A theory of cerebellar learning as a spike-based reinforcement learning in continuous time and space

Rin Kuriyama[1], Hideyuki Yoshimura[2], and Tadashi Yamazaki[1,✉]

[1]Graduate School of Informatics and Engineering, The University of Electro-Communications, Tokyo, Japan; [2]Okinawa Institute of Science and Technology Graduate University, Okinawa, Japan

The cerebellum has been considered to perform error-based supervised learning via long-term depression (LTD) at synapses between parallel fibers and Purkinje cells (PCs). Since the discovery of multiple synaptic plasticity other than LTD, recent studies have suggested that synergistic plasticity mechanisms could enhance the learning capability of the cerebellum. Indeed, we have proposed a concept of cerebellar learning as a reinforcement learning (RL) machine. However, there is still a gap between the conceptual algorithm and its detailed implementation. To close this gap, in this research, we implemented a cerebellar spiking network as an RL model in continuous time and space, based on known anatomical properties of the cerebellum. We confirmed that our model successfully learned a state value and solved the mountain car task, a simple RL benchmark. Furthermore, our model demonstrated the ability to solve the delay eyeblink conditioning task using biologically plausible internal dynamics. Our research provides a solid foundation for cerebellar RL theory that challenges the classical view of the cerebellum as primarily a supervised learning machine.

Cerebellum | Spiking Neural Network | Reinforcement Learning | Simulation

Correspondence: ***contact25@numericalbrain.org***

Theories of learning in the cerebellum have historically attracted considerable interest. One of the most widely known theories, known as the Marr-Albus-Ito model (1–3), proposes that the cerebellum acts as an error-based supervised learning (SL) machine (4) to modify output and reduce the discrepancy between actual and desired outputs. In this theory, climbing fibers (CFs) deliver error or teacher signals that can trigger long-term depression (LTD) at parallel fiber (PF) and Purkinje cell (PC) synapses (5). Since the discovery of LTD in the early 1980s (6), biological research has found many other forms of synaptic plasticity in the cerebellum (7). These findings have been extending the Marr-Albus-Ito theory gradually (8–10). Furthermore, researchers have examined the learning theory using realistic spiking network models of the cerebellum (11–17). Especially, Hausknecht et al. (13) examined the learning capability of the cerebellum through a cerebellar spiking network modeled in an SL context. The network successfully performed SL and control tasks but failed to solve RL and complex recognition tasks. Geminiani et al. (17) implemented a spiking network-based cerebellar model featuring bidirectional plasticity at PF-PC and PF-molecular layer interneurons (MLIs) synapses across multiple microzones, with several microzones more likely depression and others more likely potentiation during motor control. Their results highlight the cerebellum's capacity for error correction and adaptive learning, particularly within the context of associative learning and motor control. Both computational studies modeled the cerebellum as a conventional associative learning machine and highlighted the cerebellar learning capability, especially in the SL context.

One of the other research directions is to adapt reinforcement learning (RL) theory (18) as an alternative to conventional SL. In an RL context, there is an "agent" and an "environment". The agent takes an action affecting the state of the environment. In response to the action, the environment sends a "reward" to the agent which represents how good/bad the action was. The agent learns appropriate actions that maximize an expected future reward. One of the most notable differences between SL and RL is feedback information from the environment. An SL machine receives errors or teacher signals representing optimal actions, while an RL machine receives only reward signals.

Swain et al. (19) reported that the cerebellum fulfills three key requirements to establish reinforcement learning: it receives sensory information about the external environment and the internal environment, selects a behavior to execute, and processes evaluative feedback on the success of that behavior. Building on this idea, Yamazaki and Lennon (20) proposed the concept that the cerebellum might perform RL. Specifically, CFs deliver reward information, PCs select actions, and MLIs evaluate the current state. Synaptic weights at PF–MLI as well as PF–PC synapses are modulated by the CF signals. The researchers interpreted this structure as an RL machine, specifically as an actor-critic model (18). However, there has been no detailed spike-based implementation for the cerebellar RL concept.

In this research, based on a continuous time actor-critic framework with spiking neurons (21), we implemented a cerebellar spiking network as an actor-critic model based on the known anatomical properties of the cerebellum. In

particular, we modified a weight-updating rule of the framework for the cerebellum. We evaluated the model using a simple machine learning benchmark task, known as the mountain car task (22). We also examined the behavior of the model using delay eyeblink conditioning (23), which is a standard cerebellum-dependent motor learning task. To our knowledge, the model presented here is the first cerebellar spiking network model that can act as an RL machine.

Results

Spike-based implementation of cerebellum-style RL

In this research, we implemented a cerebellar spiking network model as an RL machine (Figure 1) while referring to the cerebellar RL concept (20) and the spiking actor-critic framework (21). Parallel fibers (PFs) excite stellate cells (SCs), basket cells (BCs), and Purkinje cells (PCs). SCs and BCs inhibit PCs. BCs and PCs inhibit each other, and also both have a self-inhibitory connection. More specifically, BCs and PCs are divided into multiple groups, which correspond to action choices as described in the section titled 'Actor and action selection'. BCs inhibit PCs in the same group and BCs in other groups, whereas PCs inhibit BCs in the same group and PCs in other groups. PCs also inhibit neurons in the deep cerebellar nuclei (DCN), which are also divided into multiple groups. CFs provide excitatory inputs to SCs and PCs. We used CF signals to trigger synaptic plasticity. Therefore, we did not implement actual CF–SC and CF–PC connections.

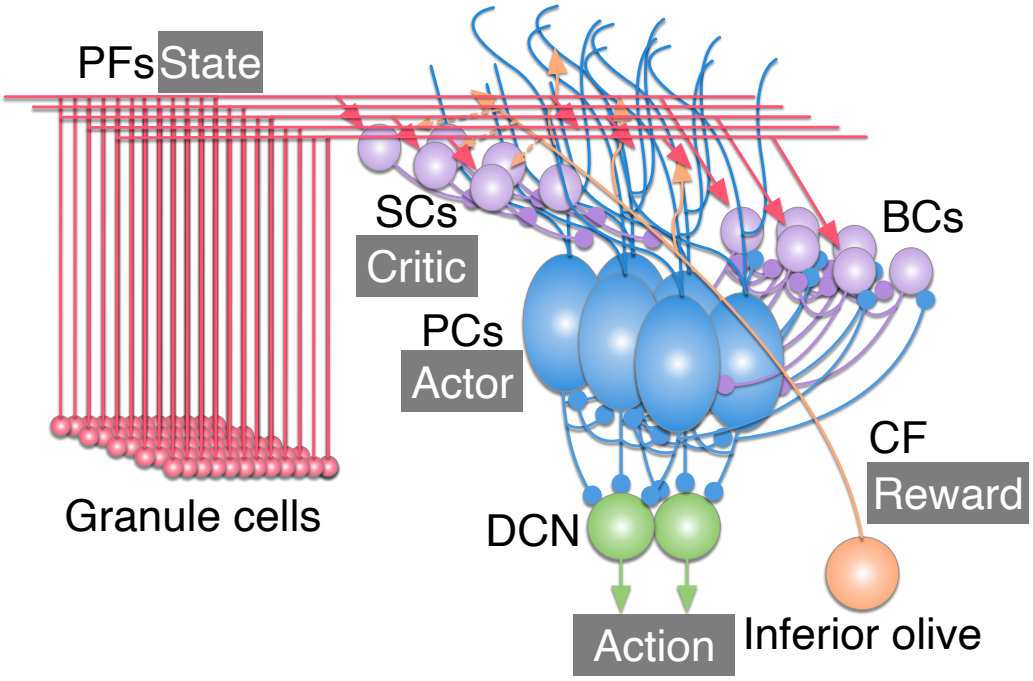

Figure 1. Structure of the present cerebellar network implemented as an actor-critic model. Red, purple, blue, green, and orange cells represent granule cells, molecular layer interneurons (SCs and BCs), PCs, DCN, and inferior olive, respectively. PFs deliver state information, CF delivers reward information, PCs serve as the actor, and SCs serve as the critic. Triangle-headed arrows and circle-headed arrows represent excitatory synapses and inhibitory synapses, respectively. Dashed lines represent excitatory signals from CF via glutamate spillover (24).

In the cerebellar spiking network model, we implemented an RL algorithm known as the actor-critic model (18) based on the previous spiking actor-critic framework (21). In RL, an agent in state $s(t)$ at time t takes action $a(t)$ following a policy π. In response to action $a(t)$, the environment changes its state and returns a reward $r(t)$ to the agent. The agent learns to find the optimal policy that maximizes the expected future reward. Specifically, the agent of the actor-critic model consists of two modules: an actor and a critic. The actor calculates preferences of actions $h_\alpha(t)$ and decides on action $a(t)$, while the critic computes the state value $V(t)$ which refers to the expected cumulative future reward. The temporal difference error (TD error) is calculated from the returned reward and the state value for updating the inner parameters of both the critic and the actor.

