Experience Adaptively Tunes the Timing Rules for Associative Plasticity

Sriram Jayabal[1*], Brandon J. Bhasin[1,2,3], Aparna Suvrathan[1,4], Jennifer DiSanto[1],

Mark S. Goldman[3*], Jennifer L. Raymond[1*]

[1]Dept. of Neurobiology, Stanford University School of Medicine, Stanford, CA, U.S.A.

[2]Dept. of Bioengineering, Stanford University, Stanford, CA, U.S.A.

[3]Center for Neuroscience, and Depts. of Neurobiology, Physiology, and Behavior, and Ophthalmology and Vision Science, University of California, Davis, CA, U.S.A.

[4]Dept. of Neurology and Neurosurgery, McGill University, Montreal, QC, Canada.

*Correspondence should be addressed to sriram.jayabal@mail.mcgill.ca or msgoldman@ucdavis.edu or jenr@stanford.edu.

ABSTRACT

A key function of the brain is to learn about the statistical relationships between events in the world. A mechanism of this learning is associative neural plasticity, controlled by the timing between neural events. Here, we show that experience can dramatically alter the timing rules governing associative plasticity to match the constraints of a particular circuit and behavior, thereby improving learning. In normal mice, the timing requirements for associative plasticity in the oculomotor cerebellum are precisely matched to the 120 ms delay for visual feedback about behavioral errors.[1] This task-specific specialization of the timing rules for plasticity is acquired through experience; in dark-reared mice that had never experienced visual feedback about oculomotor errors, plasticity defaulted to a coincidence-based rule. Computational modeling suggests two broad strategies for implementing this Adaptive Tuning of the Timing Rules for Associative Plasticity (ATTRAP), which tune plasticity to different features of the statistics of neural activity. The modeling predicts a critical role of this process in optimizing the accuracy of temporal credit assignment during learning; consistent with this, behavioral experiments revealed

a delay in the timing of learned eye movements in mice lacking experience-dependent tuning of the timing rules for plasticity. ATTRAP provides a powerful mechanism for matching the timing contingencies for associative plasticity to the functional requirements of a particular circuit and learning task, thereby providing a candidate neural mechanism for meta-learning.

MAIN TEXT

The timing between neural events controls the induction of plasticity at many synapses, providing a mechanism for learning about the correlations between events in the world.[2–4] There is considerable variation across synapses in the interval between neural events that is most effective for inducing plasticity, ranging from a few tens of milliseconds to as long as ~1 s.[1–13] In a few cases, a precise correspondence has been found between the timing contingencies for plasticity at certain synapses and specific biophysical, circuit-level or behavioral constraints on learning, suggesting that variations in the timing rules for plasticity are not random, but rather reflect functional specializations.[1,6,8,12,13] How does a given synapse "know" what timing rules are needed to support learning in the circuit in which it is embedded? Did the rules for plasticity at different synapses become specialized through evolution, to optimize learning in each neural circuit? Or are the timing rules governing synaptic plasticity themselves plastic, and adaptively tuned through experience to meet the requirements of a specific circuit and behavior?

To evaluate this, we leveraged one of the most striking specializations of the timing rules for plasticity identified to date, at the cerebellar parallel fiber-Purkinje cell synapses.[1] During learning, parallel fiber-Purkinje cell synapses that cause errors undergo associative synaptic depression, triggered by feedback about errors signaled by the single climbing fiber input to each Purkinje cell.[14] The feedback delay for climbing fibers to signal errors varies from a few tens of milliseconds for climbing fibers carrying somatosensory feedback[15] to greater than 100 ms for

visual feedback,[1,16] and likely longer for climbing fibers signaling cognitive errors.[17] To selectively weaken parallel fiber synapses that were active at the time that an error was generated, the timing contingencies for plasticity must compensate for the relevant (short or long) feedback delay for climbing fibers to signal errors. Accordingly, the optimal pairing interval between parallel fiber and climbing fiber activity for inducing associative synaptic depression varies dramatically for parallel fiber synapses onto different Purkinje cells.[1] In the cerebellar flocculus, which supports oculomotor learning, synaptic depression is very selectively induced by pairing parallel fiber activation with climbing fiber activation delayed by 120 ms,[1] which is the feedback delay between the parallel fiber activity generating an eye movement and the visual climbing fiber activity signaling a resulting oculomotor error *in vivo*.[1,16,18,19] Consistent with these results, we found that long-term depression (LTD) of parallel fiber-elicited spiking in Purkinje cells was selectively induced by this same, 120 ms pairing interval, and not by coincident pairing of parallel fiber and climbing fiber activation (**Fig. 1a-d; Extended Data Fig. 1**).

Adaptive tuning of the timing rules for plasticity

We tested the role of experience in generating this precise specialization of the timing rules for associative plasticity that compensates for the specific feedback delay in the circuit. To eliminate experience of the normal, 120 ms delay for visual climbing fiber feedback about oculomotor errors, mice were dark-reared from birth. We found that the timing requirement for associative LTD in the cerebellar flocculus was altogether different in dark-reared compared with normally reared mice. In dark-reared mice, the 120 ms parallel fiber-climbing fiber pairing interval was not at all effective at inducing LTD (**Fig. 1c,d**), indicating that the ability of delayed climbing fiber feedback to induce LTD is only acquired through the experience of that feedback delay. In

the absence of normal experience of visual feedback, LTD seems to default to a coincidence-based rule; in the dark-reared mice, coincident activation of parallel fibers and climbing fibers induced LTD (**Fig. 1b,d**), whereas in normally reared mice coincident pairing was totally ineffective at inducing LTD (**Fig. 1b,d**). Thus, the specialized timing rules for plasticity are not pre-determined through evolution; rather, experience completely rewrites the timing rules for plasticity, abolishing the ability of coincident neural activity to drive associative plasticity, and instead creating a new ability of delayed climbing fiber activation to induce plasticity. We refer to this novel capacity of the nervous system as Adaptive Tuning of the Timing Rules for Associative Plasticity (ATTRAP).

Modeling the role of ATTRAP in oculomotor learning

To evaluate the functional implications of ATTRAP, we built a simple computational model of oculomotor learning. The eye velocity response to an optokinetic visual stimulus was modeled as the sum of a fixed, non-cerebellar component and a plastic contribution from a population of Purkinje cells.[14] Each Purkinje cell received input from parallel fibers encoding the optokinetic stimulus and a climbing fiber that was driven, at a 120 ms delay, by oculomotor error, defined as the difference between the velocity of the eye and the visual stimulus (**Fig. 2a**; see **Methods**). We compared the effectiveness of learning, as measured by reduction of oculomotor errors, when the timing rules for plasticity were tuned or not tuned to the feedback delay in the circuit, as observed in the normally reared and dark-reared mice, respectively (**Fig. 2b**).

The untuned, coincidence-based plasticity rule caused climbing fiber feedback to depress not the synapses that caused the oculomotor error and associated image motion on the retina (retinal slip), but the synapses active at the later time when the climbing fibers reported the error to the flocculus. In contrast, the tuned plasticity rule enabled spikes in the climbing fibers to selectively

modify synapses that were active at the earlier time when the error was generated, thus solving

what is widely referred to as the temporal credit assignment problem.[20,21] Accordingly, the model

predicted that the timing of learned eye movements in dark-reared mice should be delayed

relative to normally reared mice (**Fig. 2c,d**). In addition, the model predicted that lack of tuned

timing rules for plasticity should prevent the dark-reared mice from achieving the same

asymptotic level of performance as normally reared mice, i.e., less effective reduction of

behavioral errors over the course of learning (**Extended Data Fig. 2a,b**).

Behavioral correlate of ATTRAP

We tested the model predictions experimentally and found a close correspondence between the

tuning of the timing requirements for associative LTD and the timing of learned behavior, as

predicted by the model. We tested optokinetic reflex (OKR) adaption, a form of oculomotor

learning supported by the cerebellar flocculus.[22,23] With one hour of training, normally reared

mice learned to increase the velocity of their eye movements to better stabilize an optokinetic

visual stimulus on the retina (**Fig. 2e,f**). Dark-reared mice were also capable of oculomotor

learning, even though the OKR adaptation training session was their first visual experience (**Fig.
2e,f**). The baseline peak amplitude of the OKR in the dark-reared mice was similar to that of

normally reared mice (**Fig. 2f**), and it increased with OKR training (**Fig. 2e,f**). To isolate the

recently learned, flocculus-dependent component of the OKR,[23] we subtracted the pre-training

eye movement response to the optokinetic stimulus from the post-training eye movement

response at the end of the one-hour training session (**Fig. 2g**). This revealed a striking difference

in the timing of the learned eye movements in the dark-reared mice compared to the normally

reared mice, which was evident in both the population average and in the timing of the learned

eye movements in individual mice (**Fig. 2h**). In dark-reared mice, the learned eye movements

were delayed by nearly 100 ms compared to normally reared mice, a time shift comparable to the altered timing requirements for associative plasticity measured at the neural level (**Fig. 1**), as predicted by the model (**Fig. 2c,d**).

Theoretical framework for implementation of ATTRAP

The experimental results suggest that experience, and the temporal patterns of experience-driven neural activity in a circuit, can adjust the timing requirements for plasticity much in the same way that the temporal patterns of activity in the circuit can adjust synaptic weights (**Fig. 3a**). Hence, ATTRAP can be thought of as a kind of meta-Spike Timing Dependent Plasticity (meta-STDP).

