A vector calculus for neural computation in the cerebellum

Mohammad Amin Fakharian, Alden M. Shoup, Paul Hage, Hisham Y. Elseweifi, and Reza Shadmehr

Laboratory for Computational Motor Control, Dept. of Biomedical Engineering
Johns Hopkins School of Medicine, Baltimore, Maryland 21205, USA

Correspondence: Reza Shadmehr or Mohammad Amin Fakharian, Johns Hopkins School of Medicine, 410 Traylor Building, 720 Rutland Ave., Baltimore, MD 21205.
Email: shadmehr@jhu.edu or mafakharian@gmail.com

Acknowledgements: The work was supported by grants from the NIH (R01-EB028156, and R37-NS128416).

Author contributions: MAF, AMS, and PH collected the data, RS conceived experiments, MAF, HYE, and RS analyzed the data, MAF and HYE performed simulations, made the figures, and performed statistical analysis, RS and MAF wrote the manuscript.

Competing interests: None.

ORCID ID: Mohammad Amin Fakharian 0000-0002-7741-8689, Alden Shoup 0000-0003-1206-6587, Paul Hage 0000-0002-7006-9877, Hisham Elseweifi 0009-0005-8008-5829, Reza Shadmehr 0000-0002-7686-2569.

When a neuron modulates its firing rate during a movement, we tend to assume that it is contributing to control of that movement. However, null space theory makes the counter-intuitive prediction that neurons often generate spikes not to cause behavior, but to prevent the effects that other neurons would have on behavior. What is missing is a direct way to test this theory in the brain. Here, we found that in the marmoset cerebellum, spike-triggered averaging identified a potent vector unique to each Purkinje cell (P-cell) along which the spikes of that cell displaced the eyes. Yet, the P-cells were active not just for saccades along their potent vector, but for all saccades. Simultaneous recordings revealed that two spikes in two different P-cells produced superposition of their potent vectors. The result was a population activity in which the spikes were canceled if their contributions were perpendicular to the intended movement. But why was one neuron producing spikes that would be canceled by another neuron? To answer this question, we recorded from the mossy fiber inputs and the interneurons in the molecular layer. We found that the mossy fibers provided a copy of the motor commands as well as the sensory goal of the movement, then the interneurons transformed these inputs so that the P-cells as a population predicted when the movement had reached the goal and should be stopped. Because this output had faster dynamics than was present in the individual neurons, the cerebellum placed the cells in subtractive competition with each other such that the downstream effects of spiking in one neuron were partially or completely eliminated because of the spiking in another neuron.

The principal theory currently employed for understanding neuronal computation in the brain is that of null spaces (*1*). This theory was developed to solve a puzzle: why are neurons in the motor regions of the cerebral cortex active not just during production of a movement, but also during its preparation and after its conclusion? The theory proposes that each neuron contributes to a movement with a single vector of weights that represents the effective connection strength of that neuron to the various muscles (*2*). A neuron may be active, but this activity does not necessarily produce a movement because the weighted combination of spikes in the population produces a sum that is zero during certain periods, resulting in cancelation. Thus, the critical elements of the theory are the weight vectors that are assigned to each neuron.

The problem is that these weights are currently found through a model that maps the neuronal activities to movements, necessarily producing cancelation when there is no movement (*2*, *3*). But how can we know whether the weights are more than just the fitted parameters of an equation? That is, how can we test the idea that in the brain, one neuron engages in apparently wasteful production of spikes despite the fact that the downstream effects of those spikes will be nullified by the spikes of another neuron? Indeed, why are the neurons producing these wasteful spikes in the first place?

If each weight vector is a reflection of functional anatomy, then two predictions must hold: 1) when each neuron's weight vector is estimated not via model fitting, but independently via spike-triggered averaging, then the resulting sum of neuronal activities in the population should cancel before and after the movement, but not during the movement, and 2) simultaneous spikes in two neurons must obey the principle of superposition, as predicted by their weight vectors.

Here, we consider a region of the cerebellum concerned with eye movements. We first find the weight vector for each Purkinje cell (P-cell) independently by relying on its climbing fiber input, which acts as a stochastic perturbation to the P-cell. Then, we use these *a priori* weights to map the computation in the network from its inputs via mossy fibers, to the molecular layer interneurons, and finally to the P-cell outputs. We find that in its input, the network receives information regarding the sensory goal of the movement alongside a copy of the motor commands, and in its output predicts when and how to stop that movement. However, this output is generated not in any one group of neurons, but via a population in which one group cancels the downstream effects of the spikes in another group. The results provide an answer to the question of why this cancelation is not wasteful but necessary: by relying on a subtractive process, the cerebellum generates fast neural dynamics in the output of the population despite having individual neurons that are burdened by slow changing firing rates.

Spike interactions organized neurons in the cerebellar cortex into discrete populations

We trained marmosets to make saccades to visual targets (Fig. 1A) and used Neuropixels and other silicon probes to record from 520 definitive P-cells, 798 putative MLIs (pMLIs), and 1716 putative mossy fibers (pMFs) from lobule VII of the vermis (Supplementary Figs. S1, S2, & S3). Fig. 1B presents an example session in which we simultaneously isolated over 100 neurons. We identified the P-cells based on their simple and complex spikes (SS and CS) and then used the distinct shapes of the CS and SS waveforms in the dendritic tree, soma, and axon of each P-cell (*4*, *5*) to identify the molecular, P-cell, and granular layers, respectively. For example, we identified the molecular layer via the downward, broad CS waveforms in the dendrite tree of the P-cells (channels 58-70, CS traces, Fig. 1B). In the granular layer we looked for spike waveforms that exhibited an "m" shape with a slow negative after-wave and labeled those cells as pMF glomeruli (*6*–*8*). In the molecular layer, we labeled neurons that inhibited the P-cells

at 1ms latency or sooner as putative MLI1s (pMLI1), and neurons that inhibited pMLI1s, excited a P-cell, and experienced excitation from CS spillover as pMLI2s (*9*–*11*).

The probe recorded from multiple folia simultaneously, allowing us to ask whether the various neurons clustered into "cliques," where each clique was composed of neurons that exhibited strong spike interactions with each other but not with the other simultaneously recorded neurons (*12*). To quantify these interactions, we computed the jitter corrected (*13*) conditional probability that one cell produced a spike at time $t + \Delta$, given that another cell fired at time t (Fig. 1D). We then formed an adjacency matrix (Fig. 1C) where the numerical value in each element of the matrix was the strength of the spike interaction between the two cells. Finally, we applied graph spectral clustering to the values in the matrix and identified boundaries that divided the neurons into small subgroups of strongly connected cells, labeled as cliques. This revealed that clique membership was largely but not exclusively due to the folium boundary. For example, most cells near the 1st P-cell layer were members of clique 1, and most cells near the 2nd P-cell layer were members of clique 2 (exceptions are noted as outliers in Fig. 1C).

