Population coding in the cerebellum and its implications for learning from error

Reza Shadmehr

Department of Biomedical Engineering
Johns Hopkins School of Medicine, Baltimore MD

Correspondence: Reza Shadmehr, 410 Traylor Building, Johns Hopkins School of Medicine, 720 Rutland Ave., Baltimore, MD 21205, USA. Email: shadmehr@jhu.edu.

Acknowledgements: This work was supported by grants from the Office of Naval Research (N00014-15-1-2312), the National Institutes of Health (R01NS078311), and the National Science Foundation (CNS-1714623).

Abstract

Artificial intelligence provides algorithms that describe how neural networks learn from error, but can the algorithms help uncover the neuronal wiring of the brain? Here, I consider the cerebellum, a major learning site that resembles a feedforward network of neurons. In the cerebellum, Purkinje cells (P-cells) organize in populations that converge onto neurons in a deep cerebellar nucleus (DCN). The outputs produced by DCN neurons are predictions that are compared to the actual observations in the inferior olive, resulting in prediction errors that are fed back strongly to the P-cells, but rather weakly to the neurons in the nucleus. Furthermore, unlike an artificial neural network, a P-cell has a very limited view of the error space: it is aware of only those errors that it can sense via its single climbing fiber. The strength disparity in the error signals, and the limited view of the error space, are critical clues that suggest how the P-cells and their targets in the nucleus should be wired into populations. As a result, the fundamental unit of computation in the cerebellum is not a single neuron, but a group of neurons that share a single teacher. To learn efficiently, the error signal organizes the P-cells, which in turn act as surrogate teachers for the neurons that they project to in the nucleus. The result is a population coding that may be responsible for a number of remarkable features of behavior, including multiple timescales, protection from erasure, and spontaneous recovery.

Ever tried. Ever failed. No Matter. Try again. Fail again. Fail better.
Samuel Beckett

Introduction

During electrophysiological recording from Purkinje cells (P-cells), listening to the sound of spikes as they are reported by an electrode, one cannot help but be impressed by the complex spike. Whereas the simple spike that the P-cell produces appears ordinary and common, a rain drop falling on the roof, the complex spike is more like lightening, a thunderous event that makes the P-cell pause its production of simple spikes. Indeed, after generating a complex spike, the P-cell requires 10-20 ms of recovery before it resumes production of simple spikes (Thach, 1967).

The significance of the complex spike as an event in the life of a P-cell is illustrated by the fact that persistent stimulation of climbing fibers, the sole source of complex spikes, can lead to excitotoxic damage of P-cells (Slemmer et al., 2005). Indeed, sustained increase or decrease in the rate of complex spikes above or below baseline damages the P-cells (O'Hearn and Molliver, 1997). Thus, survival of a P-cell depends on its ability to tame the production of complex spikes, regulating it to around baseline (De Schutter, 1995; Mauk and Donegan, 1997).

Complex spikes arise from climbing fibers that originate from cells in the inferior olive (Eccles et al., 1966). The olive cells receive inhibitory projections from deep cerebellar nucleus (DCN) neurons, cells that carry the predictions that the cerebellum is making (Kim et al., 1998; Rasmussen et al., 2015; Rasmussen and Hesslow, 2014). The olive cells receive excitatory inputs from other regions of the brain (Kralj-Hans et al., 2007), which presumably allows an olive cell to compare the predictions of the cerebellum to the actual observations (Soetedjo et al., 2009). Indeed, a complex spike is often generated after unexpected occurrence of a broad class of sensory events (Ju et al., 2019). This implies that as a first approximation, the signal in the climbing fiber is a prediction error (Kitazawa et al., 1998), that is, the difference between what was observed and what was predicted.

Prediction errors are fundamental to learning in artificial neural networks. In such networks, the error signal is conveyed to every unit that participates in computing the prediction. However, in the cerebellum, the output layer neurons in the DCN receive a rather weak signal from the olive (Lu et al., 2016), a signal that weakens even further as the animal reaches adulthood (Najac and Raman, 2018). In comparison, the middle layer neurons (the P-cells) receive an extraordinarily strong error signal, but their signal is via only a single climbing fiber. The weak error signal to the DCN raises the question of how the output layer neurons learn from their prediction errors. The single climbing fiber to a P-cell raises the concern that the teaching signal is not a fair reflection of the entire error space, but instead biased to provide only a limited view.

To illustrate an implication of these ideas, suppose you are playing basketball and shoot from the free throw line, but your ball misses the basket. If you are looking at the basket as the ball passes to one side, say to the right, the event will produce an increase in probability of complex spikes among P-cells that prefer rightward prediction errors in visual sensory space (Herzfeld et al., 2015; Soetedjo et al., 2008; Soetedjo and Fuchs, 2006). That same error will suppress probability of complex spikes in P-cells that prefer leftward prediction errors. However, the rightward error will have no consequence for the complex spikes of P-cells that prefer upward or downward errors. As a result, each P-cell receives a

highly personalized view of the error space: some care about rightward prediction errors and some care about leftward prediction errors, but none care about all parts of the error space.

Why is the error signal transmitted weakly to the nucleus, but strongly to the P-cells? Why is the error signal to each P-cell biased toward a specific part of the error space? To answer these questions, we will employ elementary mathematics from artificial intelligence to make conjectures, and then interpret the physiological and anatomical data in the computational framework.

We will suggest that the very limited distribution of prediction errors from a single olive cell to a handful of P-cells organizes the P-cells into populations that share the same teacher (De Zeeuw et al., 2011; Heck et al., 2013): P-cells that learn from the same teacher likely project together as a population to a single neuron in the DCN, thus conveying their error information to the output neuron that had the responsibility for generating the prediction that produced this error (Chaumont et al., 2013; Tang et al., 2019). In this framework, each olivary cell is the teacher to a handful of P-cells. Those P-cells, in turn, are the teachers to a handful of DCN neurons (Medina and Mauk, 1999).

Furthermore, because an error may be preferred by some P-cells, increasing their probability of complex spikes, while the same error is anti-preferred for other P-cells, decreasing their probability of complex spikes, the differing sensitivities to error, as reflected in the differing rates of plasticity associated with the presence vs. absence of a complex spike (Herzfeld et al., 2018; Yang and Lisberger, 2014a), are likely to generate multiple timescales of behavior during learning, some fast, others slow.

Finally, the personalized view of error leads to a remarkable feature of behavior: protection from erasure. That is, behavior becomes easier to learn than to unlearn. Rather than countering plasticity in the P-cells that learned from error, reversal of error engages learning in a separate group of P-cells. As a result, limiting the P-cell's view of the error space may be responsible for a fundamental feature of memory: learning followed by extinction returns behavior to baseline, but following passage of time, behavior reverts back toward that which was learned during initial training, thus producing spontaneous recovery (Criscimagna-Hemminger and Shadmehr, 2008; Kojima et al., 2004; Sarwary et al., 2018; Smith et al., 2006).

Prediction errors of the cerebellum

As a first approximation, the anatomy of the cerebellum resembles a feedforward network (Fig. 1A) (Raymond and Medina, 2018). Like an artificial neural network, inputs (mossy fibers) bring information to the first layer of neurons (granule cells), which distribute them to the second layer (P-cells), which in turn project to the output layer (DCN neurons) (Fig. 1B). To be sure, this resemblance is an approximation. The input from the granule cells to the P-cells is direct, as well as indirect (via basket cells). P-cells send a few collateral axons to neighboring P-cells (Witter et al., 2016), and at least in some lobules, P-cells also send a few collaterals to granule cells (Guo et al., 2016). As a result, there is a degree of synchrony among neighboring P-cells (Sedaghat-Nejad et al., 2019), a synchrony that remains even when chemical synapses are inactivated (Han et al., 2018). Despite these complications, a feedforward network is a useful approximation, particularly from the P-cell layer to the output layer, which is the focus of our analysis.

In a typical artificial neural network, all cells in a given layer connect to all cells in the next layer. However, in the cerebellum, about 50 P-cells converge upon a single DCN neuron (Person and Raman, 2012a). That is, P-cells organize into small populations. What determines the membership in this

population? That is, what property must a P-cell have to make it eligible to be a member of a population that projects to a single DCN neuron?

In our network, the activities of the neurons in the output layer (DCN neurons) represent the predictions of the cerebellum. Let us label the activation, i.e. firing rate, of each outward projecting DCN neuron as $a_i^{(d)}$, where the superscript d refers to DCN neurons. Some of the projecting DCN neurons are GABA-ergic and send their axons to the inferior olive. We label activity of these DCN neurons as $a_i^{(d-)}$. Another subset is non-GABA-ergic and send their axons to various other regions of the brain, e.g., brainstem, superior colliculus, red nucleus, thalamus, etc. It is with these non-GABA-ergic neurons that the cerebellum influences behavior. We label activity of the non-olive projecting DCN neurons as $a_i^{(d+)}$. Thus, in the output layer the superscript d includes GABA-ergic as well as non-GABA-ergic projecting DCN neurons: $d \in \left\{d^-, d^+\right\}$.

We assume that a cell in the olive compares the prediction $a_i^{(d-)}$ that it receives from a DCN neuron with the observed event a_i^*. For example, if the prediction is associated with visual information that should be present following conclusion of a saccadic eye movement, the actually observed visual event a_i^* is conveyed to the olive from the superior colliculus (Kojima and Soetedjo, 2018, 2017). The activity in the inferior olive neuron that projects back to the DCN and the rest of the cerebellum is a prediction error, i.e., the difference between what was predicted and what was observed. Notably, this error is only associated with activity of GABA-ergic DCN neurons $a_i^{(d-)}$. Let us label the prediction error as $e_i^{(d-)}$:

$$e_i^{(d-)} = a_i^* - a_i^{(d-)} \tag{1}$$

Thus, the error signal depends on the activity $a_i^{(d-)}$ of GABA-ergic olive-projecting DCN neurons, but not directly on the activity $a_i^{(d+)}$ of the non-GABA-ergic, non-olive projecting DCN neurons. And so we arrive at our first puzzle:

Puzzle 1. *Despite the fact that non-GABA-ergic DCN neurons are responsible for conveying predictions of the cerebellum to the rest of the brain, the errors associated with their predictions is not directly a part of the olivary signal back to the cerebellum. How do non-GABA-ergic DCN neurons receive information regarding their prediction error?*

The consequence of computing error information based only on the activity of GABA-ergic DCN neurons, and not the other output neurons, can be illustrated by an experiment by Kim et al. (1998). The authors presented rabbits a tone followed by an airpuff, teaching them to blink just before arrival of the airpuff. In naïve animals the airpuff alone produced complex spikes in P-cells. Early in training with the tone+airpuff trials the P-cells continued to produce complex spikes following the airpuff, but late in training tone+airpuff no longer produced complex spikes: as performance improved and errors were reduced, so were the complex spikes. Now the authors took the well trained animals and injected a GABA inhibitor into the inferior olive. This effectively eliminated the efficacy of the DCN input to the olive. However, blocking the DCN input re-introduced the complex spikes during tone+airpuff trials. That

is, despite the fact that the animal continued to blink predictively, the olive nevertheless signaled prediction errors to the cerebellum. This suggests that even when non-GABA-ergic DCN neurons are correctly predicting an output $a_i^{(d+)}$, if the inhibitory input to the olive is suppressed, the error signal to the cerebellum returns.

To consider this puzzle, let us further develop the mathematics associated with the problem of learning. A generic DCN neuron's activity $a_i^{(d)}$ depends partly on the inputs that it receives, and partly on its internal state. The inputs are the weighted activities of neurons that project to it. This includes inputs from P-cells (via inhibitory synapses), and inputs from mossy fibers (via excitatory synapses), as shown in Fig. 1C. We approximate the inputs to a generic DCN neuron via the following equation:

$$z_i^{(d)} = \sum_{j,k} w_{i,j}^{(d)} a_j^{(p)} + w_{i,k}^{(d)} a_k^{(m)} + b_i \qquad (2)$$

In the above equation, $a_j^{(p)}$ is the activity of P-cells that project to a given DCN neuron i, $a_k^{(m)}$ is the activity conveyed via mossy fibers to that neuron, and b_i is the internal bias of that neuron (making it possible for the neuron to fire despite have zero or negative inputs).

Notice that we have omitted the influence of projections to the DCN neuron from the olive. This is because the axons that bring olivary input to the cerebellum fire at rates that are two orders of magnitude smaller than P-cells and mossy fibers. Furthermore, in the adult animal, the excitatory post-synaptic currents (EPSCs) that are produced in a typical DCN neuron following activation of olivary axons are remarkably small (Lu et al., 2016). For example, activation of olivary axons can produce an EPSC event with an amplitude of 2 pA in a P-cell, 0.4 pA in the DCN neuron of a juvenile mouse (Najac and Raman, 2018), and only 8×10^{-3} pA in the DCN neuron of an adult mouse (Lu et al., 2016). Therefore, in the adult animal the input to the DCN from the olive is quite weak. This fact simplifies our equation, but also introduces a second puzzle:

Puzzle 2. *Given that the olive's input to the DCN is weak, how can a DCN neuron learn from its prediction error?*

To summarize, prediction errors are needed to teach the output layer of our network, the layer that represents the DCN neurons. However, projections from the olive to the DCN appear to represent the error that was made by only a subset of neurons: GABA-ergic olive projecting DCN neurons, but not the non-GABA-ergic neurons that project to the rest of the brain. Furthermore, the axons that project from the olive to the DCN, i.e., the axons that carry the error signal, appear to have weak synapses in adulthood, and are essentially silent as compared to mossy fiber and P-cell axons that converge on the same neurons. This raises the question of how can the olive, our sole source of prediction error, be an effective teacher for the output layer of our network. These puzzles will shape our discussions regarding organization of the P-cells and their downstream DCN neurons.

The problem of teaching a deep cerebellar nucleus neuron

The purpose of a teaching signal is to minimize an objective function, typically a sum of squared errors. For our network, the sum of errors is a function of all output neurons, i.e., both olive-projecting and non-olive-projecting DCN neurons:

$$J = \sum_i \left(e_i^{(d)} \right)^2 \tag{3}$$

To minimize this objective function, we need to use the error associated with each output neuron's activity to change its predictions. Thus, we need to change the activity of both the GABA-ergic DCN neurons $a_i^{(d-)}$, and the non-GABA-ergic DCN neurons $a_i^{(d+)}$. This is done by changing the weights associated with the inputs that these neurons receive, as well as their internal biases.

The activity (firing rate) of a generic neuron in our network is related to its inputs and internal biases via a non-linearity, for example a sigmoidal function:

$$a_i^{(d)} = \sigma\left(z_i^{(d)} \right) \tag{4}$$

To change the synaptic weight of a given input to a neuron in the output layer, we compute the gradient of the objective function with respect to that weight. For example, the magnitude of the weight change for a GABA-ergic DCN neuron is negatively proportional to the following gradient:

$$\begin{aligned} \frac{dJ}{dw_{i,j}^{(d-)}} &= \frac{dJ}{da_i^{(d-)}} \frac{da_i^{(d-)}}{dz_i^{(d-)}} \frac{dz_i^{(d-)}}{dw_{i,j}^{(d-)}} \\ &= -2 \frac{da_i^{(d-)}}{dz_i^{(d-)}} \frac{dz_i^{(d-)}}{dw_{i,j}^{(d-)}} e_i^{(d-)} \end{aligned} \tag{5}$$

The above expression has a simple meaning: the gradient of a DCN neuron's activation z_i with respect to its mossy fiber synaptic weights is positive (because those inputs are excitatory), and thus positive errors (e.g., olivary output above baseline) should increase the mossy fiber synaptic weights. In comparison, the gradient of z_i with respect to P-cell synaptic weights is negative (because those inputs are inhibitory), and thus positive errors should decrease those synaptic weights. Similarly, to change the internal bias of the DCN neuron, we compute the gradient with respect to that bias:

$$\begin{aligned} \frac{dJ}{db_i^{(d-)}} &= \frac{dJ}{da_i^{(d-)}} \frac{da_i^{(d-)}}{dz_i^{(d-)}} \frac{dz_i^{(d-)}}{db_i^{(d-)}} \\ &= -2 \frac{da_i^{(d-)}}{dz_i^{(d-)}} e_i^{(d-)} \end{aligned} \tag{6}$$

Aside from providing general rules for learning in the DCN, the above expressions state that the gradients of our objective function with respect to both weights and internal biases of a given DCN neuron are proportional to the prediction errors of that neuron. This means that for a DCN neuron that projects to the olive, it must receive error associated with its own predictions, not predictions of some other DCN neurons. This leads to our first conjecture:

Conjecture 1. *If a DCN neuron projects to one or more inferior olive neurons, it should receive error feedback from precisely the same olive neurons*.

Indeed, olivary projections to the DCN are strictly reciprocal: if a region in the olive projects to a region in the DCN, that DCN region also projects back to that specific region of the olive (Ruigrok and Voogd, 2000). However, even if conjecture C1 were true, we still have DCN neurons that do not project

to the olive, but rather project to some other region of the brain (puzzle P1). In principle, because these non-GABA-ergic neurons do not project to the olive, we have no error signal with which to teach them.

There are a number of ways to consider this problem. First, it is possible that as the non-GABA-ergic DNC neurons project to their destination, they synapse on GABA-ergic neurons, and those GABA-ergic neurons then project back to the olive. In this way, the olivary cell would indirectly receive the predictions of the non-GABA-ergic DCN neuron. However, this introduces a delay in communication, and importantly, imposes a timing difference in the comparison of predictions with observations for some DCN neurons (non-GABA-ergic) with respect to others (GABA-ergic). Given that timing of complex spikes depends on the precise internal state of olive neurons, which oscillate at 4-5 Hz (Khosrovani et al., 2007), and plasticity in P-cells depends on timing of the complex spikes (Herzfeld et al., 2018; Suvrathan et al., 2016), a timing difference may not be a good solution.

Another way to consider this problem is to note that a single P-cell projects to only 4-5 individual DCN neurons (Person and Raman, 2012b), and critically, this small group of DCN neurons includes both GABA-ergic, as well as non-GABA-ergic neurons (Teune et al., 1998). Therefore, for each GABA-ergic DCN neuron that projects to the olive with activity $a_i^{(d-)}$, there are one or two non-GABA-ergic DCN neurons that receive input from the same P-cell but project elsewhere with activity $a_i^{(d+)}$. In order for the non-olive projecting DCN neuron $a_i^{(d+)}$ to learn from its prediction error, it can pair its activity with a "sister" DCN neuron that projects to the olive $a_i^{(d-)}$. The two sister DCN neurons would have to receive inputs from the same P-cells and mossy fibers so that their activities are similar (Fig. 1C).