There are at least two general classes of biologically plausible mechanisms that could implement ATTRAP (**Fig. 3b**). One, which we call the timer adjustment mechanism, is a single, adjustable molecular timer that controls the eligibility of a synapse for associative plasticity. In this mechanism, a parallel fiber input starts the timer, such as the accumulation of a specific molecule, which, upon reaching a threshold, makes the synapse eligible to undergo plasticity if a climbing fiber spike arrives.[24,25] The timing requirements for plasticity could then be adjusted by changing the rate of accumulation (the speed of the clock) or by changing the threshold for eligibility (**Fig. 3b,c**; **Extended Data Fig. 3a,b**). A second potential mechanism for ATTRAP, which we call timer selection, is a bank of fixed molecular timers that are adaptively coupled to the eligibility for plasticity (**Fig. 3b,d**). One can conceive of this as an array of accumulators with different speeds or as a cascade of sequentially active components that forms a temporal basis set for constructing the eligibility window (**Extended Data Fig. 3c,d**).

For each mechanism, we constructed biochemical network models (**Extended Data Figs. 3,4; Methods**) for adjusting the timing rules for plasticity based on the statistics of neural activity,

and, in particular, the correlation structure between parallel fiber and climbing fiber activity. For the timer adjustment mechanism, we implemented an ATTRAP rule that reduced the threshold for eligibility, and thus the time of peak eligibility, if the climbing fiber input arrived before the timer reached the eligibility threshold, and increased the eligibility threshold if the climbing fiber input arrived after the timer reached threshold. This rule tuned the timing requirements for synaptic plasticity to near the median of the distribution of timing intervals between parallel fiber and climbing fiber spikes (**Fig. 3e**; **Extended Data Fig. 3e,f**; see **Supplementary Methods** for mathematical derivation).

For the timer selection mechanism, the coupling of each timer to the eligibility for plasticity was modifiable. We implemented an ATTRAP rule that strengthened the coupling of the timer that was most active at the time of the climbing fiber input and weakened the coupling of the other timers to the eligibility for plasticity (**Fig. 3b,d**; **Extended Data Fig. 3c**; see **Methods**). If the couplings were modified by fixed amounts in response to each climbing fiber spike, the ATTRAP rule tuned the timing requirement for synaptic plasticity to the mode of the distribution of intervals between parallel fiber and climbing fiber spikes (**Fig. 3f; Extended Data Fig. 3g**). This could be desirable in circuits supporting tasks like oculomotor learning, in which the neural inputs relevant for learning arrive at a characteristic delay relative to each other, and inputs at other timing delays are noise. Alternatively, if the couplings were modified by amounts proportional to their values at the time of each climbing fiber spike, the ATTRAP rule instead sculpted the eligibility time course to approximate the shape of the distribution of intervals between parallel fiber and climbing fiber spikes (**Fig. 3f**; **Extended Data Fig. 3g**; see **Supplementary Methods** for mathematical derivation). Here, learning would be driven according to the expected probability of climbing fiber inputs at each delay, which could be

advantageous in cases where a parallel fiber input leads to errors (signaled by climbing fiber input) that probabilistically arrive over a range of different time delays. The fixed and proportional modifications can be derived biochemically from zero- or first-order chemical kinetics (**Supplementary Methods; Extended Data Fig. 3c,4**); realistic networks might interpolate between these solutions. Thus, different implementations of ATTRAP have the potential to tune the timing rules for plasticity to match different features of the temporal correlation structure of their inputs, to support different learning objectives.

DISCUSSION

Together, these results demonstrate a novel, powerful adaptive capacity of the nervous system, Adaptive Tuning of the Timing Rules for Associative Plasticity (ATTRAP). Previously, experience-dependent changes in the *magnitude* of, or *threshold* for associative plasticity have been discovered at numerous synapses throughout the brain.[26–30] ATTRAP is qualitatively different, not only in the way it alters plasticity (i.e., altering the preferred pairing interval rather than the amount of plasticity), but also in its function. Whereas experience-dependent changes in the threshold for plasticity have been hypothesized to support synaptic homeostasis, synaptic tagging, and the general enhancement or impairment of learning by factors such as stress, arousal, an enriched environment or exercise,[27,31–33] ATTRAP provides a mechanism for meta-learning, or learning to learn. Through experience-dependent improvement in the accuracy of credit assignment and in the targeting of plasticity to the appropriate synapses during learning, ATTRAP provides a mechanism for increasing the efficacy of learning about events that are separated in time.

The major role of ATTRAP in the cerebellar flocculus is fairly surprising, because in this circuit the feedback delay for climbing fibers to signal oculomotor errors is highly stereotyped across individuals and across species;[1] hence, the optimal timing rule for plasticity could, in principle, have been specified through evolution. However, ATTRAP provides a more general and flexible mechanism for tuning the timing rules for plasticity, which could support meta-learning in circuits throughout the brain, including those where the relevant temporal correlations between events are not as predictable across individuals, over the lifetime of an animal (e.g., with myelination or body growth), or with changes in the environment. Notably, ATTRAP could also provide the flexibility to efficiently learn complex cognitive functions that rely on computations involving widely distributed brain circuits whose communication is subject to significant transmission and processing delays. Similar problems of uncertain propagation delays are also a challenge in engineered systems such as wireless communication networks and the internet of things,[34,35] and might benefit from ATTRAP-inspired solutions. By adaptively defining the temporal correlations between events that can be learned, ATTRAP provides a candidate neural mechanism of meta-learning, a hallmark of intelligence, whose neurobiological underpinnings have been a mystery.

METHODS

Animals

All experimental procedures were approved by the Administrative Panel on Laboratory Animal Care at Stanford University. C57Bl/6J mice were purchased from The Jackson Laboratory (stock no. 000664). Mice used in experiments were all bred at Stanford, and housed from birth either on a reversed 12-hour light/12-hour dark cycle (normally reared mice) or on a 24-hour dark cycle (dark-reared mice) in a custom-made light-proof chamber. For dark-reared mice, husbandry was

performed using night vision goggles (Wanney Group, cat. no. WNNVE-B201A) with an additional filter on the infra-red emitter to minimize emission in the visible range. All animals had *ad libitum* access to food and water. Experiments on normally reared mice were conducted during the animals' dark cycle.

Acute slice electrophysiology

Mice (24-42 days old) were anaesthetized using isoflurane (1-3%) and decapitated; their brains were extracted immediately and sliced in ice-cold artificial cerebrospinal fluid (ACSF (concentration in mM): NaCl(119), KCl(2.5), NaH_2PO_4(1), $NaHCO_3$(26.2), $MgCl_2$(1.3), $CaCl_2$(2.5), D-Glucose(10)), equilibrated with Carbogen (95% O_2, 5% CO_2, Praxair) to maintain pH at 7.3 (osmolality ~323 mOsm). 250 µm-thick coronal floccular slices were cut using a VT 1200 Vibratome (Leica) and were incubated in ACSF at 37 °C for 30-45 min and then for an hour at room temperature before electrophysiology experiments were conducted.

Slices were bathed in ACSF containing 100 µM picrotoxin to block $GABA_A$ channels, and held at 32-34 °C. Recordings were obtained in extracellular configuration from Purkinje cells identified visually using an Olympus BX51WI upright microscope. The climbing fiber input to a Purkinje cell was stimulated in the granule cell layer just beneath the Purkinje cell soma (~20-50 µm from Purkinje cell layer), and its recruitment was confirmed by the presence of complex spike responses in the Purkinje cell. Parallel fibers were stimulated by placing a stimulating electrode in the upper third of the molecular layer. Long term plasticity of parallel fiber-elicited spiking in the Purkinje cell was induced by pairing parallel fiber and climbing fiber stimulation. Each pairing consisted of a single stimulus delivered to the parallel fibers and a single stimulus to the climbing fiber, each 150 µs in duration. The interval between parallel fiber and climbing fiber stimulation (PF-CF pairing interval) was either 0 ms (coincident stimulation) or 120 ms

(parallel fibers then climbing fiber). Pairings were repeated 300 times at 1 Hz. Before and after pairing, parallel fiber-elicited spiking in the Purkinje cell was tested with parallel fiber stimulation at 0.05 Hz. All stimulation was done using bipolar electrodes pulled from theta glass (Sutter Instruments). Recording electrodes (3-5 MΩ) were pulled from borosilicate glass tubing (Harvard Apparatus) and filled with ACSF. Signals were digitized at 20 kHz using a Multiclamp 700B amplifier and low-pass filtered at 10 kHz. All reagents were obtained from Sigma-Aldrich unless otherwise indicated.

Slice physiology data were acquired using PClamp10 and analyzed using IgorPro with NeuroMatic, with statistical tests performed using JMP statistical suite. Spikes were counted in the first 50 ms (**Fig. 1**), or the first 30 or 100 ms (**Extended Data Fig. 1**) after the parallel fiber stimulus and normalized to the spontaneous spike count measured in a window of 50 ms (or 30 or 100 ms for **Extended Data Fig. 1**), starting 100 ms before each stimulus. Plasticity of the spike count was measured as the ratio of the normalized spike count measured at each time post-pairing divided by the normalized pre-induction spike count. Experimenters were blind to rearing during experimentation and data analysis.