We found that if two P-cells belonged to the same clique, then their SSs strongly synchronized: an SS in one P-cell accompanied ~14 Hz increase in the SS firing rate of another P-cell at 0 ms delay, then a ~5 Hz decrease in firing rate at 3 ms delay (Fig. 1D, 1st row). Strikingly, this synchrony was absent when the P-cells belonged to different cliques (within vs. between, mean±SEM, SS|SS: 13.6±0.5 vs. 0.4±0.04 Hz, p<1e-100, MLI|MLI: 5.2±0.11 vs 1.04±0.02, p<1e-100, MLI2|MLI2: 1.17±0.26 vs. 0.86±0.1, p=2e-1, CS|CS: 1.04±0.07 vs. 0.54±0.03, p<1e-12, SS|MLI1 @1ms: -7.38±0.17 vs. -0.26±0.03, p<1e-100, MLI1|MLI2 @1ms: -2.43±0.23 vs 0.14±0.04, p<1e-50, SS|MLI2 @3ms: 1.24±0.19 vs 0.12±0.06. P<1e-11). Similarly, if a P-cell and a pMLI1 were in the same clique, a spike in the pMLI1 produced a reduction of ~7 Hz in the firing rate of the P-cell at 1 ms delay (Fig. 1D, 3rd row). In comparison, there was a reduction of 1 Hz when the two cells were from disparate cliques. If the pMLI2 and pMLI1 were in the same cliques, a spike in the pMLI2 produced an inhibition of ~2 Hz in the pMLI1, but little or no inhibition if the two were in different cliques. Finally, the climbing fiber input produced a spillover effect upon the pMLI2s (*9*, *10*), as well as an ephaptic coupling suppression on the neighboring P-cells (*5*), but only if these neurons belonged to the same clique as the climbing fiber (Fig. 1C, 4th row).

We next asked whether clique membership also stratified the neurons based on the information that they received in their inputs from the inferior olive. To quantify this input, we relied on the fact that following presentation of a visual target at a random direction, the climbing fibers informed the P-cells regarding the direction of that visual event (*14*–*18*) (Fig. 2A). In addition, when a movement was about to be made, the climbing fibers reported the direction of the upcoming movement (*18*). For each P-cell i, the climbing fiber encoding of the visuomotor events was along a preferred axis: the CS rates were maximum when the event was in direction θ_i and minimum when the event was in direction $\theta_i + \pi$ (Fig. 2A, right plot).

The strength with which the olivary input transmitted visuomotor events differed among the cliques. In some cliques, the climbing fiber inputs exhibited a strong tuning (Fig. 1B, clique 1), while for other cliques this tuning was weak (Fig. 1B, clique 5). We represented the information content in each climbing fiber as a vector by fitting a von Mises distribution to the CS rates as a function of direction of the visuomotor event (Fig. 2A, right subplot). This produced a vector with direction θ_i and amplitude ρ_i for P-cell i. Within a clique, these climbing fiber vectors were more similar to each other than between cliques (Supplementary Fig. S4, median±MAD: within: 23.25±17.22 deg, between: 57.73±36.05deg, rank sum test p<1E-17; angle difference from center, mean±std: -1.39±29.99 deg p=0.57).

In summary, the P-cells and the interneurons formed cliques (Fig. 1E) in which the neurons strongly interacted with each other. Each clique received climbing fiber inputs that encoded visuomotor events along a vector with direction θ and amplitude ρ. Our task was to determine what these populations computed.

The climbing fiber input identified a potent vector for each P-cell

The climbing fiber input completely suppressed the P-cell's ability to produce simple spikes (Supplementary Fig. S2A), acting as a stochastic perturbation. Previous work had shown that when this suppression took place during a movement, it slightly diverted that movement (*19–21*). We performed spike-triggered averaging, comparing saccades that had the same starting position and were made toward the same target, but did or did not experience a CS-triggered SS suppression (Fig. 2B). We evaluated the difference in the eye trajectories in velocity space and observed that this difference produced a vector that throughout the duration of the saccade remained parallel to CS-on (Fig. 2C, Supplementary Fig. S5). Moreover, when two P-cells were suppressed during the same saccade, the amplitude of the disturbance more than doubled. Thus, when P-cell i was suppressed via its climbing fiber input, the result was activation of the extraocular eye muscles such that a displacement was produced in direction θ_i. We represented this displacement as a weight $\boldsymbol{w_i} = \left[w_x^i, w_y^i\right]^T$, i.e., the horizontal and vertical displacement.

To assign a magnitude to this vector, we considered the possibility that suppression of the P-cells that belonged to a clique with a strong CS tuning (i.e., large ρ) produced a greater pull on the eyes, as compared to other cliques that had a weak CS tuning (Fig. 2D). Indeed, we found that the greater the amplitude ρ, the greater the displacement caused by the SS suppression (Fig. 2D, right subplot, projection on θ, ANOVA F= 27.3, p = 5.3E-12). We assigned the following vector to each P-cell and labeled it as its potent vector:

$$\boldsymbol{w_i} = -\frac{\rho_i}{\sum \rho_i}\left[\cos(\theta_i), \sin(\theta_i)\right] \quad (1)$$

Thus, P-cell i was assigned a weight vector with direction θ_i and amplitude $\frac{\rho_i}{\sum \rho_i}$ (the negative sign indicates suppression of SS).

P-cells obeyed superposition in the space of their potent vectors

We next asked whether the downstream effects of suppressing two P-cells simultaneously could be represented as a linear sum of their potent vectors. This was a test of superposition: when P-cells i and j were suppressed, was the resulting displacement described by the vector $\boldsymbol{w_i} + \boldsymbol{w_j}$? For all pairs of simultaneously recorded P-cells, we selected pairs with $\frac{\rho_i}{\rho_j} \geq 0.5$, then measured the difference in the angles of the two potent vectors $\boldsymbol{w_i}$ and $\boldsymbol{w_j}$, i.e., $\Delta\theta_{ij} = \theta_i - \theta_j$, then used the resulting distribution of $\Delta\theta_{ij}$ to divide the pairs into three groups (Fig. 2E). For example, in group 1, $\Delta\theta_{ij} \approx 0$ (i.e., $\boldsymbol{w_i}$ and $\boldsymbol{w_j}$ were approximately parallel, whereas for group 2, $\Delta\theta_{ij} \approx 70$ deg, i.e., $\boldsymbol{w_i}$ and $\boldsymbol{w_j}$ were on average 70 degrees apart. Indeed, simultaneous suppression of two P-cells pulled the eyes along a vector with magnitude and direction that was predicted by $\boldsymbol{w_i} + \boldsymbol{w_j}$ (red vector, Fig. 2F), demonstrating superposition.