In our implementation, $s(t)$ was transmitted through PFs to PCs, BCs, and SCs. The actor and the critic were PCs and SCs, respectively. Practically, the activity of PCs represented the avoidance of actions $-h_\alpha(t)$, while the activity of SCs represented the additive inverse of the state value $-V(t)$ rather than TD error which was represented by MLIs in a previous theoretical model (20). BCs contributed to action selection by modulating the activity of PCs (25, 26). Then, the most activated DCN group decided on an actual action $a(t)$. In response to the action, reward information was transmitted through CFs to PCs and SCs.

Representation of states and rewards. States were represented by PFs following the formulation in the classical Marr-Albus-Ito model (1–3). For simplicity, we arranged the PFs on a two-dimensional grid plane composed of a number of small tiles (27). We named this plane the PF plane. Each tile had an individual receptive field that determined the PF's firing rate based on the observed state, as described in the section titled 'Implementations'.
In the present study, rewards $r(t)$ were actually considered as punishments so that $r(t) \leq 0$. CFs delivered the information as excitatory signals by flipping the sign as follows:

$$\mathrm{CF}(t) = -r(t). \tag{1}$$

Critic. We designed the temporal activity of SCs to approximate the state value. Strictly speaking, the additive inverse of the state value was calculated by SCs as follows:

$$-V(t) = \nu \mathrm{SC}(t) - V_0, \tag{2}$$

where ν, $\mathrm{SC}(t)$, and V_0 represent a scaling factor, the temporal activity of SCs, and the baseline of the state value, respectively. The baseline V_0 was set at ν times the average value of $\mathrm{SC}(t)$, which was calculated during the latter half of an episode without learning. $\mathrm{SC}(t)$ was calculated as follows:

$$\mathrm{SC}(t) = \frac{1}{N_{\mathrm{SC}}} \sum_i \mathrm{SC}_i(t), \tag{3}$$

$$\mathrm{SC}_i(t) = \frac{\mathrm{SC}_{i,\mathrm{decay}}(t) - \mathrm{SC}_{i,\mathrm{rise}}(t)}{\tau_{\mathrm{SC,\,decay}} - \tau_{\mathrm{SC,rise}}}, \tag{4}$$

$$\frac{d\mathrm{SC}_{i,l}(t)}{dt} = -\frac{\mathrm{SC}_{i,l}(t)}{\tau_{\mathrm{SC},l}} + S_i(t), \tag{5}$$

where N_{SC} and $\mathrm{SC}_i(t)$ represent the number of SCs and the temporal activity of ith SC, respectively. l is a label in $\{\mathrm{rise}, \mathrm{decay}\}$. The temporal activity of ith SC rises with a time constant $\tau_{\mathrm{SC,rise}}$ and decays to 0 with a time constant $\tau_{\mathrm{SC,decay}}$. Spike train of the ith neuron is defined as $S_i(t) = \sum_f \delta_D(t - t_i^{(f)})$ with the Dirac delta function $\delta_D(t)$ and the fth spike of the ith neuron.

Actor and action selection. PCs performed the actor and determined the action taken (28). We interpreted the temporal activity of PCs as expressing avoidance of certain actions. Specifically, the PCs in our model compute the avoidance of an action $h_{\bar{\alpha}}(t)$ instead of the preference for an action $h_\alpha(t)$. Therefore, the taken action $a(t)$ was determined by the least-avoided action in action space $\mathcal{A}$ at time t.

$$a(t) = \underset{\alpha \in \mathcal{A}}{\operatorname{argmin}}\, h_{\bar{\alpha}}(t). \tag{6}$$

The avoidance of action $\alpha \in \mathcal{A}$ was represented by the activity of PCs in $\bar{\alpha}$ group as follows :

$$h_{\bar{\alpha}}(t) = \frac{1}{N_{\mathrm{PC}_{\bar{\alpha}}}} \sum_{i \in \mathrm{PC}_{\bar{\alpha}}} \mathrm{PC}_i(t), \tag{7}$$

where $N_{\mathrm{PC}_{\bar{\alpha}}}$ is the number of neurons in the group $\mathrm{PC}_{\bar{\alpha}}$. $\mathrm{PC}_i(t)$ was the temporal activity of ith PC defined as follows:

$$\mathrm{PC}_i(t) = \frac{\mathrm{PC}_{i,\mathrm{decay}}(t) - \mathrm{PC}_{i,\mathrm{rise}}(t)}{\tau_{\mathrm{PC,decay}} - \tau_{\mathrm{PC,rise}}}, \tag{8}$$

$$\frac{d\mathrm{PC}_{i,l}(t)}{dt} = -\frac{\mathrm{PC}_{i,l}(t)}{\tau_{\mathrm{PC},l}} + S_i(t), \tag{9}$$

where l, $\tau_{\mathrm{PC},l}$, and $S_i(t)$ represent the label $l \in \{\mathrm{decay}, \mathrm{rise}\}$, the time constant for l, and the spike train of the ith PC, respectively. Practically, the actual action was generated at the downstream DCN (see the section titled 'Implementations').
Those spike activities of PCs were used for updating synaptic weights to realize which action was selected. It was vital that one PC group pauses while the other PC groups activate. We referred to this pausing behavior as a "dent", while activation of neurons in a specific area was referred to as a "bump" by Frémaux et al. (21). The inhibitory loop between BCs and PCs helped to pause the single group solely to make a dent. Although the original spiking actor-critic framework supports continuous action space, the present model only supports discrete action space with only two actions, because of the difficulty of making a single dent on PCs.

Weight-update rules. Our weight-update rule was based on TD long-term potentiation (LTP) which had been proposed in a previous spiking actor-critic model (21). However, unlike the previous model, which calculated the time derivative of the state value directly to compute the TD error, our model did not calculate it, as there was no component in the molecular layer of our model that explicitly handled the TD error.

To address this problem, we conducted an eligibility trace and transformed the weight-update rule. As a result, our weight-update rule for the state value was as follows:

$$\frac{d\Delta w_{i,j}^n}{dt} = \alpha\left(R(t)z_{i,j}(t) + \frac{\tau_r - \tau_z}{\tau_r\tau_z}V(t)z_{i,j}(t) - V(t)f_{i,j}(t)\right), \tag{10}$$

$$\frac{dz_{i,j}(t)}{dt} = -\frac{z_{i,j}(t)}{\tau_z} + f_{i,j}(t), \tag{11}$$

$$f_{i,j}(t) = S_i(t)\kappa_{i,j}(t), \tag{12}$$

$$\frac{d\kappa_{i,j}(t)}{dt} = -\frac{\kappa_{i,j}(t)}{\tau_\kappa} + S_j(t), \tag{13}$$

where $\Delta w_{i,j}^n$, α, $R(t)$, $z_{i,j}(t)$, τ_r, τ_z, $f_{i,j}(t)$, $S_i(t)$, $S_j(t)$, $\kappa_{i,j}(t)$, and τ_κ are the amount of synaptic change in nth episode, the learning rate, the negative reward, the eligibility trace for $w_{i,j}$, the reward discount time constant, the decay time constant of eligibility trace, the forcing term of eligibility trace, the postsynaptic spike train, the presynaptic spike train, the window function for spike timing dependent plasticity, and the decay time constant of the window function. According to the previous model (21), the eligibility term memorizes the history of inputs before the time of the last postsynaptic neuron spike, $\kappa_{i,j}(t)$ is reset to 0 after the postsynaptic neuron spike. The relationship of our weight-update rule to the original weight-update rule (21, 29) is described in the section S1 of SI Appendix. The synaptic weight $w_{i,j}$ was updated at the end of every episode T_{end} by $w_{i,j}^{n+1} = w_{i,j}^n + \eta\Delta w_{i,j}^n(T_{\text{end}})$ with learning rate η, and the synaptic weight was bounded between $[0,1]$. Since SCs represented the additive inverse state value and we denoted the negative reward $r(t) = -\text{CF}(t)$, the equation for the PF–SC synapses was expressed as follows:

$$\frac{d\Delta w_{i,j}^n}{dt} = \text{CF}(t)z_{i,j}(t) - \frac{\tau_r - \tau_z}{\tau_r\tau_z}V(t)z_{i,j}(t) + V(t)f_{i,j}(t). \tag{14}$$

On the other hand, the weight-update rule for PF–PC synapses was expressed as follows:

$$\frac{d\Delta w_{i,j}^n}{dt} = -\text{CF}(t)z_{i,j}(t) + \frac{\tau_r - \tau_z}{\tau_r\tau_z}V(t)z_{i,j}(t) - V(t)f_{i,j}(t). \tag{15}$$

The specific learning rates for PF–SC synapses and PF–PC synapses are denoted as η_{SC} and η_{PC}, respectively. Specifically, given that the conjunctive activation of PFs and CFs induces LTP at PF–SC synapses (30–33), our SCs learned the approximate additive inverse of the state-value. LTD at PF–SC synapses was due to PF firing, and was modulated by postsynaptic activation results in bidirectional synaptic plasticity (30–33). At PF–PC synapses, the conjunctive activation of PFs and CFs induces LTD (7), while PF activation induces LTP. These synaptic plasticities are also modulated by inhibitions of SCs (34), and we assumed that PC activity contributes to eligibility trace and bidirectional synaptic plasticity (35, 36).