If such a relationship existed between a pair of non-GABA-ergic and GABA-ergic DCN neurons such that $da_i^{(d+)} \Big/ w_{i,j}^{(d+)} = da_i^{(d-)} \Big/ dw_{i,j}^{(d+)}$, then we can compute the gradient of the objective function with respect to the weights of the non-olive projecting DCN neurons:

$$\begin{aligned} \frac{dJ}{dw_{i,j}^{(d+)}} &= \frac{dJ}{da_i^{(d-)}} \frac{da_i^{(d+)}}{dz_i^{(d+)}} \frac{dz_i^{(d+)}}{dw_{i,j}^{(d+)}} \\ &= -2 \frac{da_i^{(d+)}}{dz_i^{(d+)}} \frac{dz_i^{(d+)}}{dw_{i,j}^{(d+)}} e_i^{(d-)} \end{aligned} \tag{7}$$

The above expression implies that a non-olive-projecting DCN neuron should share the error signal with its "sister" olive-projecting DCN neuron. From this equation we arrive at our second conjecture:

Conjecture 2. *An error signal conveyed by an olive neuron must result in plasticity in at least one olive projecting GABA-ergic neuron, and one non-olive projecting, non-GABA-ergic neuron. These two DCN neurons should be coupled in the sense that their activities should be equal precisely at the time of prediction*.

At this writing, we simply do not know how the non-GABA-ergic DCN neurons receive their error information. What's more, even if conjectures C1 and C2 were true, we still have puzzle P2: while learning depends on prediction error as conveyed back to the cerebellum from the olive $e_i^{(d-)}$, in adulthood the olivary input to the DCN is weak. How can the DCN learn from its errors when it does not receive an effective signal from the olive?

The trivial answer is to suggest that perhaps DCN neurons have essentially static weights and internal biases and do not learn from their errors. But this is clearly not the case, as evidenced by the work of Mike Mauk and colleagues (Ohyama et al., 2003; Ohyama and Mauk, 2001). For example, in a classical conditioning task in which animals learn to predict the arrival of an aversive stimulus following a tone and close their eyes, training produces plasticity in the DCN, as evidence by the fact that after conclusion of training, disconnection of the P-cell input to the DCN (via GABA antagonists) produces the conditioned behavior, albeit at an earlier time with respect to the tone onset (Medina et al., 2001). Thus, prediction errors produce useful plasticity in the DCN. But given the weak olivary input, how is this possible?

The synapses that a P-cell makes upon a nuclear cell can weaken based on the history of the P-cell's firing (Telgkamp and Raman, 2002). The excitatory inputs to nuclear cells (which come predominantly from mossy fibers) can also undergo plasticity: when a period of excitation (150ms) in the DCN neuron precedes a period of strong inhibition (250ms), excitatory synapses strengthen (McElvain et al., 2010; Person and Raman, 2010; Pugh and Raman, 2008). Thus, the sequential pattern of excitation and inhibition could, in principle, be the teacher for changing the synaptic weights of the inputs onto DCN neurons (De Zeeuw et al., 2011).

However, our question is not whether the various synapses on a DCN neuron exhibit plasticity following patterned inputs from the mossy fibers and P-cells. Rather, in light of Eq. (5) and Eq. (6), puzzle P2 asks a different question: how is plasticity in a DCN neuron guided by the error that it has made in its output?

Transmitting error information to the deep nucleus via P-cells

If we assume that the error that the DCN neuron has made is not conveyed to it directly from the olive, we are left with the possibility that the nucleus neuron becomes "aware" of its error through an indirect pathway. The obvious indirect pathway is from the parent P-cells (Lang et al., 2017; Medina, 2011).

If the inferior olive is electrically (Bengtsson et al., 2011; Hoebeek et al., 2010) stimulated, at 5-10 ms latency there is occasionally a single spike in a DCN neuron. This is likely due to the direct excitatory projections from the olive to the DCN. The occasional single spike is then followed by suppression of activity that lasts around 50 ms (Fig. 2) (which can be followed by a rebound burst of activity). The suppression is a consistent feature in many DCN neurons, including output neurons that do not project to the olive (e.g., non-GABA-ergic neurons that project to the red nucleus) (Hoebeek et al., 2010). If we view olivary stimulation as an artificial means for providing error information to the cerebellum, then that error leaves its impression not via a strong EPSC in the DCN neuron at the time of error, but surprisingly, via a strong suppression of activity during a short period that follows that error.

A series of experiments by Eric Lang and colleagues (Blenkinsop and Lang, 2011; Tang et al., 2019, 2016) provided critical clues regarding how the activity in the olive may produce the suppression in the DCN. Blenkinsop and Lang (2011) inserted an electrode array into the anesthetized rat's cerebellar cortex along with a single electrode into the DCN. This allowed them to simultaneously record from 8-34 P-cells and 1-2 DCN neurons. In order to determine if one of the P-cells was connected to one of the DCN neurons, they identified complex spikes in each P-cell and used that event to align the activity in the DCN neuron. The result was a complex spike-triggered response in the DCN neuron (Fig. 3A). In about 100 cases the spike-triggered response suggested that one of the P-cells was connected to one of

the DCN neurons. In 70 of the 100 putative P-cell DCN neuron pairs the occurrence of a complex spike was followed by a suppression in the DCN, a suppression that lasted from 50-100 ms (Fig. 3A, right top subplot). In 24 pairs, the complex spike in the P-cell was coincident with a brief increase in firing rate of the DCN neuron, and then suppression (Fig. 3A, right bottom subplot). The authors interpreted the small increase as the effect of the weak projections from the olive to the nuclear neuron. Therefore, a single axon from the olive could project both to a nuclear neuron, and to the P-cell that synapsed on that nuclear neuron (Fig. 1C), thus providing the P-cell with the complex spikes that could indirectly affect the state of the downstream DCN neuron.

If a DCN neuron's prediction error is conveyed to it via its parent P-cells, and that DCN neuron also projects to the olive, then we should observe complex spikes in the parent P-cell following an error by the nucleus neuron. There is evidence for this from optogenetic stimulation of P-cells. A 50 ms period of P-cell stimulation resulted in intense production of simple spikes, but when the P-cell stimulation ended, at 80-150 ms latency the P-cells produced a complex spike (Chaumont et al., 2013) (Fig. 3B, top subplot). The interpretation is that the activated P-cells provided intense inhibition of their downstream DCN neurons, prevented those neurons from firing. This suppression of the DCN neuron's activity removed a source of inhibition at the olive. The removal of that inhibition allowed the olivary cells to reach threshold earlier and fire, thus producing a complex spike in the parent P-cell (Fig. 3B, bottom subplot). This is consistent with a closed loop in which P-cells project to a DCN neuron that project to an olive neuron that then projects back to the parent P-cell (Fig. 1C).

Lang and colleagues (Tang et al., 2019) noted that the P-cells that together projected to a single DCN neuron were not randomly distributed, but clustered along the rostra-caudal axis of the cerebellum (Fig. 3A). Furthermore, there was an important feature to the parent P-cells of a DCN neuron: they showed higher than typical levels of complex spike synchrony, suggesting that the olive inputs to the parent P-cells were not random, but closely related.

The fact that olivary projections to the DCN are much stronger in juveniles than adults led Najac and Raman (2018) to suggest that these projections are developmentally guiding the climbing fiber receiving parent P-cell to project to a specific GABA-ergic nucleus neuron that projects back to the same olive neuron that projected to the P-cell (Fig. 1C). Thus, we have potentially the following anatomy (Heck et al., 2013): the olive neuron that computes the prediction error associated with a GABA-ergic DCN neuron conveys that error via a single climbing fiber to a parent P-cell, which then projects to that specific GABA-ergic DCN neuron, which then projects to the olive neuron that computed the prediction error (Fig. 3B). This leads to our next conjecture.

Conjecture 3. *The error associated with a GABA-ergic DCN neuron's output is computed by an olive cell and then sent via climbing fibers to a few P-cells which then project back down to that specific DCN neuron.*

This sentence is the result of roughly three decades of anatomical research (Ruigrok and Voogd, 2000). A single axon from the olive projects to approximately 7 P-cells, all located at similar distances from the midline, most of which are clustered within a single lobule (Sugihara et al., 1999, 2001). Molecular markers, notably the respiratory enzyme aldolase C (commonly known as zebrin II), define about 20 longitudinal compartments in the cerebellar cortex, each receiving inputs from specific regions of the inferior olive. Labeling of P-cells that are located in the same compartment in adjacent lobules, and have similar climbing fiber receptive field characteristics, has shown that the P-cells project to the

same small area in one of the DCNs (Apps and Garwicz, 2000). Tracing studies have identified specific projections from the olive to each compartment in the cerebellar cortex (Apps and Garwicz, 2005). The olive projects to both GABA-ergic and non-GABA-ergic neurons in the DCN (De Zeeuw et al., 1997). There are no molecular markers that have thus far shown compartmentalization of the DCN. However, DCN regions are divided based on tracing of inputs from the olive (Sugihara & Shinoda, 2007). P-cells that are in a given compartment, and thus receive an input form a region in the olive, project to a region in the DCN that also receives inputs from the same olive region (Sugihara et al., 2009). Finally, as we noted in conjecture C1, olivary projections to the DCN are strictly reciprocal: if a region in the olive projects to a region in the DCN, that DCN region also project back to that specific region of the olive (Ruigrok and Voogd, 2000). Thus, we have the anatomical basis of the closed olivo-cortico-nuclear circuit that connects a longitudinal compartment of P-cells to a small subarea in the DCN (Sugihara, 2011).

It is possible that this anatomy serves a critical purpose: providing the DCN neurons with a reliable, life-long teaching signal, despite the fact that the olivary inputs to the DCN weaken with development (Najac and Raman, 2018), and may not play a significant role in adulthood (Lu et al., 2016).

Medina and Mauk (1999), building on ideas put forth by Miles and Lisberger (1981), had conjectured that whereas the climbing fiber input to the P-cell controls plasticity of parallel fiber inputs to that cell, it is the parent P-cell that controls the plasticity of mossy fiber inputs to its downstream DCN neuron. How can the P-cell accomplish this? The results in Figs. 2 and 3 demonstrate that despite the weak synapses of olivary axons upon the nucleus neurons, the prediction error conveyed to the P-cells can in fact have a dramatic effect on the firing rates of nucleus neurons: following a single complex spike, there can be a 50 ms reduction in the firing rate of the DCN neuron. To be sure, some of these results are in the anesthetized state. Do they generalize to the awake, behaving animal?

ten Brinke et al. (2017) trained mice to associate an LED with an air-puff and recorded from the interposed nucleus. As the animals learned, most task-related DCN neuron increased their discharge in response to the LED, and this increase was causally related to production of the eye-blink. The LED onset and the airpuff both tended to produce a single complex spike in a few P-cells (Fig. 4A). However, in DCN neurons, the LED onset and especially the airpuff tended to produce a transient pause in spiking (Fig. 4B). The DCN neurons that showed the airpuff-induced pause also tended to show greater facilitation in their spiking during the LED period, presumably driving the eye-blink response. Thus, during learning of a behavior, some DCN neurons showed a transient pause that appeared related to arrival of error information via a complex spike in their parent P-cells.

We are now in position to offer a partial solution to puzzle P2. De Zeeuw et al. (2011) had conjectured that the complex-spike induced suppression (and the rebound that follows in some DCN neurons) played a key role in control of DCN plasticity. Indeed, a sequential pattern of excitation then inhibition in a DCN neuron can produce LTP in the mossy fiber inputs of the same neuron (Person & Raman, 2010; Pugh & Raman, 2008). Furthermore, in the medial vestibular nucleus (a DCN analogue for some regions of the cerebellar cortex), vestibular nerve synapses (a mossy fiber analogue) on the vestibular nucleus neurons undergo LTP when nerve stimulation coincides with nucleus neuron hyperpolarization (McElvain et al., 2010). Thus, we have a plausible mechanism with which the prediction error of a olive projecting DCN neuron can cause plasticity in that specific DCN neuron: $e_i^{(d-)}$ is detected by an olive cell and sent via a climbing fiber to the parent P-cell whose complex spike

produces a long-lasting suppression of spiking in the DCN neuron. If that suppression is preceded by a period of excitation from mossy fibers, the result is an increase in the weight of the mossy fiber synapses on the DCN neuron. Notably, this learning rule is consistent with requirements of Eq. (5): a positive error results in an increase in the synaptic weight of the excitatory inputs.

According to Eq. (5), the prediction error $e_i^{(d-)}$ must be a teacher for not only the excitatory inputs from the mossy fibers, but also for the inhibitory inputs from the P-cells (as well as the cell's internal bias b_i). A positive prediction error should not only increase the weight of the excitatory synapses, it must also decrease the weight of the inhibitory synapses. Thus, the suppression of spiking in the DCN neuron following the parent P-cell's experience of a complex spike should affect the strength of the P-cell's synaptic input.

Mechanisms of plasticity in the P-cell synapses in the nucleus are poorly understood, but there is some evidence that this plasticity depends on activation patterns that resemble the suppression of spiking induced by a complex spike. Aizenman et al. (1998) found that in slices from the cerebellum of juvenile rats, following 200 ms of hyperpolarization that suppressed the nuclear cell's spiking, there was a rebound depolarization (i.e., a burst of spikes). The amount of this rebound depolarization dictated the direction of plasticity in the P-cell synapse: if the rebound was missing or weak, the synaptic strength was reduced. It is noteworthy that in vivo, generation of a complex spike in the parent P-cell not only leads to a period of suppression of spiking in the DCN neuron, this suppression is occasionally followed by a burst of spiking (Hoebeek et al., 2010). Results of Aizenman et al. (1998) imply that the magnitude of this rebound depolarization influences direction of weight change in the P-cell synapse upon the DCN neuron.

In summary, DCN neurons provide predictions of the cerebellum, and need to learn from their prediction errors. However, most of these neurons do not project to the olive where the prediction errors are computed. Furthermore, those that do project to the olive receive olivary projections that progressively weaken through development. As a result, it is puzzling how the prediction error of a DCN neuron is communicated to it in order to guide its plasticity. The possibility that we explored here is that the error signal is not conveyed to the DCN neuron from the olive directly, but indirectly through the parent P-cell. A complex spike in the parent P-cell can produce a 50 ms period of suppression in the DCN neuron. If the suppression is preceded by a period of excitation, the synaptic inputs from the mossy fibers to the DCN neuron strengthen. If the suppression is followed by rebound spiking, the synaptic inputs from the P-cells to the DCN neuron strengthen. Thus, control of plasticity at the DCN may be via the parent P-cell's complex spikes, providing an indirect pathway that communicates to the DCN a measure of its prediction error.

Effective transmission of the error signal to the nucleus via complex spike synchrony

Each olivary projection to the cerebellum represents the error signal associated with the predictions of one or more output neurons (the GABA-ergic DCN neurons). This signal needs to reach both the output layer neurons that are responsible for producing the prediction error, and the neurons in the middle layer (the P-cells) that project to those output layer neurons. The error signal seems to neglect the output layer neurons (at least in adults), and focus instead on the neurons in the middle layer, which then convey the error signal to the output layer neuron (Fig. 3). However, using the middle layer

neurons to convey error information introduces a major problem: how can an output layer neuron dissociate the error signal that it receives from the axon of a middle layer neuron, i.e. the complex spike of the P-cell, from the ordinary signals that the middle layer neuron also sends to the output layer, i.e. the simple spikes of the same P-cell?

Puzzle 3. *How does a DCN neuron sense the difference between a complex spike and a simple spike in the parent P-cell?*

A complex spike is a significant event at the soma of the P-cell, generating multiple spikelets. However, the P-cell's axon transmits this event to the DCN via ordinary spikes (Monsivais et al., 2005): typically one axonal spike at the onset of the complex spike (with near 100% probability) and another one at the final somatic spikelet (about 60% probability) (Khaliq and Raman, 2005). Thus, a complex spike is transmitted from the P-cell to the DCN neuron as a sequence of one or two spikes. Yet, a complex spike in a parent P-cell can be followed by a long-lasting reduction in the DCN neuron's activity (Fig. 3). How is this possible?

There is a remarkable similarity between the effects of a complex spike on a P-cell, and its effects on a DCN neuron: following a complex spike, both the P-cell (Fig. 5B) and the DCN neuron that it projects to (Fig. 3A) experience a long-lasting suppression. In the case of a P-cell, a complex spike is followed by 10-20 ms of total suppression of simple spikes (Fig. 5B). In the case of a DCN neuron, a complex spike in the parent P-cell is followed by a reduction in the firing rate of the nucleus neuron, lasting about 50 ms. Given that a complex spike is transmitted as roughly two ordinary spikes down the axon, after which the P-cell stops firing for 10-20 ms, it is quite surprising that there is a long-lasting inhibition in the DCN neuron.

Tang et al. (2019) noted that in anesthetized animals, the P-cells that projected to a single DCN neuron tended to have higher levels of complex spike synchrony than P-cells that did not project to the same neuron (Fig. 5A). Indeed, when they simultaneously recorded from 6 or 7 P-cells that were putatively connected to a single DCN neuron, they found that if the complex spikes in the parent P-cells were not synchronous, the complex spike had little or no effect in the DCN neuron (Fig. 5C, subplot in which 1/7 P-cells had a complex spike). However, if by chance the complex spikes in some of the parent P-cells were synchronous, the suppression deepened and its duration increased (Fig. 5C, 4/7 P-cells had a complex spike). In the rare event that all 7 of the recorded P-cells happened to produce a complex spike simultaneously, the DCN neuron's suppression was strong and long lasting (Fig. 5C, 7/7 condition).

Therefore, the simultaneous complex spikes in the parent P-cells had a far greater effect on the DCN neuron than a single complex spike in one P-cell. This provides us with a potential solution to our puzzle.

Conjecture 4. *In order for a DCN neuron to be made aware of its prediction error, there is a requirement for synchrony of complex spikes in the parent P-cells.*

Around 50 P-cells project to a single nucleus neuron (Person and Raman, 2012a). If these P-cells shared a common teacher, then they could solve an important problem: reliably transmission error information to the output layer.

Membership criterion for P-cell populations

From the perspective of learning to reduce prediction errors, organizing P-cells into small populations that have a single teacher is useful only if the error that is being transmitted to the member P-cells is the

same error that is associated with the output of the DCN neuron that the population projects to. And so we arrive at an inference regarding the membership criterion for the P-cell population: in order to transmit the error signal reliably to the DCN, there is a need for complex spike synchrony, which is more likely if the P-cell population shares a similar error signal from the olive. Furthermore, for the error signal to be a teacher for its downstream DCN neuron, it must be generated by olivary neurons that receive that DCN neuron's output.

The ideal scenario (from a computational perspective) would be the following. Suppose that a given olive neuron computes the error associated with a given DCN neuron, i.e., receives an inhibitory synapse from that DCN neuron. The olive neuron's axon to the cerebellum should split and synapse on each and every one of the 50 P-cells that project to that DCN neuron. However, a single olivary axon projects to only about 7 P-cells (Sugihara et al., 1999, 2001). Therefore, the population of P-cells that project to a nucleus neuron is bigger than what a single olive neuron can serve.