Surgery

Mice (8-12 weeks) were surgically prepared for behavioral experiments as previously described.[36] In brief, while under anesthesia with 1.5-2.5% isoflurane, a custom-made head post was attached to the top of the skull using dental acrylic cement. To enable magnetic eye tracking, two stacked neodymium magnets with a total size of 0.75 x 2 mm (diameter x height, grade N50, axially magnetized, SuperMagnetMan.com) were implanted beneath the conjunctiva on the temporal side of the left eye. An angular magnetic field sensor (HMC1512, Honeywell Inc., size: 4.8 x 5.8 mm) was soldered to an 8-pin connector and attached to the skull above the eye using

dental acrylic, in a plane parallel to horizontal (nasal-temporal) eye movements. Eye movements were measured by detecting changes in the magnetic field created by movements of the magnet implanted on the eye (for details see ref.[36]). Animals were allowed to recover after surgery for at least five days before experiments were performed. Data were collected in mice <24 weeks of age.

Behavioral experiments

Behavioral experiments were performed in a light-proof, sound-attenuated chamber. The head of the mouse was immobilized by attaching the head post to a custom-built restrainer. Visual stimuli were provided by an optokinetic drum that surrounded the animal and rotated about an earth-vertical axis. The drum was made of white translucent plastic with black and white vertical stripes, each of which subtended 7.5° of the visual field, illuminated by shutter-controlled fiber-optic lights. To elicit the optokinetic response, the optokinetic drum was sinusoidally rotated at a frequency of 1.0 Hz, with a peak velocity of ±10°/s. Training consisted of 60 repeats of 50-second long blocks of optokinetic stimulation, with 10 seconds of rest in darkness between each block, as previously described.[37] After each experiment, a calibration session was performed during which eye movements were simultaneously measured with the magnetic sensor and video-oculography, to determine the calibration factor for converting signals from the magnetic sensor (in V) to eye position (in degrees).[36]

Data were analyzed in MATLAB. Horizontal eye position was computed from the voltage measured with the surgically implanted magnetic sensor. Eye position signals were filtered using a fourth-order low pass (15 Hz) Butterworth filter. Eye velocity was calculated by differentiating the filtered eye position. Data segments containing a saccade or movement artifact (e.g., a blink, whisking), plus 50 ms before and after the artifact, were excluded from the analysis using a

threshold on the mean squared error of a sinusoidal fit to the eye velocity. The threshold was determined individually for each animal to optimize detection of artifacts with different levels of measurement noise. The experimenters were blind to the rearing condition of the mice while analyzing the data.

For each 50 s optokinetic reflex adaptation training block, data remaining after artifact removal were aligned on the visual stimulus and averaged, and the peak eye velocity and the time of peak eye velocity relative to the optokinetic stimulus were extracted from this average. Pre- and post-training values were obtained from the first minute and the average of the last three minutes of OKR training, respectively. Animals with a baseline OKR gain of less than 0.1 were excluded from further analysis (one normally reared and one dark-reared mouse).

Statistics

All statistics were performed using JMP statistical suite. A Shapiro-Wilk test was used to assess normality. Non-parametric statistical tests were used to analyze data that violated the normality assumption.

Cerebellar circuit model

We simulated oculomotor learning in a simple model of the underlying circuit consisting of parallel fiber inputs to a group of 50 similarly tuned Purkinje cells (**Fig. 2a**). The velocity of an optokinetic stimulus input $S(t)$ was a 1 Hz sinusoid. Previous modeling work has suggested that granule cell firing forms a temporal basis set that would support a wide range of learned, time-varying responses.[38–46] Consistent with this, cerebellar neurons have been shown to tile time during a trial of behavior,[47,48] as also described for neurons in other brain areas.[49] Therefore, we modeled parallel fiber responses to the optokinetic stimulus as a sequentially active bank of

1,000 rectified sinusoids with peak firing rate of 5 Hz (**Fig. 2a**, center). On each "trial", i.e., each

cycle of the sinusoidal stimulus, parallel fiber spikes were drawn from a Poisson process with

this rate. Each Purkinje cell received an independent instantiation of parallel fiber inputs, and its

firing rate $PC(t)$ was given by a weighted sum of the parallel fiber spike trains convolved with

an exponential synaptic filter,

$$PC(t) = \sum_{i=1}^{1000} w_i (PF_i(t) * e^{-t/\tau_s}),$$ (1.1)

where PF_i is the spiking activity of the i^{th} parallel fiber input, which arrives at a synapse with

weight w_i and synaptic time constant $\tau_s = 10$ ms. Weights were allowed to be negative,

representing a negative net contribution of the combination of direct, excitatory input from

parallel fibers to Purkinje cells and disynaptic inhibitory input via the molecular layer

interneurons.

We modeled the eye movement responses as resulting from a fixed non-cerebellar component

representing a baseline response with gain g_0 minus a cerebellar component proportional to the

summed response of the activities of 50 Purkinje cells, reflecting their inhibitory output:

$$E(t) = g_0 S(t) - c \cdot \sum_{j=1}^{50} PC_j(t),$$ (1.2)

where c was chosen to be 0.02 to roughly reflect the eye velocity sensitivity of floccular Purkinje

cells.[50,51] The error signal driving learning, the retinal slip (RS) of the optokinetic stimulus, was

then modeled as the difference between the stimulus and the actual eye movement output,

$$RS(t) = S(t) - E(t).$$ (1.3)

The retinal slip, delayed by 120 ms, was then used to generate the climbing fiber response to each Purkinje cell, drawn from a Poisson process with rate

$$r_{CF}(t) = F([\mathrm{RS}(t - 0.12)]_+)^3, \tag{1.4}$$

where $[\cdot]_+$ indicates positive rectification, and F is proportional to the cumulative distribution function of the retinal slip errors from the previous trial (stimulus cycle)[52] , normalized so that climbing fibers on average fired at 1 Hz.

Associative plasticity model

The trial-to-trial change in weight at each parallel fiber-Purkinje cell synapse due to climbing fiber-triggered associative LTD was computed as

$$w_i \leftarrow w_i + \Delta w_i \tag{1.5}$$

$$\Delta w_i = -\alpha \sum_k \sum_\ell e(t_{CF}^\ell - t_{\mathrm{PF},i}^k), \tag{1.6}$$

where k indexes the parallel fiber spikes and ℓ indexes the climbing fiber spikes that occurred in the trial. Each parallel fiber spike was assumed to causally create an eligibility window for plasticity $e(\tau)$ (**Fig. 2b**),

$$e(\tau > 0) = e^{-(\tau - \tau_0)^2 / 2\sigma^2},$$

where τ represents the time after the parallel fiber spike. All Purkinje cells received the same climbing fiber spikes on each trial.

The eligibility window defines the timing requirements for LTD. For the oculomotor learning simulations in **Fig. 2**, we simulated the effect of ATTRAP on these timing requirements as a change in τ_{peak}, the time of peak eligibility, from $\tau_{\mathrm{peak}} = 0$ ms for the untuned plasticity rule to

$\tau_{\text{peak}} = 120$ ms for the tuned plasticity rule. In both cases, we used a width σ of 7.5 ms, to approximate published experimental results.[1] The amplitude of the eligibility window, $\alpha = 0.005$, controlled the learning rate.

To generate **Fig. 2c** we ran 10 simulations with each eligibility (untuned, before ATTRAP and tuned, after ATTRAP), with 3,600 trials in each simulation, similar to the number of trials experienced by the animals over an hour of training. To see the asymptotic performance of learning (**Extended Data Fig. 2**) we continued the learning for an additional 3,600 trials. At the start of each simulation, the synaptic weights were all set to 0. To match the average gains measured in the dark-reared and normally reared animals before learning, we set the baseline gain g_0 to 0.41 and 0.35 in the untuned and tuned simulations respectively.

ATTRAP model

We propose two general mechanisms for implementing ATTRAP and its activity-dependent induction based on the temporal statistics of neural activity experienced: adjustment of a single timer, and selection from a bank of timers (**Fig. 3b**). We modeled synapses as receiving a large number of random parallel fiber-climbing fiber pairs, with pairing intervals randomly drawn from a specified distribution. To generate the results in **Fig. 3e,f** and **Extended Data Fig. 3e-g**, we used a bimodal distribution given by density function

$$p_{\text{CF}}(\tau) = \mathbf{P}(\tau_{\text{CF}} = \tau) \propto c_1 e^{-(\tau - \mu_1)^2/2\sigma_1^2} + c_2 e^{-(\tau - \mu_2)^2/2\sigma_2^2}, \tag{1.7}$$

where $\mu_1 = 50$ ms, $\mu_2 = 150$ ms, $\sigma_1 = 50$ ms and $\sigma_2 = 10$ ms, and $c_1 = 1$, $c_2 = 2$. The distribution was defined over intervals from 0 ms to 200 ms, corresponding to a maximal parallel fiber firing rate of 5 Hz. Additional distributions were used in **Extended Data Fig. 3g** and **Extended Data Fig. 4** as described below. Intervals were sampled by drawing uniform random

numbers on [0, 1] and evaluating the inverse of the cumulative distribution function at these values, rounded to a time step bin of 0.1 ms.