Additionally, we confirmed that our critic is capable of learning state values through a Linear Track Task simulation (see SI Appendix, section S2), which is a simplified version of the water-maze task used to isolate learning by the critic from that by the actor (21).

Simulation of Mountain Car Task

In order to evaluate the learning capability of our model as an RL machine, we applied our model to the mountain car task (37), a classic RL challenge, and evaluated its performance over 10 trials. In this task, an agent controls a car situated between two hills (Figure 2A). The goal is to have the car climb up the mountain to reach a flag positioned at the peak of the right hill. Due to the car's limited engine power and the hill's steepness, direct ascent is not possible. An agent observes a position $p_x(t)$ and a horizontal element of the car velocity $v_x(t)$, and decides the direction (either left or right) to push the car with a constant force $F(t) = \bar{F}a(t),\ \ a(t) \in \{-1, 1\}$. In early episodes, the car was confined to moving at the bottom of the valley, and failed to reach the goal (Figure 2B). However, after hundreds of episodes, our model learned to utilize the slope for acceleration, and eventually reached the goal (Figure 2C). Specifically, we found that the trajectory involved the car ascending the right hill, then the left, and finally the right hill again. The additive inverse of the state value $-V(t)$, which was approximated by SC activity, was noisy but mostly stayed around 0 at every state before learning (Figure 2B). In a successful episode after learning (Figure 2C), the value was high at the start and decreased over time as the car approached the goal. The moving average of success rate over 10 trials progressively increased and stabilized at approximately 80% after 600 episodes (Figure 2D). These results indicate that our model solved the mountain car task, and suggest that our model successfully acted as an RL machine.

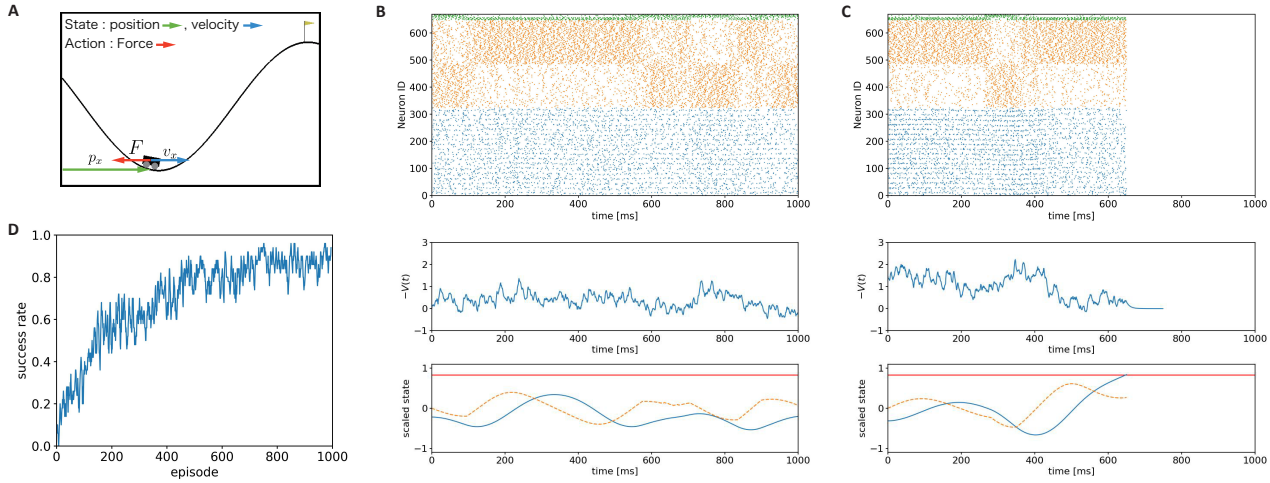

Figure 2. Simulation of the mountain car task. (**A**) Schematic of the mountain car task. p_x is the horizontal position of the cart, v_x is the horizontal velocity, and F is the external force applied to the cart by the agent. The flag marks the goal position. For clarity in the image, we adjusted the x-axis to start at 0 on the left edge. (**B,C**) Comparison of behaviors in the mountain car task. (**B**) Example behavior in early episodes. (**C**) Example behavior in later episodes. Top panels show raster plots of SCs (blue), BCs (orange) and PCs (green). The number of these neurons are 324, 324, and 20, respectively. Middle panels show trajectory of $-V(t)$ computed from the activity of SCs. Bottom panels show trajectory of scaled states. Blue line and orange dashed line represent the normalized value of a position and that of a velocity. Red line represents the goal position. When the position reaches the goal, an episode ends while a simulation lasts 100 ms (see the section titled 'Pre- and post-episode processes'). (**D**) Success rate of mountain car task. The moving average of success rate across 10 trials with window size 5 episodes.

Simulation of Delay Eyeblink Conditioning Task

To examine whether our cerebellar model shows internal dynamics consistent with the biological cerebellum, we conducted simulation of a standard cerebellum-dependent motor learning task known as the delay eyeblink conditioning task (23). In this task, an animal is presented with two types of stimuli: a conditioned stimulus (CS) and an unconditioned stimulus (US) (Figure 3A). The CS, a neutral stimulus such as a tone, initially does not trigger any response. On the other hand, the US, such as an air puff, naturally induces an unconditioned blink reflex in the animal. However, when the CS is repeatedly paired with the US, the animal begins to anticipate the air puff whenever the CS is presented. This anticipation leads the animal to blink in response to the CS alone, a learned behavior known as the conditioned response (CR). We modified the task to have two-dimensional state space: time t and eyelid position $p_e(t)$, along with two actions: opening and closing the eyelid (Figure 3B). The time axis ranges from 0 ms to 1,000 ms, corresponding to the duration of stimulus presentation, while the eyelid position is measured on a scale from fully closed (0) to fully open (1). The CS was presented from 0 ms to 1,000 ms as a time signal, and the US was introduced at 500 ms as a punishment if the eyelid was open, i.e., if $p_e(500)$ was greater than 0.1.

In early episodes, the agent preferred to open the eyelid (Figure 3C and 3B), so the agent received the US as a negative reward. As learning progressed, the anti-open PC group became active, whereas the anti-close PC group paused as if PCs expected the US onset. Thus, the eyelid started to close before the US onset. For every 10 trials, the agent successfully acquired the CR to the CS: closing its eyelid to avoid a punishment before the arrival of the US (Figure 3E). Note that we defined a successful episode as a sufficiently closed eyelid to avoid the punishment at the US onset. When focusing on the activity changes of a single PC in the anti-close group (Figure 3F), the neuron acquired a behavior to pause before US onset. These behaviors were consistent with behavioral experiments of the delayed eyeblink conditioning task (23). We also analyzed the activity of SCs during the task (Figure 3G). In early episodes, since the agent received the punishment as the US, the SC activity as the additive inverse of the state value $-V(t)$ increased exponentially toward the arrival of the US. These results are not inconsistent with biological experiments (38, 39), which show increases in MLIs' firing frequency from the time of CS onset to the time of US onset. In later episodes when the agent successfully closed the eyelid more than 80% of the time, the SC activity returned to zero. This returning behavior resulted from the disappearance of CF signals due to the successful avoidance of punishments.

We also conducted a simulation of the same task without PF–MLI synaptic plasticity. The activity of the anti-close PC group decreased as it approached the US onset (Figure 4A). However, the agent was not able to elicit a stable closing response in our settings. In a later episode, the time to close the eye was delayed (Figure 4B) compared to the that in the intact simulation (Figure 3D), so the agent could not close the eye sufficiently. These results are also consistent with experimental results where MLI–PC inhibitions were impaired (40). Furthermore, these results suggested that our model successfully solved the eyeblink task with biologically plausible internal dynamics.