A typical DCN neuron receives inputs from approximately 8 olivary axons (Najac and Raman, 2018). If we imagine that each of these axons is a collateral from an olivary axon that travels to the cerebellar cortex and becomes the climbing fiber for 7 P-cells, then the olive to nucleus convergence ratio (8:1) combined with the olive to P-cell divergence ratio (1:7) produces roughly 56 P-cells that receive information about the error made by a single nucleus neuron. In a remarkable coincidence, about 50 P-cells converge on a single nucleus (Person and Raman, 2012a). We thus arrive at our central conjecture regarding the membership criterion for P-cells that form a single population (Fig. 6A):

Conjecture 5. *The population of P-cells that project onto a nucleus neuron is composed of those P-cells that receive climbing fibers from the olivary neurons that also project to that specific nucleus neuron*.

In summary, a handful of olivary cells send axons to a single nucleus neuron, and each also send climbing fibers to a handful of P-cells. Perhaps through development, the olivary cells guide those P-cells to project to the same nucleus neuron (Najac and Raman, 2018). In this scenario, the P-cell population that projects to a single DCN neuron would consist of P-cells that receive a climbing fiber from one of those olivary cells. If most of these olivary cells also receive a projection from that DCN neuron, then the olivary cells are in position to compute the error (Eq. 1) in the DCN neuron's output and provide that information via climbing fibers to the 50 P-cells that represent the parent of the DCN neuron. The P-cells transform the error information into a complex spike, which if synchronized with other P-cells in the population, produces a reduction in the DCN neuron's activity that was responsible for the error. This in turn produces plasticity in the DCN neuron's mossy fiber and P-cell synapses. Therefore, the responsibility to teach a DCN neuron is relegated from the olive to the population of P-cells that converge on that DCN neuron.

An appropriate error signal for a P-cell

It may seem trivial to ask how we should teach the neurons in our middle layer, i.e., the P-cells. After all, each P-cell receives a climbing fiber, thus endowing it with a powerful channel to receive error information. However, if we consider this question mathematically, we arrive at a useful inference.

Due to its placement in the middle layer of our network, a P-cell does not have a direct role in making predictions, and is not responsible for any single error. Rather, it projects to a handful of output layer neurons that make predictions, and thus contribute to our cost function (Eq. 1). Indeed, because a

P-cell projects to only 4 or 5 DCN neurons (Person and Raman, 2012b), it contributes to errors made by only these output neurons, and thus must learn from their errors only.

The teaching signal for a P-cell must provide error information regarding all the output neurons that the P-cell projects to. If this was an artificial neural network, we would send a P-cell a distinct error signal associated with each of the DCN neurons it projects to. However, a P-cell has only a single climbing fiber. Thus, we have a problem.

Puzzle 5. *A P-cell contributes to activity of multiple DCN neurons, yet receives only a single climbing fiber. How can the P-cell learn to change its activity in order to reduce the errors associated with all DCN neurons it projects to?*

Suppose that P-cell k has activity $a_k^{(p)}$, where the superscript refers to the fact that this unit is in the P-cell layer. This P-cell projects to neurons $j=\{1,\cdots,N\}$ in the DCN layer, contributing to their activities $a_j^{(d)}$. As a result, the gradient of the cost with respect to activity of this P-cell is:

$$\begin{aligned}\frac{dJ}{da_k^{(p)}} &= \sum_{j=1}^{N} \frac{dJ}{da_j^{(d)}} \frac{da_j^{(d)}}{dz_j^{(d)}} \frac{dz_j^{(d)}}{da_k^{(p)}} \\ &= -2\sum_{j=1}^{N} \frac{da_j^{(d)}}{dz_j^{(d)}} \frac{dz_j^{(d)}}{da_k^{(p)}} e_j^{(d)}\end{aligned} \tag{8}$$

The above expression implies that the error that the P-cell must learn from is proportional to the weighted sum of the errors made by its downstream DCN neurons. Unfortunately, the P-cell receives only a single climbing fiber. Thus, this climbing fiber must represent the errors made by all the DCN neurons $j=\{1,\cdots,N\}$ that the P-cell projects to.

Among the 4 or 5 DCN neurons that a single P-cell projects to, some are GABA-ergic, projecting to the olive, and some are non-GABA-ergic, projecting elsewhere in the brain (Teune et al., 1998). A potential solution to our puzzle is an anatomy in which the GABA-ergic DCN neurons of a given P-cell converge their axons on a single olive neuron, and that olive neuron serves as the teacher for the parent P-cell, providing it with its single climbing fiber (Fig. 6B).

For this to be an effective solution for learning, the GABA-ergic and non-GABA-ergic DCN neurons of our single P-cell would have to be matched in their activity, which is indeed what we had conjectured earlier when we considered the problem of teaching the two groups of DCN neurons (conjecture C2). This wiring would raise the following requirement: the olive cell that provides the climbing fiber to a P-cell should be the same olive cell that receives inhibitory projections from every GABA-ergic DCN neurons that this P-cell projects to (Fig. 6, right subplot). This is consistent with the proposal that activity of the climbing fiber during development, and particularly the coincidence of complex spikes in the parent P-cell with the DCN neuron's EPSCs from the olive, may be a guide in wiring the cerebellum (Lang et al., 2017; Lu et al., 2016).

In summary, efficient learning in a P-cell requires that its climbing fiber should come from an olive neuron that is the recipient of synapses from a specific group of DCN neurons: those that are daughters of that P-cell.

A potential problem with simple spike synchrony

The axon of a P-cell transmits both the simple and the complex spike as a regular spike. While synchronous complex spikes among a population of P-cells may be a clever way with which to solve the problem of transmitting error information to a DCN neuron, it introduces a new problem: what happens if there is synchrony among the simple spikes of the parent P-cells?

One the one hand, the olive relies on the P-cell to serve as the substitute teacher for the DCN neuron, indirectly providing it with error information. On the other hand, the rest of the brain, via mossy fibers and the granule cells, rely on the P-cells to provide the DCN with a major source of information about the state of the body, thus sculpting the output of the DCN neuron, producing the cerebellum's prediction. That is, most of the brain relies on the P-cells to provide spikes that influence the DCN's predictions, while the olive relies on the P-cells to provide spikes that teach the DCN about the errors of those predictions. How can the P-cell serve both masters?

From a learning perspective, in order to disambiguate the error signal from other signals that the P-cell transmits to the DCN, there is a need for suppressing synchrony among the simple spikes. In this way, if a synchronous event is detected at the DCN neuron, it can only be due to the error signal from the olive. However, from a control perspective, if the mossy fiber input via granule cells to the P-cells are to influence the output of a DCN neuron, then there is a need for synchrony among the simple spikes (Person and Raman, 2012b). Thus, we have a puzzle:

Puzzle 4: *How does a DCN neuron disambiguate the error signal in the parent P-cells from other information that the P-cells are transmitting?*

Let us consider the control perspective. P-cells fire at a baseline rate of around 50 Hz. In a decerebrated cat this background activity produces little change in the membrane potential of the downstream DCN neuron (Bengtsson et al., 2011). Through synchrony, simple spikes of P-cells can drive the activity of a DCN neuron (Person and Raman, 2012a). For example, if the simple spikes are asynchronous, the nucleus cell is inhibited, but if the simple spikes are briefly synchronous, at 8 ms latency the DCN neuron can experience a rebound that produces a spike. Thus, through synchronization of simple spikes the downstream DCN neuron can relay the specific timing of the P-cell activity to neurons outside the cerebellum, thus affecting behavior.

Indeed in both anesthetized and behaving animals, nearby P-cells have cross-correlations that exhibit a level of synchrony among complex spikes as well as simple spikes (De Zeeuw et al., 1997; Ebner & Bloedel, 1981; Heck et al., 2007; Sedaghat-Nejad et al., 2019; Shin & De, 2006; Wise et al., 2010). For example, Heck et al. (2007) found that in rats, P-cells within a few hundred microns fired synchronously during the period of reach to grasp, but not during the period after grasp completion. Hong et al. (2016) noted evidence for simple spike synchrony around saccade onset. Simple spike synchrony exists even when chemical synapses are inactivated (Han et al., 2018), suggesting that some neighboring P-cells are electrically coupled.

However, it is unclear whether in the intact animal, control of DCN output actually relies on synchrony of simple spikes. Experimental results and simulations of Payne et al. (2019) suggest that whereas the mean rate of simple spikes is critical for control of eye velocity during a pursuit task, the timing of individual spikes plays a less significant role. Irregular simple spike timing of a population of P-cells, though produced synchronously through optogenetic stimulation, did not result in eye velocities that were different than expected from the averaged firing rates of the P-cells. If synchronous simple

spikes are present during production pursuit eye movements, they do not appear to provide an advantage for control of DCN output.

What about the idea that unambiguous transmission of error information via synchronous complex spikes might require a mechanism to suppress simple spike synchrony? Intriguingly, in the cerebellar cortex, one function of interneurons may be to reduce the synchrony of simple spikes. P-cells receive excitatory inputs from granule cells, and inhibitory inputs from small interneurons in the molecular layer, as well as Golgi cells. Hausser and Clark (1997) measured activity in P-cells and molecular layer interneurons in cerebellar slices, thus removing the influence of excitatory inputs. Without their external excitation, P-cells and the interneurons nevertheless produced spontaneous activity. Critically, the spontaneous activity of individual P-cells was highly irregular: inter-spike intervals (ISI) exhibited a long tail, indicating that the coefficient of variation of ISI was large. To measure the effect of the inhibitory interneurons on the P-cells, Hausser and Clark (1997) blocked inhibitory synapses and found that mean firing rates increased in both P-cell and the interneurons. Removal of inhibition transformed the ISI histogram into a pattern that exhibited a single narrow peak without a tail. That is, removal of inputs from the interneurons to the P-cells increased their regularity.

Do the inhibitory interneurons play a similar role in the intact animal? Brown et al. (2019) made mutant mice that lacked stellate and basket cells, the principal interneurons in the cerebellar cortex. They recorded P-cell activity in the awake mice and found that despite loss of stellate cells, there was little or no change in the mean firing rate of P-cells. Rather, the mutation affected the variability of simple spikes: without the stellate cells, the P-cell simple spike ISI histogram had a reduced tail. This argument is admittedly weak because the pattern of irregularity may be shared among P-cells, and thus irregularity by itself is not an indicator of reduced synchrony. Regardless, the available results raise the possibility that a function of the interneurons in the cerebellar cortex may be to reduce synchrony of P-cell simple spikes.

In summary, P-cells face a conflict in the two roles that they are asked to play. On the one hand, given the objective of transmitting error information from the middle layer (P-cells) to the output layer (DCN neurons), we need to encourage complex spike synchrony but avoid simple spike synchrony among the parent P-cells that converge upon a single DCN neuron. On the other hand, given the objective of driving activity of the output layer neurons and producing a prediction, it is useful to encourage simple spike synchrony. Indeed, a function of the inhibitory interneurons in the molecular layer may be to transform the spontaneous simple spike activity of P-cells into an irregular pattern. The long-tail distribution of ISIs in the P-cell simple spikes may act to reduce the probability of simple spike synchrony. However, synchrony of simple spikes has the advantage of allowing the P-cells to drive the activity of DCN neurons, thus affecting behavior. Whether simple spikes are synchronous among the P-cells that form a population (as predicted by the control hypothesis), or asynchronous (as predicted by the error signal hypothesis), remains unknown.

Population coding of P-cells

Vernon Mountcastle used to start some of his lectures by reminding his students that “if you want to know what a neuron does, figure out what it is connected to.” Understanding function can benefit enormously from a prior understanding of anatomy. That advice looms large in the motor system, where control is often via population coding: a few dozen neurons in a given region of the brain project to a

common output neuron, forming a population that encodes an aspect of behavior. In this population, what is important is not the activity of a single neuron, but the collective activity of all the neurons that are part of the population. Unfortunately, we currently have no method that can label the membership of this population in the living brain. Instead, the common technique is to collect neurons into pseudo-populations based on the statistical properties of their discharge, for example, via principal component analysis, a technique in which one assigns a weight to each neuron's output, and then labels the weighted sum of all activities as the population response.

In the cerebellum there is an alternative approach, one that benefits from an *a priori* hypothesis about the role of the teacher (inferior olive) in organizing the students (the P-cells and their downstream DCN neurons): P-cells that receive similar error signals are likely to be part of the same anatomical population (Fig. 6A). Recent work has tried this idea and found that when P-cells in the oculomotor vermis of macaques were assigned membership into populations based on their complex spike tuning with respect to visual error, the simple spikes that for individual P-cells were indecipherable, as a population became interpretable (Herzfeld et al., 2015, 2018).

Saccadic eye movements depend critically on the cerebellum, as evidenced by the deficits observed in patients with cerebellar damage (Leigh and Zee, 2015; Xu-Wilson et al., 2009) and animals with focal lesions (Barash et al., 1999; Kojima et al., 2010a). Given this evidence, it seemed likely that P-cell simple spikes should encode kinematic parameters of saccades (e.g., eye speed, amplitude, or direction). However, studies had found that individual P-cells showed little modulation of simple spikes with respect to saccade speed (Helmchen and Buttner, 1995) and direction (Thier et al., 2000). Indeed, simple spikes during saccades (Fig. 7A) presented a bewildering assortment of responses, including P-cells that produced a burst, P-cells that produced a pause, and P-cells that did both (Kojima et al., 2010b). Most puzzling was the fact that simple spikes were modulated well beyond saccade end (Thier et al., 2000), as shown in Fig. 7B. These puzzles were not unique to saccades. Both the diversity of simple spike responses among various P-cells, and the fact that modulation of discharge far outlasted the movement, were also features of P-cell discharge during wrist (Mano and Yamamoto, 1980; Tomatsu et al., 2016) and arm movements (Hewitt et al., 2015). Given that the cerebellum is critical for precisely terminating saccades (Robinson et al., 1993) as well as arm movements (Becker and Person, 2019; Chen et al., 2006; Viaro et al., 2017; Vilis and Hore, 1980), it was puzzling that the P-cells continued to modulate their activity long after the movement had ended.

The key to this puzzle was population coding. We imagined that the inferior olive divided the P-cells into groups wherein all the P-cells within a group shared a common response to error. To measure the error response of each P-cell, we relied on the observation that if a saccade concluded but the target was not on the fovea, some P-cells were likely to produce a complex spike (Junker et al., 2018; Kojima et al., 2010b; Soetedjo et al., 2008; Soetedjo and Fuchs, 2006). To measure the tuning of complex spikes, we induced visual errors. Following fixation, a target was presented in the periphery. As the monkey made a saccade toward that direction, we jumped the target to a new position (Fig. 8A). During the 50-200 ms period following saccade completion, the visual error (target position with respect to fovea) produced a complex spike with a probability that depended on direction of the error vector (Fig. 8B).

Each P-cell had a preference for a specific error direction (Fig. 8C). The direction of error that produced the largest probability of complex spikes in the post-saccadic period for a given P-cell was labeled as CS-on. The complex spike tuning for all cells is plotted in Fig. 8C. We found that if the error

was in direction CS-on, following saccade completion the probability of complex spikes peaked at around 100 ms, producing a response that was roughly twice the baseline. In contrast, if the error was in direction CS-on+180, the probability of complex spike decreased by roughly 40% below baseline. Importantly, the probability of a complex spike following saccade end was unrelated to direction of the preceding saccade, but driven by the direction of the post-saccadic error. These results were acquired in macaques and then confirmed in marmosets (Sedaghat-Nejad et al., 2019). Thus, in two different primate species, P-cells in the oculomotor vermis exhibited a strong preference for direction of the visual prediction error.

This preference for error appears to be a consistent feature of P-cells across species. In the cerebellum-like circuit of the electric fish, MG-cells are analogous to P-cells, and produce broad spikes that are analogous to complex spikes. An unexpected sensory input produces broad spikes in some MG-cells, but suppresses the broad spikes in others (Muller et al., 2019). Thus, MG-cells in the electric fish also exhibit a tuning for prediction error.

Why do the P-cells exhibit a complex spike tuning with respect to visual error? A prediction error is the difference between what is predicted and what is observed. We assume that the cerebellum predicts that at saccade end the target will be on the fovea. However, the visual stimulus is in fact elsewhere, activating neurons in an unexpected region of the visual space in the superior colliculus (Kojima and Soetedjo, 2018; Soetedjo et al., 2009). The difference between the observed location and the predicted location is the prediction error. The unexpected excitation in one region of the colliculus likely engages a few olivary neurons on the contralateral side (Saint-Cyr and Courville, 1982), which in turn engage the climbing fibers of a few P-cells. That same activation in one group of colliculus neurons inhibits collicular neurons that correspond to the polar opposite region of the visual space (Behan and Kime, 1996; Munoz and Istvan, 1998), thus removing excitation to another set of olivary neurons. As a result, a single unexpected visual event activates some collicular neurons, while simultaneously suppressing others, increasing probability of complex spikes for some P-cells, and decreasing it for others.

When we organized the P-cells in groups that shared similar complex spike tuning with respect to the visual error, a remarkable pattern emerged. The simple spikes produced by the population of P-cells predicted eye velocity in real-time (Fig. 9A). That is, whereas individual P-cells exhibited great diversity in their responses, with activity that was modulated long after the saccade had ended (Fig. 7A), as a population the simple spikes predicted the motion of the eyes precisely (Herzfeld et al., 2015). For example, when the saccade was in the anti-preferred error direction (CS-off) of the P-cells, their combined simple spikes increased before saccade onset, had a peak discharge that scaled with peak velocity of the eye (Fig. 10B), and then exhibited a suppression that coincided with deceleration of the eyes. It is remarkable that this pattern was not present in the activity of any single P-cell, but unmasked in the population, i.e., through the combined input of the bursters, pausers, and the rest.

If saccade direction aligned with the CS-on direction of the P-cells, the population still predicted eye velocity in real-time, but now with a lower gain (Fig. 9C), and without the suppression during the deceleration period. Therefore, the population response predicted real-time velocity of the eye, with a gain that multiplicatively depended on direction of motion. This encoding is termed a “gain-field”: the magnitude of the population response increased linearly with speed, and was cosine tuned in direction, with a multiplicative interaction between speed and direction. Gain fields are also employed by cells in

the posterior parietal cortex to represent position of the eye in the orbit and position of the stimulus on the fovea (Andersen et al., 1985).