The potent vector defined the axis of symmetry for the simple spikes

Some P-cells increased their SS rates around saccade onset, while others decreased their rates (Fig. 3A, top row). Given that the downstream effect of SS suppression for both groups was to pull the eyes in direction θ (Supplementary Fig. 5B) (*20*), it was puzzling that the modulations were present for all saccade directions and continued long after the movement had ended.

However, we noticed that if we represented the SS rates $r_i(t)$ as a function of saccade direction ψ with respect to the direction of the potent vector θ, then a remarkable pattern emerged: when the saccade was in direction $\psi = \theta_i + \frac{\pi}{2}$, the SS rates $r_i(\psi, t)$ were nearly identical to when the saccade was in direction $\theta_i - \frac{\pi}{2}$. As a result, a within P-cell subtraction of the rates in these two directions produced, on average, a near complete cancelation (Fig. 3B, bottom row). Indeed, this pattern held true for pairs of saccades that were at an equal angle with respect to the potent axis θ (Supplementary Fig. S6). As a result, we made the following inference regarding the within cell symmetry in the simple spikes:

$$r_i(\theta_i + \psi, t) = r_i(\theta_i - \psi, t) \quad (2)$$

Notably, Eq. (2) held true in both the burster and the pauser P-cells. Thus, just as the CS rates were symmetric with respect to θ (Fig. 2D) (a finding that is consistent in both marmosets and macaques (*15*)), so were the SS rates.

In contrast, when one saccade was in direction $\psi = \theta_i$ and another along $\psi = \theta_i + \pi$, now the within neuron difference in the SS rates produced a burst-pause pattern that crossed zero at the deceleration onset of the saccade (Fig. 3B, second column). Remarkably, this pattern held true in the bursters as well as the pausers, despite the fact that the pausers never produced an increase in their activity. Moreover, while the SS rates for saccades in all directions in both the bursters and pausers had remained modulated long after the movement had ended, in both groups the extra spikes were canceled for saccades parallel to the potent vector, producing a burst-pause pattern that returned to zero as the saccade ended (Fig. 3B).

The potent vector defined a population code for the simple spikes

According to null space theory, a neuron's contribution to a movement is described by a constant weight, which is usually found via the parameters of an equation that maps firing rates to behavior (*2*). Here, rather than fitting such a model, we had the weights *a priori:* when a movement was made in direction ψ, P-cell i produced simple spikes with the rate $r_i(\psi, t)$, and its contribution to behavior was defined by its potent vector $\boldsymbol{w_i}$:

$$\boldsymbol{p}_i = \left[w_x^i, w_y^i\right]^T r_i(\psi, t) = -\frac{\rho_i}{\sum_i \rho_i}\left[\cos(\theta_i), \sin(\theta_i)\right]^T r_i(\psi, t) \quad (3)$$

Because direction θ_i depended on the neuron's anatomical location in the cerebellum (*15*, *18*), we removed our sampling bias (Supplementary Fig. S4D) by assuming isotropy for θ, duplicating each P-cell uniformly as 8 P-cells:

$$\tilde{\boldsymbol{p}}_i(\psi, t) = -\frac{\rho_i}{\sum_i \rho_i}\sum_{n=0}^{7}\left[\cos\left(\theta_i + \frac{n\pi}{4}\right), \sin\left(\theta_i + \frac{n\pi}{4}\right)\right]^T r_i\left(\psi + \frac{n\pi}{4}, t\right) \quad (4)$$

The output of the population was simply the weighted sum of activity in all neurons:

$$\boldsymbol{p}(\psi, t) = \sum_i \tilde{\boldsymbol{p}}_i(\psi, t) \quad (5)$$

Without loss of generality, consider a movement in direction $\psi = 0$. The vector $\boldsymbol{p}(t)$ has two components: $p_x(t)$, the SS rates that move the eyes along the direction of the movement (i.e., target

direction), and $p_y(t)$, the SS rates that move the eyes perpendicular to it. We found that $p_x(t)$ exhibited a positive-negative pattern that crossed zero at the onset of the saccade's deceleration (Fig. 3C, all P-cells, left panel). In contrast, $p_y(t)$ remained near zero throughout the movement. Thus, despite the fact that in the individual neurons the SS rates were modulated for saccades in all directions, in the population the downstream effects of the SSs were canceled so that they had no effects perpendicular to the direction of the movement (Fig. 3C, right panel). Despite the fact that the SSs for individual P-cells were modulated long after the saccade ended, the downstream effects were canceled beyond the end of the movement (Fig. 3C, left panel). Finally, despite the fact that some P-cells exhibited a burst, while others exhibited a pause, in both groups $p_x(t)$ exhibited a positive-negative pattern that crossed zero at deceleration onset, and in both groups $p_y(t)$ remained near zero (Fig. 3C).

A P-cell's contribution to behavior was defined primarily by the magnitude of its potent vector, not the magnitude of its SS modulation

It is natural to assume that when a neuron's firing rates are strongly modulated during a behavior, it is contributing significantly to the control of that behavior. However, our theoretical framework makes a counter-intuitive prediction: because the downstream effects of the spikes in one neuron may be canceled by another neuron, what matters is not how much the SS rates were modulated, but how large was the potent vector. For example, consider two groups of P-cells that had roughly equal SS rate modulation during saccades, but different ρ: one group had a large ρ (i.e., large amplitude potent vector), the other a small ρ (Supplementary Fig. S7A). For each P-cell, ρ indicates the sensitivity of the CS rates for sensorimotor events in directions θ vs. $\theta + \pi$. Because the complex spikes are long-term teachers that guide sculpting of the simple spikes (*16*, *22*), a large ρ implies that the teacher dissociates these two sensorimotor events strongly, which implies that the SS rates should be different. Therefore, the theory predicts that in the population, for P-cells that have a small ρ, the SS rates $r(\theta, t)$ will be similar to $r(\theta + \pi, t)$, resulting in greater cancelation than for P-cells that have a large ρ.

To test this prediction, in our two groups of large and small ρ P-cells, we set $\rho = 1$ for all cells, and then computed the population response via Eq. (5). Whereas the two groups had roughly equal SS rate modulations as a function of movement direction, the population response of the large potent vector group was roughly twice as large as the small potent vector group (Supplementary Fig. S7B). Thus, among the P-cells with small potent vectors, despite their large modulation of simple spikes during saccades, most of these spikes were canceled.

The potent vector also defined a population code for the interneurons

In both types of MLIs, the rates in all saccade directions exhibited a burst that continued long after the movement had ended (Fig. 3D). The activities of these interneurons were difficult to understand in the context of anatomy: after all, pMLI2s inhibited pMLI1s, and pMLI1s inhibited the P-cells. Why were the interneurons bursting in all directions?