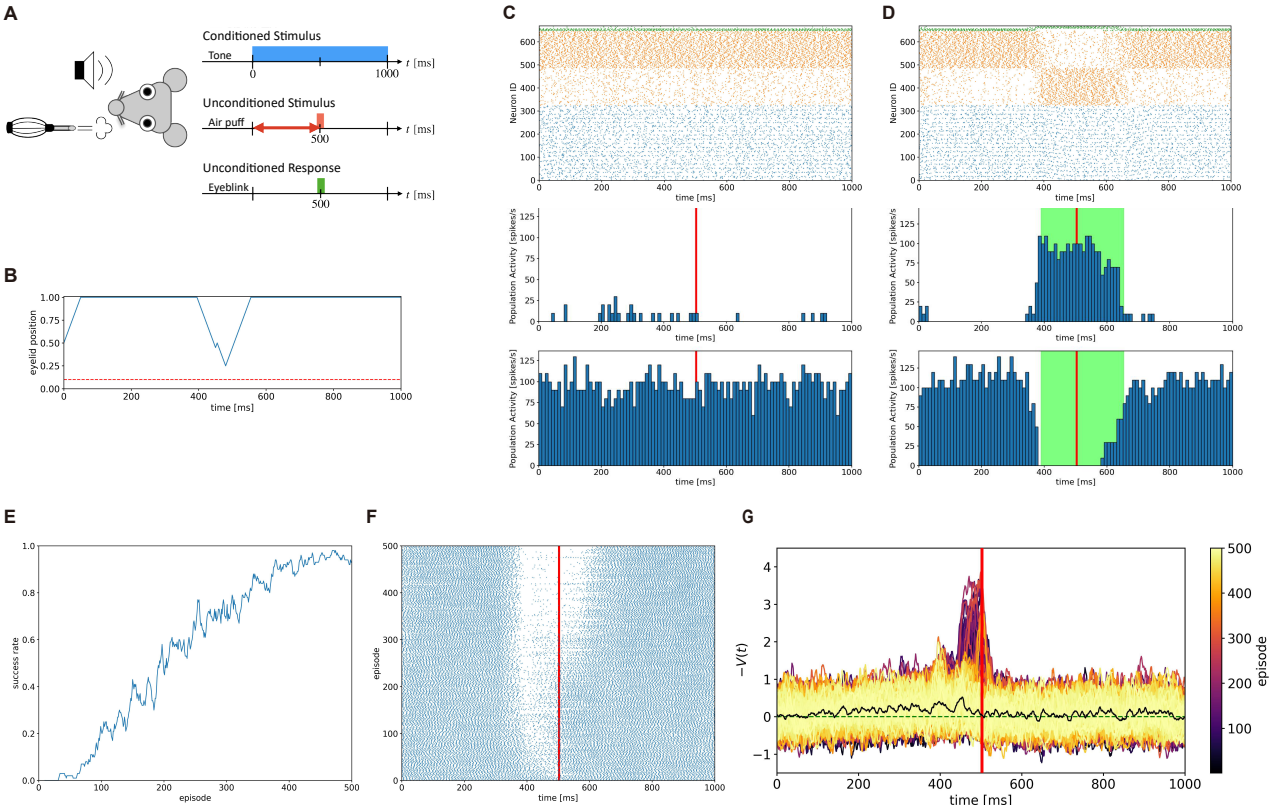

Figure 3. Simulation of the delay eyeblink conditioning task. (**A**) Schematic of the task. A neutral stimulus (CS) is presented, followed by an aversive stimulus (US) after a fixed interval. After repeated pairings, the agent closes its eyes (CR) in response to the CS alone. (**B**) Example state trajectory. The horizontal axis shows time from CS onset, and the vertical axis represents eyelid position. The dashed line marks the punishment threshold, triggered when the eyelid is above it at US onset ($t = 500$ ms). (**C,D**) Results for a failure episode (C) and a success episode (D). Top panels display SC (blue), BC (orange), and PC (green) raster plots. The number of these neurons are 324, 324, and 20, respectively. Middle and bottom panels show peristimulus time histograms for anti-open and anti-close PC population, respectively. Bin width is 10 ms. US onset is marked by a red line, and eyelid closure by green shading. (**E**) Moving average of success rate over 10 trials. (**F**) The activity changes of a single PC in the anti-close group, with dots representing spikes and red line indicating US onset. (**G**) Additive inverse of state value during conditioning, with colored traces showing stellate cell activity across 500 episodes (from dark purple in the 1st trial to light yellow in the 500th trial). The black line represents the average over the last 10 trials, green dashed line shows $-V(t) = 0$, and the red line indicates US onset.

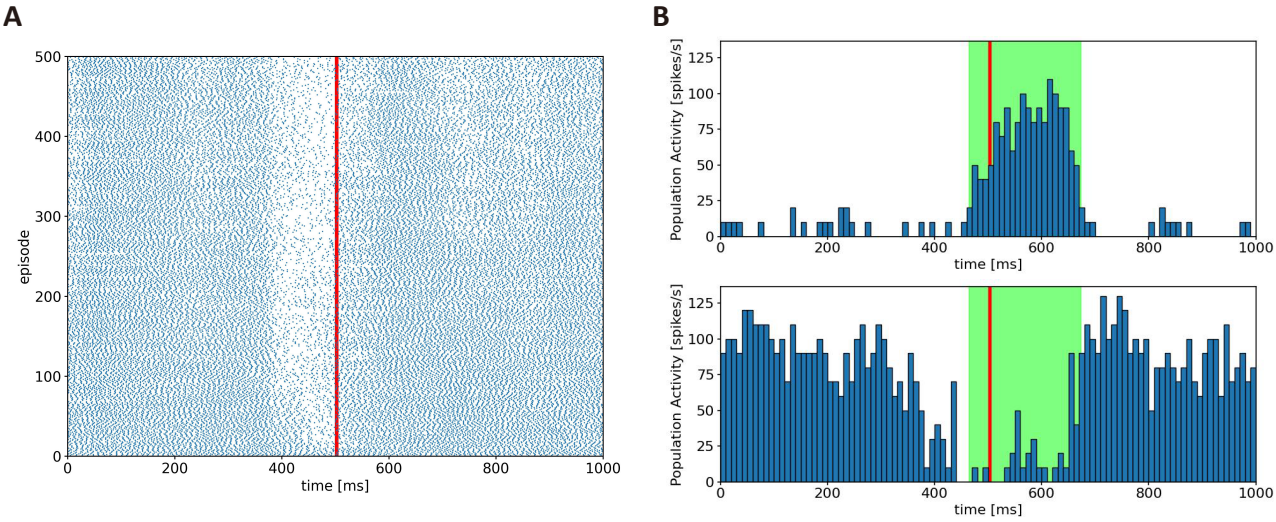

Figure 4. Results of the delay eyeblink conditioning task without PF–SC synaptic plasticity. (**A**) The activity changes of a single PC in the anti-close group. Dots represent spikes of the PC, and red line represents the time of US onset. (**B**) Example peristimulus time histograms of anti-open PC population (top) and those of anti-close PC population (bottom) in the later episode. Bin width is 10 ms. Red lines represent the time of US onset, while the green shaded intervals show when the agent closes its eyelid around the US onset.

Discussion

In this study, we implemented the cerebellar spiking network model performing RL in an actor-critic manner. Our simulations confirmed the model's ability to solve a simple RL task and demonstrated a cerebellum-dependent motor leaning task with biologically plausible internal dynamics. In contrast, Hausknecht et al. (13) reported that their cerebellar model was not able to solve RL tasks, only SL tasks. This is natural, because their model was formulated as an SL machine while ignoring PF–MLI synaptic plasticity. Recently, Geminiani et al. (17) implemented a spiking network-based cerebellar model that considers synaptic plasticity at both PF–PC and PF–MLI synapses. They reported that considering multiple synaptic plasticity mechanisms enhance learning capability in SL. Thus, those computational studies considered the cerebellum as an SL machine. Spiking network models of the cerebellum as an SL machine have been used for realtime motor control and online learning of hardware robots (14, 41). Our present cerebellar model will help those robotics studies by providing the capability of RL over SL.

In our cerebellar model, MLIs played essential roles as a critic. Recent studies have shown that MLIs actually contribute to learning in the cerebellum (38–40, 42) by their own learning capability through PF–MLI synaptic plasticity (30–32) and modulating PF–PC synaptic plasticity (34), in addition to controlling PC activities through MLI–PC inhibition (25, 26, 33, 43–46). In a behavioral experiment, genetically-modified mice lacking MLI–PC inhibition were able to acquire the CR, but the amplitude of eye closing was significantly lower than that of intact mice (40). Our results of the lesion simulation in the delay eyeblink conditioning task are consistent with the experimental results. These results suggest that the cerebellum can solve the association task even without normal SC activity, but the activity affects the performance of the acquired movement.

The present study assumed that CFs deliver negative reward information. In conventional cerebellar motor learning, CFs are considered to provide motor error signals, which are negative information. Furthermore, recent experimental results demonstrate that CFs carry not just motor-related information but also reward-related information such as reward delivery, prediction, and omission in mice (47–51). Thus, our assumption is consistent with those experimental findings. Hoang et al. (52) recapitulated the idea of negative reward information conveyed by CFs to Q-learning, which is a tabular-style reinforcement learning algorithm, and applied the Q-learning to a Go-No-go auditory discrimination task in mice. Furthermore, CF signals are provided in an all-or-none manner, and our model followed this property. However, several experiments have demonstrated that CF signals can carry graded information (53, 54), which might be useful to deliver reward information as analog values effectively.

In a standard RL algorithm called TD-learning, a reward-related signal tends to appear spontaneously at the onset of a signal that predicts the occurrence of a future reward after sufficient learning. The signal is called a reward-prediction signal or reward-prediction error. Ohmae and Medina (47) observed such reward-prediction signals conveyed by CFs at the CS onset after sufficient training of delay eyeblink conditioning in mice. Indeed, such prediction-error signals are prepared outside of the cerebellum. Therefore, in the present study, we did not consider such prediction error-type CF signals. If we incorporated them, those signals could have acted to control the balance between exploration and exploitation (18) while modulating the fluctuation of the internal network state. Another potential role would be to represent the most distant state from the goal state for reward delivery. Nevertheless, further investigation will be necessary to study the functional role of the prediction error-type CF signals.