In summary, in both macaques and marmosets, individual P-cells in the oculomotor vermis had simple spikes that were modulated for periods that lasted much longer than the saccade. The conjecture that the olive organized the P-cells into groups (Fig. 6A) presented an anatomical heuristic for population coding: P-cells that shared a similar error signal grouped together to influence a DCN neuron. To find population membership, we measured each P-cell's complex spike response to prediction error. Each P-cell had a preference for a particular direction of visual error. We assigned membership into populations based on complex spike error-dependent tuning. The resulting sum of simple spikes produced a population response that presented a consistent pattern: in direction CS-on+180, simple spikes increased before the eyes accelerated by an amount that predicted peak velocity of the eye, and then suppressed activity as the eyes decelerated. In direction CS-on, simple spikes increased by a smaller amount during acceleration, predicted the peak velocity with a smaller gain, and then returned to near baseline during deceleration. Hence, the population response seemed likely to play a role in controlling the motion of the ongoing saccade.

Actions of the DCN neurons should reflect the errors that concern the P-cell parents

We do not need the cerebellum to make a saccade (or any other movement). The superior colliculus is quite capable of activating the brainstem burst generators and the downstream motoneurons to produce a saccade. However, without the cerebellum's contributions, the saccade is dysmetric. For example, if a DCN (caudal fastigial nucleus) on one side of the cerebellum is disabled, the saccade will overshoot the targets on that same side, and undershoot the targets to the other side (Goffart et al., 2003; Guerrasio et al., 2010; Robinson et al., 1993). In particular, saccades to the side of the deactivated DCN are missing part of the deceleration commands that are needed to stop the eyes on target (Buzunov et al., 2013; Kojima et al., 2014). Thus, the cerebellum adds something to the ongoing movement. What is it adding?

To answer this question, consider that a P-cell may have a preference for a particular prediction error, but that error is not its personal responsibility. Rather, the error arises because of a miscalculation in the output of the DCN neurons that the P-cell projects to. Where do these daughter DCN neurons project to? Are the projections related to the preferred error (i.e., CS-on) of the parent P-cell?

Eq. (8) provides a potential answer: the complex spike that a P-cell generates must be associated with the miscalculation of its daughter DCN neurons. But P-cells do not respond equally to all errors. Rather, they are most concerned with a specific error, as evidenced by their complex spike tuning. Perhaps the DCN daughters should focus on producing downstream effects that will alleviate the specific error concerns of their parent P-cells.

Suppose we placed a target at location "a" on the left of the screen and made a saccade toward it (Fig. 10A). During the saccade, we erase this target and replace it with a target at "b", located a few degrees to the left of "a". Thus, the eyes do not arrive on target because the target was shifted a few degrees to the left. As the saccade ends, the cerebellum predicts that the target should be near the fovea, and thus there should be activity among the foveal-related neurons in the rostral pole of the colliculus. However, the actual sensory consequence is different: there is higher than normal activity in a group of neurons located in the right caudal superior colliculus (red region "b" in right colliculus, Fig.

10B), and lower than normal activity in a group of neurons in the left caudal superior colliculus (blue region in the left colliculus, Fig. 10B). The lower than normal collicular activity in the contralateral colliculus is because caudal neurons in one colliculus inhibit the neurons on the contralateral side (Munoz and Istvan, 1998).

Some of the affected neurons in the left and right superior colliculus project to neurons in the right and left olive, respectively (Saint-Cyr and Courville, 1982). The activity in the collicular cells is unexpected because it was not predicted by the GABA-ergic DCN neurons that inhibit these specific olive cells. The greater than expected activity in the left colliculus leads to increase in complex spike production among the P-cells parents of the right DCN neurons, whereas the smaller than expected activity in the right colliculus leads to suppression of complex spikes among the parents of the left DCN neurons.

Thus, a single event that signaled a sensory prediction error at saccade end (target is not on the fovea) is a miscalculation by two distinct groups of DCN neurons. Their miscalculations are conveyed to their parents (the P-cells). In one group of P-cells, the error results in an increase in their complex spike probability (parents of the left DCN neurons), and in another group of P-cells the same error results in a decrease in this probability. Hence, the collicular map of the visual space is partly responsible for the complex spike tuning of the P-cells in the oculomotor vermis (Kojima and Soetedjo, 2018, 2017; Soetedjo et al., 2009).

The change in complex spike probability with respect to baseline signals to the P-cell that its daughter GABA-ergic DCN neurons failed to correctly anticipate the sensory consequences of the movement. Therefore, if the daughters are to correctly predict the consequences of this saccade, they must learn to anticipate the activity that will take place in neurons which reside in the caudal region of the superior colliculus, and cancel it when it is conveyed to the inferior olive. If they did that, they would inhibit the olive neurons that these collicular neurons project to, anticipating their activity and returning the complex spikes to baseline. As a result, if a P-cell exhibits a preference for leftward visual errors, this implies that its GABA-ergic DCN daughters project to an olivary neuron that receives excitatory inputs from collicular neurons that encode a region to the left of the fovea (Fig. 10A).

However, it is not enough for the cerebellum to merely learn to correctly predict the sensory consequence of a movement. Rather, we need a mechanism with which to alter the motor commands and guide the saccade to its desired destination, placing the target on the fovea. Motor learning is not about eliminating sensory prediction errors by building better predictors (forward models). It is about learning how to make better movements. This provides us with a clue as to where the non-GABA-ergic DCN daughters of a P-cell should project to.

Conjecture 6. *There should be a correspondence between complex spike tuning of the parent P-cells, and the downstream effects of activity in the daughter non-GABA-ergic DCN neurons. The DCN daughters should project to a group of neurons that can produce an effect that will remedy the specific error that is of concern to the parent P-cells.*

For example, if at saccade end the target is to the left of the fovea (Fig. 10B), and there is a P-cell population that prefers this error, then the non-GABA-ergic DCN daughters should be able to do something constructive about eliminating this error. They could do that if their axons projected to neurons that indirectly engaged muscles that pulled the eyes horizontally to the left (Fig. 10B), for example by projecting to burst generators that act on left-pulling abducens motoneurons. Similarly, for

the P-cell population that prefers errors to the right, a constructive action would be for the daughter DCN neurons to reduce the drive for right-ward pulling abducens motoneurons. With this access to effectors, the non-GABA-ergic DCN neurons could profitably influence the saccade's trajectory in a way that would help restore the complex spikes in their parent P-cells toward baseline.

Now let us consider how presence of this leftward error at saccade end should produce trial-to-trial learning. The leftward error increases probability of complex spikes for P-cells that prefer that error, slightly suppressing their simple spikes on the next trial (push more to the left). The same error decreases the probability of complex spikes for P-cells for which this error is CS-on+180, slightly increasing their simple spikes on the next trial (push less to the right). As a result, on the next trial, as the target is shown at location "a" and the saccade is made, the cerebellum via its non-GABA-ergic DCN neuron adds a small leftward force to the ongoing motor commands, guiding it toward location "b". If such an anatomy existed, the non-GABA-ergic neurons would be able to alter the saccade's trajectory, reducing the sensory prediction error.

The correspondence between the error that concerns the parents and the downstream influence of the daughter helps us understand why the population the P-cells show a suppression of activity when the saccade is decelerating toward the CS-on+180 direction, but not when it is decelerating toward direction CS-on (Fig. 9A). In the oculomotor vermis of the cerebellum, P-cells that are on the right side project to the right caudal fastigial nucleus. These P-cells tend to prefer errors that are to the left (Herzfeld et al., 2015), and thus have their CS-on+180 to the right. According to conjecture 6, their DCN daughters indirectly affect effectors that produce a leftward force. When a rightward saccade is made, these P-cells tend to show suppression during the deceleration period, which may increase activity of the daughter DCN neurons, thus stopping the saccade.

In summary, the error that is transmitted to a P-cell must reflect the miscalculations that are made by only the DCN neurons that it projects to. P-cells exhibit a preference for a specific error. To remedy this error, it would be useful if the DCN daughters produced actions that could remedy the error concerns of the parent P-cells. In the context of saccades and visual errors, this means that the DCN neurons should project to a location where they can indirectly engage motoneurons that pull the eyes in a direction that is parallel to the preferred visual error vector of the parent P-cells.

Effect of complex spikes on changing behavior

Is there anatomical evidence for conjecture 6? Ekerot et al. (1995) examined the part of the interpositus nucleus that received inputs from P-cells which responded with complex spikes when the cat's forelimb was touched or pinched on the skin. To determine the complex spike tuning of the parent P-cells, in the decerebrated cat they recorded local field potentials from the interpositus neurons and noted that providing a noxious pinch on a specific forelimb skin area produced an excitatory response. They used this as a proxy for the complex spike tuning of the parent P-cells. They then electrically stimulated these interpositus neurons and observed multi-joint movements of the forelimb. The movements were likely generated because stimulation engaged DCN neurons that projected to the red nucleus, from where neurons projected to the spinal cord via the rubrospinal tract.

When they compared the complex spike tuning of the parent P-cells (as inferred from the field potentials in the DCN) and the stimulation results, there was a general correspondence: the movement that was evoked from stimulation of a particular DCN region pulled the limb away, protecting the part of

the skin which was pinched to produce the complex spike in the parent P-cells. For example, if the skin on the ventral forepaw was touched, generating potentials in a region of DCN that received input from P-cells that had a complex spike receptive field for that part of the skin, stimulation of that DCN region activated a palmar flexor of the wrist, producing a withdrawal of the forepaw away from the noxious stimulus.

Conjecture 6 predicts that in the case of saccadic eye movements, P-cells that have a tuning for leftward visual errors should project to non-GABA-ergic DCN neurons that in turn project to neurons that through their activity produce a leftward force on the eyes. There is an indirect way to test this prediction in the awake, behaving animal.

Complex spikes are random events that may or may not occur after a given error. For example, if an error occurs in trial *n*, and that error is in the preferred direction of a P-cell, in some fraction of trials that P-cell will generate a complex spike, whereas in most trials it will not. Conjecture 6 predicts that if the complex spike did occur, it signaled the preferred error, which in turn produced plasticity in the P-cell and its daughter DCN neurons. On the next trial, behavior should be changed: on trial *n*+1, the eyes should be pulled slightly in the direction specified by the CS-on vector of the parent P-cell.

We tested this idea in Herzfeld et al. (2018). After saccade completion, an error was present or absent, and the P-cell had produced a complex spike, or not. We measured trial-to-trial change in behavior via the difference in motor output (eye velocity) from the trial in which the error was experienced to the subsequent trial in which the same target was presented.

We began with trials in which after saccade completion, the error was in the CS-on direction of the P-cell (Fig. 11A). In some trials, the P-cell did not produce a complex spike in the post-saccadic period (Fig. 11A, CS absent). In other trials, the P-cell did produce a complex spike (Fig. 11A, CS present). In both cases, in the subsequent trial the motor commands pulled the eyes more in the direction of error, which in this case happened to coincide with CS-on of the P-cell under study (Fig. 11A, left panel). The actual errors were similar in trials that the complex spike had or had not occurred. Yet, the change in the motor output was significantly larger, i.e., the pull was stronger, following production of a complex spike in that single P-cell (Fig. 11A, left panel, difference). A similar observation has been made in a different task (pursuit) when visual error was in direction CS-on of P-cells in the flocculus region of the cerebellum (Yang and Lisberger, 2014a).

Sometimes the movement was perfect: a saccade took place and the target was precisely on the fovea (Fig. 11B). Despite this lack of error, on some fraction of trials the P-cell nevertheless generated a complex spike in the post-saccadic period. Remarkably, in the subsequent trial the eyes were again pulled in the CS-on direction of that P-cell (Fig. 11B, CS present). The precise amount of change in behavior was the same whether or not an error was present (the difference curves are the same in Fig. 11A and Fig. 11B). Therefore, even without an error, the presence of a post-saccadic complex spike in a single P-cell was followed in the subsequent trial by a change in behavior. The eyes were pulled in the CS-on direction of that P-cell.

These results hinted that if a P-cell produced a post-saccadic complex spike, in the subsequent trial that P-cell (indirectly) influenced a specific group of motoneurons, those that produced force along the P-cell's CS-on direction. To test this hypothesis, we focused on trials in which error was in direction CS-on+180^{o} (i.e., error was opposite the preferred error direction of the P-cell). Following experience of this error, the behavior in the subsequent trial changed in the direction of that error (Fig. 11C, left

panel). Remarkably, the trial-to-trial change in behavior was significantly smaller if the P-cell had produced a complex spike. As a result, the difference between the trial-to-trial change in the motor commands that took place with and without a complex spike was always in direction CS-on, regardless of whether error was in direction CS-on (Fig. 11A), CS-on+180^{o} (Fig. 11C), or absent altogether (Fig. 11B).

We therefore made our analysis blind to the error that the animal had actually experienced and instead labeled each trial based on whether the P-cell had or had not produced a post-saccadic complex spike. If the P-cell produced a complex spike, then the trial-to-trial change in saccade velocity was entirely in the CS-on direction of that P-cell (CS present, Fig. 11D), with no component along CS-on+90^{o}. In contrast, if the P-cell did not produce a complex spike, the trial-to-trial change was in direction CS-on+180^{o} of that P-cell (CS absent, Fig. 11D).

In summary, following production of a post-saccadic complex spike in a given P-cell, in the subsequent trial the eyes were pulled along a vector that was exactly parallel to the CS-on direction of that P-cell. Without a complex spike, in the subsequent trial the eyes were pushed away along that same vector. Importantly, this pattern was independent of the error direction; it occurred even when there were no errors at the end of the saccade. Thus, there appeared to be a correspondence between the complex spike tuning of a P-cell, and the direction of action of the daughter DCN neurons (Apps and Garwicz, 2005; Ekerot et al., 1995). If a P-cell had a complex spike tuning that preferred a particular direction of error, its daughter DCN neurons appeared to project to neurons that produced actions that could mitigate the specific error vector that most concerned the parent P-cells.

Multiple timescales of adaptation

The fact that a P-cell has a preference for a specific error vector, coupled with the conjecture that its DCN daughters mirror that preference in their anatomical projections, gives us the possibility to consider why behavior during learning exhibits a number of remarkable features: multiple timescales, resistance to erasure, and spontaneous recovery.

We noted that presence of a complex spike in a P-cell was followed by a change in behavior, but the absence of a complex spike was also followed by a change in behavior, one that was much smaller and in the opposite direction (Fig. 11D). This is because both the presence and absence of a complex spike leads to P-cell plasticity, but in different directions, and with different rates. For example, Yang and Lisberger (2014a) trained monkeys to visually pursue a target and induced an unexpected change in the motion of the target. Sometimes the error was in direction CS-on for the P-cell, and sometimes it was in direction CS-on+180. If the error took place and a complex spike occurred, on the next trial the simple spike rate during pursuit exhibited a decrease (Fig. 12A, pursuit). If the complex spike did not occur, the simple spike rate during the subsequent trial exhibited a small increase. We observed roughly the same pattern during saccades (Herzfeld et al., 2018): if at saccade completion a complex spike was present, on the next trial the P-cell showed a reduced rate of simple spikes (Fig. 12A, saccade). (The effect of complex spike absence on the simple spikes was more difficult to see in the saccade data.)

Yang and Lisberger (2014a) measured activity changes in P-cells over the course of 100 trials. They found that if the error was consistently in the CS-on direction of the P-cell, the simple spike rate during the pursuit declined steadily (Fig. 12B, pursuit). If the error was consistently in the CS-on+180 direction, the simple spike rate increased. Notably, the rate of increase in the simple spikes was slower than the rate of decrease during pursuit. We observed a similar pattern during saccades (Herzfeld et al.,

2018): the simple spikes produced by the population of P-cells decreased for CS-on errors, and increased for CS-on+180 errors, but the decrease occurred faster than the increase.

The differing sensitivities of simple spikes to presence or absence of a complex spike is an important feature because it suggests that a single error engages mechanisms that naturally exhibit multiple timescales, some fast (CS-on), others slow (CS-on+180). Let us show how this fact can contribute to the multiple timescales of behavior during adaptation.

Suppose we have two P-cell populations, one with a complex spike tuning to the left, and the other to the right (as in Fig. 10). These P-cells are distinct populations that projects to two different DCN neurons. On trial *n* we make a saccade and during that saccade population 1 produces simple spike response $x_1^{(n)}$. On the same trial population 2 produces response $x_2^{(n)}$. Following saccade completion we experience a post-saccadic error $e^{(n)}$: the target is to the left of our fovea. This error is in direction CS-on for population 1, and CS-on+180 for population 2. The error produces plasticity in our two populations (and their DCN daughters). On the next trial, the responses in the two populations are as follows:

$$\begin{aligned} x_1^{(n+1)} &= a_{on} x_1^{(n)} + b_{on} e^{(n)} \\ x_2^{(n+1)} &= a_{off} x_2^{(n)} + b_{off} e^{(n)} \end{aligned} \tag{9}$$

The term b_{on} refers to the trial-to-trial change in the simple spike response when the error is in direction CS-on. The term a_{on} refers to the trial-to-trial change due to passage of time between the two trials (a decay toward baseline). Let us show that we can infer information about these parameters from the available data.

Following experience of error $e^{(n)}$, the saccade is corrected by a small amount q on the subsequent trial. The adaptive response q is due to the contributions of these two populations. Suppose that if a population produced a complex spike in response to the error, it affects behavior on the next trial in the direction of its CS-on by amount y_1. If it does not produce a complex spike, the effect is in the direction of CS-on+180 by amount y_0. Fig. 8C provides the probabilities associated with production of a complex spike. From these probabilities, we form the following two equations:

$$\begin{aligned} q = &\Pr\left(CS=1 \middle| CSon=\leftarrow, err=\leftarrow\right) y_1 - \Pr\left(CS=0 \middle| CSon=\leftarrow, err=\leftarrow\right) y_0 \\ &-\Pr\left(CS=1 \middle| CSon=\rightarrow, err=\leftarrow\right) y_1 + \Pr\left(CS=0 \middle| CSon=\rightarrow, err=\leftarrow\right) y_0 \\ 0 = &\Pr\left(CS=1\right) y_1 - \Pr\left(CS=0\right) y_0 \end{aligned} \tag{10}$$

The first equation provides the probabilistic contributions of the two populations to learning from error. For example, the term $\Pr\left(CS=1 \middle| CSon=\leftarrow, err=\leftarrow\right)$ specifies the probability of a complex spike in population 1, given that the CS-on direction is leftward, and error direction is leftward as well. The term $\Pr\left(CS=1 \middle| CSon=\rightarrow, err=\leftarrow\right)$ specifies the probability of a complex spike in population 2, given that the CS-on direction is rightward, but the error direction is leftward. The values for each probability term can be acquired from Fig. 8C. The second equation provides the constraint associated with homeostatic regulation; sum total of plasticity, weighted by the overall probability of complex spikes, is zero.

The solution to Eq. (10) is $y_1 = 5.6q$ and $y_0 = 0.6q$. This means that following an error, the population with a CS-on is responsible for 79% of the adaptive response, whereas the population with CS-on+180 is responsible for 21%. Therefore, given the leftward error, both populations learn from error, one by increasing the simple spike rate, and the other by suppressing it. However, because population 1 prefers the error, it will contribute roughly 3.5 times more to the adaptive response than population 2.

