MATLAB's zero-phase filter function "filtfilt". Last, interpolated points were removed such that the filtered marker positional data reflected only data captured during the experimental session.

To segment continuous data into reaches, we found instances of the marker passing the 1-cm positional threshold in the outward direction and clipped 10-s segments centered on this time point. We defined outreach as the segment of this data from the time before threshold crossing that the hand exceeded 2 cm/s in outward and upward velocity to the time after threshold crossing where the hand stopped moving in the positive outward direction (outward velocity < 0 cm/s). Occasionally, the marker

would become obscured behind the pellet holder during reach or spurious detection of the nose would jump the marker position to the nose and be detected as a reach. Therefore, to prevent against analyzing reaches that had large segments of data missing, any threshold crossings where the marker dropped greater than 25% of points between the start and end of outreach were not considered for further analysis.

Reach velocity and acceleration were calculated using the numerical gradient between position timepoints in each dimension. To produce aligned reach position curves, we interpolated data at 10 ms centered on the time the hand passed the 1 cm positional threshold crossing in outward direction. The effect of stimulation was assessed by measuring changes in stimulation and washout reaches (early, middle, and late) relative to the last 5 baseline reaches in the 50-ms interval following the end of stimulation. To assess the unadapted effect of stimulation or washout, early reaches were defined as the first reach in each block; middle and late reaches were the middle 5 and last 5 reaches of reach block, respectively. To align random-stimulation position reaches, we found the positional threshold of stimulation on each reach, then aligned stimulation reaches and baseline reaches to the time they crossed this boundary during outreach, averaged across reaches, then measured the difference in these curves, yielding the stimulation-aligned positional difference between end baseline and stimulation reaches. For washout reaches in random-position stimulation experiments, reaches were aligned to the time of the threshold crossing at 1 cm such that the aftereffect could be compared to fixed-position stimulation experiments. To account for varying effects of stimulation seen across animals (hypermetric and hypometric movements), the direction of positional change in early stimulation reaches relative to baseline for each animal in random- or fixed-position stimulation experiments was defined as the positive direction and the opposing direction as negative for that animal in each paradigm, allowing us to group data across animals with diverging effects. To assess the time course of stimulation effects within individual animals, we measured differences in position at each timepoint between the early stimulation reaches and baseline reaches using a Wilcoxon signed rank test. Reach endpoint probability density heatmaps (Fig. 2) were made by filtering endpoint positions with a two-dimensional gaussian filter (sd 0.35 mm).

Electrophysiology

Craniotomies were made over the cerebellum ipsilateral to the reaching arm of in fully trained animals. A custom-made recording chamber was implanted over the craniotomy, the brain was covered with triple-antibiotic cream (Globe), and the recording chamber was sealed with Quik-sil silicone (World Precision Instruments) such that it could be preserved for multiple recordings.

Single electrode recordings

Single-electrode recordings were performed with 3-5 MOhm platinum/tungsten optrodes (Thomas Recording). Once animals were headfixed, the electrode was targeted to -6.25 mm anterior-posterior, 1.9 mm medial-lateral (measured from bregma) then lowered into the brain up to a depth of 1.8 mm with a motorized micromanipulator (Newport Motion Controller, Model 861). Signals were band-pass filtered at 300-5000 Hz, amplified with a MultiClamp 700A amplifier (Axon Instruments), then digitized (CED Power3 1401) and recorded with Spike2 software (CED). Once a putative PC was isolated, the brain tissue was allowed to relax for 15 minutes. Cell sorting was performed offline using Psort[76].

Neuropixel recordings

Neuropixels were lowered into the brain using a motorized micromanipulator (Sensapex uMp micromanipulator). Once the electrode shank spanned the putative PC layer, the tissue was allowed to relax for 15 minutes. Electrophysiology data was acquired using an OpenEphys system (https://open-ephys.org/gui). Data were sorted offline in Kilosort2[77] and manually curated in phy (https://github.com/cortex-lab/phy).

Neural data analysis

Following sorting, isolated units were analyzed offline using custom-written MATLAB code. In well-isolated single-electrode units, simple spikes and identifiable Cspks were sorted using Psort. To identify Cspks in Neuropixel recordings, we cross-correlated cells with high firing rates in the cortex with adjacent low-firing-rate clusters and looked for the presence of a Cspk-aligned simple spike pause and characteristic simple spike and Cspk waveforms. In many cells Cspks could not be identified across the length of the experiment. In these cases, we identified PCs based on cortical location and electrophysiological criteria using the firing rate, CV2, and median absolute difference from the median interspike interval (MAD)[78]. Cerebellar cortical cells with a firing rate > 40 spikes/s, CV2 > 0.20, MAD < 0.008 were labeled as PCs (Fig. S2). Instantaneous firing rates for PCs were calculated by taking the inverse of the ISI between spikes, convolving with a 20-ms gaussian, then sampling at 10-ms intervals. In Neuropixel recording adaptation experiments we analyzed reach modulated PCs, defined as exhibiting a firing rate change during the reach epoch ≥ 1 standard deviation of the mean firing rate of the cell. To find the time and magnitude of firing rate pause in grouped population PC data, we found the first local minimum in smoothed firing rate traces that occurred after reach onset.

In pontocerebellar stimulation experiments, to assure that observed simple spike adaptation was not the result of changing unit isolation across the experiment, we assessed unit stability with two metrics: waveform correlation and unit displacement across the experiment[79]. To assess waveform correlation, we isolated the template waveforms for each unit on the electrode with the greatest spike amplitude and the 32 surrounding electrodes (33 total). We averaged 1000 randomly selected spike waveforms for each channel from the baseline block and the washout block, concatenated waveform templates across the 33 channels, then correlated the concatenated waveforms from the baseline and washout blocks (Pearson correlation). As a shuffled control, we correlated concatenated templates from neighboring units in the baseline and washout block. Neighboring units were defined as those whose 32 surrounding electrodes overlapped with the unit of interest. PCs whose across experiment waveform correlation did not exceed the 99th percentile (0.735) of the across-unit shuffled control correlation were excluded from further analysis.

To assess cell displacement across the experiment we calculated the position of unit (x, y) using

$$(x, y) = \left(\frac{\sum_{i=1}^{N} x_i a_i^2}{\sum_{i=1}^{N} a_i^2}, \frac{\sum_{i=1}^{N} y_i a_i^2}{\sum_{i=1}^{N} a_i^2} \right)$$

where N is the number of electrodes, x_i and y_i are the lateral and upward position of the electrode, and a_i is the peak-to-peak spike waveform amplitude on the i^{th} electrode. Unit displacement was defined as the Euclidean distance between unit positions in the baseline and washout blocks. As a shuffled control, the displacement between neighboring units (as defined above) across the experiment was calculated. PCs whose displacement was above the 1st percentile (2.172 μm) of shuffled control were excluded from further analysis.

LASSO regression

To quantify the variance of PC simple spike firing rate that could be explained by reach kinematics, we used least absolute shrinkage and selection operator (LASSO) regression[34]. LASSO has the advantage of performing both regressor selection and regularization, producing a sparse model of many correlated kinematic regressors. 23 kinematic variables were used as regressors, including position, velocity, and acceleration in the upward, outward, and lateral directions, speed, unsigned acceleration, with each velocity and acceleration term additionally broken into positive and negative components. A full list of regressors is included in Fig. S3. Data for each reach were clipped into 2-s segments centered at the time of a 1-cm threshold crossing in the outward direction. Regression was performed using a custom-written MATLAB code using the “lasso” function. All kinematic data were standardized to have a mean

of 0 and a variance of 1, and regression was performed with a 10-fold cross validation to avoid overfitting. To find the appropriate offset of firing rate and kinematics, instantaneous simple spike firing rates for each reach were offset by lags from 0 ms to -300 ms (firing rate leading kinematics) in 10-ms steps. The lag that minimized the mean squared error (MSE) of the regression was selected for each cell. To calculate the variance of firing rate explained, the predicted firing rates from the best fit regression were calculated from the kinematic data and compared to empirical data. R^2 was calculated using:

$$R^2 = 1 - \frac{SS_{res}}{SS_{tot}}$$

where SS_{res} is the sum of squared residuals and SS_{tot} is the total sum of squares.

For the spike shuffled control, spike times on individual trials were shuffled in time so that each reach epoch had the same mean firing rate, then converted to instantaneous firing rates as described above. For the reach shuffled control, reaches were assigned to firing rates recorded on different reaches. For both controls regressions were performed at the lag that minimized the MSE for empirical data and repeated 100 times; R^2 values of each shuffled control were taken as the average of these 100 regressions. To assess the unique contribution of individual kinematic regressors to the fraction of variance explained in the empirical data regression, each regressor was time shuffled independently and regressions were repeated. The change in R^2 value between the regressor shuffled regression compared to the complete empirical data model is the fraction of unique contribution to total variance explained for each kinematic variable[40].

Cerebellar model

The cerebellar model in the paper was derived from a previously published model[62] and written using custom code in Python. A major difference between our paper's model and the cited model is the assumption of a continuous temporal input of parallel fiber activity distributed across a hypothetical 400-ms movement, rather than a single parallel fiber input trial over trial. The model PC was fed 1000 parallel fibers that either increased or decreased their activity from baseline for 30 ms during a 400-ms interval, mimicking hypothesized temporal basis sets produced by the granule cell layer[80–82]. PC firing at time t on the n^{th} trial was calculated as the sum of the weighted contribution of all parallel fibers at time *t*:

$$PC_n(t) = PC_0(t) + \sum_{i}^{1000} w_n^i PF_n^i(t)$$

Here, w_n^i is the weight of parallel fiber *i* on the n^{th} trial and PC_0 is the baseline firing rate of the PC. Parallel fiber weights were adjusted following each trial according to two parameters: the probability of a Cspk as a function of trial error $\beta P(CS|E_n)$ where β dictates the strength of synaptic depression in response to a Cspk, and a decay term, α_{PF}, that relaxes parallel fiber weights back to their initial value w_0^i:

$$w_{n+1}^i = \begin{cases} w_n^i - (1-\alpha_{PF})\left(w_n^i - w_0^i\right) - \beta P(CS(t)|E_n(t)) & \text{if } PF^i(t) > 0 \\ w_n^i - (1-\alpha_{PF})\left(w_n^i - w_0^i\right) & \text{if } PF^i(t) < 0 \end{cases}$$

The probability of a Cspk is a function of *t*, where greater deviations in the PC rate from baseline at time *t* lead to greater probabilities of Cspks. Specifically, we took the mean error in the preceding 25-ms interval ($E_n\,(t)$) to calculate the probability of a Cspk at each time in the movement interval:

$$P(CS(t)|E_n(t)) = \frac{a}{1+e^{-\tau E_n(t)}} - \frac{a}{2}$$

To obtain values for the parameters a and τ, we fit a curve to the change in position of early, middle, and late stimulated reaches in fixed-position stimulation experiments, then took the derivative of this curve to obtain the error correction (trial-over-trial positional change) for a given error magnitude.