Timer adjustment mechanism

For the timer adjustment mechanism (**Fig. 3c**), each synapse contains a single timer, started by the parallel fiber, that determines the window of eligibility (**Fig. 3a**, $e(\tau)$) for the synapse to undergo climbing fiber-driven plasticity. The timer, and hence the timing requirement for plasticity, is adjusted according to the temporal statistics of parallel fiber and climbing fiber activity. If a climbing fiber spike occurs before eligibility has reached its peak, then the timer is too slow, and is adjusted so that eligibility peaks earlier. Conversely, if the climbing fiber spike occurs after the peak eligibility, then the timer is too fast and is adjusted so that eligibility peaks later.

We implemented this mechanism by constructing an accumulation-to-bound timer,[24,25] for which a biochemical reaction diagram is shown in **Extended Data Fig. 3a**. Briefly, parallel fiber input causes a molecular species A to be quickly and maximally activated, which is required for the production of a molecular species B* that builds up in the synapse. Once species B* reaches a threshold, species C* is produced, which controls the eligibility (i.e., the time-varying concentration of species C* defines the spike-timing dependent plasticity eligibility window). Species C* causes the production of molecular species D*, which resets the timer by causing species A* to return to an inactivated form, thus leading to inactivation of species B* and then species C* (**Extended Data Fig. 3b**). Each of the four species can exist in an active or inactive form, with total concentration conserved by all reactions. The dynamical evolution of the fractional concentrations of these species is given by

$$\frac{dA^*}{d\tau} = -k_A^- \frac{D^{*10}}{K_{DA}^{10} + D^{*10}} A^*$$

$$\frac{dB^*}{d\tau} = k_B^+ A^* (1 - B^*) - k_B^- B^* (1 - A^*)$$

$$\frac{dC^*}{d\tau} = k_C^+ \Theta(B^* > \theta)(1 - C^*) - k_C^- C^*$$

(1.8)

$$\frac{dD^*}{d\tau} = k_D^+ \frac{C^{*20}}{K_{CD}^{20} + C^{*20}} (1 - D^*) - k_D^- D^*,$$

where asterisks indicate the active form of each species, $\Theta(\cdot)$ is a step function that evaluates to 1

when its argument is true, and θ is a tunable threshold for activation of species C. Parameter

values are listed in **Extended Data Table 1**. The large Hill coefficients used here are in the

range of what could be produced by ultrasensitive signaling.[24,53,54]

In this accumulation-to-bound timer, eligibility is dependent on the time that species B* takes to

reach threshold, which can be controlled either by the rate of production of B*, k_B^+, or the value

of the threshold itself. Here, for simplicity, we chose to move the threshold, because linear

changes in threshold lead to approximately linear changes in the time of activation of the

eligibility molecule C, while changes in k_B^+ are inversely (and thus nonlinearly) related to

changes in the activation time of C.

In order to adjust the timing of the eligibility window based on the experienced distribution of

intervals between parallel fiber and climbing fiber spikes, we constructed a timer adjustment

window, which can be thought of as a meta-STDP window (**Fig. 3c**). Just as the eligibility

window defines how the synaptic weight should increase or decrease as a function of the time of

the climbing fiber spike, the timer adjustment window analogously defines how the threshold θ

should increase or decrease at that time. If species B* is present but below threshold (so C* is

not present) when a climbing fiber spike occurs, then the eligibility timing is too slow, and the

threshold for activation of C is decreased; similarly, if D* is present and A has already been shut off when the climbing fiber spike occurs, then the eligibility timing is too fast and the threshold is increased (**Extended Data Fig. 3a**). Mathematically, the timer adjustment window was defined as a function of the time after the parallel fiber spike by

$$\Delta\theta(\tau) = \frac{X^5}{K_{X\theta}{}^5 + X^5} - \frac{Y^5}{K_{Y\theta}{}^5 + Y^5} A, \tag{1.9}$$

where

$$X = \frac{(1-A^*)^{20}}{K_{DX}{}^{20} + (1-A^*)^{20}} D^*$$

$$Y = \Theta(C^* < C_{min}^*)B^*, \tag{1.10}$$

(see **Extended Data Table 1** for parameter values). Because of the approximately linear relationship between them, the change in threshold, $\Delta\theta(\tau)$ in equation (1.9), tightly corresponds to the shift in peak eligibility time $\Delta\tau_{peak}(\tau)$ in **Fig. 3c** (*Timer adjustment window*). We ran 10 simulations of 5,000 pairs of parallel fiber and climbing fiber spikes. For each pair, equations (1.8) were simulated using a fourth/fifth-order Runge-Kutta solver (solve_ivp in the SciPy ODE integration package), with initial concentration $A^*(0) = 1$, and all other initial concentrations set to 0 (**Extended Data Fig. 3b**). A pairing interval yielding a climbing fiber spike time τ_{CF} was then drawn from the distribution p_{CF}, equation (1.7), and used to calculate $\Delta\theta$, which was added to the current value of the threshold,

$$\theta \leftarrow \theta + \beta \cdot \Delta\theta(\tau_{CF}), \tag{1.11}$$

with learning rate $\beta = 10^{-3}$. The timing of the eligibility window approached an equilibrium level around which it fluctuated in steady-state (**Extended Data Fig. 3e**), so we generated the

eligibility window shown in **Fig. 3e** by averaging the value of the threshold across the last 4,000 iterations of all the simulations.

To better understand the behavior of the timer adjustment mechanism of ATTRAP, we also considered a simple idealized model in which each climbing fiber input shifted a fixed-shape eligibility window by a constant amount earlier or later, depending on whether it arrived too early or too late, respectively. For a stationary distribution of intervals between parallel fiber and climbing fiber spikes, we found that this idealized timer adjustment ATTRAP mechanism shifts the window of eligibility for plasticity so that its peak is at the median of distribution of intervals between parallel fiber and climbing fiber spikes (**Extended Data Fig. 3f**; see **Supplementary Methods** for derivation).

We note that accumulation-to-bound timers place a potential sparsity constraint on the rate of spikes in a given parallel fiber. If a synapse receives two parallel fiber spikes, with the second arriving before the accumulator has reached threshold, this might cause the accumulator to cross threshold prematurely. Alternatively, the synapse could ignore one of the two inputs. Little is known about the spike rates in parallel fibers *in vivo*; however, they have been hypothesized to fire at low (temporally sparse) rates.[55–58] It is also possible that ATTRAP could be selectively triggered by a burst of parallel fiber spikes that occurs relatively rarely.[58]

Timer selection mechanism

For the timer selection mechanism of ATTRAP (**Fig. 3d**), the eligibility window for plasticity is determined by a weighted sum of activations $\{X_m{}^*\}$ of a bank of M molecular timers, with eligibility coupling weights $\{u_m\}$,

$$e(\tau) = \sum_{m=1}^{M} u_m X_m^{*}(\tau).$$ (1.12)

The time course of eligibility and hence the timing requirements for plasticity are altered by changing the coupling weights $\{u_m\}$. The selection mechanism itself is quite general, and does not depend on the particular implementation of the individual timers, so long as the bank of timers are started by arrival of the parallel fiber input and generate peaks that approximately tile time (**Fig. 3d**, *Molecular timers*).

We consider two methods for updating the eligibility coupling weights $\{u_m\}$ following the arrival of each climbing fiber input, termed "fixed" and "proportional", as detailed below. For both methods, eligibility coupling weights are increased or decreased additively following the climbing fiber spike,

$$u_m \leftarrow u_m + \Delta u_m.$$ (1.13)

If the timer activations each have the same integrated area, both methods keep the total integrated area under the eligibility trace constant by increasing the eligibility coupling weight u_n of the timer that is most active at the time of the climbing fiber input and decreasing the eligibility coupling weights of the other timers.

For the "fixed" update method, we subtract a fixed amount $\delta > 0$ (up to thresholding to keep the coupling weights non-negative) from the eligibility coupling weights of all the timers except the one that is most activated at the time of the climbing fiber input,

$$\Delta u_{m \neq n} = \begin{cases} -\delta & \text{if } \Delta u_{m \neq n} > \delta \\ -\Delta u_{m \neq n} & \text{otherwise} \end{cases}.$$ (1.14)

The coupling of the most activated timer is then increased by an amount

$$\Delta u_n = \sum_{m \neq n} -\Delta u_m, \tag{1.15}$$

so that the sum total coupling change is 0. This could arise biochemically from competition for a limited pool of molecular resources for coupling the timers to the eligibility for plasticity: these limited resources unbind at the time of the climbing fiber spike from cofactors that couple out-of-sync timers to eligibility and are transferred to cofactors that couple the most in-sync timer(s) to eligibility (**Extended Data Fig. 3c**; **Extended Data Fig. 4a,b**; see **Supplementary Methods**). These fixed updates cause a sharpening of the tuning of plasticity for the most likely parallel fiber-climbing fiber pairing intervals, as out-of-sync timers have their coupling to eligibility reduced to zero.

For the "proportional" update method, we increase the eligibility coupling weight u_n of the timer that is most active at the time of the climbing fiber spike by an amount proportional to its difference from the maximal coupling weight (here, defined to be 1), while decreasing the coupling of the other timers by an amount proportional to their current coupling weight,

$$\begin{aligned} \Delta u_n &= \gamma(1 - u_n) \\ \Delta u_{m \neq n} &= -\gamma u_m, \end{aligned} \tag{1.16}$$

where $\gamma > 0$ is a learning rate. This update automatically preserves the sum of the eligibility coupling weights and assures that all coupling weights remain within the range [0, 1]. The proportional coupling update method can be biochemically implemented in a manner similar to the fixed coupling update method (**Extended Data Fig. 4c**; see **Supplementary Methods**). For ease of simulation, for all simulation results shown in **Fig. 3** and **Extended Data Fig. 3** we used the instantaneous updates described in equations (1.14), (1.15) and (1.16).