The next critical feature of the P-cell's response to error is that between the trial in which the error is experienced and the subsequent trial, the response decays with passage of time. Yang and Lisberger (2014a) found that the trial-to-trial change in a P-cell's adaptive response depended not only on production of a complex spike, but also on the passage of time between trials. The passage of time produced decay in the error induced change in the simple spike response.

We do not know whether the plasticity induced by production of a complex spike decays differently than the plasticity induced by lack of a complex spike. But we can make a guess by examining the data in Fig. 12B. After a long series of trials, the change in simple spikes reaches an asymptote. Remarkably, this asymptote is not different for CS-on errors that encourage production of complex spikes, and CS-on+180 errors that suppress it. However, the trial-to-trial rate of change is faster when the errors are in direction CS-on.

A constant asymptote but differing rates of convergence require that trial-by-trial forgetting be smaller for the process that learns more rapidly. To see this, consider that a process that learns from error e by amount b but suffers from forgetting between trials by amount a (as in Eq. 9), exhibits a steady state $x^{(ss)}$ that at asymptote is defined by:

$$x^{(ss)} = \frac{be^{(ss)}}{1-a} \tag{11}$$

Eq. (10) and the data in Fig. 12B imply that when the error is in direction CS-on, the P-cell population that prefers this error will exhibit an error sensitivity that is larger than the population for which this error is in direction CS-on+180. That is, $b_{on} > b_{off}$. The data in Fig. 12B implies that $x_1^{(ss)} = x_2^{(ss)}$. Hence, we infer that $a_{on} < a_{off}$ and $b_{on} > b_{off}$.

In summary, when an error is aligned with CS-on, sensitivity to that error in the P-cells and their downstream DCN neurons is greater than when the same error is aligned with CS-on+180. However, if the error is aligned to CS-on, the resulting adaptation decays at a faster rate as time passes from one trial to the next, as compared to when error is aligned to CS-on+180. Thus, given a single error, one population of P-cells and their DCN neurons learns a great deal (CS-on), but also forgets at a faster pace, whereas another population of P-cells learns less (CS-on+180), but retains a greater amount.

Resistance to erasure and spontaneous recovery of memory

In numerous paradigms, from fear conditioning to motor adaptation, learning exhibits a remarkable property: acquisition of a novel behavior followed by its extinction does not erase the acquired memory. Rather, following passage of time, behavior reverts toward the originally learned pattern. That is, unlike an artificial neural network, in many biological systems one cannot unlearn a behavior by reversing the

sign of the prediction errors (Pekny et al., 2011). Our framework provides clues as to what may be the cause of this ubiquitous feature of learning.

In the field of classical conditioning, bees can learn to associate an odor with nectar, extending their proboscis upon presentation of the odor (Stollhoff et al., 2005). They will extinguish this response if the odor is presented without the nectar. However, following passage of time, the bees once again extend their proboscis when they are presented with the odor. In the field of fear conditioning, a stimulus can be associated with a shock, inducing fear. This fear can be extinguished if the stimulus is presented without the shock (Schiller et al., 2010). However, fear of the stimulus returns following passage of time. In the field of motor learning, people and other animals respond to a perturbation by modifying their motor commands through learning from error. If the perturbation changes direction, reversing the error vector, behavior returns to baseline. However, with passage of time the behavior spontaneously reverts back: subjects reproduce the motor commands that they had originally learned (Criscimagna-Hemminger and Shadmehr, 2008; Kojima et al., 2004; Sarwary et al., 2018; Smith et al., 2006).

The central theme in all of these experiments is resistance to erasure, leading to spontaneous recovery. A mathematical model (Kording et al., 2007; Smith et al., 2006) had suggested that during learning, changes in behavior appear to be supported by two adaptive processes: a fast adaptive process that has high sensitivity to error along with poor retention, and a slow adaptive process that has poor sensitivity to error along with robust retention. In this model, learning is due to changes in both the fast and the slow processes. When error reverses direction, behavior returns to baseline, but the memory is not erased. Rather, the slow process retains what it had learned, but that learning is masked by learning in the fast process. With passage of time, the fast process decays, inducing spontaneous recovery of the previously learned behavior.

In our model of the cerebellum, training starts with errors that are in a constant direction, thus producing rapid changes in the population with CS-on aligned with that error, and slower change in other populations. During extinction training, the error reverses direction. The populations that had learned slowly now learn rapidly because the error is aligned with their CS-on. In contrast, the populations that had learned rapidly now learn slowly because the error is aligned with their CS-on+180. Critically, reversal of error direction produces adaptation that is approximately 3.5 times slower in the CS-on+180 direction as compared to CS-on. Thus, reversal of error cannot easily erase the history of past trials.

Behavior exhibits apparent extinction not because the populations have returned to their baseline state, but because of a balance has been reached between the changes in the CS-on and CS-on+180 populations of P-cells. With passage of time, the previously rapidly adapting population exhibits decay toward baseline, resulting in the spontaneous recovery of behavior. Thus, a consequence of preference for error in the cerebellum may be that behavior is more easily learned than unlearned: unlearning requires 4 or more times the amount of training than the amount that it took to acquire the original behavior.

Discussion

Activity of single neurons in the cerebral cortex and the cerebellum are often difficult to decode as a function of the ongoing behavior. In the cerebral cortex, the current approach is to pool the neurons

into pseudo-populations in which a weight is associated with each neuron's response based on a statistical measure of its activity (e.g., principal components), and then sum the weighted activities to produce a population response. In the cerebellum, however, P-cells actually organize into small groups, each converging on a single DCN neuron (Person and Raman, 2012a). Is there a principled way to understand population coding in the cerebellum?

Here, we began with the assumption that activities of DCN neurons are the cerebellum's predictions, and the inputs from the inferior olive are their prediction errors. A theoretical analysis of the learning problem highlighted three key features of the anatomy. First, only some DCN neurons project to the olive, making it unclear how the prediction error is computed for the non-olive projecting DCN neurons. Second, olivary inputs to the DCN neurons (i.e., output layer) are weak (Lu et al., 2016), but the same inputs are remarkably strong to the P-cells (i.e., middle layer). This disparity makes it unclear how the output layer neurons learn from their prediction error. Finally, a P-cell projects to 4 or 5 DCN neurons (Person and Raman, 2012b), thus requiring the error signal that it receives to include information from this specific group of output neurons. Yet, a P-cell is endowed with only a single climbing fiber, thus limiting its view of the error space.

The idea that we considered is that the error associated with a DCN neuron's output is computed by an olive cell and then sent via climbing fibers to a few P-cells which then project back to that specific DCN neuron (Chaumont et al., 2013). That is, the error signal is conveyed to a DCN neuron indirectly through its parent P-cells (Heck et al., 2013). However, a complex spike is effective in influencing a DCN neuron's activity only if that it is synchronized with complex spikes that are generated in other P-cells in the same population (Tang et al. 2019). Thus, control of plasticity at the DCN may be via synchronous complex spikes in the parent P-cells, providing an indirect pathway that communicates to the DCN neuron its prediction error. However, the requirement for complex spike synchrony implies that the parent P-cells of a DCN neuron cannot be selected randomly. Rather, the olive must select the P-cells so that they share the same preference for error.

This framework implies that the population of P-cells that project onto a nucleus neuron is likely composed of those P-cells that receive climbing fibers from the olivary neurons that also project to that nucleus neuron (Fig. 6A). Furthermore, efficient learning in a P-cell requires that its single climbing fiber should come from a specific olivary neuron: the olivary neuron that computes the prediction errors associated with the daughter DCN neurons (Fig. 6B).

This framework has proved useful in decoding simple spike activity of P-cells during voluntary movements (Herzfeld et al., 2018, 2015). During a saccade, individual P-cells exhibited simple spike rates that increased, decreased, or did both, with a modulation that outlasted the movement by 50 ms or more (Kojima et al., 2010b; Thier et al., 2000). However, when the P-cells were organized into populations based on their complex spike response to prediction error, the sum total of simple spikes presented a clear pattern: in direction CS-on+180, simple spikes increased before the eyes accelerated and then suppressed activity as the eyes decelerated. In direction CS-on, simple spikes increased by a smaller amount during acceleration but returned to near baseline during deceleration. Thus, the population response during saccades appeared related to the acceleration needed to move the eyes, and the deceleration needed to stop it.

But P-cells are not the output of the cerebellum, and so they do not directly control any aspect of behavior. To understand their role in control of behavior, it is useful to ask where the DCN neurons of

a given population of P-cells project to. Trial-by-trial analysis of complex spikes (Herzfeld et al., 2018; Yang and Lisberger, 2014a) has provided evidence for the conjecture that there is a correspondence between the complex spike tuning of a P-cell, and the direction of action of the daughter DCN neuron (Apps and Garwicz, 2005; Ekerot et al., 1995). The direction of error that a P-cell is concerned with (CS-on) appears to be aligned with the direction of action of its DCN neuron. Thus, the predictions that a DCN neuron makes, and the actions it influences, are aligned with the specific prediction errors that its parent P-cell is concerned with.

There may be a consequence of this anatomy on how we learn. During the past two decades, numerous paradigms have quantified behavior during motor adaptation and noted a few consistencies. First, during a single session, errors decline as training proceeds, but there are at least two rates of learning, one fast, and the other slow. Second, if the sensory prediction errors reverse direction, thus causing "unlearning", behavior returns to baseline (termed extinction training). However, reversal of the errors does not erase the memory, as evidenced by the fact that with passage of time, there is spontaneous recovery of the initially learned behavior.

It seems possible that the multiple timescales of learning are related to the fact that a performance error engages multiple P-cell populations, some that produce a complex spike in response to that error, and others that suppress this production. Both induce plasticity that affects behavior, but the effects are asymmetric: simple spike rates change more after experience of a complex spike, and may also decay more with passage of time. In addition, the conjecture that the P-cells act as substitute teachers for their DCN daughters (Medina and Mauk, 1999) raises the possibility that even slower timescales of adaptation arise from plasticity in the DCN neurons (Herzfeld et al., 2020). The asymmetric response that a P-cell population exhibits to experience of error may be responsible for the fact that reversal of the prediction error during extinction training does not reverse the effects of past learning. Instead, it produces spontaneous recovery with passage of time.

However, our central conjecture that complex spikes signal a prediction error is not without problems. Complex spikes occur at onset of salient events that by themselves may have no obvious interpretation as a prediction error (Welsh et al., 1995). Indeed, the probability of a complex spike depends on many factors beyond the presence or absence of a prediction error. For example, while sudden occurrence of stimulus by itself may not produce a complex spike in a P-cell of a naïve animal, it will do so if the animal has learned to associate that stimulus with an aversive event (Ohmae and Medina, 2015), with a greater reward (Heffley and Hall, 2019; Larry et al., 2019), or a previous performance error (Junker et al., 2018). Withholding of reward when it was expected can also modulate complex spike rates (Heffley et al., 2018). These results suggest that the signal in the climbing fiber is not merely the difference between what was predicted and what was observed, but also a measure of value of that difference.

It is noteworthy that complex spikes, rare events that occur approximately once per second, appear poorly matched to reflect prediction errors of DCN neurons, neurons which fire at rates of tens of spikes per second. The key appears to be complex spike timing. Small errors produce a complex spike that has high temporal variability, while large errors produce a complex spike that has low temporal jitter (Herzfeld et al., 2018; Najafi et al., 2014), synchronized with other P-cells (Najafi et al., 2014). Furthermore, complex spikes that coincide with large errors tend to have more somatic spikelets. Both the timing (Herzfeld et al., 2018; Suvrathan et al., 2016) and the number of spikelets (Rasmussen et al.,

2013; Yang and Lisberger, 2014b) affect not only the amount of plasticity at the P-cell, but also the strength of suppression in the downstream DCN neuron, thus potentially affecting its plasticity.

Much of this evidence appears consistent with the conjecture that the olive relies on synchrony of complex spikes within a population of P-cells to convey error information to a DCN neuron. However, synchrony is a double edge sword because it forces P-cells to play two conflicting roles. In order to transmit error information to the DCN neurons, the population needs to encourage synchrony among the complex spikes (Tang et al., 2019) but not the simple spikes. However, in order to drive the output of a DCN neuron, there is a need for simple spike synchrony (Person and Raman, 2012a). It is possible that a function of the molecular layer inhibitory interneurons is to control synchrony of simple spikes (Brown et al., 2019; Hausser and Clark, 1997). The answer to this puzzle requires simultaneous recording from P-cells that have similar complex spike tuning.

We have focused on an elementary movement (saccades), but there is no denying that the vast majority of cerebellar output is destined to the cerebral cortex. These outputs affect many aspects of behavior, including cortical activity that supports decision-making. For example, fastigial (Gao et al., 2018) or dentate nucleus (Chabrol et al., 2019) stimulation biases the delay period activity in the motor cortex as an animal evaluates the evidence associated with choosing a movement. Disruption of the interpositus nucleus eliminates working memory related activity in the prefrontal cortex (Siegel and Mauk, 2013). In humans, damage along the vermis has been associated with autism. Children who suffer from autism spectrum disorder, a developmental disorder that leads to impairments in social and communication skills, exhibit anatomic abnormalities in their cerebellum, including reduced number of P-cells (Whitney et al., 2008), particularly along the vermis (Marko et al., 2015; Scott et al., 2009). How are these P-cells contributing to social and communications skills? The answer may be revealed once we understand the nature of the prediction errors that are sent from the olive to these P-cells.

Aizenman, C.D., Manis, P.B., Linden, D.J., 1998. Polarity of Long-Term Synaptic Gain Change Is Related to Postsynaptic Spike Firing at a Cerebellar Inhibitory Synapse. Neuron 21, 827–835. https://doi.org/10.1016/S0896-6273(00)80598-X

Andersen, R.A., Essick, G.K., Siegel, R.M., 1985. Encoding of spatial location by posterior parietal neurons. Science 230, 456–458.

Apps, R., Garwicz, M., 2005. Anatomical and physiological foundations of cerebellar information processing. Nat.Rev.Neurosci. 6, 297–311. https://doi.org/10.1038/nrn1646

Apps, R., Garwicz, M., 2000. Precise matching of olivo-cortical divergence and cortico-nuclear convergence between somatotopically corresponding areas in the medial C1 and medial C3 zones of the paravermal cerebellum. Eur.J.Neurosci. 12, 205–214. https://doi.org/10.1046/j.1460-9568.2000.00897.x

Barash, S., Melikyan, A., Sivakov, A., Zhang, M., Glickstein, M., Thier, P., 1999. Saccadic dysmetria and adaptation after lesions of the cerebellar cortex. J.Neurosci. 19, 10931–10939.

Becker, M.I., Person, A.L., 2019. Cerebellar Control of Reach Kinematics for Endpoint Precision. Neuron 103, 335–348. https://doi.org/10.1016/j.neuron.2019.05.007

Behan, M., Kime, N.M., 1996. Intrinsic circuitry in the deep layers of the cat superior colliculus. Vis.Neurosci. 13, 1031–1042.

Bengtsson, F., Ekerot, C.-F., Jörntell, H., 2011. In Vivo Analysis of Inhibitory Synaptic Inputs and Rebounds in Deep Cerebellar Nuclear Neurons. PLoS One 6. https://doi.org/10.1371/journal.pone.0018822

Blenkinsop, T.A., Lang, E.J., 2011. Synaptic action of the olivocerebellar system on cerebellar nuclear spike activity. J.Neurosci. 31, 14708–14720. https://doi.org/10.1523/JNEUROSCI.3323-11.2011

Brown, A.M., Arancillo, M., Lin, T., Catt, D.R., Zhou, J., Lackey, E.P., Stay, T.L., Zuo, Z., White, J.J., Sillitoe, R.V., 2019. Molecular layer interneurons shape the spike activity of cerebellar Purkinje cells. Sci.Rep. 9, 1742. https://doi.org/10.1038/s41598-018-38264-1

Buzunov, E., Mueller, A., Straube, A., Robinson, F.R., 2013. When during horizontal saccades in monkey does cerebellar output affect movement? Brain Res. 1503, 33–42. https://doi.org/10.1016/j.brainres.2013.02.001

Chabrol, F.P., Blot, A., Mrsic-Flogel, T.D., 2019. Cerebellar Contribution to Preparatory Activity in Motor Neocortex. Neuron 103, 506–519. https://doi.org/10.1016/j.neuron.2019.05.022

Chaumont, J., Guyon, N., Valera, A.M., Dugue, G.P., Popa, D., Marcaggi, P., Gautheron, V., Reibel-Foisset, S., Dieudonne, S., Stephan, A., Barrot, M., Cassel, J.C., Dupont, J.L., Doussau, F., Poulain, B., Selimi, F., Lena, C., Isope, P., 2013. Clusters of cerebellar Purkinje cells control their afferent climbing fiber discharge. Proc.Natl.Acad.Sci.U.S.A. 110, 16223–16228. https://doi.org/10.1073/pnas.1302310110

Chen, H., Hua, S.E., Smith, M.A., Lenz, F.A., Shadmehr, R., 2006. Effects of human cerebellar thalamus disruption on adaptive control of reaching. Cereb.Cortex 16, 1462–1473.

Criscimagna-Hemminger, S.E., Shadmehr, R., 2008. Consolidation patterns of human motor memory. J.Neurosci. 28, 9610–9618.

De Schutter, E., 1995. Cerebellar long-term depression might normalize excitation of Purkinje cells: a hypothesis. Trends in Neurosciences 18, 291–295. https://doi.org/10.1016/0166-2236(95)93916-L

De Zeeuw, C.I., Hoebeek, F.E., Bosman, L.W., Schonewille, M., Witter, L., Koekkoek, S.K., 2011. Spatiotemporal firing patterns in the cerebellum. Nat.Rev.Neurosci. 12, 327–344. https://doi.org/10.1038/nrn3011

De Zeeuw, C.I., Koekkoek, S.K., Wylie, D.R., Simpson, J.I., 1997. Association between dendritic lamellar bodies and complex spike synchrony in the olivocerebellar system. J.Neurophysiol. 77, 1747–1758. https://doi.org/10.1152/jn.1997.77.4.1747

De Zeeuw, C.L., Van Alphen, A.M., Hawkins, R.K., Ruigrok, T.J., 1997. Climbing fibre collaterals contact neurons in the cerebellar nuclei that provide a GABAergic feedback to the inferior olive. Neuroscience 80, 981–986. https://doi.org/10.1016/s0306-4522(97)00249-2

Ebner, T.J., Bloedel, J.R., 1981. Correlation between activity of Purkinje cells and its modification by natural peripheral stimuli. J.Neurophysiol. 45, 948–961. https://doi.org/10.1152/jn.1981.45.5.948

Eccles, J.C., Llinás, R., Sasaki, K., 1966. The excitatory synaptic action of climbing fibres on the Purkinje cells of the cerebellum. The Journal of Physiology 182, 268–296. https://doi.org/10.1113/jphysiol.1966.sp007824

Ekerot, C.F., Jorntell, H., Garwicz, M., 1995. Functional relation between corticonuclear input and movements evoked on microstimulation in cerebellar nucleus interpositus anterior in the cat. Exp.Brain Res. 106, 365–376.