We ran the simulations mimicking the experimental block structure used for empirical data, including a baseline block with no perturbation, an experimental block with a perturbation on every trial, and a washout block with the perturbation removed. For perturbation trials, we added activity to a random subset of 100 parallel fibers at t = 200 ms for 50 ms that, when combined, drove 40 simple spikes/second in PCs at their initial weights. After 20 such trials, the perturbation was removed, and the model was run for an additional 20 washout trials. To simulate random-position perturbation experiments, the time of perturbation was changed on every trial.

Acknowledgements

We thank the members of the Person lab for critical feedback on the manuscript; the Neurotechnology Center at the University of Colorado Anschutz Medical campus for use of core facilities including the Advanced Light Microscopy Core and the Optogenetics and Neural Engineering Core. We thank Mr. Michael Spindle and Ms. Elena Judd for technical assistance. Work was supported by F31 NS113395-01 to DJC; and NS114430, NSF CAREER 1749568, and by a grant from the Simons Foundation as part of the Simons-Emory International Consortium on Motor Control to ALP.

639

640 **Contributions**

641 DJC and ALP designed experiments, interpreted data, and wrote the manuscript. DJC performed
642 experiments, built the computational model, made figures, and analyzed data. MIB developed the closed-
643 loop system, edited the manuscript, and contributed to early conceptualization of the project.

644

Figures:

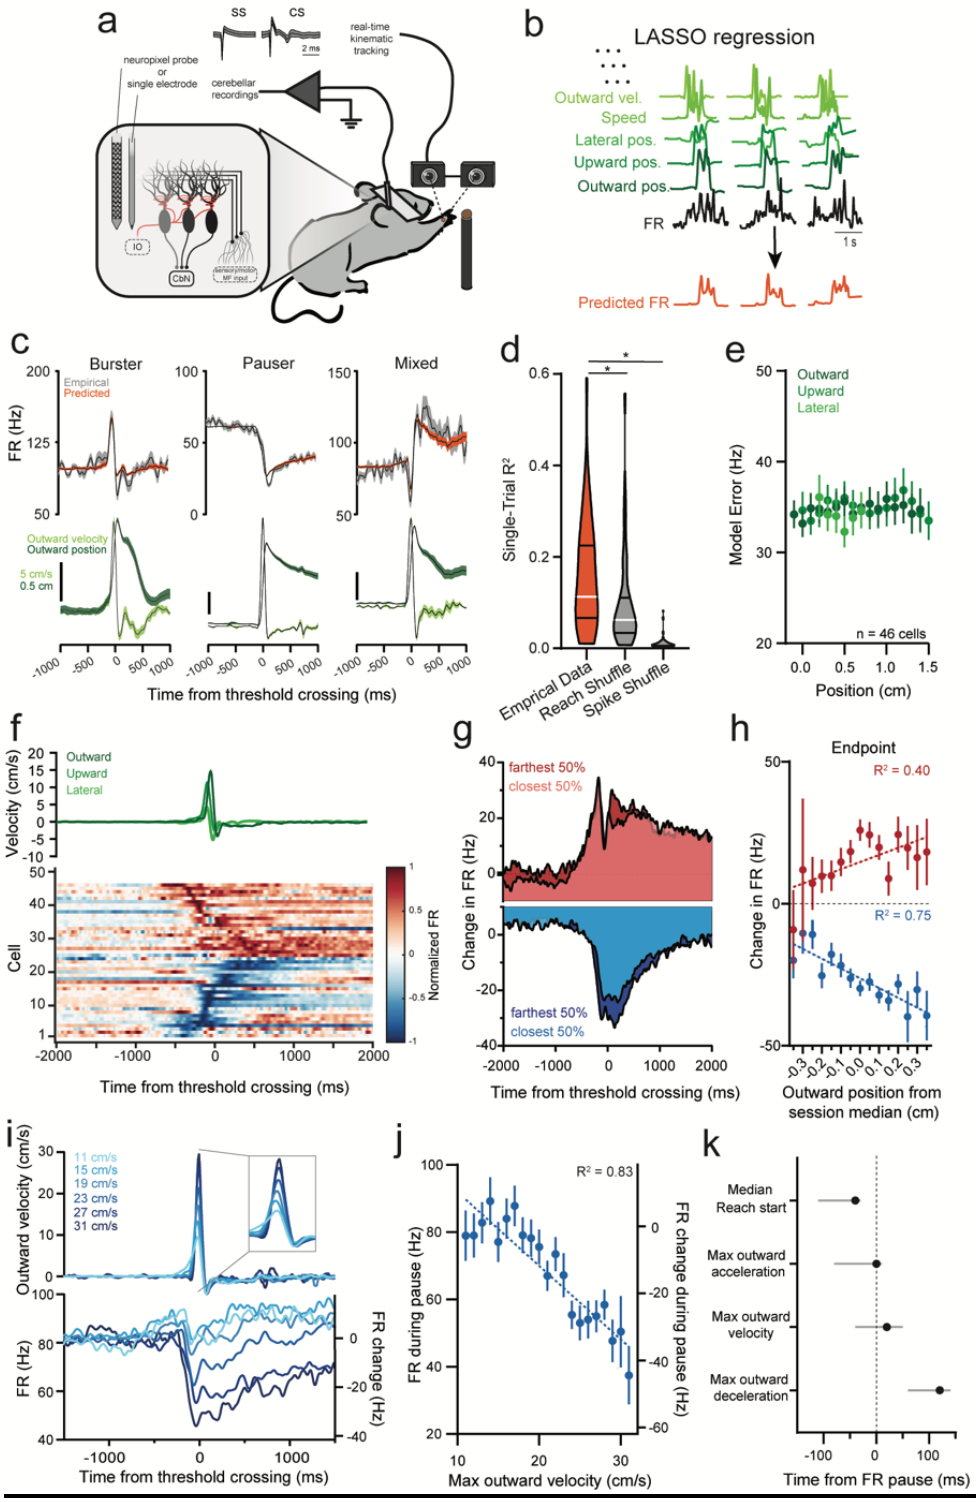

Figure 1. Net population activity in PCs predicts reach velocity.

- **a.** PCs of the deep central sulcus were recorded with either single electrodes or Neuropixel probes while the reaching hand was tracked in real-time with high-speed cameras.
- **b.** Kinematic regressors in multilinear LASSO regression were used to model firing rates on individual reaches across sessions.
- **c.** Examples of 3 PCs fit with LASSO regression. Top: trial-averaged empirical and LASSO predicted firing rates. Bottom: outward position and velocity aligned to firing rate at optimal lag.
- **d.** Modest single-trial R^2 for single cells in empirical, reach shuffled, and spike shuffled LASSO regressions.
- **e.** Stable continuous encoding of reach kinematics across reach epoch seen with invariant absolute value of the difference in empirical firing rate and predicted firing rate (model error) across outward, upward, and lateral positions binned at 0.1 cm.
- **f.** During reach (green), PCs group roughly into bursters (red) and pausers (blue), aligned to the time the hand passed 'threshold', 1-cm in the outward direction.
- **g.** Depth of modulation scales with reach distance for both bursters (red) and pausers (pausers). 50% farthest and 50% closest reaches grouped.
- **h.** Change in firing rate for bursters (red) and pausers (blue) during outreach for endpoints binned at 0.05 cm in the outward direction relative to the median endpoint for each session.
- **i.** Pooling burst and pause PC populations reveals net pausing activity that scales with reach velocity. Top: Reach velocity curves binned at 4 cm/s for the population of reaches across all recorded PCs. Bottom: PC population firing rate corresponding to each bin (left y axis), firing rate change relative to 1-2 s before threshold crossing (right y axis).
- **j.** Firing rate during the pause in population activity was strongly negatively related to reach velocity binned at 1 cm/s.
- **k.** Time of population pause corresponds to time of peak outward acceleration and precedes deceleration. Plot relates the median timing of reach start, peak outward acceleration, peak outward velocity, and peak outward deceleration to the time of population simple spike pauses for all reach velocity bins (n = 21 bins, 1 cm/s).

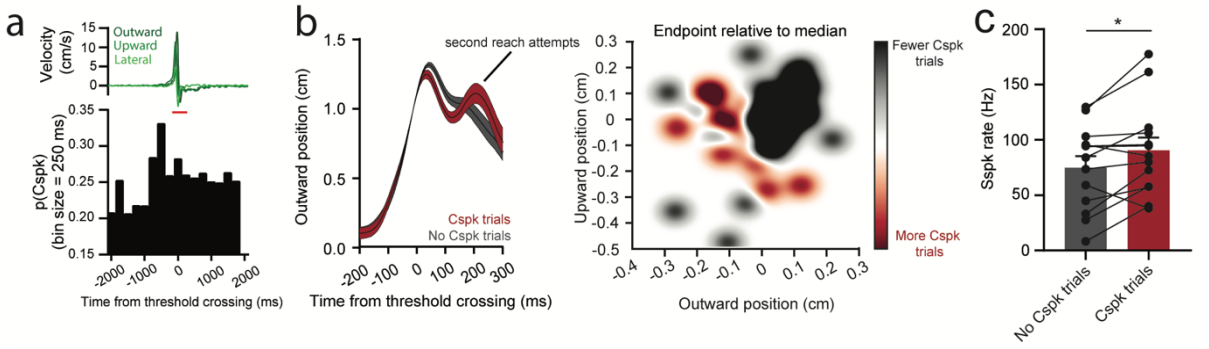

Figure 2. Reaches with Cspks have erroneous kinematics and elevated simple spike rates.

a. Positive modulation of Cspks around the time of reach. Top: Mean velocity of reaches with Cspks recorded. Bottom: PETH of Cspk activity relative to the time of threshold crossing. Window used to define trials with Cspks spanned the outreach +250 ms (red bar).

b. Kinematics from an example session separated into reaches with (red) and without Cspks (black). Left: positional profiles. Right: probability density heatmap of upward and outward endpoints.

c. PC simple spike (Sspk) rates were significantly higher during outreach in trials with Cspks (n = 13 cells).