The timer selection mechanism of ATTRAP can also be thought of as updating the eligibility window according to a meta-STDP window (**Fig. 3d**, *Selection window*). However, in this case, because there is a whole bank of timers to update in each update step, the selection window corresponds to a 2-D matrix describing how the entire vector of eligibility coupling weights is updated when a climbing fiber arrives at any possible time delay τ_{CF} following a parallel fiber input. That is, the vector of coupling weight changes $\Delta \boldsymbol{u}$ occurring due to a climbing fiber input at time τ_{CF} after a parallel fiber input corresponds to a slice of the two-dimensional selection window, and matrix multiplying the 2-D matrix of all possible slices by the shape of the climbing fiber input gives the eligibility update (up to a possible rescaling of the eligibility updates for the "fixed" method, denoted by the function f in **Fig. 3d**, to account for eligibility coupling weights needing to be non-negative). Similarly, a slice of the two-dimensional window for a fixed value of τ gives a selection window for how the value of the eligibility window at time τ will change as a function of the climbing fiber input time τ_{CF} (**Extended Data Fig. 3h**).

For an idealized set of molecular timers with identical, narrowly peaked, symmetrical activations that tile time and a stationary distribution of intervals between parallel fiber and climbing fiber spikes, simulations suggest that the fixed eligibility coupling weight updates result in an eligibility window whose maximum value occurs near the mode of the distribution (up to the width of the timer activations), while the proportional updates result in an eligibility window that can be shown (see **Supplementary Methods**) to approximate the shape of the distribution (**Extended Data Fig. 3g**). For the fixed updates, in addition to the distribution defined in equation (1.7), we used distributions with $\mu_1 = 50$ ms, $\sigma_1 = 10$ ms, $\mu_2 = 150$ ms and $\sigma_2 = 5$ ms, but where we varied which of the two bimodal peaks contained the mode, with either $c_1 = 1.2$ and $c_2 = 1$, or $c_1 = 1$ and $c_2 = 1.2$ (**Extended Data Fig. 3g**, left, bottom two panels

respectively). The eligibility windows in **Extended Data Fig. 3g** are examples from simulations

of 10 million parallel fiber-climbing fiber pairs with $\delta = 10^{-5}$ for the fixed update method and 5

million parallel fiber-climbing fiber pairs with $\gamma = 10^{-5}$ for the proportional update method.

The eligibility windows in **Fig. 3f** were generated using a bank of timers derived from a

biochemical implementation of a molecular cascade. For both fixed and proportional coupling

weight updates, we ran 5 million parallel fiber-climbing fiber pairs with $\gamma = \delta = 10^{-4}$.

We modeled the bank of molecular timers as a cascade of sequentially activated molecules, each

of which is activated by the previous element in the chain, with the first element corresponding

to the parallel fiber input. A reaction diagram is shown in **Extended Data Fig. 3c**. Briefly, the

model contains three types of molecular species: species A, which is maximally activated by a

parallel fiber input and is continuously degraded; a set of species $\{X_m\}$ that form the temporal

basis for eligibility; and a set of species $\{D_m\}$ that are used to produce a delayed activation

between subsequent X_m's (**Extended Data Fig. 3d**). The total concentration of active and

inactive molecules was conserved by all reactions. The dynamical evolution of the fractional

concentrations of these species are mathematically described by

$$
\begin{aligned}
\frac{\mathrm{dA}}{\mathrm{d}\tau} &= -k_A^- \mathrm{A} \\
\frac{\mathrm{dX}_1^*}{\mathrm{d}\tau} &= k_{X,1}^+ \frac{\mathrm{A}^{20}}{K_A^{20} + \mathrm{A}^{20}}(1 - \mathrm{X}_1^*) - k_{X,1}^- \mathrm{X}_1^* \\
\frac{\mathrm{dD}_m^*}{\mathrm{d}\tau} &= k_{D,m}^+ (1 - \mathrm{D}_m^*)\mathrm{X}_{m-1}^* - k_{D,m}^- \mathrm{D}_m^* \\
\frac{\mathrm{dX}_{m+1}^*}{\mathrm{d}\tau} &= k_{X,m+1}^+ \frac{\mathrm{D}_m^{*20}}{K_m^{20} + \mathrm{D}_m^{*20}}(1 - \mathrm{X}_{m+1}^*) - k_{X,m+1}^- \mathrm{X}_{m+1}^*,
\end{aligned}
\tag{1.17}
$$

where asterisks indicate the active form of each species X_m and D_m. The cascade could be of

arbitrary length M, which we chose here to be 26, spanning >200 ms (24 shown in **Fig. 3d** for

visualization purposes). Parameter values are listed in **Extended Data Table 2**. Similarly to the accumulation-to-bound timer implemented above, equations (1.8), the Hill coefficients are comparable to those which could result from an ultrasensitive process. The dynamics were simulated in the same manner as the timer adjustment mechanism (see above).

A bank of timers constructed from such a cascade places a lesser sparsity constraint on parallel fiber input than the accumulation-to-bound timer described for the timer adjustment mechanism (**Fig. 3c**, **Extended Data Fig. 3a,b**). Whereas the time during which subsequent parallel fiber inputs cause interference is related to the time to threshold in the accumulation-to-bound timer, here that window is defined only by the time interval in which the first state X_1 is most active, effectively just tens of milliseconds in our implementation.

Alternatively, the timer selection method could use a bank of accumulation-to-bound timers with different, fixed timings. These might exist within a single synapse, or each of the ~100,000 parallel fiber-Purkinje cell synapses could contain a single timer with a different timing. Each of these particular implementations fall into the same general class of mechanism in which the eligibility trace is built up from a temporal basis set defined by a bank of timers and their eligibility coupling weights.

Simulation methods

All simulations were run in Python 3.7, using the SciPy ecosystem (NumPy, SciPy, Matplotlib, and Jupyter/IPython).[59–61]

REFERENCES

1. Suvrathan, A., Payne, H. L. & Raymond, J. L. Timing Rules for Synaptic Plasticity Matched to Behavioral Function. *Neuron* **92**, 959–967 (2016).

2. Roberts, P. D. & Bell, C. C. Spike timing dependent synaptic plasticity in biological systems. *Biological cybernetics* **87**, 392–403 (2002).

3. Hebb, D. O. *The organization of behavior: A neuropsychological theory*. (Psychology Press, 2005).

4. Markram, H., Gerstner, W. & Sjöström, P. J. A history of spike-timing-dependent plasticity. *Frontiers in synaptic neuroscience* **3**, 4 (2011).

5. Bell, C. C., Han, V. Z., Sugawara, Y. & Grant, K. Synaptic plasticity in a cerebellum-like structure depends on temporal order. *Nature* **387**, 278 (1997).

6. Dudman, J. T., Tsay, D. & Siegelbaum, S. A. A role for synaptic inputs at distal dendrites: instructive signals for hippocampal long-term plasticity. *Neuron* **56**, 866–879 (2007).

7. Shouval, H. Z., Wang, S. S.-H. & Wittenberg, G. M. Spike timing dependent plasticity: a consequence of more fundamental learning rules. *Frontiers in computational neuroscience* **4**, 19 (2010).

8. Cho, J.-H. *et al.* Coactivation of thalamic and cortical pathways induces input timing–dependent plasticity in amygdala. *Nature neuroscience* **15**, 113 (2012).

9. Feldman, D. E. The spike-timing dependence of plasticity. *Neuron* **75**, 556–571 (2012).

10. Basu, J. *et al.* A cortico-hippocampal learning rule shapes inhibitory microcircuit activity to enhance hippocampal information flow. *Neuron* **79**, 1208–1221 (2013).

11. Mehaffey, W. H. & Doupe, A. J. Naturalistic stimulation drives opposing heterosynaptic plasticity at two inputs to songbird cortex. *Nature neuroscience* **18**, 1272 (2015).

12. Bittner, K. C., Milstein, A. D., Grienberger, C., Romani, S. & Magee, J. C. Behavioral time scale synaptic plasticity underlies CA1 place fields. *Science* **357**, 1033–1036 (2017).

13. Leroy, F., Brann, D. H., Meira, T. & Siegelbaum, S. A. Input-timing-dependent plasticity in the hippocampal CA2 region and its potential role in social memory. *Neuron* **95**, 1089-1102. e5 (2017).

14. Ito, M., Yamaguchi, K., Nagao, S. & Yamazaki, T. Long-term depression as a model of cerebellar plasticity. in *Progress in brain research* vol. 210 1–30 (Elsevier, 2014).

15. Jörntell, H., Ekerot, C.-F., Garwicz, M. & Luo, X.-L. Functional organization of climbing fibre projection to the cerebellar anterior lobe of the rat. *The Journal of physiology* **522**, 297–309 (2000).

16. Raymond, J. L. & Lisberger, S. G. Neural learning rules for the vestibulo-ocular reflex. *Journal of Neuroscience* **18**, 9112–9129 (1998).

17. Schmahmann, J. D., Guell, X., Stoodley, C. J. & Halko, M. A. The theory and neuroscience of cerebellar cognition. *Annual review of neuroscience* **42**, (2019).