Gao, Z., Davis, C., Thomas, A.M., Economo, M.N., Abrego, A.M., Svoboda, K., De Zeeuw, C.I., Li, N., 2018. A cortico-cerebellar loop for motor planning. Nature 563, 113–116. https://doi.org/10.1038/s41586-018-0633-x

Goffart, L., Chen, L.L., Sparks, D.L., 2003. Saccade dysmetria during functional perturbation of the caudal fastigial nucleus in the monkey. Ann.N.Y.Acad.Sci. 1004, 220–228.

Guerrasio, L., Quinet, J., Buttner, U., Goffart, L., 2010. Fastigial oculomotor region and the control of foveation during fixation. J.Neurophysiol. 103, 1988–2001. https://doi.org/10.1152/jn.00771.2009

Guo, C., Witter, L., Rudolph, S., Elliott, H.L., Ennis, K.A., Regehr, W.G., 2016. Purkinje Cells Directly Inhibit Granule Cells in Specialized Regions of the Cerebellar Cortex. Neuron 91, 1330–1341. https://doi.org/10.1016/j.neuron.2016.08.011

Han, K.S., Guo, C., Chen, C.H., Witter, L., Osorno, T., Regehr, W.G., 2018. Ephaptic Coupling Promotes Synchronous Firing of Cerebellar Purkinje Cells. Neuron 100, 564–578. https://doi.org/10.1016/j.neuron.2018.09.018

Hausser, M., Clark, B.A., 1997. Tonic synaptic inhibition modulates neuronal output pattern and spatiotemporal synaptic integration. Neuron 19, 665–678. https://doi.org/10.1016/s0896-6273(00)80379-7

Heck, D.H., De Zeeuw, C.I., Jaeger, D., Khodakhah, K., Person, A.L., 2013. The neuronal code(s) of the cerebellum. J.Neurosci. 33, 17603–17609. https://doi.org/10.1523/JNEUROSCI.2759-13.2013

Heck, D.H., Thach, W.T., Keating, J.G., 2007. On-beam synchrony in the cerebellum as the mechanism for the timing and coordination of movement. Proc.Natl.Acad.Sci.U.S.A 104, 7658–7663. https://doi.org/10.1073/pnas.0609966104

Heffley, W., Hall, C., 2019. Classical conditioning drives learned reward prediction signals in climbing fibers across the lateral cerebellum. eLife 8:e46764.

Heffley, W., Song, E.Y., Xu, Z., Taylor, B.N., Hughes, M.A., McKinney, A., Joshua, M., Hull, C., 2018. Coordinated cerebellar climbing fiber activity signals learned sensorimotor predictions. Nat Neurosci 21, 1431–1441. https://doi.org/10.1038/s41593-018-0228-8

Helmchen, C., Buttner, U., 1995. Saccade-related Purkinje cell activity in the oculomotor vermis during spontaneous eye movements in light and darkness. Exp.Brain Res. 103, 198–208.

Herzfeld, D.J., Hall, N.J., Tringides, M., Lisberger, S.G., 2020. Principles of operation of a cerebellar learning circuit. eLife 9, e55217. https://doi.org/10.7554/eLife.55217

Herzfeld, D.J., Kojima, Y., Soetedjo, R., Shadmehr, R., 2018. Encoding of error and learning to correct that error by the Purkinje cells of the cerebellum. Nat.Neurosci. 21, 736–743. https://doi.org/10.1038/s41593-018-0136-y

Herzfeld, D.J., Kojima, Y., Soetedjo, R., Shadmehr, R., 2015. Encoding of action by the Purkinje cells of the cerebellum. Nature 526, 439–442.

Hewitt, A.L., Popa, L.S., Ebner, T.J., 2015. Changes in Purkinje cell simple spike encoding of reach kinematics during adaption to a mechanical perturbation. J.Neurosci. 35, 1106–1124. https://doi.org/10.1523/JNEUROSCI.2579-14.2015

Hoebeek, F.E., Witter, L., Ruigrok, T.J.H., Zeeuw, C.I.D., 2010. Differential olivo-cerebellar cortical control of rebound activity in the cerebellar nuclei. PNAS 107, 8410–8415. https://doi.org/10.1073/pnas.0907118107

Hong, S., Negrello, M., Junker, M., Smilgin, A., Thier, P., De, S.E., 2016. Multiplexed coding by cerebellar Purkinje neurons. eLife 5. https://doi.org/10.7554/eLife.13810

Ju, C., Bosman, L.W.J., Hoogland, T.M., Velauthapillai, A., Murugesan, P., Warnaar, P., Genderen, R.M. van, Negrello, M., Zeeuw, C.I.D., 2019. Neurons of the inferior olive respond to broad classes of sensory input while subject to homeostatic control. The Journal of Physiology 597, 2483–2514. https://doi.org/10.1113/JP277413

Junker, M., Endres, D., Sun, Z.P., Dicke, P.W., Giese, M., Thier, P., 2018. Learning from the past: A reverberation of past errors in the cerebellar climbing fiber signal. PLOS Biology 16, e2004344. https://doi.org/10.1371/journal.pbio.2004344

Khaliq, Z.M., Raman, I.M., 2005. Axonal Propagation of Simple and Complex Spikes in Cerebellar Purkinje Neurons. J. Neurosci. 25, 454–463. https://doi.org/10.1523/JNEUROSCI.3045-04.2005

Khosrovani, S., Giessen, R.S.V.D., Zeeuw, C.I.D., Jeu, M.T.G.D., 2007. In vivo mouse inferior olive neurons exhibit heterogeneous subthreshold oscillations and spiking patterns. PNAS 104, 15911–15916. https://doi.org/10.1073/pnas.0702727104

Kim, J.J., Krupa, D.J., Thompson, R.F., 1998. Inhibitory cerebello-olivary projections and blocking effect in classical conditioning. Science 279, 570–573.

Kitazawa, S., Kimura, T., Yin, P.B., 1998. Cerebellar complex spikes encode both destinations and errors in arm movements. Nature 392, 494.

Kojima, Y., Iwamoto, Y., Yoshida, K., 2004. Memory of learning facilitates saccadic adaptation in the monkey. J.Neurosci. 24, 7531–7539.

Kojima, Y., Robinson, F.R., Soetedjo, R., 2014. Cerebellar fastigial nucleus influence on ipsilateral abducens activity during saccades. J.Neurophysiol. 111, 1553–1563. https://doi.org/10.1152/jn.00567.2013

Kojima, Y., Soetedjo, R., 2018. Elimination of the error signal in the superior colliculus impairs saccade motor learning. Proc.Natl.Acad.Sci.U.S.A 115, E8987–E8995. https://doi.org/10.1073/pnas.1806215115

Kojima, Y., Soetedjo, R., 2017. Change in sensitivity to visual error in superior colliculus during saccade adaptation. Sci.Rep. 7, 9566. https://doi.org/10.1038/s41598-017-10242-z

Kojima, Y., Soetedjo, R., Fuchs, A.F., 2010a. Effects of GABA agonist and antagonist injections into the oculomotor vermis on horizontal saccades. Brain Res. 1366, 93–100. https://doi.org/10.1016/j.brainres.2010.10.027

Kojima, Y., Soetedjo, R., Fuchs, A.F., 2010b. Changes in simple spike activity of some Purkinje cells in the oculomotor vermis during saccade adaptation are appropriate to participate in motor learning. J.Neurosci. 30, 3715–3727. https://doi.org/10/3715

Kording, K.P., Tenenbaum, J.B., Shadmehr, R., 2007. The dynamics of memory as a consequence of optimal adaptation to a changing body. Nat.Neurosci. 10, 779–786.

Kralj-Hans, I., Baizer, J.S., Swales, C., Glickstein, M., 2007. Independent roles for the dorsal paraflocculus and vermal lobule VII of the cerebellum in visuomotor coordination. Exp.Brain Res. 177, 209–222. https://doi.org/10.1007/s00221-006-0661-x

Lang, E.J., Apps, R., Bengtsson, F., Cerminara, N.L., De Zeeuw, C.I., Ebner, T.J., Heck, D.H., Jaeger, D., Jörntell, H., Kawato, M., Otis, T.S., Ozyildirim, O., Popa, L.S., Reeves, A.M.B., Schweighofer, N., Sugihara, I., Xiao, J., 2017. The Roles of the Olivocerebellar Pathway in Motor Learning and

Motor Control. A Consensus Paper. Cerebellum 16, 230–252. https://doi.org/10.1007/s12311-016-0787-8

Larry, N., Yarkoni, M., Lixenberg, A., Joshua, M., 2019. Cerebellar climbing fibers encode expected reward size. eLife 8. https://doi.org/10.7554/eLife.46870

Leigh, R.J., Zee, D.S., 2015. The neurology of eye movements. Oxford University Press.

Lu, H., Yang, B., Jaeger, D., 2016. Cerebellar Nuclei Neurons Show Only Small Excitatory Responses to Optogenetic Olivary Stimulation in Transgenic Mice: In Vivo and In Vitro Studies. Front Neural Circuits 10. https://doi.org/10.3389/fncir.2016.00021

Mano, N., Yamamoto, K., 1980. Simple-spike activity of cerebellar Purkinje cells related to visually guided wrist tracking movement in the monkey. J.Neurophysiol. 43, 713–728.

Marko, M.K., Crocetti, D., Hulst, T., Donchin, O., Shadmehr, R., Mostofsky, S.H., 2015. Behavioural and neural basis of anomalous motor learning in children with autism. Brain 138, 784–797. https://doi.org/10.1093/brain/awu394

Mauk, M.D., Donegan, N.H., 1997. A model of Pavlovian eyelid conditioning based on the synaptic organization of the cerebellum. Learn.Mem. 4, 130–158.

McElvain, L.E., Bagnall, M.W., Sakatos, A., du Lac, S., 2010. Bidirectional Plasticity Gated by Hyperpolarization Controls the Gain of Postsynaptic Firing Responses at Central Vestibular Nerve Synapses. Neuron 68, 763–775. https://doi.org/10.1016/j.neuron.2010.09.025

Medina, J.F., 2011. The multiple roles of Purkinje cells in sensori-motor calibration: to predict, teach and command. Current Opinion in Neurobiology, Sensory and motor systems 21, 616–622. https://doi.org/10.1016/j.conb.2011.05.025

Medina, J.F., Garcia, K.S., Mauk, M.D., 2001. A mechanism for savings in the cerebellum. J.Neurosci. 21.

Medina, J.F., Mauk, M.D., 1999. Simulations of cerebellar motor learning: computational analysis of plasticity at the mossy fiber to deep nucleus synapse. J.Neurosci. 19, 7140–7151.

Miles, F.A., Lisberger, S.G., 1981. Plasticity in the vestibulo-ocular reflex: a new hypothesis. Annu.Rev.Neurosci. 4, 273–299. https://doi.org/10.1146/annurev.ne.04.030181.001421

Monsivais, P., Clark, B.A., Roth, A., Häusser, M., 2005. Determinants of Action Potential Propagation in Cerebellar Purkinje Cell Axons. J. Neurosci. 25, 464–472. https://doi.org/10.1523/JNEUROSCI.3871-04.2005

Muller, S.Z., Zadina, A.N., Abbott, L.F., Sawtell, N.B., 2019. Continual Learning in a Multi-Layer Network of an Electric Fish. Cell. https://doi.org/10.1016/j.cell.2019.10.020

Munoz, D.P., Istvan, P.J., 1998. Lateral inhibitory interactions in the intermediate layers of the monkey superior colliculus. J.Neurophysiol. 79, 1193–1209.

Najac, M., Raman, I.M., 2018. Synaptic excitation by climbing fibre collaterals in the cerebellar nuclei of juvenile and adult mice. The Journal of Physiology 6703–6718. https://doi.org/10.1113/JP274598@10.1002/(ISSN)1469-7793(CAT)VirtualIssues(VI)EditorsChoice

Najafi, F., Giovannucci, A., Wang, S.S., Medina, J.F., 2014. Coding of stimulus strength via analog calcium signals in Purkinje cell dendrites of awake mice. eLife 3, e03663. https://doi.org/10.7554/eLife.03663

O'Hearn, E., Molliver, M.E., 1997. The Olivocerebellar Projection Mediates Ibogaine-Induced Degeneration of Purkinje Cells: A Model of Indirect, Trans-Synaptic Excitotoxicity. J. Neurosci. 17, 8828–8841. https://doi.org/10.1523/JNEUROSCI.17-22-08828.1997

Ohmae, S., Medina, J.F., 2015. Climbing fibers encode a temporal-difference prediction error during cerebellar learning in mice. Nat.Neurosci. 18, 1798–1803. https://doi.org/10.1038/nn.4167

Ohyama, T., Mauk, M.D., 2001. Latent acquisition of timed responses in cerebellar cortex. J.Neurosci. 21, 682–690.

Ohyama, T., Nores, W.L., Mauk, M.D., 2003. Stimulus generalization of conditioned eyelid responses produced without cerebellar cortex: implications for plasticity in the cerebellar nuclei. Learn.Mem. 10, 346–354.

Payne, H.L., French, R.L., Guo, C.C., Nguyen-Vu, T.B., Manninen, T., Raymond, J.L., 2019. Cerebellar Purkinje cells control eye movements with a rapid rate code that is invariant to spike irregularity. eLife 8. https://doi.org/10.7554/eLife.37102

Pekny, S.E., Criscimagna-Hemminger, S.E., Shadmehr, R., 2011. Protection and expression of human motor memories. J.Neurosci. 31, 13829–13839. https://doi.org/10.1523/JNEUROSCI.1704-11.2011

Person, A.L., Raman, I.M., 2012a. Purkinje neuron synchrony elicits time-locked spiking in the cerebellar nuclei. Nature 481, 502–505. https://doi.org/10.1038/nature10732

Person, A.L., Raman, I.M., 2012b. Synchrony and neural coding in cerebellar circuits. Front Neural.Circuits. 6, 97. https://doi.org/10.3389/fncir.2012.00097

Person, A.L., Raman, I.M., 2010. Deactivation of L-type Ca Current by Inhibition Controls LTP at Excitatory Synapses in the Cerebellar Nuclei. Neuron 66, 550–559. https://doi.org/10.1016/j.neuron.2010.04.024

Pugh, J.R., Raman, I.M., 2008. Mechanisms of Potentiation of Mossy Fiber EPSCs in the Cerebellar Nuclei by Coincident Synaptic Excitation and Inhibition. J. Neurosci. 28, 10549–10560. https://doi.org/10.1523/JNEUROSCI.2061-08.2008

Rasmussen, A., Hesslow, G., 2014. Feedback control of learning by the cerebello-olivary pathway. Prog.Brain Res. 210, 103–119. https://doi.org/10.1016/B978-0-444-63356-9.00005-4

Rasmussen, A., Jirenhed, D.A., Zucca, R., Johansson, F., Svensson, P., Hesslow, G., 2013. Number of spikes in climbing fibers determines the direction of cerebellar learning. J.Neurosci. 33, 13436–13440. https://doi.org/10.1523/JNEUROSCI.1527-13.2013

Rasmussen, A., Zucca, R., Johansson, F., Jirenhed, D.A., Hesslow, G., 2015. Purkinje cell activity during classical conditioning with different conditional stimuli explains central tenet of Rescorla-Wagner model [corrected]. Proc.Natl.Acad.Sci.U.S.A 112, 14060–14065. https://doi.org/10.1073/pnas.1516986112

Raymond, J.L., Medina, J.F., 2018. Computational Principles of Supervised Learning in the Cerebellum. Annual Review of Neuroscience 41, 233–253. https://doi.org/10.1146/annurev-neuro-080317-061948

Robinson, F.R., Straube, A., Fuchs, A.F., 1993. Role of the caudal fastigial nucleus in saccade generation. II. Effects of muscimol inactivation. J.Neurophysiol. 70, 1741–1758.

Ruigrok, T.J.H., Voogd, J., 2000. Organization of projections from the inferior olive to the cerebellar nuclei in the rat. Journal of Comparative Neurology 426, 209–228. https://doi.org/10.1002/1096-9861(20001016)426:2<209::AID-CNE4>3.0.CO;2-0

Saint-Cyr, J.A., Courville, J., 1982. Descending projections to the inferior olive from the mesencephalon and superior colliculus in the cat. An autoradiographic study. Exp.Brain Res. 45, 333–348.

Sarwary, A.M.E., Wischnewski, M., Schutter, D.J.L.G., Selen, L.P.J., Medendorp, W.P., 2018. Corticospinal correlates of fast and slow adaptive processes in motor learning. Journal of Neurophysiology 120, 2011–2019. https://doi.org/10.1152/jn.00488.2018

Schiller, D., Monfils, M.H., Raio, C.M., Johnson, D.C., LeDoux, J.E., Phelps, E.A., 2010. Preventing the return of fear in humans using reconsolidation update mechanisms. Nature 463, 49–53.

Scott, J.A., Schumann, C.M., Goodlin-Jones, B.L., Amaral, D.G., 2009. A comprehensive volumetric analysis of the cerebellum in children and adolescents with autism spectrum disorder. Autism Res. 2, 246–257. https://doi.org/10.1002/aur.97

Sedaghat-Nejad, E., Herzfeld, D.J., Hage, P., Karbasi, K., Palin, T., Wang, X., Shadmehr, R., 2019. Behavioral training of marmosets and electrophysiological recording from the cerebellum. J.Neurophysiol. 122, 1502–1517.

Shin, S.L., De, S.E., 2006. Dynamic synchronization of Purkinje cell simple spikes. J.Neurophysiol. 96, 3485–3491. https://doi.org/10.1152/jn.00570.2006

Siegel, J.J., Mauk, M.D., 2013. Persistent activity in prefrontal cortex during trace eyelid conditioning: dissociating responses that reflect cerebellar output from those that do not. J.Neurosci. 33, 15272–15284. https://doi.org/10.1523/JNEUROSCI.1238-13.2013

Slemmer, J.E., De Zeeuw, C.I., Weber, J.T., 2005. Don't get too excited: mechanisms of glutamate-mediated Purkinje cell death, in: Progress in Brain Research, Creating Coordination in the Cerebellum. Elsevier, pp. 367–390. https://doi.org/10.1016/S0079-6123(04)48029-7

Smith, M.A., Ghazizadeh, A., Shadmehr, R., 2006. Interacting adaptive processes with different timescales underlie short-term motor learning. PLoS.Biol. 4, e179.

Soetedjo, R., Fuchs, A.F., 2006. Complex spike activity of purkinje cells in the oculomotor vermis during behavioral adaptation of monkey saccades. J Neurosci. 26, 7741–7755.