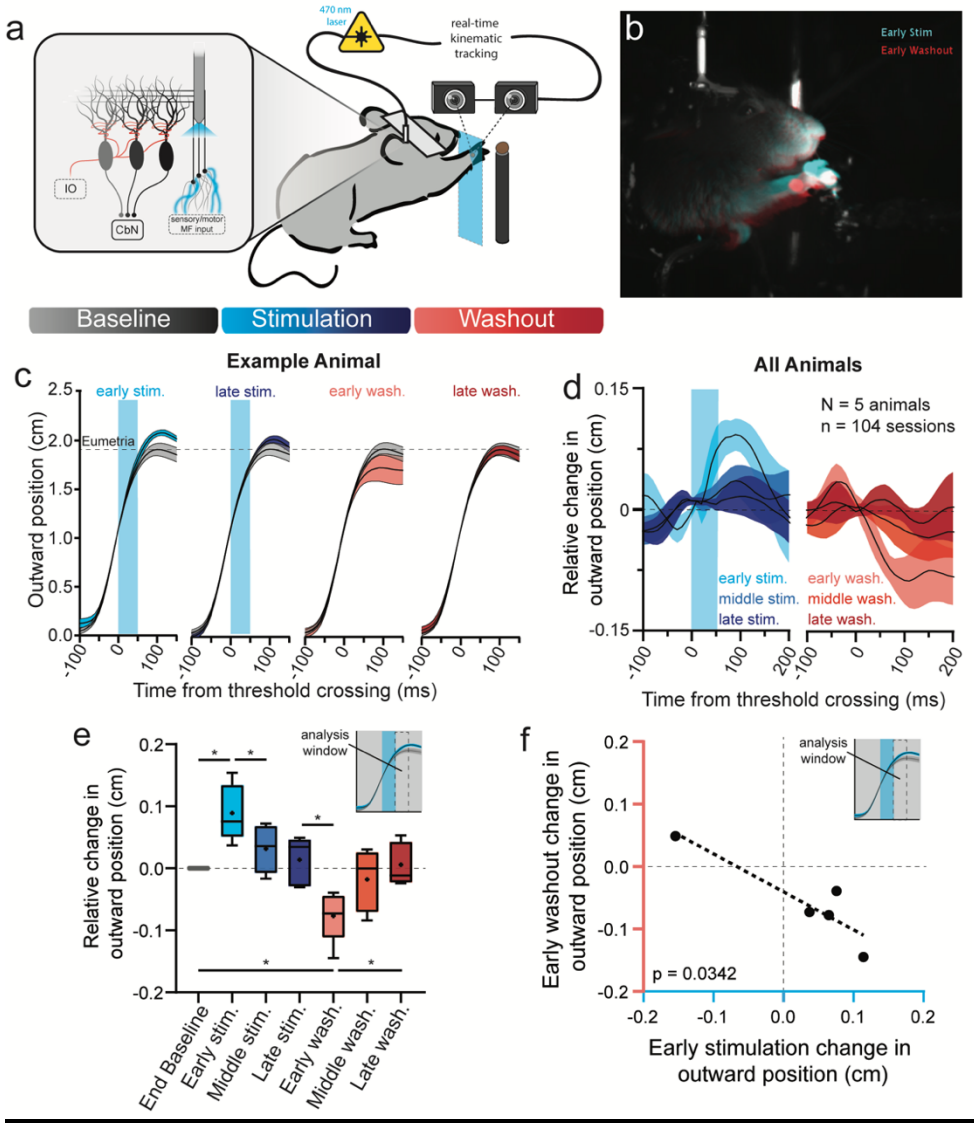

Figure 3. Adaptation to mossy fiber stimulation during reach.

- **a.** Headfixed mice expressing ChR2 in pontocerebellar mossy fibers were trained to reach for food pellets while the hand was tracked with high-speed cameras. On laser trials, light directed to cerebellar primary fissure through an implanted fiber was triggered in closed loop after the hand crossed a plane 1 cm outward from reach start position. Bottom: Perturbation schedule followed canonical adaptation structure, with a baseline (no-stimulation) block, stimulation block with stimulation on every reach, followed by a washout block with stimulation omitted.
- **b.** Hand position 100 ms after threshold crossing in the first stimulated (blue) and washout (red) reaches heading to the target (white).
- **c.** Hand position during baseline (grey), compared to hand position measured across the adaptation and washout blocks in an example mouse (n = 20 sessions). Blue shading denotes the time of mossy fiber stimulation.
- **d.** Summary of stimulation-induced kinematic effects, which decay over the adaptation block and show opposing aftereffects. Baseline subtracted hand position, rectified relative to the direction of kinematic effect of stimulation, is shown for reaches in the early (first reach), middle (middle 5), and late (last 5) phases for both stimulation (blue) and washout (red) blocks. N=5 mice; 104 sessions.
- **e.** Summary of adaptation effects across animals and sessions. Relative change in outward position was assessed in the 50-ms window following the end of stimulation. Asterisks indicate statistically significant differences between blocks (p values reported in main text).
- **f.** The magnitude and direction of early stimulation effect was related to aftereffects. Plot shows linear regression relating the magnitude of the early stimulation effect and early washout effect compared to baseline reaches.

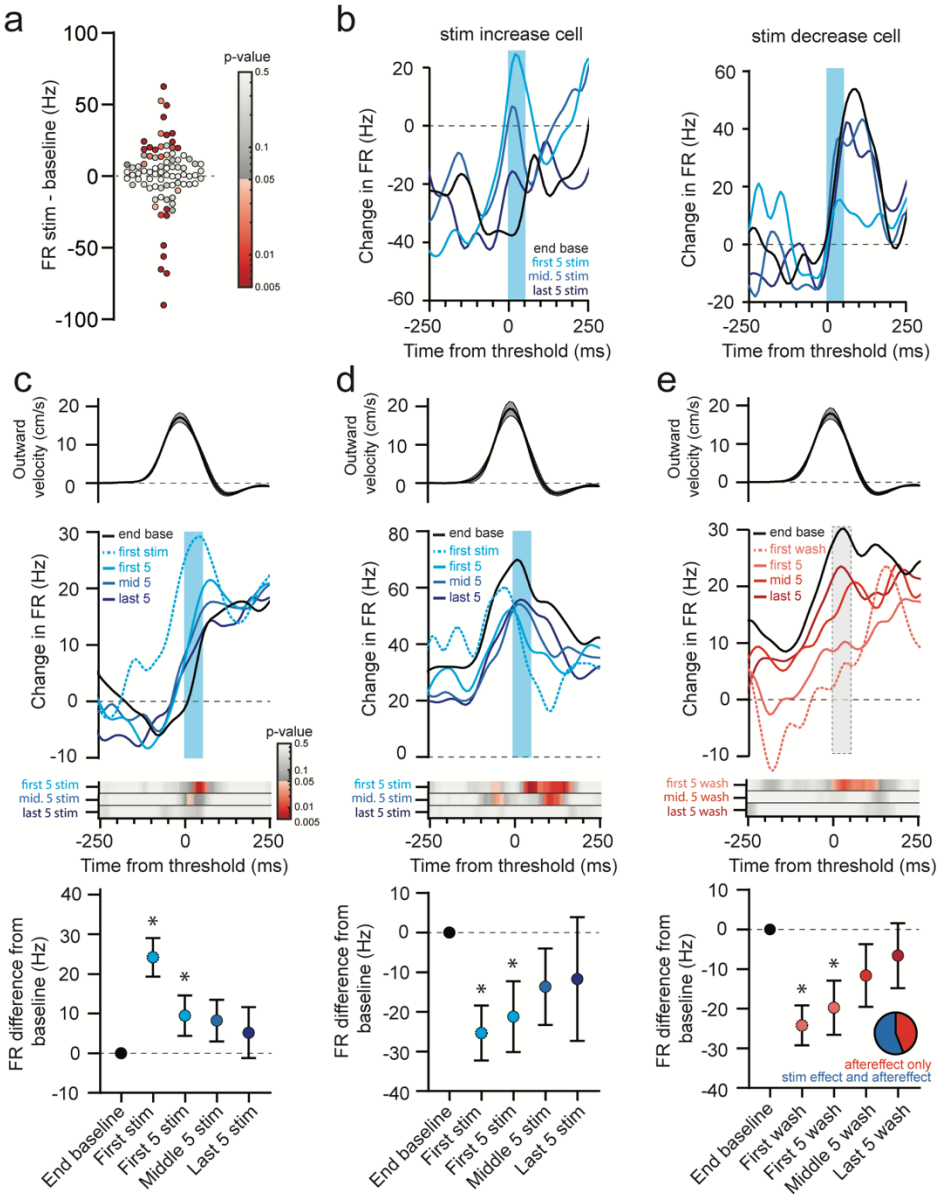

Figure 4. PCs show electrophysiological correlates of behavioral adaptation over the stimulation and washout blocks.

- **a.** Mossy fiber stimulation effect during reach of all reach-modulated PCs. The difference in simple spike rate during the stimulation window is compared to the same epoch during baseline reaches (n = 85 cells). Significant differences are denoted by the color map on the right.
- **b.** Two example cells showing stimulation responses over the adaptation block. (Left) A cell that was positively modulated by stimulation. Colors denote firing rates computed at different points in the session. Black, end of baseline; Blue, light-dark first 5; middle 5, and last 5 reaches of the stimulation block. (Right) Same as left but for a PC that was negatively modulated by stimulation. Blue shading denotes the time of the 50-ms stimulation window.
- **c.** Population summary of activity of PCs firing rate adaptation over stimulation block for all PCs positively modulated by stimulation. Top: mean reach velocity for all sessions. Middle: Average baseline subtracted simple spike rates for the last 5 baseline reaches and the first, first 5, middle 5, and last 5 stimulated reaches. The heat map below indicates the point-by-point significance of differences of the first 5, middle 5, and last 5 stimulated reaches relative to the last 5 baseline reaches (10-ms bins). Bottom: Quantification of changes of PC rates compared between stimulation and baseline reaches, measured during the time window of the stimulation epoch. n = 19 cells.
- **d.** Same as in **c.** but for the population of PCs negatively modulated by stimulation. n = 10 cells
- **e.** Same as in **c.** but for the population of cells that showed significant decreases in washout block relative to baseline reaches, measured during the stimulation epoch. The pie chart indicates the number of “aftereffect cells” that also responded to stimulation. n = 25 cells