Each clique had P-cells with a potent vector $\boldsymbol{w}$. We used this vector for the parent clique to analyze the firing patterns of the MLIs. We found that the potent axis of the P-cells was inherited by the MLIs such that the axis of symmetry roughly held true for the interneurons. That is, when one saccade was in direction $\psi = \theta$ and another in direction $\psi = \theta + \pi$, the within-cell difference in the rates was a pause-burst pattern in the pMLI1s, and a burst pattern in pMLI2s (Fig. 3E). Moreover, when the two saccades were in directions $\psi = \theta + \pi/2$ and $\psi = \theta - \pi/2$, the difference in the two rates was nearly zero for pMLI2s and missing the burst in pMLI1s.

Thus, to estimate the population response of the pMLIs, we applied the same analysis as for the P-cells: we used Eqs. (4) and (5) and assigned to each pMLI the potent vector of its clique. Now we found that in the population of pMLI1s, in the direction of the target the response was a pause-burst pattern (Fig. 3F, left panel), i.e., roughly the opposite of the pattern in the P-cells. In the pMLI2s, the population response was a burst that returned to zero at deceleration onset. Plotting these population rates of the various neurons against one another (Fig. 3G) unmasked a negative correlation, consistent with the inhibitory influence that the pMLI2s have on pMLI1s, and the inhibitory influence that the pMLI1s have on the P-cells.

Importantly, the population response in pMLI2s perpendicular to the target remained nearly zero during the saccade (Fig. 3F, right panel). However, this was not the case for pMLI1s, an inconsistency that may indicate that pMLI1s were composed of unknown subtypes (*23*).

In summary, when we used the climbing fiber input to assign a potent vector to the interneurons of each clique, the results revealed computations consistent with inhibitory influence of pMLI2s on pMLI1s, and pMLI1s on the P-cells.

The weight vectors revealed the computations performed in this region of the cerebellum

The burst-pause pattern in the P-cells implied a computation associated with predicting when to stop the saccade. For this to be possible, this region of the cerebellum should receive two kinds of information: a copy of the motor commands in muscle coordinates, and a copy of the goal location in visual coordinates. To examine this hypothesis, we recorded from n=1716 pMFs (Supplementary Fig. S8) and found one group of inputs that provided information regarding the goal of the movement in visual coordinates, termed goal pMF (Fig. 4A, and Supplementary Fig. S9), and another group that encoded the motor commands, termed state pMF (Fig. 4D, and Supplementary Fig. S10).

The goal pMFs generally exhibited buildup activity before saccade onset with a relatively narrow spatial response that encoded the amplitude of the visual target (Fig. 4B, left subfigure). However, the goal information was present only if a reward relevant visual target was the destination of the saccade. When a similar saccade was made spontaneously without this target, the goal pMFs exhibited a muted response (Fig. 4A and 4B). In contrast, the state pMFs reported a faithful copy of the motor commands regardless of whether the movement was reward relevant or not (Fig. 4D). Indeed, a simple integration of the number of spikes in the state pMFs was an excellent predictor of the saccade amplitude in all conditions (Fig. 4C).

These results suggested a working model (Fig. 4F). When the saccade was aimed at a reward relevant target, the cerebellar cortex was made aware of the target location x_g via the goal pMFs. The cerebellar network integrated the spikes in the state pMFs to a bound set by the goal pMFs. Once the bound was reached, the P-cells signaled to the nucleus via disinhibition to produce forces that would aid in stopping the movement. The model predicted that when the goal information was withheld from the cerebellum (i.e., reward irrelevant saccades), the integration would still occur, but the P-cells would not be able to provide a stopping signal.

To test this prediction, we binned the movements based on saccade peak velocity and acceleration durations and found that regardless of saccade velocity, amplitude, or acceleration duration, in reward relevant saccades the SS rates transitioned from burst to pause at deceleration onset (Fig. 4G, first column, Fig. 4H & 4I, Supplementary Fig. S12). The one exception was for very long saccade acceleration periods. In these saccades, the velocity during the acceleration phase was double-peaked, and both the goal and state pMFs showed a double peak in their rate profiles, indicating rare saccades in

which mid-flight of the movement, the goal was changed (Supplementary Fig. S12). For these "double saccades", the SS rates crossed zero before deceleration onset, but for all other reward relevant saccades the zero-crossing remained locked to deceleration onset. Moreover, regardless of saccade velocity, amplitude, or acceleration duration, if the target was reward relevant then the firing rates of pMLI1s reached a constant peak value (Fig. 4H, first column, and Fig. 4I & 4J).

Critically, these features were missing for the irrelevant saccades. For example, in the irrelevant saccades the SS rates exhibited an increase as a function of velocity but did not cross zero (Fig. 4H, right column, Supplementary Figs. S12 & S13). Also, the firing rates of pMLI1s reached half the peak for reward irrelevant saccades (Fig. 4G, second column and Supplementary Figs. S12 & S13).

Taken together, the results suggest a conceptual model in which the interneurons in this region of the cerebellum integrated the motor commands $u(t)$ that were present in the state MFs to predict displacement $\hat{x}(t)$, which was then compared to the target location x_g provided by the goal MFs (Fig. 4F). The fact that the firing rates of pMLI1s reached a constant peak for all reward relevant saccades suggested that pMLI1s played a dominant role in setting the bound, directing the P-cells to disinhibit the nucleus and stop the movement.

Discussion

We uncovered the logic of computation in a region of the cerebellum that is concerned with control of eye movements, finding evidence that neurons often produced spikes not to affect behavior, but to prevent the effects that other neurons would have on behavior.

The P-cells, the interneurons, and the climbing fiber inputs exhibited spike interactions that grouped them into discrete networks, called cliques. Using the climbing fiber input, we found that each P-cell pulled the eyes along a specific 2D vector, called the potent vector. However, the P-cells and the interneurons were modulated not just for saccades along their potent vector, but for all saccades. Indeed, their modulation continued long after the movement had ended. Critically, when two P-cells with different potent vectors were simultaneously suppressed, the downstream effect on behavior was summation of the individual vectors. This raised the possibility that some neurons were modulated not to contribute to a movement, but to eliminate another neuron's unwanted contributions to that movement.

A critical clue emerged in the symmetry that was present in the P-cell SS rates: the modulations were nearly equal when two saccades had the same absolute value of angular displacement with respect to the cell's potent vector. The purpose of this symmetry was revealed in the population. During a saccade, as each P-cell generated its simple spikes, pulling the eyes along its potent vector, the superposition of the potent vectors in the population removed the component that would have pulled the eyes perpendicular to the direction of movement. What remained was a vector in the direction of the movement, exhibiting a fast-changing burst-pause pattern that, regardless of saccade amplitude or velocity, consistently crossed zero at the onset of deceleration. Thus, it appeared that the P-cells as a population produced an output that specified when and how to stop the movement.

For a neural circuit to make such a prediction it would need to be provided in its inputs at least two kinds of information: the goal location in sensory coordinates, and the real-time commands in motor coordinates. Indeed, the mossy fiber inputs separately provided each piece of information, but only if the planned movement was reward relevant. For the irrelevant saccades, the state mossy fibers continued to provide accurate motor command information, but the goal information was missing. Now, without the goal, the population output of the P-cells was a burst without a pause. As a result, when the

cerebellum was denied information about the goal of the movement, the P-cells were unable to predict when to stop it.