If both the cerebellum and the basal ganglia perform RL, how do these regions cooperate for learning? One possible interpretation is hierarchical RL (HRL) (18) as discussed in our previous research (10). HRL considers two RL machines organized hierarchically. The higher RL machine breaks down a complex task into simpler, smaller subtasks. The lower RL machine tries to solve these subtasks one by one. Eventually, HRL involves creating a hierarchy of policies that makes the learning process more efficient and scalable (18, 55, 56). If the basal ganglia and the cerebellum cooperate to perform HRL, the resulting network would demonstrate powerful learning capabilities.

To our knowledge, our model is the first to implement cerebellar RL with spiking neurons. Our research provides an implementation that supports cerebellar RL theory to challenge the traditional view of the cerebellum as primarily a SL machine. Furthermore, our interpretation of a potential cooperative role of the basal ganglia and the cerebellum as HRL sheds light on how the multiple brain regions act together synergistically as a whole.

Materials and Methods

Implementations

First, we discretized a state space and mapped it on a PF plane with grid size $N_x \times N_y$. A tile positioned at (i, j) on the PF plane had its own responsible state $\mathbf{o}_{i,j}$ as follows:

$$\mathbf{o}_{i,j} = [i\Delta s_x, j\Delta s_y]^T, \tag{16}$$

$$\Delta s_x = \frac{s_{x,\text{max}} - s_{y,\text{min}}}{N_x}, \quad \Delta s_y = \frac{s_{y,\text{max}} - s_{y,\text{min}}}{N_y} \tag{17}$$

Table 1. Neuron-specific parameters. These parameters were adjusted from those used in our previous model (15).

Type	C_m [pF]	g_{L} [mS]	E_L [mV]	U_{th} [mV]	U_{r} [mV]	I_{ext} [pA]	ζ [pA]
SC	107	2.32	-68.0	-55.0	-70.0	30.0	5.0
BC	107	2.32	-68.0	-55.0	-70.0	30.0	1.0
PC	107	2.32	-65.0	-55.0	-66.0	30.0	5.0

where x and y are the axes of the PF plane, and $(s_{x,\mathrm{min}}, s_{x,\mathrm{max}})$ and $(s_{y,\mathrm{min}}, s_{y,\mathrm{max}})$ represent the ranges of state x and y, respectively. When the agent observed state $\mathbf{o}(t) = [x(t), y(t)]^T$, PFs on a tile at (i, j) fired as an inhomogeneous Poisson spike generator with a firing rate $\rho_{i,j}(\mathbf{o}(t))$ defined as follows :

$$\rho_{i,j}(\mathbf{o}(t)) = \frac{\rho_{\mathrm{PF}}}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{||\mathbf{o}(t) - \mathbf{o}_{i,j}||^2}{2\sigma^2}\right), \tag{18}$$

where $\rho_{\mathrm{PF}} = 50$ Hz is a scaling factor of firing rate, and $\sigma = 0.5$ is a standard deviation. To maintain the total activity of all PFs when the agent is at the edge of a state space, we added "margin" tiles as an outer frame. Although we modeled PFs as Poisson spike generators for simplicity, more realistic implementations are reported (57).
We utilized a Leaky Integrate-and-Fire model (LIF) to SCs, BCs, and PCs. A LIF model was implemented as follows.

$$C_m \frac{du_i}{dt} = -g_{\mathrm{L}}(u_i - E_{\mathrm{L}}) + I_{i,\mathrm{syn}}(t) + I_{\mathrm{ext}} + I_{i,\mathrm{noise}}(t), \tag{19}$$

$$t_i^{(f)} : u_i\left(t_i^{(f)}\right) = U_{\mathrm{th}}, \tag{20}$$

$$\lim_{t_i^{(f)} \to t,\ t > t_i^{(f)}} u_i(t) = U_{\mathrm{reset}}, \tag{21}$$

where C_m, u_i, t, g_{L}, E_{L}, $I_{i,\mathrm{syn}}$, I_{ext}, and $I_{i,\mathrm{noise}}$ represent the membrane capacitance, the membrane potential of the ith neuron, time, the leak conductance, the resting potential, the synaptic current to the ith neuron, the external input current, and the noise current to the ith neuron, respectively. When the membrane potential of the ith neuron reaches the threshold potential U_{th}, that time is denoted as $t_i^{(f)}$, and the neuron fires the fth spike and resets its membrane potential to the reset potential U_{reset}. A synaptic current was defined as follows:

$$I_{i,\mathrm{syn}}(t) = \sum_c s_c w_{i,j} I_{i,j}(t), \tag{22}$$

$$\frac{dI_{i,j}(t)}{dt} = -\frac{I_{i,j}(t)}{\tau_c} + S_j(t), \tag{23}$$

where c is a synapse label, s_c, $w_{i,j}$, τ_c, and $S_j(t)$ are the scaling factor, synaptic weight between the jth presynaptic neuron and the ith postsynaptic neuron, decay time constant, and spike train of the jth presynaptic neuron, respectively. Whether it is an inhibitory or excitatory current is determined by the sign of the scaling factor. Noise current $I_{i,\mathrm{noise}}(t)$ was modeled as follows:

$$\frac{dI_{i,\mathrm{noise}}}{dt} = -\frac{I_{i,\mathrm{noise}}}{\tau_{\mathrm{noise}}} + \zeta n(t), \tag{24}$$

where τ_{noise} is a time constant, and ζ is an amplitude factor. $n(t)$ is a uniform random variable in the range $[-1, 1]$. Parameters of all neurons were set as listed in Table 1, and parameters of all synapses were set as listed in Table 2.

Table 2. Synaptic parameters for each connection type. A hyphen in a weight cell means it will be modified because of synaptic plasticity.

Connection Type	scaling factor	weight	τ_c [ms]	probability
PF–SC	200.0	-	8.3	0.1
PF–BC	10.0	1.0	8.3	0.5
PF–PC	5.0	-	8.3	0.5
SC–PC	-0.05	0.03	10.0	0.5
BC–BC	-2.0	1.0	10.0	0.3
BC–PC	-1.0	1.0	10.0	0.3
PC–BC	-2.0	1.0	10.0	0.8
PC–PC	-19.8	1.0	5.0	1.0

As mentioned above, PFs were distributed on the two-dimensional grid plane, and the grid size was determined by each task. Each tile contained 100 PFs. SCs were prepared 1 per tile. Each SC positioned at $(i\Delta s_x, j\Delta s_y)$ could

receive excitatory inputs from PFs located in $s_{x,\min} \leq x \leq s_{x,\max}$ and $(j-1)\Delta s_y \leq y \leq (j+1)\Delta s_y$. The number of BCs was the same as that of SCs, and the number of PCs was 20. Both BCs and PCs could receive excitatory PF inputs from all tiles. The BC–BC, BC–PC, PC–BC, and PC–PC connections were constructed based on the grouping manner mentioned in the section titled 'Spike-based implementation of cerebellum-style RL'. The actual synapses were formed stochastically (Table 2). In this study, we divided BCs and PCs into two groups for all tasks. Additionally, for the sake of simplicity, we assumed that all CFs synchronized their activity so that our model had a single CF. Synaptic plasticity was applied to PF–SC and PF–PC synapses. The details of the weight-update rules are described in Equations Eq. (14) and Eq. (15).

Finally, the outcome of actions were generated by DCN cells prepared in an equal number of PCs. Each DCN cell was inhibited by its counterpart PC. This configuration means that the activity of each DCN cell, denoted as DCN_i, was modulated by the corresponding PC activity $\mathrm{PC}_i(t)$ defined by Equation 8, as follows:

$$\frac{d\mathrm{DCN}_i(t)}{dt} = -\frac{\mathrm{DCN}_i(t) - \mathrm{DCN}_0}{\tau_{\mathrm{DCN}}} - \nu \mathrm{PC}_i(t), \tag{25}$$

where τ_{DCN}, DCN_0, and ν represent the decay time constant, baseline activity of DCN, and a scaling factor, respectively.

Since PCs in the group $\mathrm{PC}_{\bar{\alpha}}$ represented the avoidance of action α and inhibited DCN cells, the activity of DCN group DCN_α represented the preference of actions $h_\alpha(t)$ as follows:

$$h_\alpha(t) = \frac{1}{N_{\mathrm{DCN}_\alpha}} \sum_{i \in \mathrm{DCN}_\alpha} \mathrm{DCN}_i(t), \tag{26}$$

where N_{DCN_α} is the number of neurons in the group DCN_α, and $\mathrm{DCN}_i(t)$ is the activity of ith DCN in the group DCN_α corresponding to action α. Therefore, the most activated DCN group represented the most preferred action so that the final outcome of actions was chosen as follows:

$$a(t) = \underset{\alpha \in \mathcal{A}}{\operatorname{argmax}}\, h_\alpha(t). \tag{27}$$

Task Definitions

Pre- and post-episode processes. As a pre-episode process, at the start of each episode, a free run for 200 ms was carried out to discard initial transient activity that could affect learning. Then, at the end of each episode, a simulation continued for 100 ms while stopping the firing of PFs to update synaptic weights. These pre- and post-episode processes were performed in all tasks.