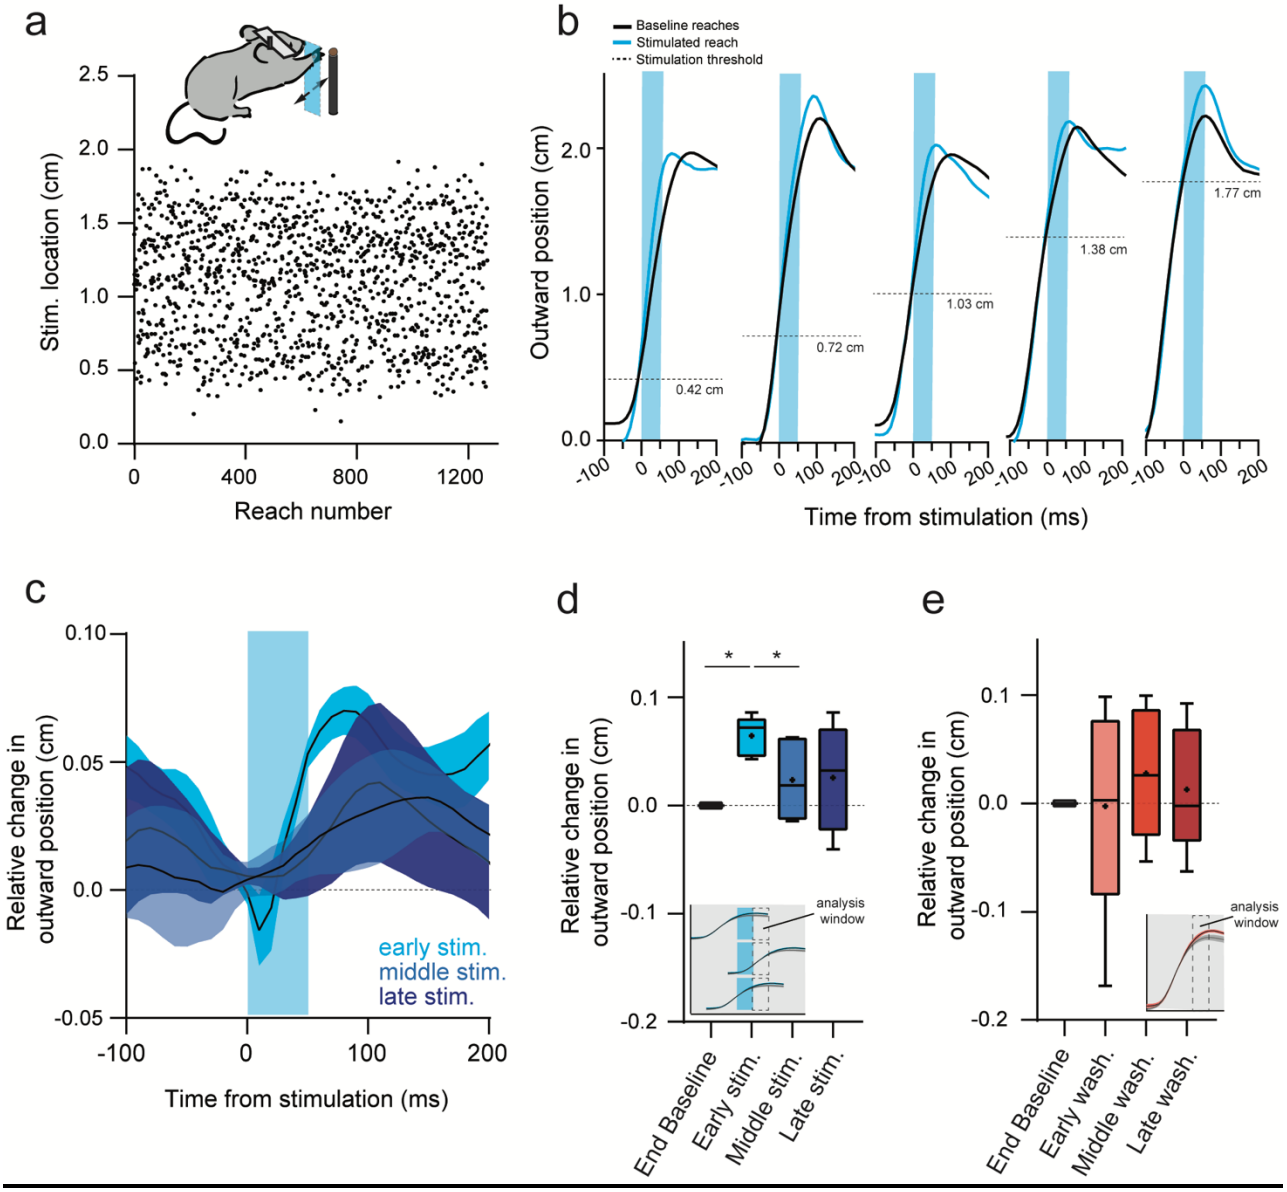

Figure 5. Dissociation of adaptation and aftereffects with randomized stimulation position.

- **a.** Stimulation location during outreach was distributed pseudorandomly between 0.3 and 1.8 cm in the outward direction during the stimulation block.
- **b.** Examples of reaches stimulated at 5 different locations during outreach. Each stimulated reach is compared to the last 5 baseline reaches of each session. The horizontal dashed line indicated the threshold crossing that triggered stimulation.
- **c.** Summary data of relative change in outward position for stimulation reaches in the early, middle, and late block. N=5 mice; 60 sessions.
- **d.** Quantification of stimulation effect on outward position across adaption block. For each reach, the analysis window was the 50-100 ms after stimulation onset aligned to the time of threshold crossing for each reach (inset).
- **e.** Quantification aftereffects on outward position during washout block. Here, the analysis window is the 50-100 ms after crossing the 1-cm threshold for each reach – the same as the analysis in fixed-position stimulation experiments.

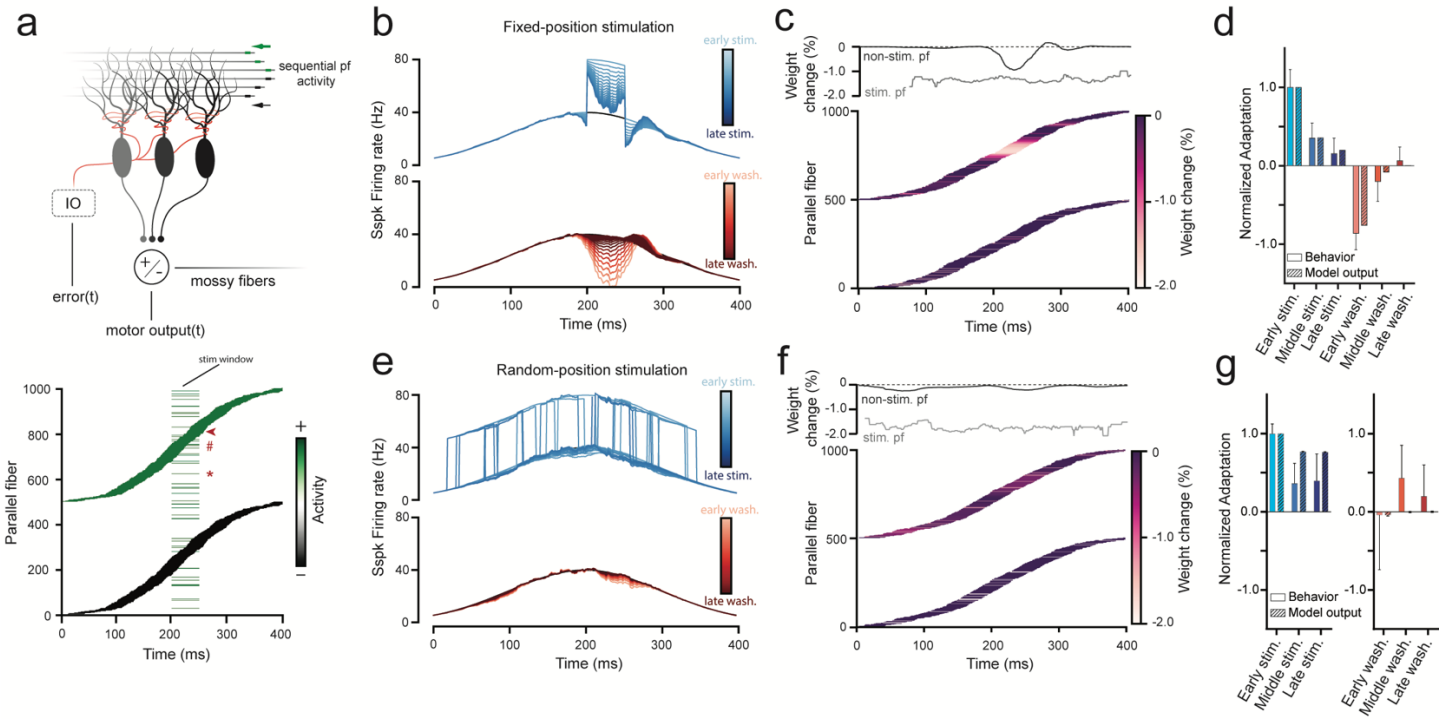

Figure 6. A cerebellar model accounts for adaptation and aftereffect dissociation.

a. Schematic diagram of the temporal cerebellar-learning model. The model input is 1000 parallel fibers, divided into 2 balanced populations, activated during a brief window during a simulated 400-ms movement. The output of the PC module that receives this information is compared to the input in the cerebellar nuclei. At equilibrium, the 2 parallel fiber populations are perfectly balanced (one increases, and one decreases activity; bottom) and the PC module outputs an activity curve (gaussian) that spans the movement. Deviations from this normal curve (errors) lead to mismatch in the nuclei and subsequent activation of the inferior olive, which reduces the weights of parallel fibers active shortly before the error. To simulate optogenetic perturbation experiments (bar-code like pattern at 200 ms), a step of activity was added to a subset of parallel fibers for 50 ms in the center of the movement (fixed stim) or randomized across the block (random stim). Note that stimulation can either activate a cell twice (e.g. parallel fiber 617 *) or overlap with endogenous activity (e.g. 756 #), and non-stimulated neurons can be endogenously active during the stimulus window (e.g. 805 arrow).

b. PC simple spike activity during the stimulation block (top, blue) and washout block (bottom, red) showing progressively adapting response magnitudes during the adaptation block and progressively decaying aftereffects during washout.

c. Parallel fiber weights at the end of the fixed-position stimulation block. Top: change in weights of “artificially” stimulated and non-stimulated parallel fibers plotted by time of endogenous activation. Bottom: heatmap of parallel fiber weight changes. Note population weight change concentrated at time of stimulation, seen in both artificially stimulated and unstimulated fibers during stimulation epoch.

d. Comparison of model output to empirical observations for fixed-position stimulus conditions (Fig.3). Model closely matches behavioral adaptation.

e. Same as **b.** but here the stimulation window is randomized across the reach.

f. Same as **c.** but for random position stimulation experiments. Note the absence of clustered weight changes in unstimulated parallel fibers.

g. Comparison of model output to empirical observations for random-position stimulus conditions (Fig.5) showing that both model and empirical observations show adaptation but not directional aftereffects.

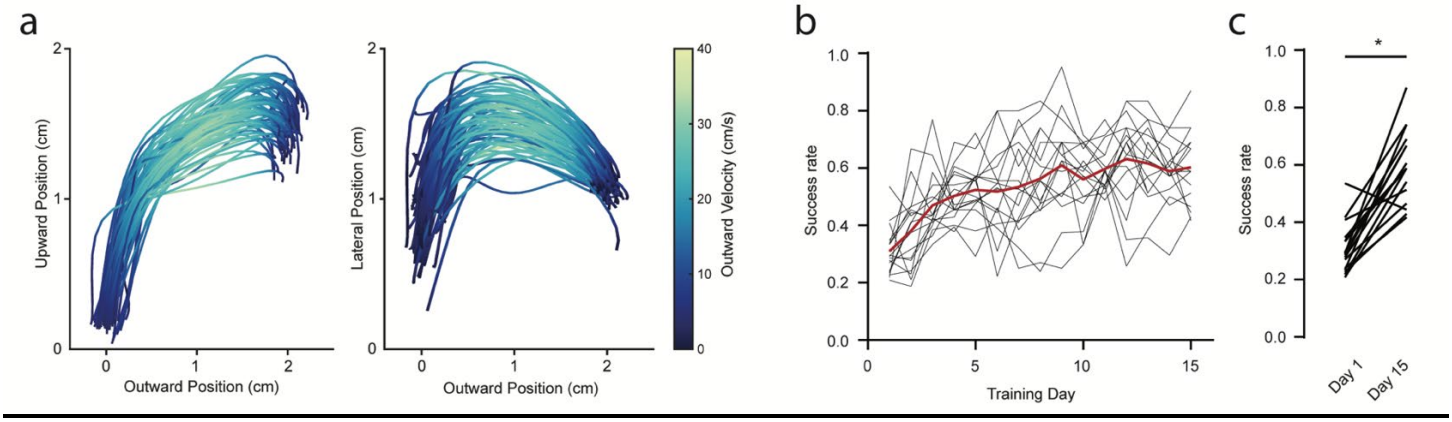

Figure S1. Reach tracking and reach performance over sessions.

a. The right hand was tracked with high-speed cameras as mice reached upwards and outwards towards a food pellet. Positional outreach trajectories from a single session viewed are shown from a lateral (left) or bottom-up (right) vantage point with traces colored by the magnitude of outward velocity.

b. Mice were trained for a minimum of 15 days on the reaching task. Pellet retrieval success was tracked throughout training for each mouse, mean is shown in red (n = 17 mice).

c. Quantification of success rate on day 1 of training and day 15.