Thus, the population output critically depended on competitive spike cancelation among the neurons. But why should a neuron produce spikes if its downstream consequences will be nulled by another neuron? The answer appears to be in the difference between what the population produced, and what the individual neurons produced: because the population output was a fast-changing function of time, the much slower individual P-cells appeared to be placed in subtractive competition with each other, eliminating each other's spikes in order to produce the fast dynamics of the output. The same pattern was present in the pMLIs: their population response exhibited a much faster dynamic than those of the individual neurons.

However, this explanation alone is not sufficient to account for all of our data. Notably, we found that all of the SS modulations produced by one group of P-cells for saccades in directions perpendicular to their potent vector were canceled by another group of P-cells with the opposite potent vector. The rationale for a neuron to produce spikes despite complete cancelation by another neuron remains unknown, but perhaps is indicative of a reserve for learning, for example, in case there is unilateral weakness in the extraocular muscles (*24*). A more likely explanation is that gaze control normally includes eye, head, and body movements, and the complete cancelation that we observed here would not occur if other body parts were allowed to move and thus contribute to directing the gaze.

Finally, while our results provide significant evidence in support of the null space theory of neural computation (*2*), it also raises a serious concern. Here, we used an *a priori* estimate of the potent vector for each neuron and then relied on anatomical information to distribute the vectors uniformly in order to compute a population response. In most other approaches the potent vectors are found via a model that maps neuronal activity to behavior. This means that if the recorded neurons are a biased sample of the underlying distribution of potent vectors, the model-estimated potent vectors will not resemble the actual one.

Methods

Data were collected from two marmosets, *Callithrix Jacchus*, 350-370 g, subjects 132F (Charlie), and 65F (Barney), during a 2-year period. The marmosets were born and raised in a colony that Prof. Xiaoqin Wang has maintained at the Johns Hopkins School of Medicine since 1996. Our procedures were approved by the Johns Hopkins University Animal Care and Use Committee in compliance with the guidelines of the United States National Institutes of Health.

Data acquisition

Following recovery from head-post implantation surgery, the animals were trained to make saccades to visual targets and rewarded with a mixture of applesauce and lab diet (*25*). Visual targets were presented on an LCD screen with 500Hz refresh rate and low latency (Dell AW2524H). Binocular eye movements were tracked at 2000 Hz using VPIXX eye tracking system. Tongue movements were tracked with a 522 frame/sec Sony IMX287 FLIR camera, with frames captured at 100 Hz.

We performed MRI and CT imaging on each animal and used this data to design an alignment system that defined trajectories from the burr hole to various locations in the cerebellar vermis (*25*), including points in lobule VI and VII. We used 3D Slicer software to first align the T2 MRI to the marmoset atlas (*26*). We then aligned the CT scan image to the transferred T2 MRI (Supplementary Fig. S1). We used a piezoelectric, high precision microdrive (0.5 micron resolution) with an integrated absolute encoder (M3-LA-3.4-15 Linear smart stage, New Scale Technologies) to advance the electrode. In order to reach the cerebellar cortex with Neuropixels NHP 1.0 and Neuropixels Ultra probes, we planned a posterior burr hole and avoided the confluence of sinuses using MRI T2 images. This approach enabled us to record from multiple folia in lobules VI and VII simultaneously.

We recorded from the cerebellum using Neuropixels 1.0 NHP probes, Neuropixels Ultra probes, as well as 64-channel checkerboard or linear high-density silicon probes (M1 and M2 probes, Cambridge Neurotech). Because the conductive coating on the Cambridge probes degraded after each insertion into the brain, we re-coated the probes and restored their low impedance after every 3-4 recording sessions (*27*). For the M1 and M2 probes we connected each electrode to a 64-channel head stage amplifier and digitizer (RHD2132 and RHD2164, Intan Technologies, USA), then connected the head stage to a communication system (RHD2000 Evaluation Board, Intan Technologies, USA). For Neuropixels we used the National Instrument acquisition system (NI PXIe-1071). Data were sampled at 30 kHz and aligned to the eye tracking system time as the reference time using random TTL signals.

Behavioral protocol

Each trial began with fixation of a center target after which a primary target appeared at one of 8 randomly selected directions at a distance of 5-6.5 deg. As the subject made a saccade to this primary target, that target was erased, and a secondary target was presented at a distance of 2.5-3.5 deg, also at one of 8 randomly selected directions. The subject was rewarded if following the primary saccade it made a corrective saccade to the secondary target, landed within a 1.5-2.0 deg radius of the target center, and maintained fixation for at least 200 ms. The reward was food that was provided in two small tubes (4.4 mm diameter), one to the left and the other to the right of the animal. A successful trial produced a food increment in one of the tubes and would continue to do so for 50 consecutive trials, then switch to the other tube. Because the food increment was small, the subjects naturally chose to work for a few consecutive trials, tracking the visual targets and allowing the food to accumulate, then stopped tracking and harvested the food via a licking bout (*21*, *28*).

We analyzed eye movements using a deep neural network that detected all saccades and microsaccades (*29*). The pre-trained networks for human and macaque monkeys did not perform well for marmosets. Hence, we designed a custom Matlab GUI-based program to curate the saccades (https://github.com/ShadmehrLCMC/SACCURATE). We then retrained the network through transfer learning for each individual animal using seven 30-45 minutes of recording. To prevent any potential false positive saccades, we pruned the saccades by fitting a bivariable Gaussian distribution to two different feature spaces defined as biologically relevant metrics (*30*) and removed saccades that were outliers (less than 1% chance of belonging to the distribution). The two feature spaces were the log-log main-sequence plots (maximum velocity vs. amplitude) and the log-log acceleration time to deceleration time ratio vs. amplitude of saccade. We then analyzed and found valid fixations among candidate fixations after or before each detected saccade. We measured and used data-driven thresholds on steady eye position criteria including maximum displacement, dispersion on x and y axes, fixation duration as well as maximum velocity. We discarded fixation candidates with lost eye signals due to instability of eye tracking or blinking.

We detected the onset and offset time of each saccade using a trained neural network, as described above. We then low pass filtered the eye position traces with a 3rd order Butterworth filter with 100 Hz cut-off frequency. We then calculated the saccade velocity by differentiating the eye position trace and found peak velocity times and values. Using the onset, offset and maximum velocity times we formed the acceleration and deceleration duration of each saccade.

Our data included rare saccades with extremely long acceleration durations that had a double-step velocity profile (Supplementary Fig. S12). We labeled these as double saccades and analyzed them separately, finding that the goal information provided in the mossy fiber inputs appeared to change mid-saccade.