Mountain Car Task. The dynamics of the car are described by the following differential equations :

$$\frac{dv_x(t)}{dt} = v_x(t) + \bar{F}a(t) - g\cos(3p_x(t)) \tag{28}$$

$$\frac{dp_x(t)}{dt} = v_x(t), \tag{29}$$

where $\bar{F} = 0.001$ is a force amplitude, $a(t) \in \{-1, 1\}$ represents an action as the force direction, $g = 0.0025$ represents the gravity. Observation spaces of the velocity and position are bounded by $[-1.2, 0.6]$ and $[-0.07, 0.07]$, respectively. When the car collides with the wall, its velocity is immediately reset to 0 without any negative reward associated with the collision. An agent receives constant negative reward at every step in the original settings. However, the maximum firing frequency of CF, which is approximately 10 Hz (58), was not enough to represent the constant negative reward. Therefore, in this study, we modeled a new negative reward function for CF. CF fired when the car was far from the goal and the speed was slow. Thus, the firing rate $\nu(t)$ was defined as follows:

$$\nu(t) = \bar{\nu}\left(1 - \frac{p_{\mathrm{goal}} - p_x(t)}{p_{\mathrm{goal}} - p_{\mathrm{wall}}}\right)\left(1 - \frac{|v_x(t)|}{v_{\max}}\right), \tag{30}$$

where $\bar{\nu}$, p_{goal}, p_{wall}, and $v_{\max}$ are max firing rate of 5 Hz, goal position, left side wall position, and limit of speed respectively. Then a negative reward function is defined as follows:

$$R(t) = \begin{cases} -5 & \text{when CF fires at time } t, \\ 0 & \text{otherwise.} \end{cases} \tag{31}$$

The car starts from the bottom of the valley with a velocity of 0 in every episode. The action $a(t)$ represents the direction in which the agent pushes the car. An episode ends when the car reaches the goal or when the elapsed time exceeds 1,000 ms, with no special negative reward given.
The grid size of the PF plane was 16×16. The parameters of the state value function, as defined in Equation Eq. (2), were $\nu = 200$ and $V_0 = 2.48$. The reward discount time constant τ_r, the decay time constant of eligibility trace τ_z, and the decay time constant of the window function τ_κ were set to 100 ms, 20 ms, and 20 ms, respectively. Initial synaptic weights of PF–SC and PF–PC were 0.05 and 0.8, respectively. Both learning rates for the critic and the agent were 0.4.
We conducted 10 trials of 1,000 episodes each. We calculated the changes in success rate for each trial using a moving average of results (success=1, failure=0) with a window size of 10 episodes, and then calculated the average across trials.

Delay Eyeblink Conditioning Task. The dynamics of the eyelid is described as follows:

$$\frac{dp_e(t)}{dt} = 0.01a(t), \tag{32}$$

where $a(t) \in \{-1, 1\}$ represents the action the agent takes at time t. Also, negative reward is defined as follows:

$$r(t) = \begin{cases} -5 & \text{when } t = 500 \text{ and } p_e(t) > 0.1, \\ 0 & \text{otherwise.} \end{cases} \tag{33}$$

The grid size of the PF plane was 16×16. The parameters of the state value function, which defined in Equation Eq. (2), were $\nu = 200$ and $V_0 = 2.48$. The reward discount time constant τ_r, the decay time constant of eligibility trace τ_z, and the decay time constant of the window function τ_κ were set to 100 ms, 20 ms, and 20 ms, respectively. To prioritize the 'open' action, we initially set synaptic weights for PF–PC$_{\text{anti close}}$ to 0.8, while PF–PC$_{\text{anti open}}$ was set to 0.7. We conducted the simulation for 500 episodes with learning rates $\eta_{\text{SC}} = 0.05$ and $\eta_{\text{PC}} = 0.3$.
The whole code of our model and all environments were written in C++. All ordinary differential equations were solved numerically with a forward Euler method with a temporal resolution of 1.0 ms.

Conflict of interest statement

The authors declare no competing interest

Author contributions

T.Y. conceived and designed the research. R.K. and H.Y. contributed the formulation and the implementation of the algorithms. R.K. wrote the code, performed simulation, and analyzed the data. R.K. and T.Y. wrote the original draft. R.K., H.Y., and T.Y. discussed the draft and revised it. All authors contributed to the article and approved the submitted version.

Funding

This work was supported by JSPS KAKENHI Grant numbers JP22KJ1372 and JP22H05161. Part of this study was supported by the MEXT Program for Promoting Researches on the Supercomputer Fugaku hp220162.

Acknowledgment

We would like to thank Profs. Kazuo Kitamura, Jun Igarashi, Shogo Ohmae, and Kenji Doya for their fruitful discussions and helpful insights.

Bibliography

1. D. Marr, "A theory of cerebellar cortex," The Journal of Physiology, vol. 202, no. 2, pp. 437–470, 1969.
2. M. Ito, "Neurophysiological aspects of the cerebellar motor control system.," International journal of neurology, vol. 7, no. 2, pp. 162–76, 1970.
3. J. S. Albus, "A theory of cerebellar function," Mathematical Biosciences, vol. 10, no. 1-2, pp. 25–61, 1971.
4. F. Rosenblatt, "The perceptron: A probabilistic model for information storage and organization in the brain," Psychological Review, vol. 65, no. 6, pp. 386–408, 1958.
5. M. Ito, "Cerebellar Long-Term Depression: Characterization, Signal Transduction, and Functional Roles," Physiological Reviews, vol. 81, no. 3, pp. 1143–1195, 2001.
6. M. Ito and M. Kano, "Long-lasting depression of parallel fiber-Purkinje cell transmission induced by conjunctive stimulation of parallel fibers and climbing fibers in the cerebellar cortex," Neuroscience Letters, vol. 33, no. 3, pp. 253–258, 1982.
7. E. D'Angelo, The Organization of Plasticity in the Cerebellar Cortex: From Synapses to Control, vol. 210 of Progress in Brain Research, pp. 31–58. Elsevier, 2014.