Soetedjo, R., Fuchs, A.F., Kojima, Y., 2009. Subthreshold activation of the superior colliculus drives saccade motor learning. J.Neurosci. 29, 15213–15222. https://doi.org/10.1523/JNEUROSCI.4296-09.2009

Soetedjo, R., Kojima, Y., Fuchs, A.F., 2008. Complex spike activity in the oculomotor vermis of the cerebellum: a vectorial error signal for saccade motor learning? J.Neurophysiol. 100, 1949–1966.

Stollhoff, N., Menzel, R., Eisenhardt, D., 2005. Spontaneous recovery from extinction depends on the reconsolidation of the acquisition memory in an appetitive learning paradigm in the honeybee (Apis mellifera). J.Neurosci. 25, 4485–4492.

Sugihara, I., 2011. Compartmentalization of the Deep Cerebellar Nuclei Based on Afferent Projections and Aldolase C Expression. Cerebellum 10, 449–463. https://doi.org/10.1007/s12311-010-0226-1

Sugihara, I., Fujita, H., Na, J., Quy, P.N., Li, B.-Y., Ikeda, D., 2009. Projection of reconstructed single purkinje cell axons in relation to the cortical and nuclear aldolase C compartments of the rat cerebellum. Journal of Comparative Neurology 512, 282–304. https://doi.org/10.1002/cne.21889

Sugihara, I., Shinoda, Y., 2007. Molecular, Topographic, and Functional Organization of the Cerebellar Nuclei: Analysis by Three-Dimensional Mapping of the Olivonuclear Projection and Aldolase C Labeling. J. Neurosci. 27, 9696–9710. https://doi.org/10.1523/JNEUROSCI.1579-07.2007

Sugihara, I., Wu, H., Shinoda, Y., 1999. Morphology of single olivocerebellar axons labeled with biotinylated dextran amine in the rat. J.Comp.Neurol. 414, 131–148. https://doi.org/10.1002/(SICI)1096-9861(19991115)414:2<131::AID-CNE1>3.0.CO;2-F

Sugihara, I., Wu, H.-S., Shinoda, Y., 2001. The Entire Trajectories of Single Olivocerebellar Axons in the Cerebellar Cortex and their Contribution to Cerebellar Compartmentalization. J. Neurosci. 21, 7715–7723. https://doi.org/10.1523/JNEUROSCI.21-19-07715.2001

Suvrathan, A., Payne, H.L., Raymond, J.L., 2016. Timing Rules for Synaptic Plasticity Matched to Behavioral Function. Neuron 92, 959–967. https://doi.org/10.1016/j.neuron.2016.10.022

Tang, T., Blenkinsop, T.A., Lang, E.J., 2019. Complex spike synchrony dependent modulation of rat deep cerebellar nuclear activity. eLife 8. https://doi.org/10.7554/eLife.40101

Tang, T., Suh, C.Y., Blenkinsop, T.A., Lang, E.J., 2016. Synchrony is Key: Complex Spike Inhibition of the Deep Cerebellar Nuclei. Cerebellum. 15, 10–13. https://doi.org/10.1007/s12311-015-0743-z

Telgkamp, P., Raman, I.M., 2002. Depression of inhibitory synaptic transmission between Purkinje cells and neurons of the cerebellar nuclei. J.Neurosci. 22, 8447–8457.

ten Brinke, M.M., Heiney, S.A., Wang, X., Proietti-Onori, M., Boele, H.-J., Bakermans, J., Medina, J.F., Gao, Z., De Zeeuw, C.I., 2017. Dynamic modulation of activity in cerebellar nuclei neurons during pavlovian eyeblink conditioning in mice. eLife 6, e28132. https://doi.org/10.7554/eLife.28132

Teune, T.M., Burg, J. van der, Zeeuw, C.I. de, Voogd, J., Ruigrok, T.J.H., 1998. Single Purkinje cell can innervate multiple classes of projection neurons in the cerebellar nuclei of the rat: A light microscopic and ultrastructural triple-tracer study in the rat. Journal of Comparative Neurology 392, 164–178. https://doi.org/10.1002/(SICI)1096-9861(19980309)392:2<164::AID-CNE2>3.0.CO;2-0

Thach, W.T., 1967. Somatosensory receptive fields of single units in cat cerebellar cortex. J.Neurophysiol. 30, 675–696. https://doi.org/10.1152/jn.1967.30.4.675

Thier, P., Dicke, P.W., Haas, R., Barash, S., 2000. Encoding of movement time by populations of cerebellar Purkinje cells. Nature 405, 72–76.

Tomatsu, S., Ishikawa, T., Tsunoda, Y., Lee, J., Hoffman, D.S., Kakei, S., 2016. Information processing in the hemisphere of the cerebellar cortex for control of wrist movement. J.Neurophysiol. 115, 255–270. https://doi.org/10.1152/jn.00530.2015

Viaro, R., Bonazzi, L., Maggiolini, E., Franchi, G., 2017. Cerebellar Modulation of Cortically Evoked Complex Movements in Rats. Cereb.Cortex 27, 3525–3541. https://doi.org/10.1093/cercor/bhw167

Vilis, T., Hore, J., 1980. Central neural mechanisms contributing to cerebellar tremor produced by limb perturbations. J.Neurophysiol. 43, 279–291.

Welsh, J.P., Lang, E.J., Suglhara, I., Llinas, R., 1995. Dynamic organization of motor control within the olivocerebellar system. Nature 374, 453–457. https://doi.org/10.1038/374453a0

Whitney, E.R., Kemper, T.L., Bauman, M.L., Rosene, D.L., Blatt, G.J., 2008. Cerebellar Purkinje cells are reduced in a subpopulation of autistic brains: a stereological experiment using calbindin-D28k. Cerebellum. 7, 406–416. https://doi.org/10.1007/s12311-008-0043-y

Wise, A.K., Cerminara, N.L., Marple-Horvat, D.E., Apps, R., 2010. Mechanisms of synchronous activity in cerebellar Purkinje cells. J.Physiol 588, 2373–2390. https://doi.org/10.1113/jphysiol.2010.189704

Witter, L., Rudolph, S., Pressler, R.T., Lahlaf, S.I., Regehr, W.G., 2016. Purkinje Cell Collaterals Enable Output Signals from the Cerebellar Cortex to Feed Back to Purkinje Cells and Interneurons. Neuron 91, 312–319. https://doi.org/10.1016/j.neuron.2016.05.037

Xu-Wilson, M., Chen-Harris, H., Zee, D.S., Shadmehr, R., 2009. Cerebellar contributions to adaptive control of saccades in humans. J.Neurosci. 29, 12930–12939.

Yang, Y., Lisberger, S.G., 2014a. Role of Plasticity at Different Sites across the Time Course of Cerebellar Motor Learning. J. Neurosci. 34, 7077–7090. https://doi.org/10.1523/JNEUROSCI.0017-14.2014

Yang, Y., Lisberger, S.G., 2014b. Purkinje-cell plasticity and cerebellar motor learning are graded by complex-spike duration. Nature. 510, 529–532. https://doi.org/10.1038/nature13282

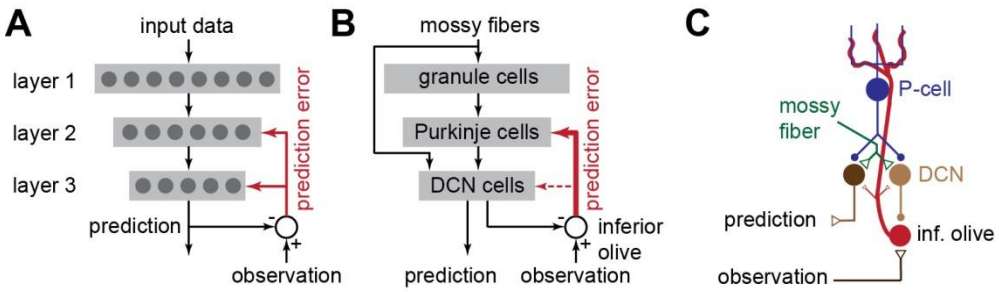

Figure 1. A feedforward network as a simplified model of the cerebellum. **A**. In an artificial neural network with three layers, the third layer provides the predictions, which are compared to observations and then fed back via an error signal to units in layers 2 and 3. **B**. In the cerebellum, the 3-layers are comprised of granule cells, Purkinje cells, and deep cerebellar nucleus neurons. The predictions of the cerebellum are conveyed via GABA-ergic DCN neurons to the inferior olive, which in turn provides the cerebellum with an error signal. This signal is conveyed strongly to the P-cells, but weakly to the DCN neurons (dashed line). The predictions of the cerebellum are also conveyed via non-GABA-ergic DCN neurons to the rest of the brain. **C**. A single P-cell projects to both GABA-ergic and non-GABA-ergic DCN neurons. However, the error signal sent from the olive to the cerebellum depends directly on the GABA-ergic DCN neurons, but not on the non-GABA-ergic neurons. Filled circles are inhibitory synapses, triangles are excitatory synapses.

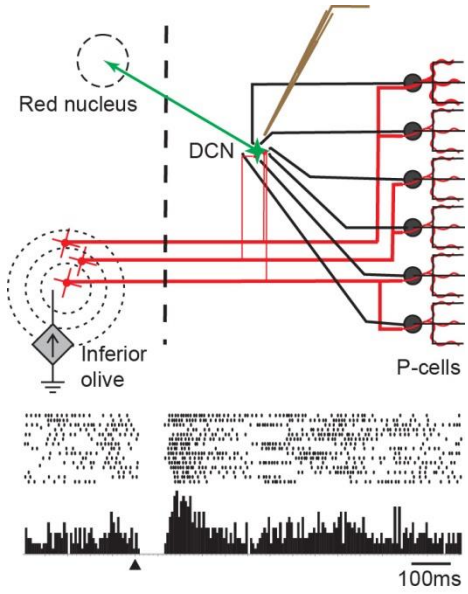

Figure 2. Effect of inferior olive stimulation on DCN neurons (interposed nucleus) in an anesthetized rat. Neurons in the inferior olive were stimulated while an electrode recorded activity of an excitatory neuron that projected to the red nucleus. Stimulation of the inferior olive produced a 50 ms pause in the activity of the DCN neuron. From Hoebeek et al. (2010), used by permission.

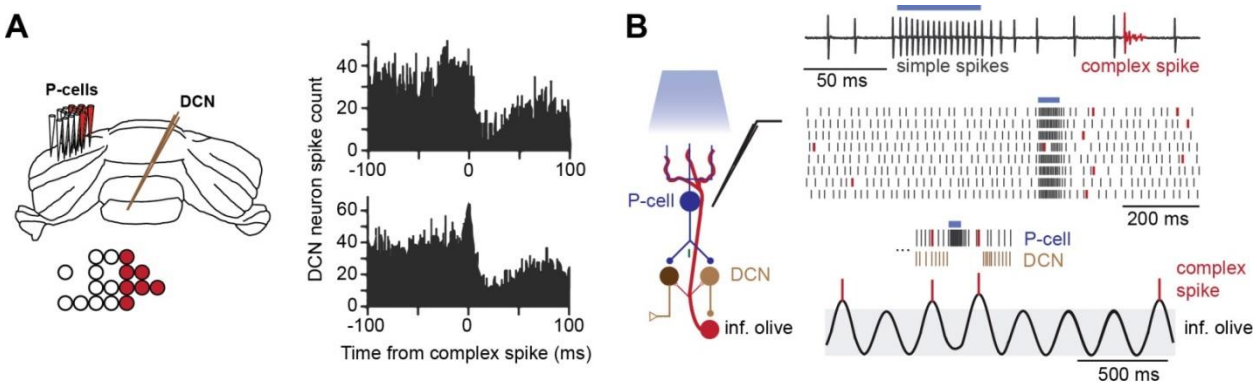

Figure 3. An olivary neuron that projects to a P-cell receives input from the DCN daughter neuron of that P-cell. **A**. Electrode array recordings from multiple P-cells and a single DCN neuron in an anesthetized rat. To determine if a P-cell projected to the DCN neuron, the complex spike in the P-cell was used as the trigger to align the spikes in the DCN neuron. The data on the right are from two instances in which the spike-triggered averaging suggested a P-cell to DCN projection. In 70% of the pairs the complex spike in the P-cell was followed by a suppression in the DCN neuron (top right). In 21% of the pairs there was an initial spike in the DCN neuron followed by suppression. The P-cells that appeared connected to a single DCN neuron were usually clustered together along the rostra-caudal axis of the cerebellum, as shown on the left with red-filled electrodes in the array. From Blenkinsop and Lang (2011) and Tang et al. (2019), used by permission. **B**. Optogenetic stimulation of P-cells is followed by production of a complex spike. P-cells were optogenetically stimulated while activity was recorded with an electrode. Stimulation resulted in intense production of simple spikes, which was then followed at a latency of about 100 ms with a complex spike. A model suggests that P-cell stimulation strongly suppresses the DCN neuron, which in turn removes inhibition from the olivary neuron. The olivary neuron has a membrane potential that oscillates. Removal of inhibition allows the potential to reach threshold sooner, resulting in climbing fiber activity in the parent P-cell, and thus a complex spike. From Chaumont et al. (2013), used by permission.

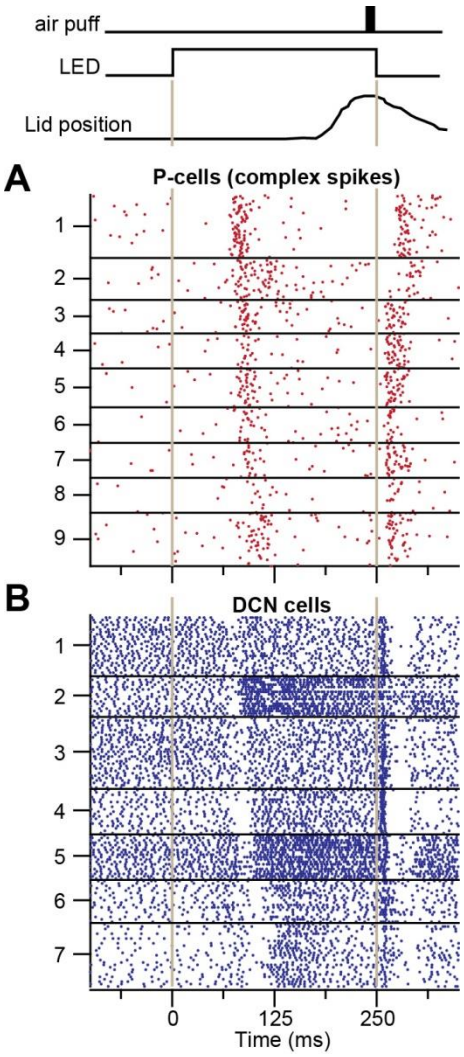

Figure 4. During learning of a behavior, some DCN neurons show a transient pause that appear related to arrival of error information via a complex spike in their parent P-cells. Recordings are from P-cells and neurons in the interposed nucleus as mice learned to associate an LED with an airpuff. The LED onset and the airpuff both tended to produce complex spikes in P-cells, which coincided with a transient pause in DCN neurons. From ten Brinke et al. (2017), used with permission.

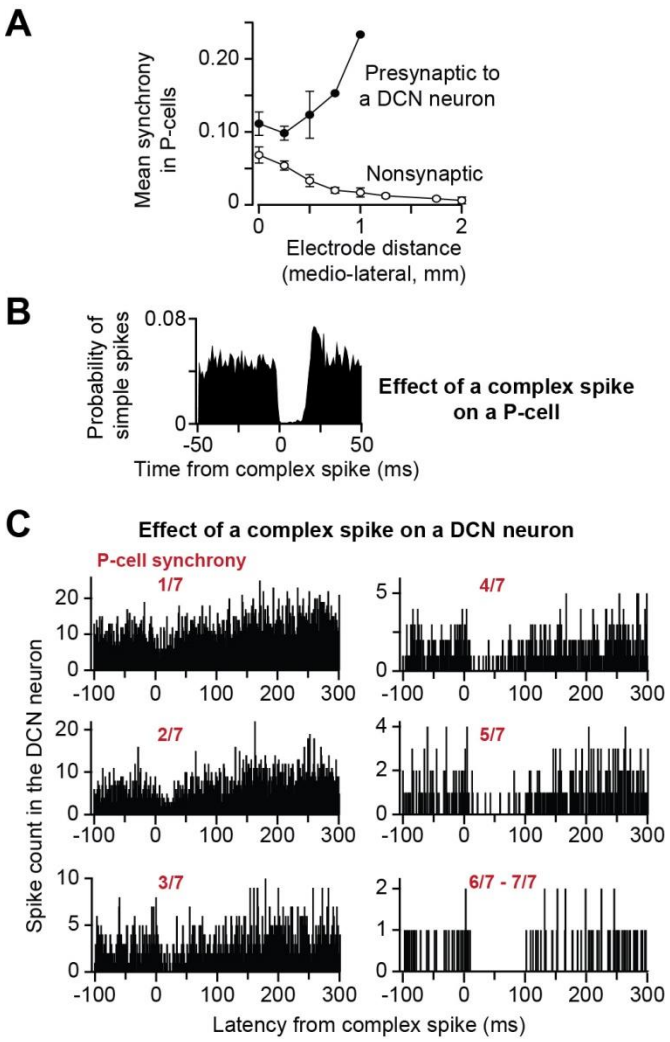

Figure 5. Effective transmission of the error signal to the nucleus requires complex spike synchrony. **A**. Groups of P-cells that project to a single DCN neuron tend to have higher levels of complex spike synchrony than groups of P-cells that do not project to the same neuron. From Tang et al. (2019), used with permission. **B**. Following a complex spike, there is a 10-20 ms period of total simple spike suppression in the P-cell. Data are from awake behaving marmoset. From Sedaghat-Nejad et al. (2019), used with permission. **C**. Effect of a complex spike is greater if the event is synchronized among the parent P-cells. From Tang et al. (2019), used with permission.