18. Maekawa, K. & Simpson, J. L. Climbing fiber responses evoked in vestibulocerebellum of rabbit from visual system. *Journal of Neurophysiology* **36**, 649–666 (1973).

19. Stone, L. S. & Lisberger, S. G. Detection of tracking errors by visual climbing fiber inputs to monkey cerebellar flocculus during pursuit eye movements. *Neuroscience letters* **72**, 163–168 (1986).

20. Dayan, P. & Abbott, L. F. *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*. (MIT Press, 2001).

21. Sutton, R. S. & Barto, A. G. *Reinforcement learning: An introduction*. (MIT press, 2018).

22. Katoh, A., Kitazawa, H., Itohara, S. & Nagao, S. Dynamic characteristics and adaptability of mouse vestibulo-ocular and optokinetic response eye movements and the role of the flocculo-olivary system revealed by chemical lesions. *Proceedings of the National Academy of Sciences* **95**, 7705–7710 (1998).

23. Shutoh, F., Ohki, M., Kitazawa, H., Itohara, S. & Nagao, S. Memory trace of motor learning shifts transsynaptically from cerebellar cortex to nuclei for consolidation. *Neuroscience* **139**, 767–777 (2006).

24. Goldbeter, A. A minimal cascade model for the mitotic oscillator involving cyclin and cdc2 kinase. *Proceedings of the National Academy of Sciences* **88**, 9107–9111 (1991).

25. Rensing, L., Meyer-Grahle, U. & Ruoff, P. Biological timing and the clock metaphor: oscillatory and hourglass mechanisms. *Chronobiol. Int.* **18**, 329–369 (2001).

26. Abraham, W. C. & Bear, M. F. Metaplasticity: the plasticity of synaptic plasticity. *Trends in Neurosciences* **19**, 126–130 (1996).

27. Abraham, W. C. Metaplasticity: tuning synapses and networks for plasticity. *Nature Reviews Neuroscience* **9**, 387 (2008).

28. Guo, Y. *et al.* Dark exposure extends the integration window for spike-timing-dependent plasticity. *Journal of Neuroscience* **32**, 15027–15035 (2012).

29. Kuśmierz, Ł., Isomura, T. & Toyoizumi, T. Learning with three factors: modulating Hebbian plasticity with errors. *Current opinion in neurobiology* **46**, 170–177 (2017).

30. Gerstner, W., Lehmann, M., Liakoni, V., Corneil, D. & Brea, J. Eligibility traces and plasticity on behavioral time scales: experimental support of neohebbian three-factor learning rules. *Frontiers in neural circuits* **12**, 53 (2018).

31. Kim, J. J., Foy, M. R. & Thompson, R. F. Behavioral stress modifies hippocampal plasticity through N-methyl-D-aspartate receptor activation. *Proceedings of the National Academy of Sciences* **93**, 4750–4753 (1996).

32. Van Praag, H., Christie, B. R., Sejnowski, T. J. & Gage, F. H. Running enhances neurogenesis, learning, and long-term potentiation in mice. *Proceedings of the National Academy of Sciences* **96**, 13427–13431 (1999).

33. Duffy, S. N., Craddock, K. J., Abel, T. & Nguyen, P. V. Environmental enrichment modifies the PKA-dependence of hippocampal LTP and improves hippocampus-dependent memory. *Learning & Memory* **8**, 26–34 (2001).

34. Mahmood, A., Exel, R., Trsek, H. & Sauter, T. Clock synchronization over IEEE 802.11—A survey of methodologies and protocols. *IEEE Transactions on Industrial Informatics* **13**, 907–922 (2016).

35. Mani, S. K., Durairajan, R., Barford, P. & Sommers, J. An architecture for IoT clock synchronization. in *Proceedings of the 8th International Conference on the Internet of Things* 1–8 (2018).

36. Payne, H. L. & Raymond, J. L. Magnetic eye tracking in mice. *eLife* **6**, e29222 (2017).

37. Inoshita, T. & Hirano, T. Occurrence of long-term depression in the cerebellar flocculus during adaptation of optokinetic response. *Elife* **7**, e36209 (2018).

38. Fujita, M. Adaptive filter model of the cerebellum. *Biol Cybern* **45**, 195–206 (1982).

39. Buonomano, D. V. & Mauk, M. D. Neural Network Model of the Cerebellum: Temporal Discrimination and the Timing of Motor Responses. *Neural Computation* **6**, 38–55 (1994).

40. Medina, J. F. & Mauk, M. D. Computer simulation of cerebellar information processing. *nature neuroscience* **3**, 1205 (2000).

41. Medina, J. F., Garcia, K. S., Nores, W. L., Taylor, N. M. & Mauk, M. D. Timing Mechanisms in the Cerebellum: Testing Predictions of a Large-Scale Computer Simulation. *The Journal of Neuroscience* **20**, 5516–5525 (2000).

42. Yamazaki, T. & Tanaka, S. Neural Modeling of an Internal Clock. *Neural Computation* **17**, 1032–1058 (2005).

43. Dean, P., Porrill, J., Ekerot, C.-F. & Jörntell, H. The cerebellar microcircuit as an adaptive filter: experimental and computational evidence. *Nature Reviews Neuroscience* **11**, 30–43 (2010).

44. Clopath, C., Badura, A., De Zeeuw, C. I. & Brunel, N. A Cerebellar Learning Model of Vestibulo-Ocular Reflex Adaptation in Wild-Type and Mutant Mice. *Journal of Neuroscience* **34**, 7203–7215 (2014).

45. Kennedy, A. *et al.* A temporal basis for predicting the sensory consequences of motor commands in an electric fish. *Nature neuroscience* **17**, 416 (2014).

46. Rössert, C., Dean, P. & Porrill, J. At the Edge of Chaos: How Cerebellar Granular Layer Network Dynamics Can Provide the Basis for Temporal Filters. *PLOS Computational Biology* **11**, e1004515 (2015).

47. Li, J. X., Medina, J. F., Frank, L. M. & Lisberger, S. G. Acquisition of Neural Learning in Cerebellum and Cerebral Cortex for Smooth Pursuit Eye Movements. *J. Neurosci.* **31**, 12716–12726 (2011).

48. Sauerbrei, B. A., Lubenov, E. V. & Siapas, A. G. Structured Variability in Purkinje Cell Activity during Locomotion. *Neuron* **87**, 840–852 (2015).

49. Paton, J. J. & Buonomano, D. V. The Neural Basis of Timing: Distributed Mechanisms for Diverse Functions. *Neuron* **98**, 687–705 (2018).

50. Katoh, A., Shin, S.-L., Kimpo, R. R., Rinaldi, J. M. & Raymond, J. L. Purkinje cell responses during visually and vestibularly driven smooth eye movements in mice. *Brain Behav* **5**, e00310–e00310 (2015).

51. Payne, H. L. *et al.* Cerebellar Purkinje cells control eye movements with a rapid rate code that is invariant to spike irregularity. *Elife* **8**, e37102 (2019).

52. Laughlin, S. A simple coding procedure enhances a neuron's information capacity. *Z. Naturforsch., C, Biosci.* **36**, 910–912 (1981).

53. Goldbeter, A. & Koshland, D. E. An amplified sensitivity arising from covalent modification in biological systems. *Proceedings of the National Academy of Sciences* **78**, 6840–6844 (1981).

54. Ferrell, J. E. & Ha, S. H. Ultrasensitivity part I: Michaelian responses and zero-order ultrasensitivity. *Trends in Biochemical Sciences* **39**, 496–503 (2014).

55. Marr, D. A theory of cerebellar function. *J Physiol* **202**, (1969).

56. Albus, J. S. A theory of cerebellar function. *Mathematical Biosciences* **10**, 25–61 (1971).

57. Sawtell, N. B. Multimodal integration in granule cells as a basis for associative plasticity and sensory prediction in a cerebellum-like circuit. *Neuron* **66**, 573–584 (2010).

58. Powell, K., Mathy, A., Duguid, I. & Häusser, M. Synaptic representation of locomotion in single cerebellar granule cells. *Elife* **4**, e07290 (2015).

59. Virtanen, P. *et al.* SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nature Methods* 1–12 (2020) doi:10.1038/s41592-019-0686-2.

60. Hunter, J. D. Matplotlib: A 2D graphics environment. *Computing in Science & Engineering* **9**, 90–95 (2007).

61. Pérez, F. & Granger, B. E. IPython: a System for Interactive Scientific Computing. *Computing in Science and Engineering* **9**, 21–29 (2007).

Acknowledgements:

We thank Maxwell Gagnon for help with codes for data acquisition and analysis, and Pragna Gaddam, James Dang, Macarena Martínez Rey, and Jaydev Bhateja for technical assistance. We also thank Benjamin Lankow, Alireza Alemi, and members of the Raymond and Goldman labs for helpful discussions. **Funding:** Supported by NIH R01 DC04154 (J.L.R.), R01 NS072406 (J.L.R.), R01 NS104926 (M.S.G.), the Simons Collaboration on the Global Brain 543031 (J.L.R.) and 542989 (M.S.G.), a Stanford School of Medicine Dean's Fellowship and a Wu Tsai Neurosciences Institute Interdisciplinary Scholar Award (S.J.), and a Stanford Center for Mind, Brain, Computation and Technology Traineeship and a Stanford Interdisciplinary Graduate Fellowship (B.J.B.).