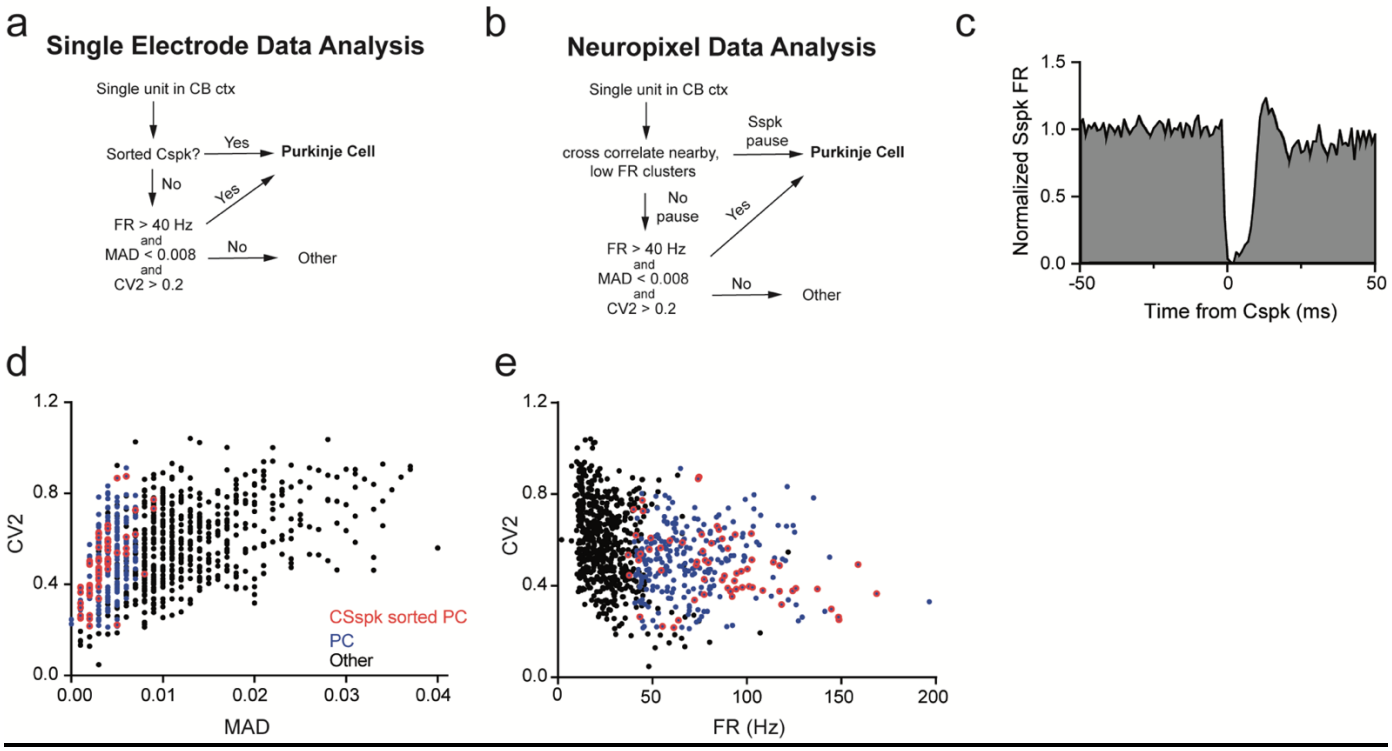

Figure S2. PC identification by firing rate characteristics.

a. Cerebellar recordings using single electrodes were first anatomically targeted to cerebellar cortex. If a recorded cell had visible Cspks they were classified as PCs. Otherwise, if cells had a firing rate > 40 Hz, a median absolute difference firing rate from the median interspike interval (MAD) < 0.008, and a CV2 > 0.2, they were classified as PCs[78].

b. Neuropixel-recorded single units were cross correlated with nearby (<200 microns) low firing rate (<5 Hz) single units. If this cross correlation exhibited the characteristic firing rate pause seen in PC simple spikes after a Cspk, these units were classified as the simple spikes and Cspks of a single PC. If no pause was seen, cells that exhibited the same firing rate, MAD, and CV2 profile described in **a.** were classified as PCs.

c. Example simple spike pause aligned to the time of a Cspk from a Neuropixel recording.

d. MAD and CV2 values plotted for all recorded cells. Colors indicate whether a cell was a PC classified with Cspks (red), a PC classified by firing rate criteria (blue), or cells not classified as PCs (black).

e. Same as in **d.** but plotting CV2 and firing rate for each cell.

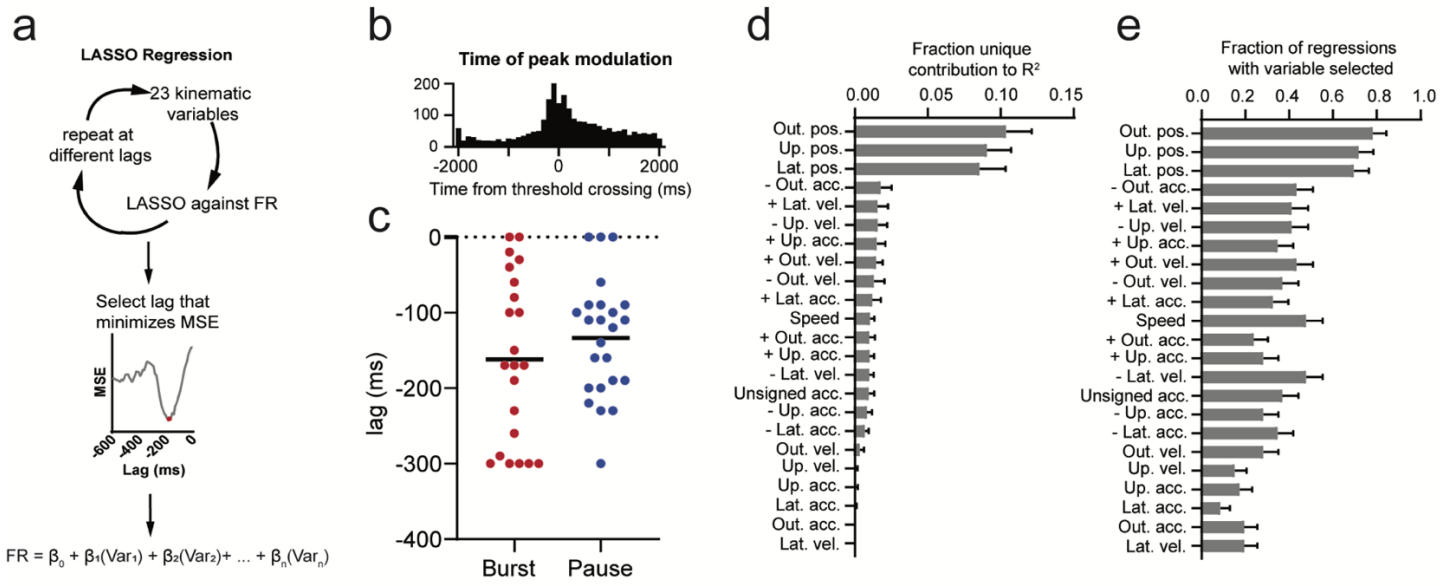

Figure S3. Lasso regression details.

a. Schematic of lasso regressions. 23 kinematic variables were regressed against firing rate at different lags from 0 to -300 ms. The lag that minimized the mean squared error (MSE) of the regressions was selected.

b. Peak modulation time of all cells across all reaches (n = 2391 reaches, 46 cells).

c. Optimal lags of the LASSO regression for each cell.

d. Fraction of the unique contribution to total variance explained for each regressor.

e. Fraction of regressions with each variable selected.

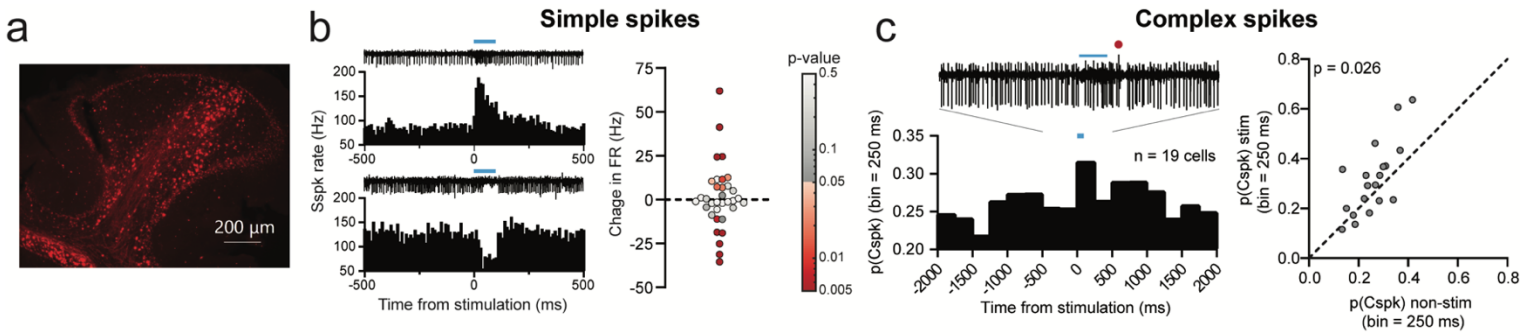

Figure S4. Changes in PC firing during optogenetic stimulation of mossy fibers.

a. Mossy fiber boutons expressing hSyn-ChR2-mCherry in the cerebellar cortex.

b. Simple spike responses to mossy fiber stimulation. Left: examples of single-cell simple spike responses to mossy fiber stimulation. Right: quantification of simple spike responses to all recorded cells. Significance of differences are indicated by the color and corresponding p-value map.

c. Cspk responses to mossy fiber stimulation. Left: PSTH of the population of recorded cells with Cspks binned at 250 ms. A single trace showing a Cspk after stimulation is shown above. Right: Quantification of Cspk probability in the stimulated and non-stimulated epochs for each cell.