Neurophysiological analysis

We used OpenEphys (*31*) for electrophysiology data acquisition, and then used Kilosort 2.0 and Phy 2.0 (*32*) to manually identify and curate the spikes. Each recording was curated twice by two experienced neurophysiologists. We controlled for contamination in the cross-correlograms in P-cell simple and complex spikes and removed the spikelets and double detection of CSs as SSs by aligning the two waveforms to each other and removing the copies in SS units (*33*). We further tracked the sorted units across three to seven 30-45 minute recordings using a semi-supervised custom written MATLAB GUI software. To do so, we used electrophysiology properties of the cells including waveform and location on the probe, auto-correlogram and raster plots aligned to saccade onsets. We discarded unstable cells with varying baseline firing rates across time (commonly due to the physical pressure of the probe). We performed manual cell type identification through identification of the layers. First, we identified the P-cells via suppression of SS following a CS. As P-cells are large cells, their spike waveforms are present on multiple channels of the silicon probes (Fig. 1B). Next, using the CS waveforms on the channels we identified the orientation of the P-cell's axon and dendrites, thus identifying the molecular, Purkinje, and granular layers (*4*, *5*, *33*).

Quantifying the quality of isolation of each neuron. A summary of the measures used to verify the quality of the neurophysiological data is provided in Supplementary Fig. S2A. To measure the isolation quality of each neuron, we computed the conditional probability $\Pr(s(t+\Delta)=1|s(t)=1)$, that is, the probability of a spike at time delay Δ, given that the neuron produced a spike at time zero. We then multiplied this probability by 1000 (bin size 1 ms) and plotted the results as firing rate. For well-

isolated cells, we expected a clean refractory period. To measure the noise rate in each neuron's spiking, we quantified this conditional probability at $\Delta = 1$ms, except for complex spikes, for which we used a 5ms period. The data for all cell types are shown in Supplementary Fig. S2A, second row. For example, to quantify the quality of the P-cells, we measured the noise rate and found that for the simple spikes, this rate was 0.58±0.41 Hz (median ± median absolute deviation, MAD). For the complex spikes, the noise rate was 0±0 Hz.

Identification of the MLIs. Following the identification of the molecular, P-cell, and granule layers, we identified the putative MLI1 and MLI2s (labeled as pMLI1 and pMLI2) based on their spike interactions with each other and the P-cells (*11*). These interactions were extracted and measured through cross-correlograms that were then corrected via spike jittering: for each cell pair, the interaction was measured via a cross-correlation, then corrected for chance interaction due to firing rate fluctuations using interval jittering (*13*), as shown in Fig. 1D. pMLI1s were identified using their milliseconds inhibition of P-cells as well as their synchrony with each other, while pMLI2s were identified through their millisecond suppression of pMLI1s and lack of interaction or later disinhibition of P-cells (*11*). We relied on the interactions of MLIs with P-cells to identify them, then confirmed that their waveform, baseline rates, and auto-correlograms were similar to previous findings in definitively identified neurons in mice (*11*).

Identification of the mossy fibers: We identified the mossy fibers based on their unique "m" shaped spike waveform (Fig. 1B) in the granule layer. We did this by relying on our Neuropixels Ultra recordings to acquire a very high resolution "picture" of the mossy fiber spike waveforms (Supplementary Fig. S8). As previously reported, mossy fibers exhibit a triphasic waveform with a negative after wave or a fast narrow spike (*8*, *34*–*38*). We used an approach similar to (*39*) to extract spatio-temporal waveforms of each cell by concatenating spike waveforms on different channels to form a 2D image, then centered the absolute peak to the center of image in time and space (Supplementary Fig. S8C). In order to get a consistent result for the checkerboard and linear probes we picked two columns out of four columns of the checkerboard probe with highest waveform values for each cell and linearized it. We then centered the absolute peak of the waveform in time and space (see Supplementary Fig. S8C). Finally, we used UMAP to cluster the mossy fibers from other granular layer interneurons (Supplementary Fig. S8D). UMAP on a linearized feature space of these spatio-temporal waveforms revealed two groups of mossy fibers and a potential Golgi cell group. As the mossy fibers exhibited different spike waveforms based on the location of recording (Supplementary Fig. S8A, Neuropixel Ultra recording), we found one UMAP group with downward fast spiking waveform and another UMAP group with upward triphasic waveform with a negative after wave. The GLI group, on the other hand, had larger signature in space and showed wider spike waveforms (Supplementary Fig. S8E).

Identifying mossy fibers that encoded goal and state mossy fibers. The oculomotor region of the vermis receives mossy fibers mainly from the pontine nuclei, the nucleus reticularis tegmenti pontis (NRTP), and the paramedian pontine reticular formation (PPRF) (*40*) (Supplementary Fig. S8F). NRTP neurons receive information regarding the goal of the movement from the superior colliculus (*41*, *42*). The PPRF region houses the horizonal burst generator neurons (*43*). This suggested that the mossy fibers in this region of the cerebellum should separately encode the goal of the movement, and a copy of the motor commands (Supplementary Fig. S9-S11).

We measured the firing rates of mossy fibers during fixation, saccades, and after presentation of the visual cue. One group of mossy fibers had responses similar to neural integrator neurons in the brainstem, exhibiting tonic rates that encoded fixation position of the eyes. To identify these neurons, we

computed the firing rates during task-irrelevant fixations (>200ms fixation) in a 50ms window centered during the fixation period. We then fitted the projected position of the eye on the preferred direction of the mossy fiber and used R-squared larger than 0.25 to identify fixation encoding mossy fibers. We removed these mossy fibers from the population. Next, we removed mossy fibers with baseline rates higher than 20Hz as NRTP and PPRF neurons usually have low baseline values. We removed mossy fibers with less than 500 targeted saccades toward their preferred direction and less than 100 Hz of increase in their response aligned to deceleration onset of targeted saccades. Among 345 remaining mossy fibers we used a normalized distance measure between the responses to reward relevant and irrelevant saccades while controlling for kinematics (direction, amplitude, peak velocity). We used UMAP to cluster the mossy fibers into two groups using these distances as well as features including normalized amplitude and direction tuning of the cells (Supplementary Fig. S11B).

Forming cliques based on neuron-neuron spike interactions

After acquiring the cross-correlograms for each pair of neurons, we clustered them into small networks, called cliques. The procedure began by forming an adjacency matrix (Fig. 1C), where each neuron was a row and a column in the matrix. We then made the adjacency matrix symmetric by keeping the maximum absolute value of the CCCG during the 0-3 ms period for each cell pair. Next, we used spectral clustering to divide the adjacency matrix into groups. Spectral clustering is a graph clustering method which uses the adjacency matrix of the graph and represents the nodes (each cell) in a 2D or 3D subspace of eigenvectors of the Laplacian matrix (spectral space). In the spectral space the interconnected nodes have shorter distance from each other than the weakly connected nodes (*44*). To illustrate clique membership, we used the weights to assign line thickness and then connected the cell's location on the probe to its clique (Fig. 1B, left subplot).