Author contributions: S.J., A.S., and J.L.R. designed the experiments; S.J., and J.D. performed the experiments and analyzed the data, B.J.B., M.S.G., and J.L.R. designed the modeling study with input from S.J., B.J.B. and M.S.G. ran the simulations and performed mathematical analyses, and S.J., B.J.B., M.S.G., and J.L.R. made the figures and wrote the manuscript with input from all authors.

Competing interests: The authors declare no competing interests.

Data and materials availability: All the data used to support the conclusions of the paper are presented within the paper and its supplementary materials, and are available upon request.

Supplementary Information is available for this paper.

Fig. 1. Adaptive Tuning of the Timing Rules for Associative Plasticity (ATTRAP).

(**a**) Plasticity was induced by pairing stimulation of the parallel fiber (PF) and climbing fiber (CF) inputs to a Purkinje cell (300x at 1 Hz) in slices of the cerebellar flocculus of normally and dark-reared mice.

(**b**) Coincident PF-CF pairing. ***Top***, PF-elicited spiking in sample Purkinje cells from normally reared (***left***) and dark-reared mice (***right***). ***Bottom***, Time course of changes in PF-elicited spiking, in individual cells (*grey*) and the group averages from dark-reared (*dark green*) and normally reared (*light green*) mice.

(**c**) 120 ms PF-CF pairing. ***Top***, PF-elicited spiking in sample Purkinje cells. ***Bottom***, Time course of changes in PF-elicited spiking in normally reared mice (***left***, *light blue*) and dark-reared mice (***right***, *dark blue*).

(**d**) Summary of long-term depression (LTD) of spike count, measured 30 min after pairing. ***Left***, In normally reared mice, LTD of PF-elicited Purkinje cell spiking was induced by a PF-CF pairing interval of 120-ms (**$P = 0.001$, single-sample t-test, $n = 11$ cells) but not 0 ms ($P = 0.863$; single-sample t-test, $n = 11$ cells). ***Right***, In dark-reared mice, LTD of PF-elicited spiking was induced by a PF-CF pairing interval of 0 ms (*$P = 0.013$, single-sample t-test, $n = 9$ cells) but not 120 ms ($P = 0.226$, single-sample t-test, $n = 9$ cells), demonstrating the role of normal experience in the Adaptive Tuning of the Timing Rules for Associative Plasticity (ATTRAP) to match the feedback delay in the circuit. Error bars in all panels are S.E.M.

Fig. 2. Experience eliminates delay of learned eye movements.

(**a**) *Left*, During optokinetic reflex (OKR) adaptation, mice learn to more accurately track the motion of a striped visual stimulus (*grey arrow*) with their eye movements (*pink arrows*). **Right**, Model architecture used to simulate oculomotor learning. The model received input from a sinusoidal optokinetic stimulus, which it attempted to match with eye movements, modeled as the sum of Purkinje cell population activity and a fixed, non-cerebellar contribution. Each Purkinje cell was driven by 1,000 parallel fibers (PFs) with firing rates that tiled time during the stimulus half-cycle, with spikes drawn randomly on each trial. Oculomotor errors, defined as motion of the optokinetic stimulus relative to the eye (retinal slip) probabilistically drove climbing fiber (CF) spikes at a 120 ms delay, which in turn drove associative LTD of the weights, **w**, of parallel fiber-Purkinje cell synapses.

(**b**) The model was run using plasticity rules with different timing requirements. Dark-reared mice were simulated with an untuned plasticity rule, where peak eligibility for undergoing climbing fiber-triggered associative plasticity was coincident with a PF spike (*green trace*), so that a PF-CF pairing interval of 0 ms, but not 120 ms, induced LTD of the PF synaptic weight w (*green bar*). Normally reared mice were simulated with a plasticity rule tuned to the feedback delay in the circuit, where eligibility for plasticity peaked 120 ms after the PF spike (*blue trace*), so that a PF-CF pairing interval of 120 ms, but not 0 ms, induced LTD of the PF synaptic weight (*blue bar*).

(**c**) ***Top,*** Simulated eye movement responses in one tuned (*left*) and untuned (*right*) model run, before (Pre, *black)* and after (Post, *blue, green)* training with 3,600 sinusoidal stimulus cycles. ***Bottom,*** Mean simulated eye movement responses (± 1 standard deviation), for 10 tuned (*left*) and untuned (***right***) model simulations. The learned component of the eye movement response (*shaded grey regions*) was isolated by subtracting the simulated pre-training (*black traces*) eye movement

response to the optokinetic visual stimulus (*grey trace*) from the simulated post-training eye movement response (*blue* or *green trace*).

(**d**) ***Top***, Learned component of the simulated eye movement response to a stimulus (*grey*) after simulating 3,600 sinusoidal stimulus cycles (60 min) using the untuned (*green*) or tuned (*blue*) plasticity rule. Mean ± 1 standard deviation across 10 simulations. ***Bottom***, Timing of the learned component of the eye movement response (time of peak learned eye velocity relative to peak optokinetic stimulus velocity, calculated for each simulation run, *circles*) was delayed in models with untuned plasticity (*green*) relative to models with tuned plasticity (*blue*). *Solid vertical lines*: average timing across simulations; shaded region: ± 1 standard deviation.

(**e**) Raw traces showing horizontal eye velocity responses to a sinusoidally moving visual stimulus (i.e., the OKR) at the beginning (Pre) and end (Post) of the 60-min training session, in representative normally reared (***left***) and dark-reared (***right***) mice.

(**f**) One hour of OKR training induced a learned increase in the peak eye velocity response to the optokinetic visual stimulus in both normally reared (***left***; ******P*** < 0.0001, matched pairs *t*-test) and dark-reared (***right***; ***P* $= 0.025$, matched pairs *t*-test) mice (*dashed black lines*). *Grey*: individual mice. The amplitude of the OKR pre-training was similar in normally and dark-reared mice ($P = 0.3087$, independent samples *t*-test).

(**g**) ***Top***, Average eye movement responses (averaged across 10-45 stimulus cycles) in representative normally reared (***left***) and dark-reared (***right***) mice, pre- and post-training. ***Bottom***, The learned component of the eye movement response (*shaded grey regions*) was isolated by subtracting the pre-training (*black traces*) response to the optokinetic visual stimulus (*grey trace*) from the post-training eye movement response (*blue* or *green* trace). Traces show the mean ± S.E.M. across animals, for normally reared (***left***; $N = 14$) and dark-reared (***right***; $N = 10$) mice.

(**h**) ***Top***, Learned component of the eye movement response to an optokinetic visual stimulus (*grey*) in dark-reared mice (*green*) and normally reared mice (*blue*). Mean ± S.E.M. across animals. ***Bottom***, Timing of the learned eye movement (time of peak learned eye velocity relative to peak optokinetic stimulus velocity, calculated for each mouse, *circles*) was delayed in dark-reared mice (*green)* relative to normally reared mice (*blue*) (**P = 0.0002, independent samples *t*-test), and relative to the visual stimulus (P = 0.0007, single-sample *t*-test). *Solid vertical lines*: average peak learned timing across animals; *shaded region*: ± S.E.M.

Fig. 3. Theoretical framework for implementing Adaptive Tuning of the Timing Rules for Associative Plasticity (ATTRAP)

(**a**) Schematic of spike timing-dependent associative LTD at parallel fiber-Purkinje cell synapses. A parallel fiber (PF) input spike at time $\tau = 0$ starts a timer that creates a time-varying eligibility window $e(\tau)$ for plasticity. The change in weight at that PF synapse is proportional to the value of the eligibility ($e(\tau)$; intensity of shading in stopwatch) at the time τ_{CF} of subsequent climbing fiber (CF) input.

(**b, c, d**) Schematics of two possible mechanisms of ATTRAP: adjustment of a single timer, and selection from a bank of timers. Schematics were generated from biochemical network implementations of the ATTRAP algorithms shown in **Extended Data Figs. 3,4**. In the timer adjustment mechanism (**b**, *top*; **c**), an eligibility window $e(\tau)$ for spike timing-dependent plasticity is generated by a single molecular timer initiated by PF input, such as an accumulator reaching a threshold. This eligibility window is itself plastic and can be shifted earlier or later according to the value of an ATTRAP timer adjustment window. The value of the window at the time of the subsequent CF input determines the change in the time of peak eligibility for plasticity. If the CF spikes before peak eligibility, the time of peak eligibility is shifted earlier, e.g., by decreasing the eligibility threshold (*green*), but if the CF spikes after peak eligibility, the time of peak eligibility is shifted later (*blue*) relative to the PF input (change in the location of the *pink wedge* indicating eligibility for plasticity in **b**). In the timer selection model (**b**, *bottom*; **d**), eligibility for plasticity results from a weighted sum of activities of a bank of molecular timers that peak at different times after PF input. Tuning of the timing requirements for plasticity is accomplished by changing the contribution of each molecular timer to the eligibility for plasticity by increasing the coupling to eligibility of the timer most active during CF input and decreasing the couplings of the other timers (changes in the shade of the *pink wedges* in **b**). This effectively changes the level of eligibility for plasticity at every point in time according to the correlation of the CF input with an ATTRAP selection window.

(**e**) ATTRAP by timer adjustment shifts the eligibility window and hence the timing requirements for associative plasticity (*black*) to contain the median (*red*) of the distribution of intervals between PF and CF spikes (*grey*).