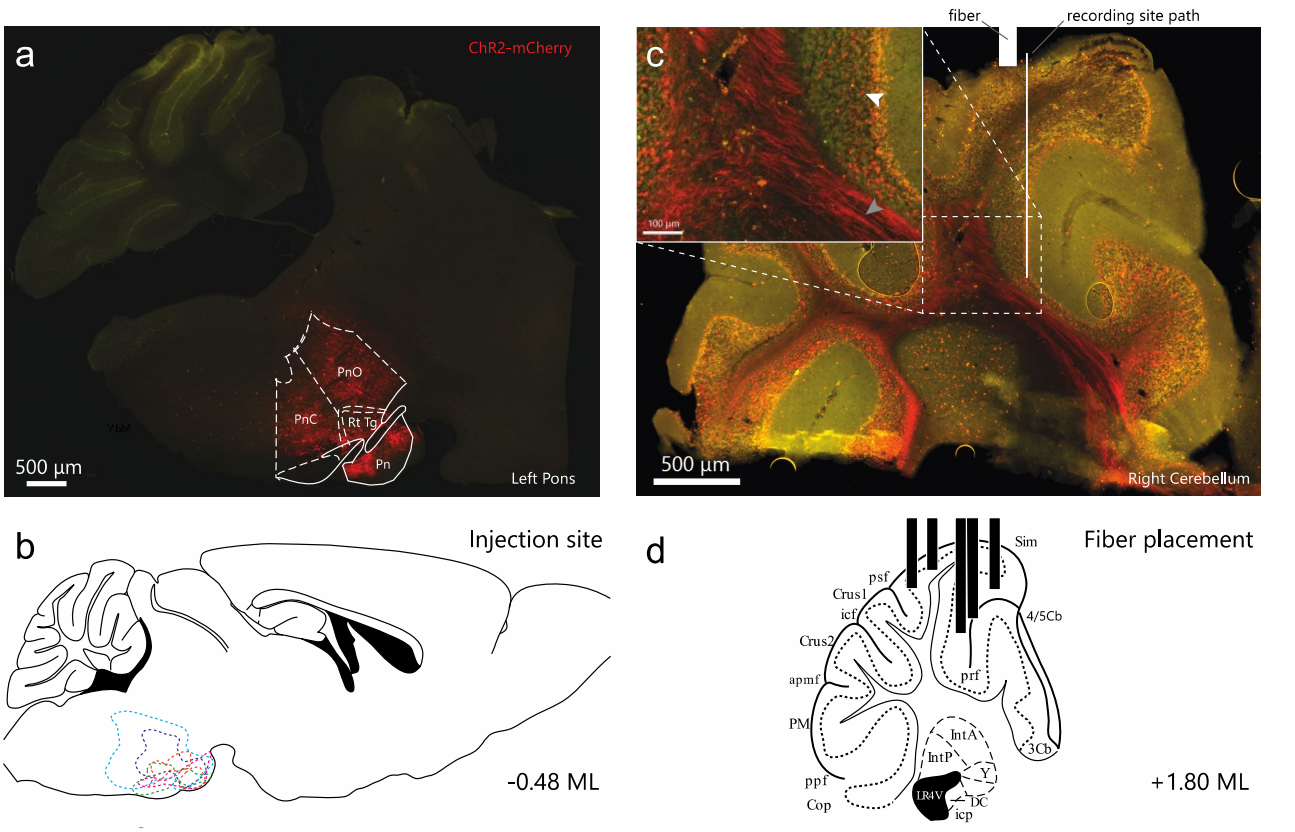

Figure S5. Opsin expression for mice in behavioral experiments.

a. Histological section showing ChR2-mCherry expression at the injection site in the left pontine nuclei (Pn: pontine nuclei; RtTg: reticulotegmental nuclei; PnO: pontine reticular nuclei, oral part; PnC: pontine reticular nuclei, caudal part).

b. Contours of ChR2 expression in the pontine nuclei for mice used in behavioral experiments.

c. Right cerebellum of the animal shown in **a.** Mossy fiber axons (grey arrow) and boutons (white arrow) can be seen expressing ChR2 in the cerebellar cortex. The approximate location of the optical fiber and recording site path are shown in white.

d. Location of fiber placement in a representative section for animals used in behavioral experiments.

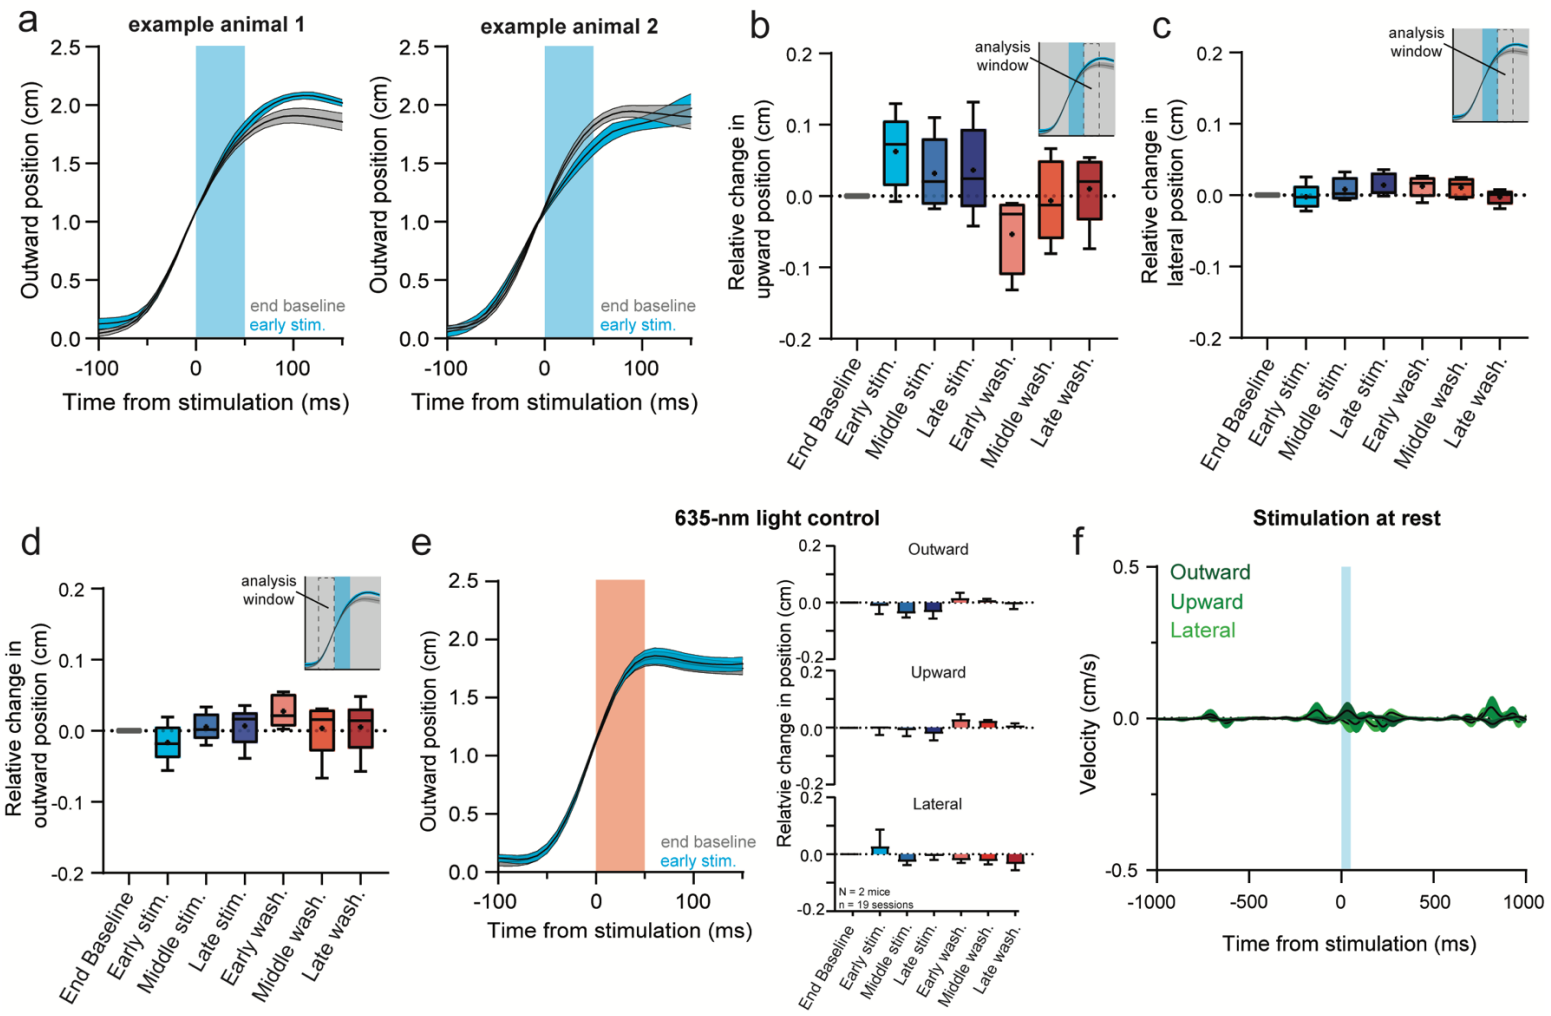

Figure S6. Fixed-position stimulation supplemental data

a. Two example mice with differing effects of stimulation on early reaches in the stimulation block. To account for diverging effects we define the direction of deviation with stimulation as positive and the opposing direction as negative.
b. Summary of the relative change in upward position for the same data shown in Fig. 3e. Relative change in upward position was assessed in the 50-ms window following the end of stimulation.
c. Summary of the relative change in lateral position for the same data shown in Fig. 3e.
d. Summary of the relative change in outward position for in the 50-ms window before stimulation.
e. Stimulating with 635-nm light did not cause deviations in position or adaptation profiles.
f. Stimulating while the mouse was awake with its hand at rest on the bar produced virtually no movement.