Computing the potent vector

We used the properties of the climbing fiber input that the P-cells of a clique received to assign a potent vector to all the neurons that were a member of the clique. The climbing fibers reported to the cerebellum both the direction of the visual event, and then independently, the direction of the planned saccade (*18*). The strength of this encoding differed among the various P-cells. We quantified this encoding by measuring the CS rate as a function of the direction of the visuomotor event during two different time windows: 40 to 85 ms with respect to the visual cue onset, and -70 to +30 ms window from saccade onset (Fig. 2A, Supplementary Fig. S4A). We then fitted a Von Mises distribution to the resulting rates as a function of angle (*45*). This produced two vectors, one that quantified the strength and direction of the CS response to the visual events, and a second that quantified the strength and direction of the CS response to the motor events. Each vector had an angle θ and amplitude ρ, where the amplitude identified the sharpness of tuning. Thus, the amplitude of the vector was bounded by 0 (not tuned), and 1 (maximum sharpness). This produced two distinct vectors for each P-cell: one describing the tuning with respect to visual events, and one describing the tuning with respect to motor events (Supplementary Fig. S4B). However, encoding of these two vectors were very similar: the magnitude of the vector for the visual event was highly correlated with the magnitude of the vector for the motor event (Supplementary Fig. S4B). Hence, for each P-cell, we used the vector with the larger amplitude and labeled it as its *potent vector*.

Testing the validity of the potent vector

The potent vector was defined based on the climbing fiber information that each P-cell received. However, we conjectured that the direction and magnitude of the potent vector were predictors of the downstream influence that the P-cells had on the eye movements. To examine this conjecture, we performed three tests.

First, we tested our estimate of the direction of the potent vector θ. To do this, we quantified the effects that a CS-triggered SS suppression had on the trajectory of a saccade (Fig. 2B). We observed that the downstream effects, on average, were focused on the direction of the potent vector, with little or no contributions in the direction perpendicular to this vector (Fig. 2C). Moreover, when two P-cells with potent vectors that were nearly parallel were simultaneously suppressed, the effects in the direction of the potent vector roughly tripled in direction θ with little or no contribution in direction $\theta + \frac{\pi}{2}$ (Fig. 2C). These results implied that our estimate of the direction of the potent vector for each P-cell was an unbiased estimate of the actual direction of the weight vector $\boldsymbol{w}$ that functionally defined the connection between that P-cell and the eye muscles.

Second, we tested our estimate of the magnitude of the potent vector. To do this, we divided our population of P-cells into 3 groups based on the amplitude ρ [edges: $\rho = 0, 0.15, 0.3$]. We observed that as ρ increased between the groups, so did the downstream effects of SS suppression (Fig. 2D). Next, we considered the effects of simultaneous suppression of two P-cells, finding that the size of the displacement produced grew linearly with the size of the amplitude of the sum of the two potent vectors (Supplementary Fig. S5C). These results implied that our estimate of the amplitude of the potent vector for each P-cell was an unbiased estimate of the actual amplitude of the weight vector that functionally defined the connection between that P-cell and the eye muscles.

Third, we asked whether simultaneous suppression of two P-cells during a single saccade produced a displacement of the eyes in a direction and amplitude that was predicted by the superposition of the two potent vectors of the P-cells. We observed results that were consistent with this prediction (Fig. 2E).

Assigning a potent vector to each clique

We organized the P-cells and the neighboring MLIs into neuronal cliques, and then assigned a single potent vector to all the neurons in that clique. To do this, we considered cliques that had at least one P-cell (both SS and CS). We observed that the climbing fibers that projected to the same clique carried information that was much more similar to each other than between cliques (Supplementary Fig. 4C). For example, the visuomotor tuning distance, measured as the difference in the strength of tuning ρ of visual vs. motor events (distance of two points in Supplementary Fig. S4B, as shown in Supplementary Fig. S4C), was much less within a clique than between cliques (rank sum test, p=1.5 E-26). Similarly, the difference in the angle θ of the potent vectors was much less within a clique than between (rank sum test, p=1.12 E -18). To compute the potent vector for a given clique, we performed a vector average of the potent vectors of the P-cells that belonged to that clique, then assigned that single potent vector to all the P-cells and interneurons in that clique. Individual climbing fiber's preferred direction differences from the clique's potent direction was centered at zero, as was the difference in the amplitude (Supplementary Fig. S4D & S4E).

Perturbation of saccades following a CS-triggered SS suppression

For this analysis we used the spike-triggered averaging techniques introduced in (*20*). We focused on the saccades to the primary target as they had a fixed amplitude and started from the center fixation point.

For each P-cell, we divided saccades that were made from the same initial position to the same target into two groups: (1) saccades in which the P-cell did not experience a CS during the time window -60 to +40ms with respect to saccade onset (NOCS: no cs), (2) saccades in which the P-cell experienced a single CS during the time window -30 to +30ms with respect to saccade onset (WCS: with cs). Then, we took steps to ensure that the initial start positions of the saccades were comparable between WCS and NOCS saccades. For each target direction, we computed the initial position of each WCS and NOCS saccade with respect to the fixation point and projected this vector onto the potent vector. We removed the minimum number of NOCS saccades required to match their mean projected start position to within ±0.03 degrees of the mean projected start position of the WCS saccades. We then computed the difference in saccade velocity vectors in each direction as a function of time between WCS and NOCS saccades, projected onto the potent vector, and the orthogonal vector. For each target direction, this produced two scalar quantities, i.e., change in velocity, as a function of time. We then averaged the projected changes in velocity across all saccade directions. The null distribution was formed through bootstrapping of the populations of CS with randomized tagging of trials, labeling them as saccades with or without a CS (number of bootstraps=1000).

As a sanity check, we checked that the CS-triggered SS suppression displaced the eyes only if the suppression event was within a short window centered on the saccade itself. Indeed, the SS suppression events that occurred longer than 50 ms before saccade onset had no effects on the saccade (Supplementary Fig. S5A).

Testing superposition of the potent vectors

Neuropixels allowed us to record from multiple folia simultaneously, producing cliques that had diverse potent vectors. For example, in Fig. 1B, clique 1 has a large amplitude potent vector, whereas clique 5 has a very small potent vector. Moreover, the direction of the potent vector for clique 1 is near 135°, whereas for clique 3 it is near 180°. Because the climbing fiber inputs to all of these cliques were recorded simultaneously, we performed spike-triggered averaging on simultaneous complex spikes. This allowed us to ask whether simultaneous suppression of two P-cells produced a displacement that obeyed superposition in the vector space of the two potent vectors.