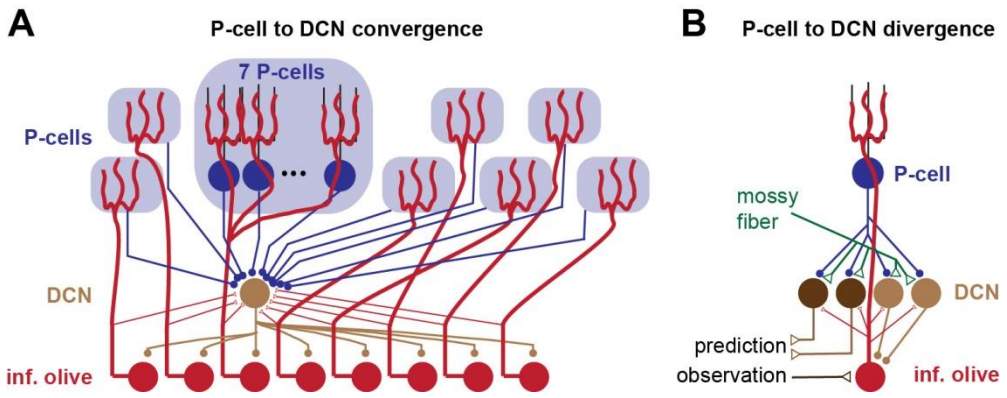

Figure 6. Membership criterion for population coding in the cerebellum. **A**. The population of P-cells that project onto a nucleus neuron is composed of those P-cells that receive climbing fibers from the olivary neurons that also project to that specific nucleus neuron. **B**. A single P-cell projects to only a few DCN neurons, which include both GABA-ergic and non-GABA-ergic cells. Efficient learning requires that the error signal that the P-cell receives represent these errors only. Here, the GABA-ergic DCN daughters of a given P-cell converge their axons on a single olive neuron, and that olive neuron serves as the teacher for the parent P-cell, providing it with its single climbing fiber. The olive cell in term sends excitatory projections to all the DCN daughters of the P-cell.

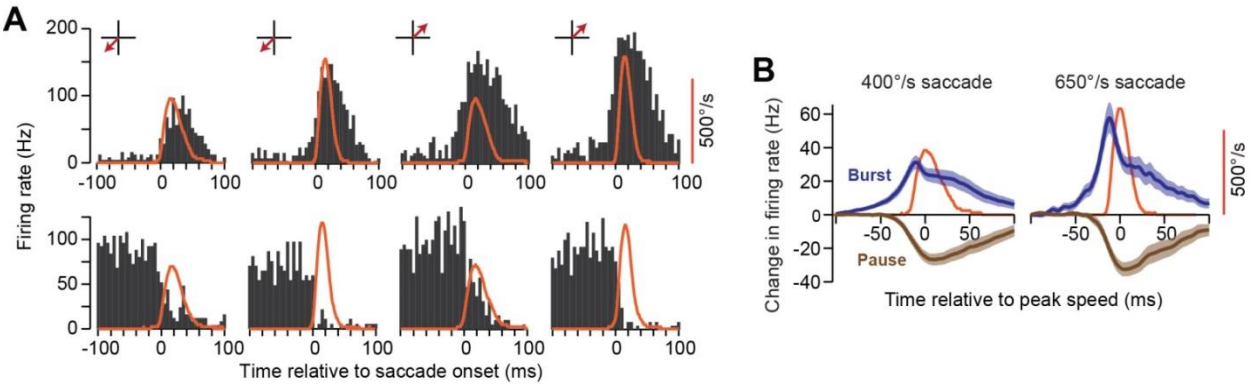

Figure 7. Activity of P-cells during a saccade. **A**. Saccades were made to various directions with various amplitudes and velocities. The red trace shows the eye velocity, and black bars indicate instantaneous firing rates of two P-cells. Each row is a single P-cell. Some P-cells exhibited a burst during the saccade, while others exhibited a pause. However, the duration of modulation tended to outlast the saccade by 50 ms or more. **B**. Averaged activity of burst and pause type P-cells during a saccade. From Herzfeld et al. (2015), used with permission.

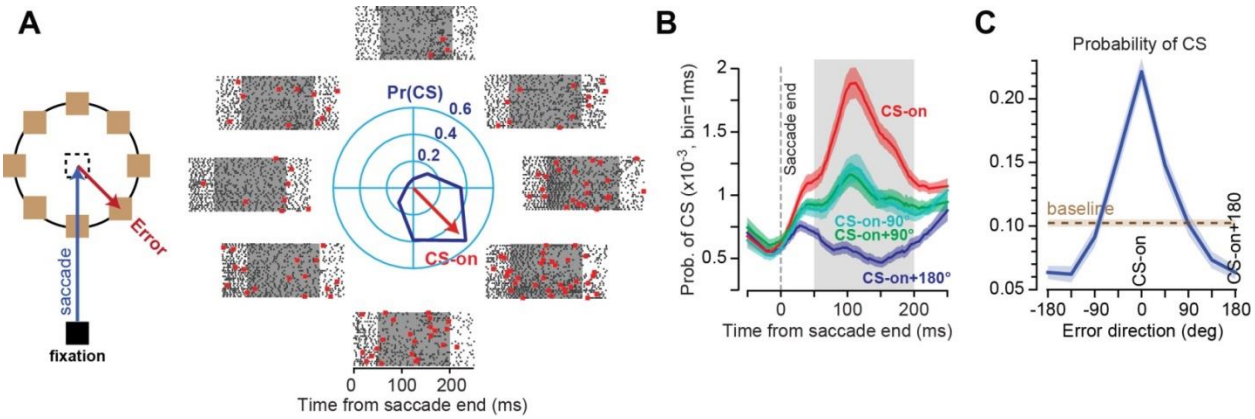

Figure 8. P-cells in the oculomotor vermis exhibit complex spike tuning with respect to visual error. **A**. Animals made a saccade to a target location and experienced a visual error following conclusion of the movement. The data show the simple (black dots) and complex spikes (red dots) of a P-cell in the period following saccade end. The center polar plot describes the probability of complex spike in the 50-200 ms period following saccade end. The direction of error that induced the greatest probability of complex spikes is noted by the red vector labeled CS-on. **B**. Probability of complex spike during the time following saccade end. Data are averaged from 72 P-cells. **C**. Probability of complex spike was averaged across 72 P-cells, aligned to CS-on. From Herzfeld et al. (2015), used with permission.

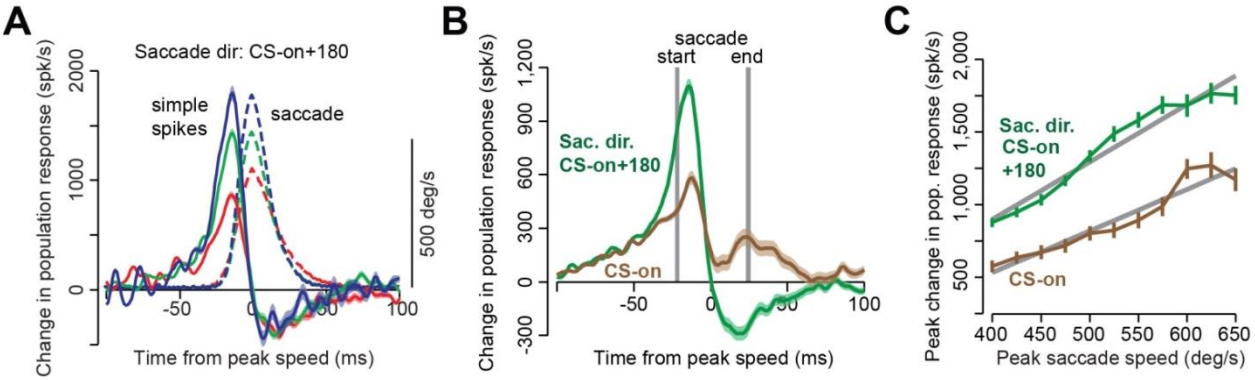

Figure 9. Simple spike population response of P-cells during a saccade. **A**. P-cells were organized into 50 member populations and their simple spike rates were summed as a saccade was made in direction CS-on+180. Saccades associated with three different peak velocities are shown. **B**. Simple spike population response are shown across all saccades made in direction CS-on and CS-on+180. C. Peak simple spike population response increased linearly with peak saccade velocity, but with a higher gain when the saccade was in direction CS-on+180. From Herzfeld et al. (2015), used with permission.

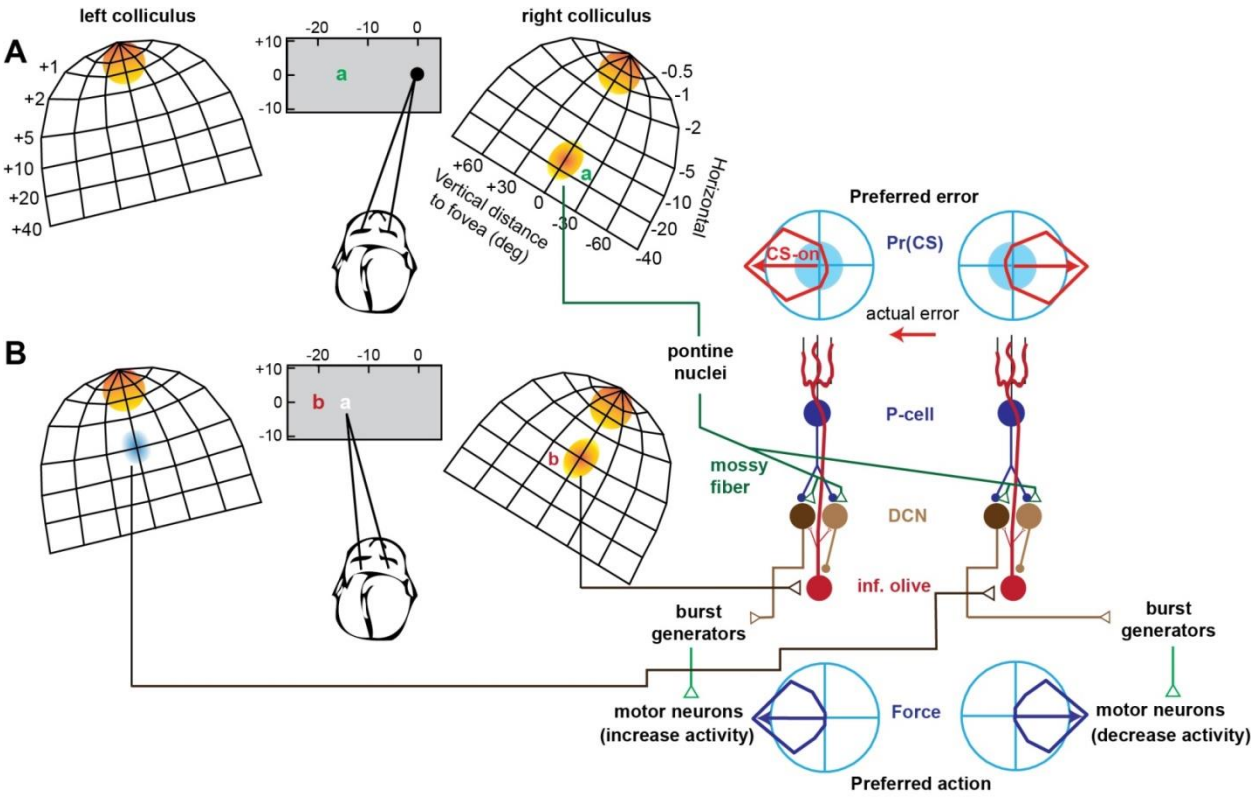

Figure 10. Hypothetical relationship between preferred error of a P-cell and the projections of its DCN daughters. **A**. A model of activity in the superior colliculus as a subject views a fixation point and a target is presented to the left at location "a". During fixation, foveal related neurons in the rostral part of both colliculi are active. Presentation of the target activates neurons in the caudal region of the right colliculus (orange, labeled "a"), which results in a saccade. A copy of the saccade related activity is sent via mossy fibers to the cerebellum. **B**. At saccade end, the target is expected to be on the fovea, but is in fact at location "b". This produces activity in a region labeled "b" in the right colliculus, which in turn inhibits a region on the contralateral colliculus. The unexpected activity on the right colliculus engages olivary neurons on the contralateral side, which produce a complex spike among P-cells the left of the oculomotor vermis. The unexpected suppression on the left colliculus also affected the contralateral olive, which reduces the complex spike probability for P-cells to the right of the vermis. Presence of complex spike for the left P-cells slightly suppresses the parallel fiber synapses that were activated during the preceding saccade. The lack of complex spike for the right P-cells slightly increases the weight of the parallel fiber synapse. The P-cells on the left project to DCN neurons that project to burst generators that activate motoneurons that pull the eyes to the left. On the next saccade, presentation of the same mossy fiber input now produces a slight suppression of P-cells on the left (with respect to the previous trial), which results in greater force production to the left. Similarly, the P-cells on the right side of the vermis produce slightly more simple spikes, and this results in production of slightly less rightward force. Hence, if there is a correspondence between the preferred action of the DCN neurons and the preferred error of their P-cell parents, a prediction error results in learning that correctly compensates for the experienced error. That is, the DCN daughters should project to a group of neurons that can produce an effect that will remedy the specific error that is of concern to the parent P-cells.

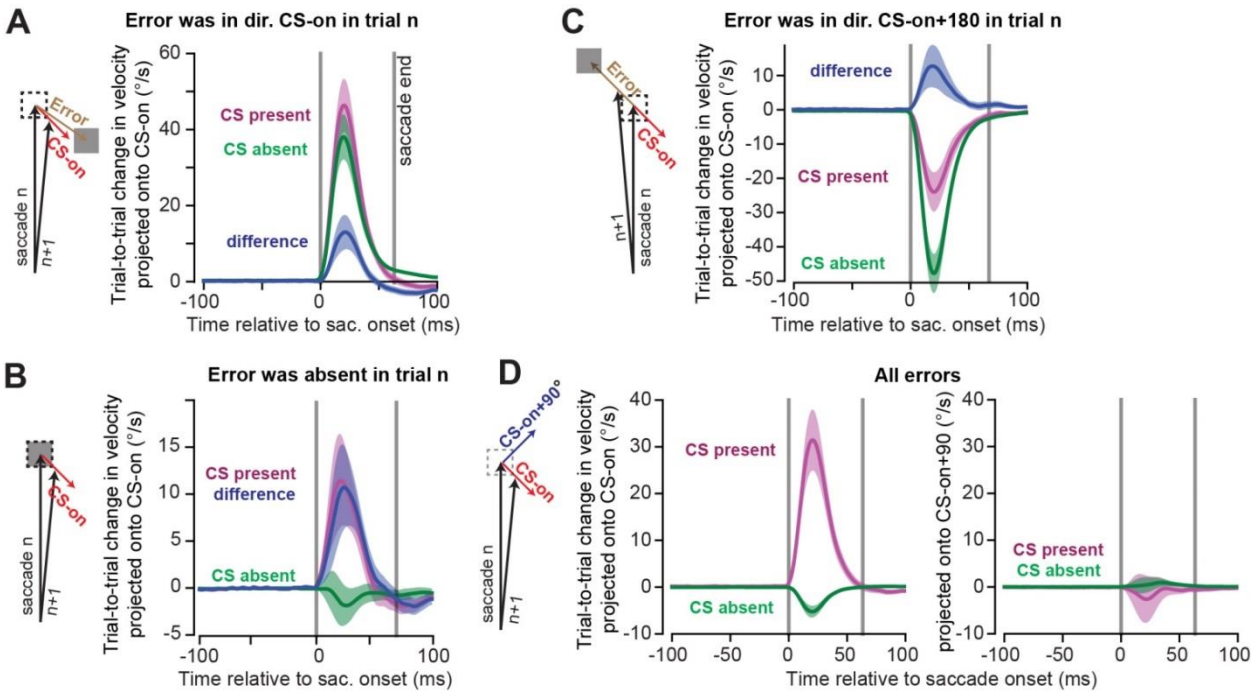

Figure 11. The influence of P-cells on motor output. **A.** Analysis of trials in which the error was in the CS-on direction of the P-cell. Eye trajectory (velocity vectors as a function of time) was measured during the saccade in the trial in which the error was experienced, as well as the trajectory in the subsequent trial in which saccade was made to the same target. To measure change in behavior, the trial-to-trial difference in the 2D sequence of velocity vectors was projected onto the CS-on direction of the P-cell. Following experience of an error, the next trial exhibited an increase in the velocity vector along the direction of error. However, if the P-cell produced a complex spike, velocity on the next trial was larger in direction CS-on compared to when a complex spike was absent. **B.** Same as (A) except for trials in which there was no post-saccadic error (|error| < 0.25°). Even without an error, presence of a complex spike led to increased motor output along direction CS-on of the P-cell that had produced the complex spike. **C.** Same as (A), except for trials in which errors were in CS-on+180° of the P-cell under study. If the error was opposite the preferred direction, presence of a complex spike still biased behavior in CS-on direction of that cell. **D.** Trial-to-trial change following any error. The trial-to-trial change in velocity was projected onto direction CS-on (left) and CS-on+90° (right) of the P-cell. When a CS was present, change in motor output was only in direction CS-on. From Herzfeld et al. (2018), used by permission.

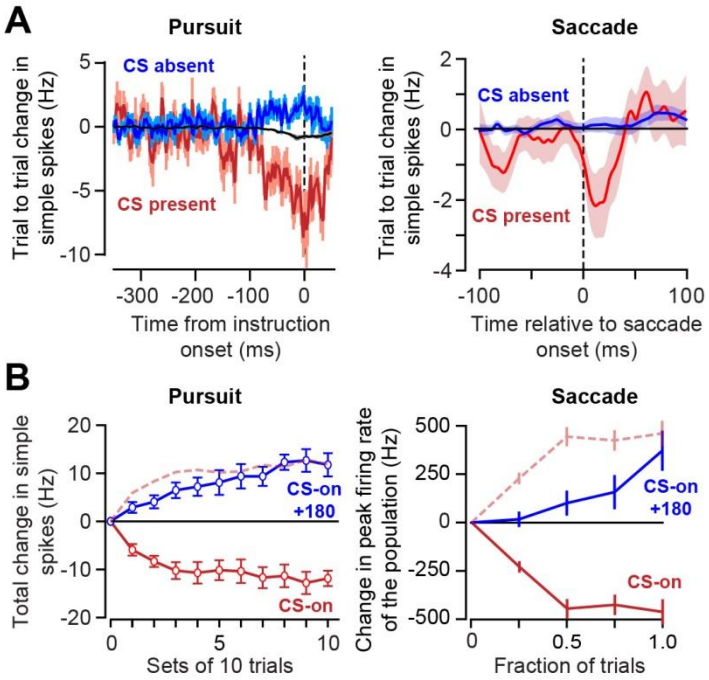

Figure 12. Change in simple spikes following presence or absence of a complex spike. Data are from pursuit and saccade tasks. **A**. Trial-to-trial change in simple spikes. In the pursuit task, the inter-trial interval was 2.5 sec. Time zero indicates the time during pursuit in which an unexpected change in the visual stimulus occurred. The black trace is the baseline relationship during pursuit. In the saccade task, change in simple spikes are plotted with respect to saccade onset. **B**. Total amount of change in simple spikes when the errors are in CS-on or CS-on+180 direction. For the pursuit task, the change is the expected value for a single P-cell, averaged across all P-cells. For the saccade task, the change is the sum for 50 P-cells that make an average population. The dashed lines are the CS-on data mirrored along the x-axis. It appears that the rate of change in simple spikes is faster when errors are in direction CS-on. From Yang and Lisberger (2014), and Herzfeld et al. (2018), used by permission.