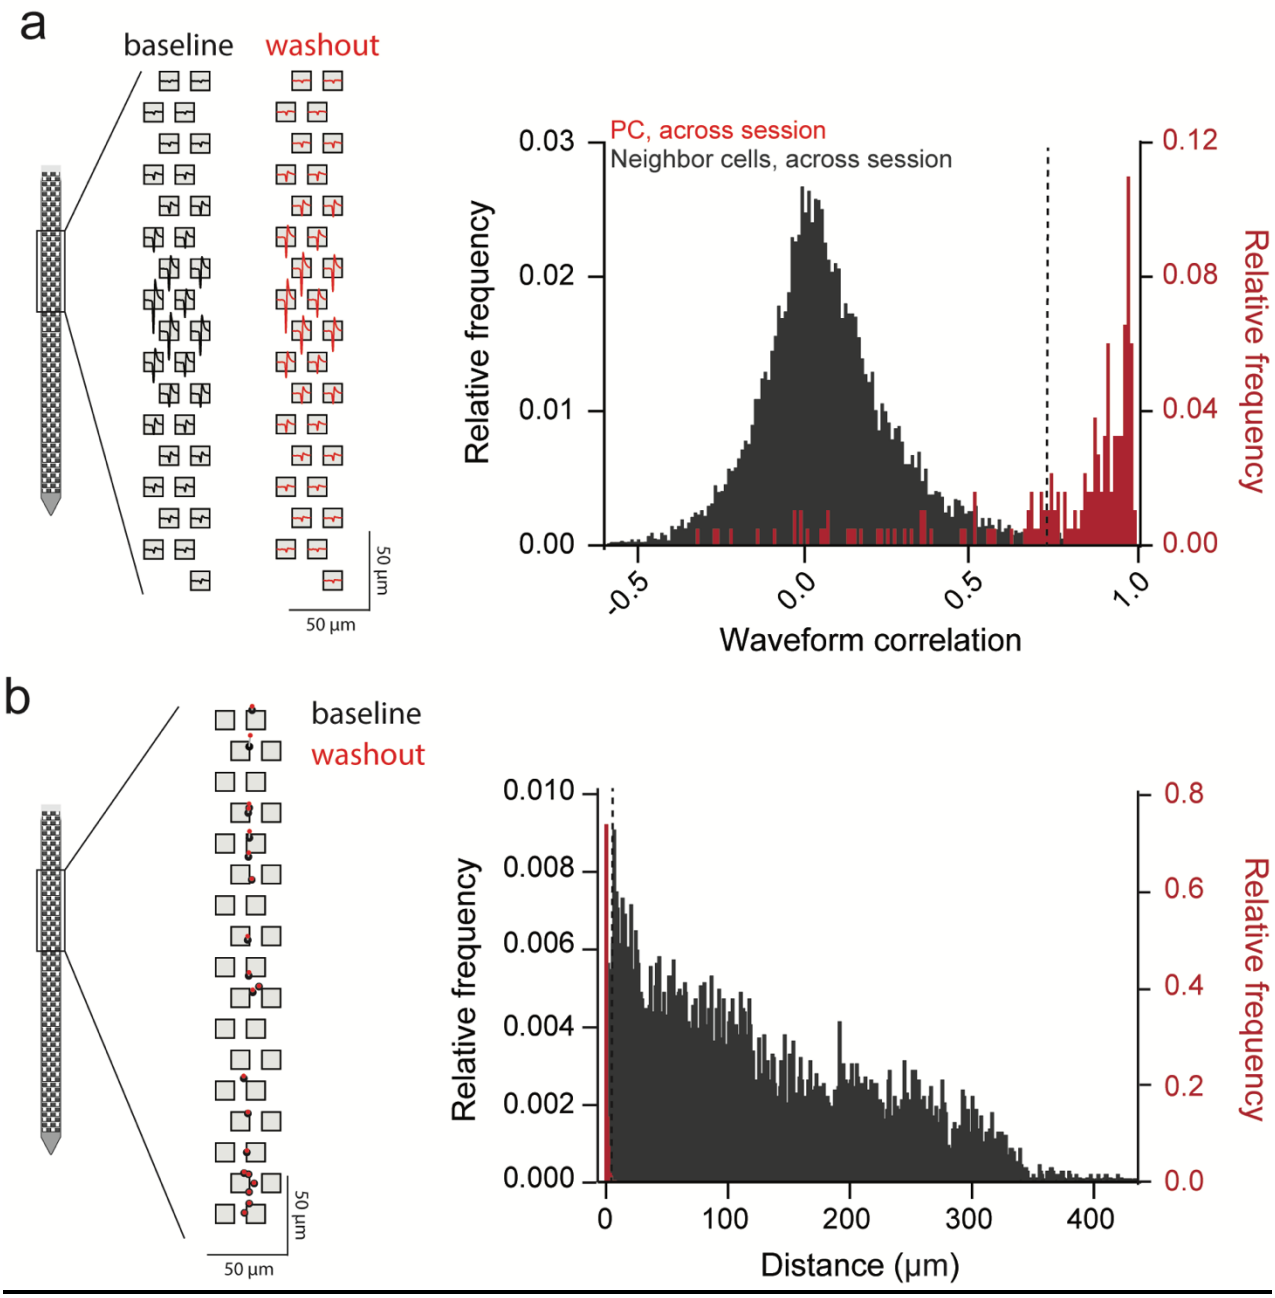

Figure S7. Assessing unit stability across recording sessions.

a. Left: Waveforms templates detected on each Neuropixel electrode for a cell during baseline and during washout. Right: Histogram of waveform correlation of PCs across sessions (red) and of mismatched neighboring cells, across the session (shuffled control, grey). PCs with an across-session waveform correlation that fell below the 99th percentile of the shuffled control (dashed line) were excluded from further analysis.

b. Left: Unit displacement for cells across a session. Baseline unit position is shown in grey and washout position is shown in red. Right: Histogram of unit displacement of PCs across sessions (red) and of mismatched neighboring cells, across the session (shuffled control, grey). PCs with an across-session displacement that fell below the 1st percentile of the shuffled control (dashed line) were excluded from further analysis.

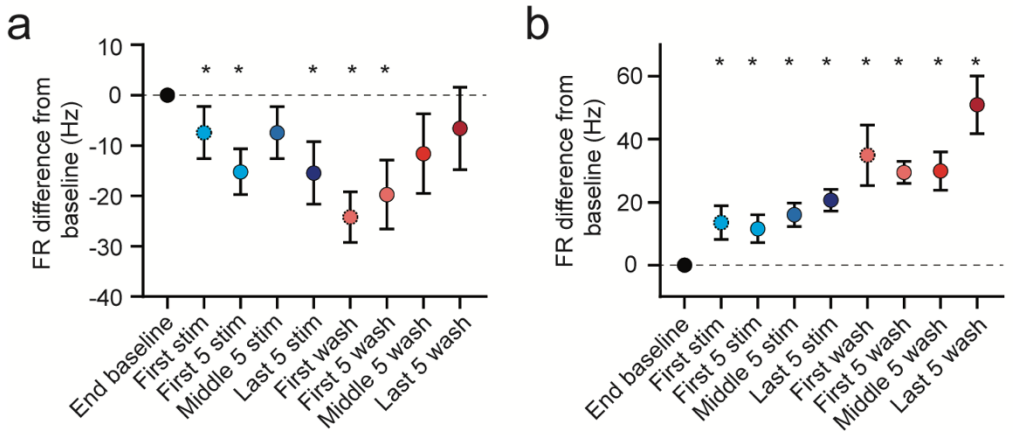

Figure S8. Firing rate profiles of aftereffect cells in stimulation and washout.

a. Cells with a decreased firing rate in washout compared to baseline showed a progressive decline in the stimulation block and recovery during washout (same as in Fig. 4e).

b. Cells with an increased firing rate through stimulation and washout.

References:

1. Holmes, G. The symptoms of acute cerebellar injuries due to gunshot injuries. *Brain* **40**, 461–535 (1917).
2. Becker, M. I. & Person, A. L. Cerebellar Control of Reach Kinematics for Endpoint Precision. *Neuron* **103**, 335-348.e5 (2019).
3. Lang, C. E. & Bastian, A. J. Cerebellar subjects show impaired adaptation of anticipatory EMG during catching. *J. Neurophysiol.* **82**, 2108–2119 (1999).
4. Baizer, J. S., Kralj-Hans, I. & Glickstein, M. Cerebellar lesions and prism adaptation in macaque monkeys. *J. Neurophysiol.* **81**, 1960–1965 (1999).
5. Tseng, Y. W., Diedrichsen, J., Krakauer, J. W., Shadmehr, R. & Bastian, A. J. Sensory prediction errors drive cerebellum-dependent adaptation of reaching. *J. Neurophysiol.* **98**, 54–62 (2007).
6. Morton, S. M. & Bastian, A. J. Cerebellar contributions to locomotor adaptations during splitbelt treadmill walking. *J. Neurosci.* **26**, 9107–9116 (2006).
7. Smith, M. A. & Shadmehr, R. Intact ability to learn internal models of arm dynamics in Huntington's disease but not cerebellar degeneration. *J. Neurophysiol.* **93**, 2809–2821 (2005).
8. Miall, R. & Wolpert, D. Forward Models for phyiological motor control. *Neural Networks* **9**, 1265–1279 (1996).
9. Kawato, M. & Gomi, H. The cerebellum and VOR / OKR learning models. *Trends Neurosci.* **15**, 445–453 (1992).
10. Ohyama, T., Nores, W. L., Murphy, M. & Mauk, M. D. What the cerebellum computes. *Trends Neurosci.* **26**, 222–227 (2003).
11. Shadmehr, R. & Krakauer, J. W. A computational neuroanatomy for motor control. *Exp. Brain Res.* **185**, 359–381 (2008).
12. Shadmehr, R., Smith, M. A. & Krakauer, J. W. Error correction, sensory prediction, and adaptation in motor control. *Annu. Rev. Neurosci.* **33**, 89–108 (2010).
13. Lavond, D. G., Knowlton, B. J., Steinmetz, J. E. & Thompson, R. F. Classical Conditioning of the Rabbit Eyelid Response With a Mossy-Fiber Stimulation CS: II. Lateral Reticular Nucleus Stimulation. *Behav. Neurosci.* **101**, 676–682 (1987).
14. Mauk, M. D., Steinmetz, J. E. & Thompson, R. F. Classical conditioning using stimulation of the inferior olive as the unconditioned stimulus. *Proc. Natl. Acad. Sci. U. S. A.* **83**, 5349–5353 (1986).
15. Coesmans, M., Weber, J. T., De Zeeuw, C. I. & Hansel, C. Bidirectional parallel fiber plasticity in the cerebellum under climbing fiber control. *Neuron* **44**, 691–700 (2004).
16. Ito, M. & Kano, M. Long-lasting depression of parallel fiber-Purkinje cell transmission induced by conjunctive stimulation of parallel fibers and climbing fibers in the cerebellar cortex. *Neurosci. Lett.* **33**, 253–258 (1982).
17. Lee, K. H. *et al.* Circuit mechanisms underlying motor memory formation in the cerebellum. *Neuron* **86**, 529–540 (2015).
18. Ekerot, C. F., Garwicz, M. & Schouenborg, J. Topography and nociceptive receptive fields of climbing fibres projecting to the cerebellar anterior lobe in the cat. *J. Physiol.* **441**, 257–274 (1991).
19. Bloedel, J. R. & Bracha, V. On the cerebellum, cutaneomuscular reflexes, movement control and the elusive engrams of memory. *Behav. Brain Res.* **68**, 1–44 (1995).
20. Ito, M. Cerebellar control of the vestibulo-ocular reflex--around the flocculus hypothesis. *Annu. Rev. Neurosci.* **5**, 275–297 (1982).
21. Blazquez, P. M., Hirata, Y., Heiney, S. A., Green, A. M. & Highstein, S. M. Cerebellar Signatures of Vestibulo-Ocular Reflex Motor Learning. *J. Neurosci.* **23**, 9742–9751 (2003).
22. Rowan, M. J. M. *et al.* Graded Control of Climbing-Fiber-Mediated Plasticity and Learning by

Inhibition in the Cerebellum. *Neuron* 1–17 (2018). doi:10.1016/j.neuron.2018.07.024