8. M. Fujita, "A theory of cerebellar cortex and adaptive motor control based on two types of universal function approximation capability," Neural Networks, vol. 75, pp. 173–196, 2016.
9. J. L. Raymond and J. F. Medina, "Computational Principles of Supervised Learning in the Cerebellum," Annual Review of Neuroscience, vol. 41, no. 1, pp. 233–253, 2018.
10. T. Yamazaki, "Evolution of the marr-albus-ito model." in *Cerebellum as a CNS Hub*, H. Mizusawa, S. Kakei, Eds. (Springer, 2021), pp. 239–255, 2021.
11. J. F. Medina, K. S. Garcia, W. L. Nores, N. M. Taylor, and M. D. Mauk, "Timing Mechanisms in the Cerebellum: Testing Predictions of a Large-Scale Computer Simulation," The Journal of Neuroscience, vol. 20, no. 14, pp. 5516–5525, 2000.
12. M. Gosui and T. Yamazaki, "Real-World-Time Simulation of Memory Consolidation in a Large-Scale Cerebellar Model," Frontiers in Neuroanatomy, vol. 10, p. 21, 2016.
13. M. Hausknecht, W.-K. Li, M. Mauk, and P. Stone, "Machine Learning Capabilities of a Simulated Cerebellum," IEEE Transactions on Neural Networks and Learning Systems, vol. 28, no. 3, pp. 510–522, 2017.
14. I. Abadía, F. Naveros, E. Ros, R. R. Carrillo, and N. R. Luque, "A cerebellar-based solution to the nondeterministic time delay problem in robotic control," Science Robotics, vol. 6, no. 58, p. eabf2756, 2021.
15. R. Kuriyama, C. Casellato, E. D'Angelo, and T. Yamazaki, "Real-Time Simulation of a Cerebellar Scaffold Model on Graphics Processing Units," Frontiers in Cellular Neuroscience, vol. 15, p. 623552, 2021.
16. Y. Shinji, H. Okuno, and Y. Hirata, "Artificial cerebellum on FPGA: realistic real-time cerebellar spiking neural network model capable of real-world adaptive motor control," Frontiers in Neuroscience, vol. 18, p. 1220908, 2024.
17. A. Geminiani, C. Casellato, H.-J. Boele, A. Pedrocchi, C. I. D. Zeeuw, and E. D'Angelo, "Mesoscale simulations predict the role of synergistic cerebellar plasticity during classical eyeblink conditioning," PLOS Computational Biology, vol. 20, no. 4, p. e1011277, 2024.
18. R. S. Sutton and A. G. Barto, Reinforcement learning: an introduction. Adaptive computation and machine learning series, Cambridge, Massachusetts: The MIT Press, second edition ed., 2018.
19. R. A. Swain, A. L. Kerr, and R. F. Thompson, "The Cerebellum: A Neural System for the Study of Reinforcement Learning," Frontiers in Behavioral Neuroscience, vol. 5, p. 8, 2011.
20. T. Yamazaki and W. Lennon, "Revisiting a theory of cerebellar cortex," Neuroscience Research, vol. 148, pp. 1–8, 2019.
21. N. Frémaux, H. Sprekeler, and W. Gerstner, "Reinforcement Learning Using a Continuous Time Actor-Critic Framework with Spiking Neurons," PLoS Computational Biology, vol. 9, no. 4, p. e1003024, 2013.
22. A. W. Moore, Efficient Memory-Based Learning for Robot Control. PhD thesis, University of Cambridge, 1990.
23. D.-A. Jirenhed, F. Bengtsson, and G. Hesslow, "Acquisition, Extinction, and Reacquisition of a Cerebellar Cortical Memory Trace," Journal of Neuroscience, vol. 27, no. 10, pp. 2493–2502, 2007.
24. G. Szapiro and B. Barbour, "Multiple climbing fibers signal to molecular layer interneurons exclusively via glutamate spillover," Nature Neuroscience, vol. 10, no. 6, pp. 735–742, 2007.
25. M. J. Dizon and K. Khodakhah, "The Role of Interneurons in Shaping Purkinje Cell Responses in the Cerebellar Cortex," The Journal of Neuroscience, vol. 31, no. 29, pp. 10463–10473, 2011.
26. A. M. Brown, M. Arancillo, T. Lin, D. R. Catt, J. Zhou, E. P. Lackey, T. L. Stay, Z. Zuo, J. J. White, and R. V. Sillitoe, "Molecular layer interneurons shape the spike activity of cerebellar Purkinje cells," Scientific Reports, vol. 9, no. 1, p. 1742, 2019.
27. J. S. Albus, "A new approach to manipulator control: The cerebellar model articulation controller (CMAC)," Journal of Dynamics Systems, Measurement, and Control, vol. 97, no. 3, pp. 220–227, 1975.
28. M. Ito, The Cerebellum and Neural Control. New York: Raven Press, 1984.
29. K. Doya, "Reinforcement Learning in Continuous Time and Space," Neural Computation, vol. 12, no. 1, pp. 219–245, 2000.
30. H. Jörntell and C.-F. Ekerot, "Reciprocal Bidirectional Plasticity of Parallel Fiber Receptive Fields in Cerebellar Purkinje Cells and Their Afferent Interneurons," Neuron, vol. 34, no. 5, pp. 797–806, 2002.
31. H. Jörntell and C.-F. Ekerot, "Receptive Field Plasticity Profoundly Alters the Cutaneous Parallel Fiber Synaptic Input to Cerebellar Interneurons In Vivo," The Journal of Neuroscience, vol. 23, no. 29, pp. 9620–9631, 2003.
32. D.-A. Jirenhed, F. Bengtsson, and H. Jörntell, "Parallel fiber and climbing fiber responses in rat cerebellar cortical neurons in vivo," Frontiers in Systems Neuroscience, vol. 7, p. 16, 2013.
33. S. J. Liu, "Stellate cells: Synaptic processing and plasticity." in *Handbook of the Cerebellum and Cerebellar Disorders*, M. Manto, D. L. Gruol, J. D. Schmahmann, N. Koibuchi, F. Rossi, Eds. (Springer, 2013), vol. 2, pp. 809–828, 2013.
34. M. J. Rowan, A. Bonnan, K. Zhang, S. B. Amat, C. Kikuchi, H. Taniguchi, G. J. Augustine, and J. M. Christie, "Graded Control of Climbing-Fiber-Mediated Plasticity and Learning by Inhibition in the Cerebellum," Neuron, vol. 99, no. 5, pp. 999–1015.e6, 2018.
35. A. Bonnan, M. M. J. Rowan, C. A. Baker, M. M. Bolton, and J. M. Christie, "Autonomous Purkinje cell activation instructs bidirectional motor learning through evoked dendritic calcium signaling," Nature Communications, vol. 12, no. 1, p. 2153, 2021.
36. N. T. Silva, J. Ramírez-Buriticá, D. L. Pritchett, and M. R. Carey, "Climbing fibers provide essential instructive signals for associative learning," Nature Neuroscience, vol. 27, pp. 940–951, 2024.
37. M. Towers, J. K. Terry, A. Kwiatkowski, J. U. Balis, G. d. Cola, T. Deleu, M. Goulão, A. Kallinteris, A. KG, M. Krimmel, R. Perez-Vicente, A. Pierré, S. Schulhoff, J. J. Tai, A. T. J. Shen, and O. G. Younis, "Gymnasium," Mar. 2023.
38. M. ten Brinke, H.-J. Boele, J. Spanke, J.-W. Potters, K. Kornysheva, P. Wulff, A. IJpelaar, S. Koekkoek, and C. De Zeeuw, "Evolving Models of Pavlovian Conditioning: Cerebellar Cortical Dynamics in Awake Behaving Mice," Cell Reports, vol. 13, no. 9, pp. 1977–1988, 2015.
39. M. Ma, G. L. Futia, F. M. S. d. Souza, B. N. Ozbay, I. Llano, E. A. Gibson, and D. Restrepo, "Molecular layer interneurons in the cerebellum encode for valence in associative learning," Nature Communications, vol. 11, no. 1, p. 4217, 2020.
40. H.-J. Boele, S. Peter, M. M. T. Brinke, L. Verdonschot, A. C. H. IJpelaar, D. Rizopoulos, Z. Gao, S. K. E. Koekkoek, and C. I. D. Zeeuw, "Impact of parallel fiber to Purkinje cell long-term depression is unmasked in absence of inhibitory input," Science Advances, vol. 4, no. 10, p. eaas9426, 2018.
41. A. Antonietti, A. Geminiani, E. Negri, E. D'Angelo, C. Casellato, and A. Pedrocchi, "Brain-inspired spiking neural network controller for a neurorobotic whisker system," Frontiers in Neurorobotics, vol. 16, 2022.
42. F. Johansson, D.-A. Jirenhed, A. Rasmussen, R. Zucca, and G. Hesslow, "Memory trace and timing mechanism localized to cerebellar Purkinje cells," Proceedings of the National Academy of Sciences, vol. 111, no. 41, pp. 14930–14934, 2014.
43. P. Wulff, M. Schonewille, M. Renzi, L. Viltono, M. Sassoè-Pognetto, A. Badura, Z. Gao, F. E. Hoebeek, S. v. Dorp, W. Wisden, M. Farrant, and C. I. D. Zeeuw, "Synaptic inhibition of Purkinje cells mediates consolidation of vestibulo-cerebellar motor learning," Nature Neuroscience, vol. 12, no. 8, pp. 1042–1049, 2009.
44. K. Lee, P. Mathews, A. Reeves, K. Choe, S. Jami, R. Serrano, and T. Otis, "Circuit Mechanisms Underlying Motor Memory Formation in the Cerebellum," Neuron, vol. 86, no. 2, pp. 529–540, 2015.
45. T. Hirano, Y. Yamazaki, and Y. Nakamura, "LTD, RP, and Motor Learning," The Cerebellum, vol. 15, no. 1, pp. 51–53, 2016.
46. M. Jelitai, P. Puggioni, T. Ishikawa, A. Rinaldi, and I. Duguid, "Dendritic excitation–inhibition balance shapes cerebellar output during motor behaviour," Nature Communications, vol. 7, no. 1, p. 13722, 2016.
47. S. Ohmae and J. F. Medina, "Climbing fibers encode a temporal-difference prediction error during cerebellar learning in mice," Nature Neuroscience, vol. 18, no. 12, pp. 1798–1803, 2015.
48. W. Heffley, E. Y. Song, Z. Xu, B. N. Taylor, M. A. Hughes, A. McKinney, M. Joshua, and C. Hull, "Coordinated cerebellar climbing fiber activity signals learned sensorimotor predictions," Nature Neuroscience, vol. 21, no. 10, pp. 1431–1441, 2018.
49. W. Heffley and C. Hull, "Classical conditioning drives learned reward prediction signals in climbing fibers across the lateral cerebellum," eLife, vol. 8, p. e46764, 2019.
50. D. Kostadinov, M. Beau, M. Blanco-Pozo, and M. Häusser, "Predictive and reactive reward signals conveyed by climbing fiber inputs to cerebellar Purkinje cells," Nature Neuroscience, vol. 22, no. 6, pp. 950–962, 2019.
51. K. Ikezoe, N. Hidaka, S. Manita, M. Murakami, S. Tsutsumi, Y. Isomura, M. Kano, and K. Kitamura, "Cerebellar climbing fibers multiplex movement and reward signals during a voluntary movement task in mice," Communications Biology, vol. 6, no. 1, p. 924, 2023.
52. H. Hoang, S. Tsutsumi, M. Matsuzaki, M. Kano, K. Toyama, K. Kitamura, and M. Kawato, "Negative reward-prediction errors of climbing fiber inputs for cerebellar reinforcement learning algorithm." bioRxiv [Preprint], 2023. (accessed 9 September 2023).
53. F. Najafi and J. F. Medina, "Beyond "all-or-nothing" climbing fibers: graded representation of teaching signals in Purkinje cells," Frontiers in Neural Circuits, vol. 7, p. 115, 2013.
54. Y. Zang and E. De Schutter, "Climbing Fibers Provide Graded Error Signals in Cerebellar Learning," Frontiers in Systems Neuroscience, vol. 13, p. 46, 2019.
55. J. Morimoto and K. Doya, "Acquisition of stand-up behavior by a real robot using hierarchical reinforcement learning," Robotics and Autonomous Systems, vol. 36, no. 1, pp. 37–51, 2001.
56. T. D. Kulkarni, K. R. Narasimhan, A. Saeedi, and J. B. Tenenbaum, Hierarchical deep reinforcement learning: integrating temporal abstraction and intrinsic motivation, p. 3682–3690. NIPS'16, Red Hook, NY, USA: Curran Associates Inc., 2016.
57. C. Rössert, S. Solinas, E. D'Angelo, P. Dean, and J. Porrill, "Model cerebellar granule cells can faithfully transmit modulated firing rate signals," Frontiers in Cellular Neuroscience, vol. 8, 2014.
58. C. I. D. Zeeuw, C. C. Hoogenraad, S. Koekkoek, T. J. Ruigrok, N. Galjart, and J. I. Simpson, "Microcircuitry and function of the inferior olive," Trends in Neurosciences, vol. 21, no. 9, pp. 391–400, 1998.