(**f**) ATTRAP by timer selection can be implemented with fixed (*top*) or proportional (*bottom*) updates to the eligibility coupling weights as a result of each CF input. Fixed updates alter the eligibility window and timing requirements for associative plasticity to contain the mode of the PF-CF pairing interval distribution. Proportional updates tune the eligibility window to approximate the shape of the distribution. Colors same as (**e**).

Extended Data Fig. 1. Similar plasticity of Purkinje cell spiking in the first 30, 50, and 100 ms after a parallel fiber stimulus.

(**a,b**) Plasticity of parallel fiber-elicited spiking in representative Purkinje cells (**a**) and the population averages (**b**) from normally reared mice, with spikes counted in the 30, 50 or 100 ms after parallel fiber (PF) stimulation. The spike counts in these different measurement windows underwent similar changes after pairing parallel fiber (PF) and climbing fiber (CF) stimulation, at pairing intervals of 0 ms (***top***) or 120 ms (***bottom***). Results obtained using the 50 ms window for measuring spike count (*green, blue*) are reproduced from **Fig. 1c.**

(**c,d**) Same as in (**a,b**), but for dark-reared mice. Similar results were obtained when PF-elicited

Purkinje cell spiking was measured in the first 30, 50, or 100 ms after PF stimulation.

Extended Data Fig. 2. Model of changes in behavior over the course of training.

(**a**) Root mean squared error (RMSE) of eye movement responses as a function of the number of training trials (i.e., sinusoidal stimulus cycles), in models with untuned (*green*) and tuned (*blue*) plasticity rules. Mean ± 1 standard deviation across 10 simulations. Following the experimental results (**Fig. 2f**), we assumed a slightly higher baseline gain and therefore lower initial error for simulations with the untuned rule. Model responses shown in **Fig. 2c,d** are from 3,600 trials of learning (*dashed line*).

(**b**) Oculomotor error, as measured by failure to stabilize the optokinetic stimulus on the retina (retinal slip) in models with untuned (*green*) and tuned (*blue*) plasticity rules after extended training of 7,200 trials. Whereas oculomotor errors are nearly eliminated in the tuned model, errors persist in the untuned model.

Extended Data Fig. 3. Biologically plausible ATTRAP mechanisms extract features of the input timing distribution.

(**a**) Reaction diagram for a tunable accumulation-to-bound timer. ***Top,*** The timer consists of four molecular species, each of which can be in an active (*marked with asterisk)* or inactive form. Parallel fiber (PF) input starts a cascade that leads to the accumulation of B* which, when above threshold, activates C. Eligibility for plasticity is controlled by the time-varying concentration of C*. ***Bottom,*** the threshold for activation of C, and hence the time of peak eligibility, is altered depending on the concentrations of the four species at the time of a climbing fiber (CF) input.

(**b**) Dynamics of the four active molecular components of the accumulation-to-bound timer. In this example, the threshold for activation of C is marked by the *solid grey line*, and the time B* crosses

the threshold is marked by the *dotted black line*. The eligibility, controlled by species C*, is shown in *black*.

(**c**) Implementation of the timer selection method for ATTRAP. ***Top,*** Reaction diagram for a cascade of sequentially active molecules coupled to eligibility. PF input starts the cascade, which leads to the sequential, transient activation of timer species X_m and delay elements D_m. The timer elements, by interacting with active eligibility coupling cofactors $U_m{}^*$, produce a time-varying eligibility window. Alternatively, timer elements could be a bank of accumulation-to-bound timers, as in (**a**), similarly coupled to eligibility. ***Bottom,*** The eligibility coupling weight is increased for the timer that is most active at the time of CF input, while all others are decreased. Biochemically, this corresponds to a situation in which the activator molecules are a limited resource with a fixed concentration. The out-of-sync cofactors inactivate and unbind their activator molecules, which are then taken up by the cofactor corresponding to the most active timer.

(**d**) Dynamics of two example sequentially active timer molecules (*black*) and corresponding delay elements (*blue*). A timer molecule X_m is activated when the concentration of the preceding delay element $D_{m-1}{}^*$ reaches a threshold level (*grey dotted lines*).

(**e**) The timer adjustment mechanism extracts the median of the input timing distribution. ***Top,*** Meta-learning dynamics of the biophysical implementation of the timer adjustment mechanism using the accumulation-to-bound timer in (**a**). The boundaries of the full-width-half-maximum extent of the eligibility window for each parallel fiber-climbing fiber (PF-CF) pair presentation averaged across 10 simulations is shown. The median of the PF-CF pairing interval is shown in *red*. ***Middle,*** The eligibility window generated by averaging the threshold for eligibility over the last 4,000 pairings across 10 simulations (*black*) contains the median PF-CF pairing interval (*red*). The shape of the pairing interval distribution is shown in *grey* behind the eligibility window.

Bottom, Cumulative distribution function of PF-CF intervals. The median timing is the cumulative distribution function evaluated at 0.5 (*dotted grey line*). The full-width-half-maximum extent of the eligibility window (*black trace, **middle***) is superimposed on the cumulative probability in *light blue*.

(**f**) As in (**e**), for an idealized model in which the eligibility window is Gaussian and climbing fiber spikes occurring before or after the peak shift the peak by a fixed amount earlier or later respectively. ***Top,*** Mean ± 1 standard deviation of the peak for each PF-CF pair presentation across 10 simulations. ***Middle,*** The eligibility window generated by averaging the peak timing over the last 4,000 pairings across 10 simulations (*black*).

(**g**) The two types of the timer selection mechanism, with fixed (***left***) or proportional (***right***) updates of coupling weights, generate eligibility windows (*black traces*) that extract the mode (***top left***, *red line; **bottom**, red arrows*) of the input timing distribution (*grey*) or that approximate the shape of the distribution (***right***), respectively. Eligibility windows were generated using an idealized set of timers with narrow, Gaussian activations, and 10 million (*fixed*) or 5 million (*proportional*) PF-CF spike time intervals. ***Middle***, Eligibility coupling weight of the mode-containing timer (*blue*, mean ± 1 standard deviation across 10 simulations), for the example distribution in the ***top.***

(**h**) Slices of the 2-D selection window (**Fig. 3d**) for two given time points $\tau = 40$ ms and 120 ms indicate how the height of the eligibility window for plasticity at those points changes as a function of the time of climbing fiber input τ_{CF}. *Blue* shading indicates a positive change, and *green* indicates a negative change.

Extended Data Fig. 4. Biochemical implementations of timer selection mechanism for ATTRAP.

(**a**) Simple bank of six timers for biochemical implementations of the timer selection ATTRAP mechanism. Peak times were spread evenly across 200 ms after parallel fiber (PF) input.

(**b**) Biochemical implementation of the timer selection mechanism with fixed updates from zero-order kinetics. *Left,* The changes in concentrations of the eligibility coupling cofactors (***top***), and the limited resource (activator) molecule (***bottom***; **Extended Data Fig. 3c**) are shown as a result of 5 parallel fiber-climbing fiber (PF-CF) pair presentations. Colors correspond to timer activations in (**a**). *Right, Top,* An example eligibility window after 5,000 PF-CF pair presentations (*black*) picks out the mode (*red*) of the distribution (*grey*). *Right, Bottom,* Mean $\pm$ 1 standard deviation over 10 simulations of the concentration of eligibility coupling cofactor U_3^*, corresponding to the timer whose activation contains the mode.

(**c**) Biochemical implementation of the timer selection mechanism with proportional updates from first-order kinetics. *Top,* An example eligibility window after 50,000 PF-CF pair presentations (*black*) that approximates the shape of the PF-CF pairing interval distribution (*grey*). *Bottom,* Changes in concentrations of the eligibility coupling cofactors as a result of 5 PF-CF pair presentations. Colors correspond to timer activations in (**a**).

Extended Data Table. 1. | Reaction rates and threshold parameters for biochemical implementation of the accumulation-to-bound timer

Parameter	Value
k_A^-	1,000
k_B^+	5
k_B^-	100
k_C^+	1,500
k_C^-	250
k_D^+	200
k_D^-	20
K_{DA}	0.3
K_{CD}	0.9
$K_{X\theta}$	0.025
$K_{Y\theta}$	0.025
K_{DX}	0.9
C_{min}^*	0.001

Extended Data Table. 2. | Reaction rates and threshold parameters for biochemical implementation of the cascade of timers

Parameter	Value	Parameter	Value
k_A^-	100	K_9	0.3965
$k_{X,m}^+$	15,000	K_{10}	0.3965
$k_{X,m}^-$	200	K_{11}	0.3965
$k_{D,1}^+$	100	K_{12}	0.3964
$k_{D,1}^-$	100	K_{13}	0.3964
$k_{D,m}^+$ for $m > 1$	80	K_{14}	0.3964
$k_{D,m}^-$ for $m > 1$	80	K_{15}	0.3964
K_A	0.6000	K_{16}	0.3964
K_1	0.3800	K_{17}	0.3966
K_2	0.3670	K_{18}	0.3967
K_3	0.3864	K_{19}	0.3961
K_4	0.3931	K_{20}	0.3966
K_5	0.3953	K_{21}	0.3965
K_6	0.3960	K_{22}	0.3968
K_7	0.3963	K_{23}	0.3986
K_8	0.3964	K_{24}	0.4007
		K_{25}	0.4016