23. Yang, Y. & Lisberger, S. G. Interaction of plasticity and circuit organization during the acquisition of cerebellum-dependent motor learning. *Elife* **2013**, 1–19 (2013).
24. Herzfeld, D. J., Kojima, Y., Soetedjo, R. & Shadmehr, R. Encoding of error and learning to correct that error by the Purkinje cells of the cerebellum. *Nat. Neurosci.* **21**, (2018).
25. Raymond, J. L., Lisberger, S. G. & Mauk, M. D. The cerebellum: A neuronal learning machine? *Science (80-.).* **272**, 1126–1131 (1996).
26. Gandolfo, F., Li, C. S. R., Benda, B. J., Padoa Schioppa, C. & Bizzi, E. Cortical correlates of learning in monkeys adapting to a new dynamical environment. *Proc. Natl. Acad. Sci. U. S. A.* **97**, 2259–2263 (2000).
27. Li, C. S. R., Padoa-Schioppa, C. & Bizzi, E. Neuronal correlates of motor performance and motor learning in the primary motor cortex of monkeys adapting to an external force field. *Neuron* **30**, 593–607 (2001).
28. Mathis, M. W., Mathis, A. & Uchida, N. Somatosensory Cortex Plays an Essential Role in Forelimb Motor Adaptation in Mice. *Neuron* **93**, 1493-1503.e6 (2017).
29. Proville, R. D. *et al.* Cerebellum involvement in cortical sensorimotor circuits for the control of voluntary movements. *Nat. Neurosci.* **17**, 1233–1239 (2014).
30. Hewitt, A. L., Popa, L. S. & Ebner, T. J. Changes in Purkinje Cell Simple Spike Encoding of Reach Kinematics during Adaption to a Mechanical Perturbation. *J. Neurosci.* **35**, 1106–1124 (2015).
31. Low, A. Y. T. *et al.* Precision of Discrete and Rhythmic Forelimb Movements Requires a Distinct Neuronal Subpopulation in the Interposed Anterior Nucleus. *Cell Rep.* **22**, 2322–2333 (2018).
32. Guo, J.-Z. *et al.* Disrupting cortico-cerebellar communication impairs dexterity. *Elife* **10**, 1–31 (2021).
33. Becker, M. I., Calame, D. J., Wrobel, J. & Person, A. L. Online control of reach accuracy in mice. *J. Neurophysiol.* **124**, 1637–1655 (2020).
34. Tibshirani, R. Regression Shrinkage and Selection via the Lasso. *J. R. Stat. Soc.* **58**, 267–288 (1996).
35. Popa, L. S., Streng, M. L. & Ebner, T. J. Long-Term Predictive and Feedback Encoding of Motor Signals in the Simple Spike Discharge of Purkinje Cells. **4**, (2017).
36. Streng, M. L., Popa, L. S. & Ebner, T. J. Climbing Fibers Control Purkinje Cell Representations of Behavior. *J. Neurosci.* **37**, 1997–2009 (2017).
37. Coltz, J. D., Johnson, M. T. & Ebner, T. J. Cerebellar Purkinje cell simple spike discharge encodes movement velocity in primates during visuomotor arm tracking. *J. Neurosci.* **19**, 1782–803 (1999).
38. Hewitt, A. L., Popa, L. S., Pasalar, S., Hendrix, C. M. & Ebner, T. J. Representation of limb kinematics in Purkinje cell simple spike discharge is conserved across multiple tasks. *J. Neurophysiol.* **106**, 2232–2247 (2011).
39. Roitman, A. V., Pasalar, S., Johnson, M. T. V. & Ebner, T. J. Position, direction of movement, and speed tuning of cerebellar Purkinje cells during circular manual tracking in monkey. *J. Neurosci.* **25**, 9244–9257 (2005).
40. Musall, S., Kaufman, M. T., Juavinett, A. L., Gluf, S. & Churchland, A. K. Single-trial neural dynamics are dominated by richly varied movements. *Nat. Neurosci.* **22**, 1677–1686 (2019).
41. Herzfeld, D. J., Kojima, Y., Soetedjo, R. & Shadmehr, R. Encoding of action by the Purkinje cells of the cerebellum. *Nature* **526**, 439–441 (2015).
42. Brown, S. T. & Raman, I. M. Sensorimotor Integration and Amplification of Reflexive Whisking by Well-Timed Spiking in the Cerebellar Corticonuclear Circuit. *Neuron* 1–12 (2018). doi:10.1016/j.neuron.2018.06.028
43. Person, A. L. & Raman, I. M. Purkinje neuron synchrony elicits time-locked spiking in the cerebellar nuclei. *Nature* **481**, 502–505 (2012).
44. Thier, P., Dicke, P. W., Haas, R. & Barash, S. Encoding of movement time by populations of

cerebellar Purkinje cells. *Nature* **405**, 72–76 (2000).
45. Medina, J. F. & Lisberger, S. G. Links from complex spikes to local plasticity and motor learning in the cerebellum of awake-behaving monkeys. *Nat. Neurosci.* **11**, 1185–1192 (2008).
46. Kimpo, R. R., Rinaldi, J. M., Kim, C. K., Payne, H. L. & Raymond, J. L. Gating of neural error signals during motor learning. *Elife* **3**, 1–23 (2014).
47. Avraham, G., Taylor, J. A., Breska, A., Ivry, R. B. & McDougle, S. D. Contextual effects in motor adaptation adhere to associative learning rules. *bioRxiv* 2020.09.14.297143 (2021).
48. Wagner, M. J. *et al.* A neural circuit state change underlying skilled movements. *Cell* 3731–3747 (2021). doi:10.1016/j.cell.2021.06.001
49. Kitazawa, S., Kimura, T. & Yin, P. B. Cerebellar complex spikes encode both destinations and errors in arm movements. *Nature* **392**, 494–497 (1998).
50. Gaffield, M. A. & Christie, J. M. The cerebellum encodes and influences the initiation and termination of discontinuous movements. *bioRxiv* 1–31 (2021).
51. Miall, R. C., Keating, J. G., Malkmus, M. & Thach, W. T. Simple spike occurrence of complex spikes in cerebellar Punkinje cells. *Nat Neurosci* **1**, 13–15 (1998).
52. Medina, J. F., Nores, W. L. & Mauk, M. D. Inhibition of climbing fibres is a signal for the extinction of conditioned eyelid responses. *Nature* **416**, 330–333 (2002).
53. Kim, O. A., Ohmae, S. & Medina, J. F. sufficient to cause extinction of associative motor learning. *Nat. Neurosci.* **23**, 1550–1554 (2021).
54. Burroughs, A. *et al.* The dynamic relationship between cerebellar Purkinje cell simple spikes and the spikelet number of complex spikes. *J. Physiol.* **595**, 283–299 (2017).
55. Brodal, A. & Jansen, J. The ponto-cerebellar projection in the rabbit and cat. *J. Comp. Neurol.* **84**, 31–118 (1946).
56. Brodal, P. & Bjaalie, J. G. Organization of the pontine nuclei. *Neurosci. Res.* **13**, 83–118 (1992).
57. Leergaard, T. B. *et al.* Three-dimensional topography of corticopontine projections from rat sensorimotor cortex: Comparisons with corticostriatal projections reveal diverse integrative organization. *J. Comp. Neurol.* **478**, 306–322 (2004).
58. Sillitoe, R. V., Fu, Y. H. & Watson, C. *Cerebellum*. *The Mouse Nervous System* (Elsevier Inc., 2012). doi:10.1016/B978-0-12-369497-3.10011-1
59. Huang, C. C. *et al.* Convergence of pontine and proprioceptive streams onto multimodal cerebellar granule cells. *Elife* **2013**, 1–17 (2013).
60. Kim, J. & Augustine, G. J. Molecular Layer Interneurons: Key Elements of Cerebellar Network Computation and Behavior. *Neuroscience* **462**, 22–35 (2021).
61. Rasmussen, A., Jirenhed, D. A., Wetmore, D. Z. & Hesslow, G. Changes in complex spike activity during classical conditioning. *Front. Neural Circuits* **8**, 1–13 (2014).
62. Herzfeld, D. J., Hall, N. J., Tringides, M. & Lisberger, S. G. Principles of operation of a cerebellar learning circuit. *Elife* **9**, 1–30 (2020).
63. Mauk, M. D. & Buonomano, D. V. The neural basis of temporal processing. *Annu. Rev. Neurosci.* **27**, 307–340 (2004).
64. Shimansky, Y., Wang, J. J., Bauer, R. A., Bracha, V. & Bloedel, J. R. On-line compensation for perturbations of a reaching movement is cerebellar dependent: Support for the task dependency hypothesis. *Exp. Brain Res.* **155**, 156–172 (2004).
65. Albergaria, C., Silva, N. T., Pritchett, D. L. & Carey, M. R. Locomotor activity modulates associative learning in mouse cerebellum. *Nat. Neurosci.* **21**, 725–735 (2018).
66. Donchin, O., Francis, J. T. & Shadmehr, R. Quantifying generalization from trial-by-trial behavior of adaptive systems that learn with basis functions: Theory and experiments in human motor control. *J. Neurosci.* **23**, 9032–9045 (2003).
67. Nashef, A. *et al.* Cerebellar Shaping of Motor Cortical Firing Is Correlated with Timing of Motor Actions Article Cerebellar Shaping of Motor Cortical Firing Is Correlated with Timing of Motor Actions. *CellReports* **23**, 1275–1285 (2018).
68. Shinoda, Y., Kakei, S., Futami, T. & Wannier, T. Thalamocortical organization in the cerebello-

thalamo-cortical system. *Cereb. Cortex* **3**, 421–429 (1993).
69. Shinoda, Y., Futami, T. & Kano, M. Input-output organization of the ventrolateral nucleus of the thalamus. *Stereotact Funct Neurosurg* **60**, 17–31 (1993).
70. Shinoda, Y., Futami, T. & Kano, M. Synaptic organization of the cerebello-thalamo-cerebral pathway in the cat. II. Input-output organization of single thalamocortical neurons in the ventrolateral thalamus. *Neurosci. Res.* **2**, 157–180 (1985).
71. Hadjiosif, A. M., Krakauer, J. W. & Haith, A. M. Did we get sensorimotor adaptation wrong? implicit adaptation as direct policy updating rather than forward-model-based learning. *J. Neurosci.* **41**, 2747–2761 (2021).
72. Lee, K. H. *et al.* Circuit mechanisms underlying motor memory formation in the cerebellum. *Neuron* **86**, 529–540 (2015).
73. White, J. J. *et al.* An optimized surgical approach for obtaining stable extracellular single-unit recordings from the cerebellum of head-fixed behaving mice. *J. Neurosci. Methods* **262**, 21–31 (2016).
74. TestArduino. (2015). Available at: https://github.com/cortex-lab/phy.
75. Yu, B., Gabriel, D., Noble, L. & An, K. N. Estimate of the optimum cut- off frequency for the Butterworth low-pass digital filter. *J. Appl. Biomech.* **15**, 318–329 (1999).
76. Sedaghat-Nejad, E. *et al.* P-sort: An open-source software for cerebellar neurophysiology. *J. Neurophysiol.* **126**, 1055–1075 (2021).
77. Pachitariu, M., Steinmetz, N., Kadir, S., Carandini, M. & Kenneth D., H. Kilosort: realtime spike-sorting for extracellular electrophysiology with hundreds of channels. *bioRxiv* 061481 (2016). doi:10.1101/061481
78. Hensbroek, R. A. *et al.* Identifying Purkinje cells using only their spontaneous simple spike activity. *J. Neurosci. Methods* **232**, 173–180 (2014).
79. Schoonover, C. E., Ohashi, S. N., Axel, R. & Fink, A. J. P. Representational drift in primary olfactory cortex. *Nature* **594**, 541–546 (2021).
80. Kennedy, A. *et al.* A temporal basis for predicting the sensory consequences of motor commands in an electric fish. *Nat. Neurosci.* **17**, 416–422 (2014).
81. Lanore, F., Cayco-Gajic, N. A., Gurnani, H., Coyle, D. & Silver, R. A. Cerebellar granule cell axons support high-dimensional representations. *Nat. Neurosci.* **24**, 1142–1150 (2021).
82. Wagner, M. J. *et al.* Shared cortex-cerebellum dynamics in the execution and learning of a motor task. *Cell* **177**, 669–682 (2019).