Supplementary Information

A. Relationship between our weight-update rule and the previous weight-update rule

The weight-update rule and eligibility trace in our model are defined as follows:

$$\frac{d\Delta w^n_{i,j}(t)}{dt} = R(t)z_{i,j}(t) + \frac{\tau_r - \tau_z}{\tau_r \tau_z} V(t) z_{i,j}(t) - V(t) f_{i,j}(t), \tag{S1}$$

$$\frac{dz_{i,j}(t)}{dt} = -\frac{z_{i,j}(t)}{\tau_z} + f_{i,j}(t). \tag{S2}$$

By incorporating the definitions, $\frac{d\Delta w_{i,j}(t)}{dt}$ can be transformed as follows (Equations):

$$\frac{d\Delta w^n_{i,j}(t)}{dt} = R(t)z_{i,j}(t) + \frac{\tau_r - \tau_z}{\tau_r \tau_z} V(t) z_{i,j}(t) - V(t) f_{i,j}(t) \tag{S3}$$

$$= R(t)z_{i,j}(t) - \frac{1}{\tau_r} V(t) z_{i,j}(t) + \frac{1}{\tau_z} V(t) z_{i,j}(t) - V(t) f_{i,j}(t) \tag{S4}$$

$$= R(t)z_{i,j}(t) - \frac{1}{\tau_r} V(t) z_{i,j}(t) + V(t) \left\{ \frac{1}{\tau_z} z_{i,j}(t) - f_{i,j}(t) \right\} \tag{S5}$$

$$= R(t)z_{i,j}(t) - \frac{1}{\tau_r} V(t) z_{i,j}(t) - V(t) \frac{dz_{i,j}(t)}{dt}. \tag{S6}$$

The total weight change over an episode $\Delta w_{i,j}(T_{\text{end}})$ is expressed as follows:

$$\Delta w_{i,j}(T_{\text{end}}) = \Delta w_{i,j}(0) + \int_0^{T_{\text{end}}} \frac{d\Delta w^n_{i,j}(t)}{dt} dt \tag{S7}$$

$$= \Delta w_{i,j}(0) + \int_0^{T_{\text{end}}} \left\{ R(t)z_{i,j}(t) - \frac{1}{\tau_r} V(t) z_{i,j}(t) - V(t) \frac{dz_{i,j}(t)}{dt} \right\} dt \tag{S8}$$

$$= \Delta w_{i,j}(0) - [V(t)z(t)]_0^{T_{\text{end}}} + \int_0^{T_{\text{end}}} \left\{ R(t)z_{i,j}(t) - \frac{1}{\tau_r} V(t) z_{i,j}(t) + \dot{V}(t) z_{i,j}(t) \right\} dt \tag{S9}$$

$$= \Delta w_{i,j}(0) - [V(t)z(t)]_0^{T_{\text{end}}} + \int_0^{T_{\text{end}}} \left\{ R(t) - \frac{1}{\tau_r} V(t) + \dot{V}(t) \right\} z_{i,j}(t) dt. \tag{S10}$$

Here, by assuming $\Delta w_{i,j}(0) = 0$, $z(0) = 0$, and $V(T_{\text{end}}) = 0$, the following equation is obtained:

$$\Delta w_{i,j}(T_{\text{end}}) = \int_0^{T_{\text{end}}} \left\{ R(t) - \frac{1}{\tau_r} V(t) + \dot{V}(t) \right\} z_{i,j}(t) dt. \tag{S11}$$

The expression within the curly brackets is equivalent to the continuous time TD error defined by (1), which was used by the previous spiking actor-critic framework (2).

B. Simulation of Linear Track Task

In order to evaluate whether our critic could approximate $-V(t)$, we conducted a simulation of the linear track task (2). In this task, the environment is a narrow rectangular plane (Figure S1A), on which the agent consistently moves from the leftmost side (start) to the rightmost side (goal) at a fixed speed. Thus, the agent always reaches the goal at the same time, and receives the same negative reward across all episodes. Once the time constant of the reward discount τ_r is determined, the theoretical value can be solved (Figure S1A).

In this task, the plane size was 40×1, and a fixed starting position was $s = 5$. The agent ran with a fixed speed of 0.05 per ms toward a goal area $(s > 39)$. When the agent reached the goal, a negative reward of -5 was given. The parameters of the state value function, which defined in Equation 2, were $\nu = 300$ and $V_0 = 12.4$. The learning rate and the initial weight of PF–SC were 0.04, and 0.05. The reward discount time constant τ_r, the decay time constant of eligibility trace τ_z, and the decay time constant of the window function τ_κ were set to 100 ms, 20 ms, and 20 ms, respectively.

On the first run, the SC represented a noisy, around zero value function (Figure S1B). Over 100 trials, the SCs learned the value function progressively. Then the SC activity increased exponentially toward the goal, as if it expected the given negative reward (Figure S1B). The averaged value in 20 late trials, as plotted with a black line, nicely matches the theoretical value function. These results were consistent with previous research (2), and suggest that our weight-update rules were successful and that the SCs were able to approximate the additive inverse of the state value $-V(t)$.

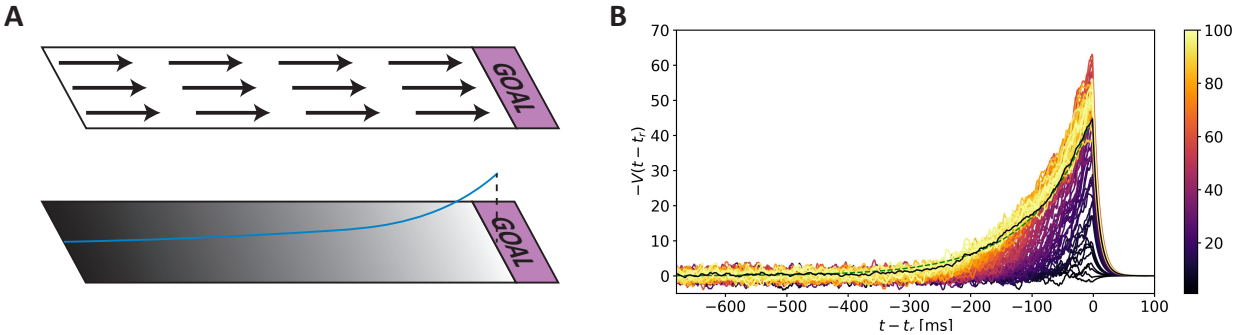

Figure S1. Simulation of Linear Track Task. (**A**) Environment of the linear track task. Top panel shows the environment and fixed policy. The agent starts at the left edge, and moves straight to the goal shown as a pink area. Bottom panel shows theoretical state value. When the agent follows a fixed policy, the theoretical state value is uniquely determined. (**B**) Additive inverse of the state value learned by the critic. Each colored trace shows the value represented by the stellate cell neurons activity against time in 100 first simulation trials (from dark purple in the 1st trial to light yellow in the 500th trial). The plots are aligned with the timing of the negative reward delivery t_r. The black line shows the average of over last 10 trials. The green dashed line shows the theoretical value function.

Bibliography

1. Doya, K. (2000). Reinforcement Learning in Continuous Time and Space. Neural Computation, 12(1):219–245. doi: 10.1162/089976600300015961.
2. Frémaux, N., Sprekeler, H., and Gerstner, W. (2013). Reinforcement Learning Using a Continuous Time Actor-Critic Framework with Spiking Neurons. PLoS Computational Biology, 9(4): e1003024. doi: 10.1371/journal.pcbi.1003024.