For two simultaneously recorded P-cells, we had two potent vectors, with properties (ρ_1, θ_1) and (ρ_2, θ_2). Next, we selected pairs of P-cells in which the size of the two potent vectors exceeded a minimum threshold: $\frac{\rho_2}{\rho_1} \geq 0.5$, where $\rho_1 > \rho_2$. This ensured that the P-cell with the small potent vector ρ_2 would not be completely overwhelmed by the P-cell with the larger potent vector ρ_1. We then divided the saccades that were made from the same initial position to the same target into two groups: (1) saccades during which neither of the two P-cells experienced a CS (NOCS, -60 to +40ms window with respect to saccade onset), (2) saccade during which both of the P-cells experienced a CS within 30ms of each other (WCS, [-30 to +30]ms window). As in the case of single P-cells, we ensured that the starting positions of the two saccades were exquisitely matched. We then tested the hypothesis that the CS-triggered change in saccade trajectory was the vector summation of the two CS-on vectors, i.e., $\rho_{net}\angle\theta_{net} = \rho_1\angle\theta_1 + \rho_2\angle\theta_2$.

Computing population responses

For clique j, we computed the potent vector as the average vector defined by its climbing fiber inputs (only P-cells with both SS and CS):

$$\rho_j \angle \theta_j = \frac{1}{n_i} \sum_{i \in j} \rho_i \angle \theta_i \quad (6)$$

For each of the cell types in a clique we first binned the saccades based on the movement angle to 8 directions (0:45:360 deg). We then averaged cell responses to saccade directions with respect to θ (Fig. 3A-B). Instantaneous firing rates were calculated from peri-event time histograms with 1 ms bin size. We used a Savitzky–Golay filter (2nd order, 31 datapoints) to smooth the traces for visualization purposes. Finally, to compute a population response, we used Eqs. (3-5). The result was a 2D vector of spikes as a function of time.

We found that in the population response for the P-cells, the rates parallel to the direction of the movement exhibited a burst-pause pattern that crossed zero at deceleration onset. We calculated the confidence intervals for this zero-crossing using bootstrapping. To do so, we calculated the first zero-crossing in a -30ms to 30ms window around deceleration onset for bootstrapped population (10,000 times with replacement). We repeated the same procedure for finding MLI1 population peak time in a -30ms to 75ms window.

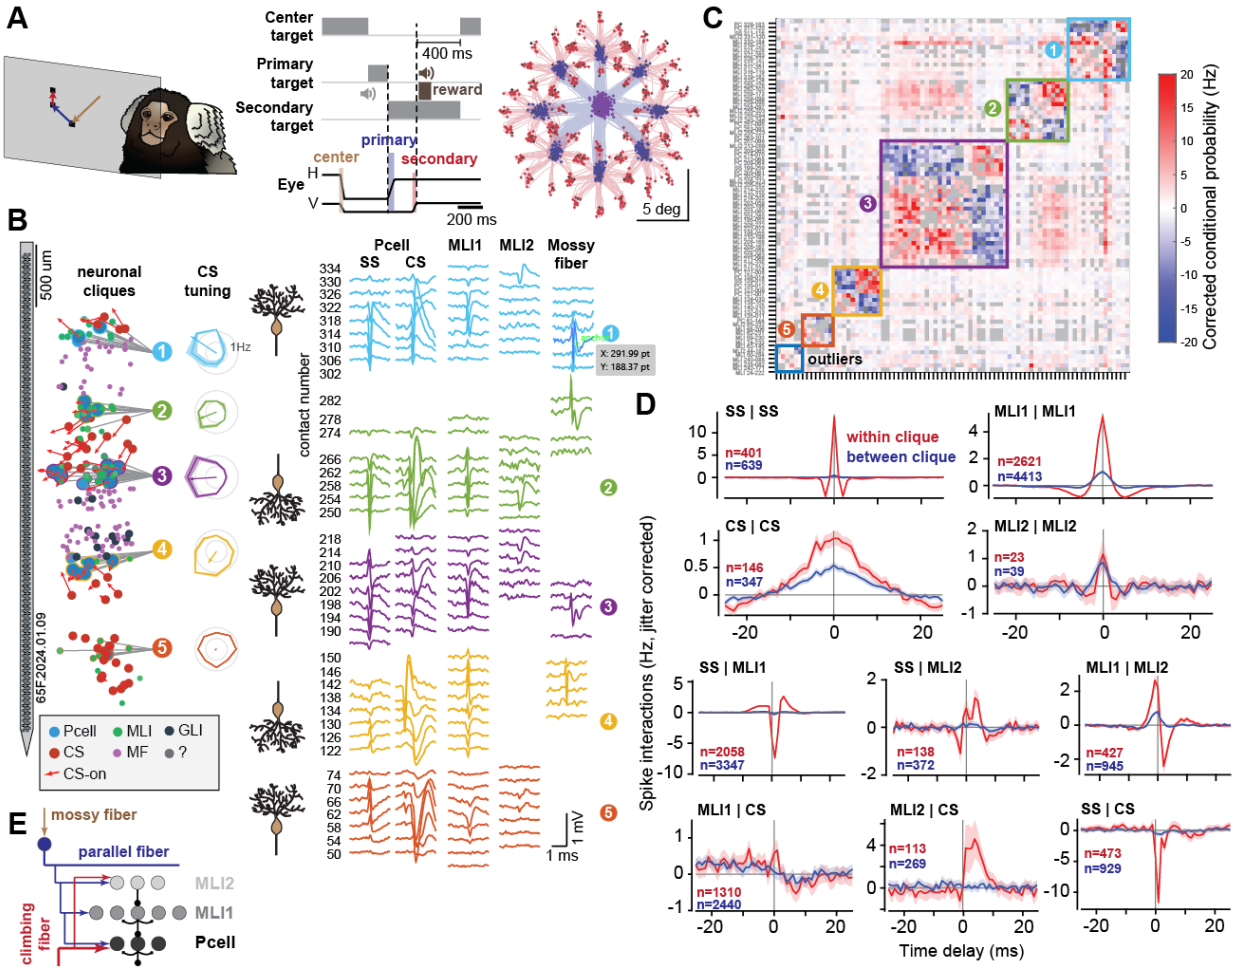

Figure 1. Spike interactions clustered neurons in the cerebellar cortex into discrete populations. **A**. Random corrective saccade paradigm. **B**. Example Neuropixels session with the location of 5 neuronal cliques. Spike waveforms are shown for example neurons. CS tuning indicates the average response in the climbing fiber inputs to each clique as a function of direction of visual events. P-cell dendrite and axon orientations are drawn to show approximate location of the molecular and granular layers. **C**. Adjacency matrix. Rows and columns are neurons recorded simultaneously, with the color indicating the strength of between neuron spike interactions (jitter-corrected conditional probability). Cliques and outliers are identified by the boundaries. **D**. The jitter-corrected probability that a spike occurred in one neuron, given that a spike had occurred in another neuron at a given time delay, for neuron pairs within and between cliques. **E**. Within-clique architecture of the network, with MLI2 inhibiting MLI1, and MLI1 inhibiting P-cells. Error bars are SEM.

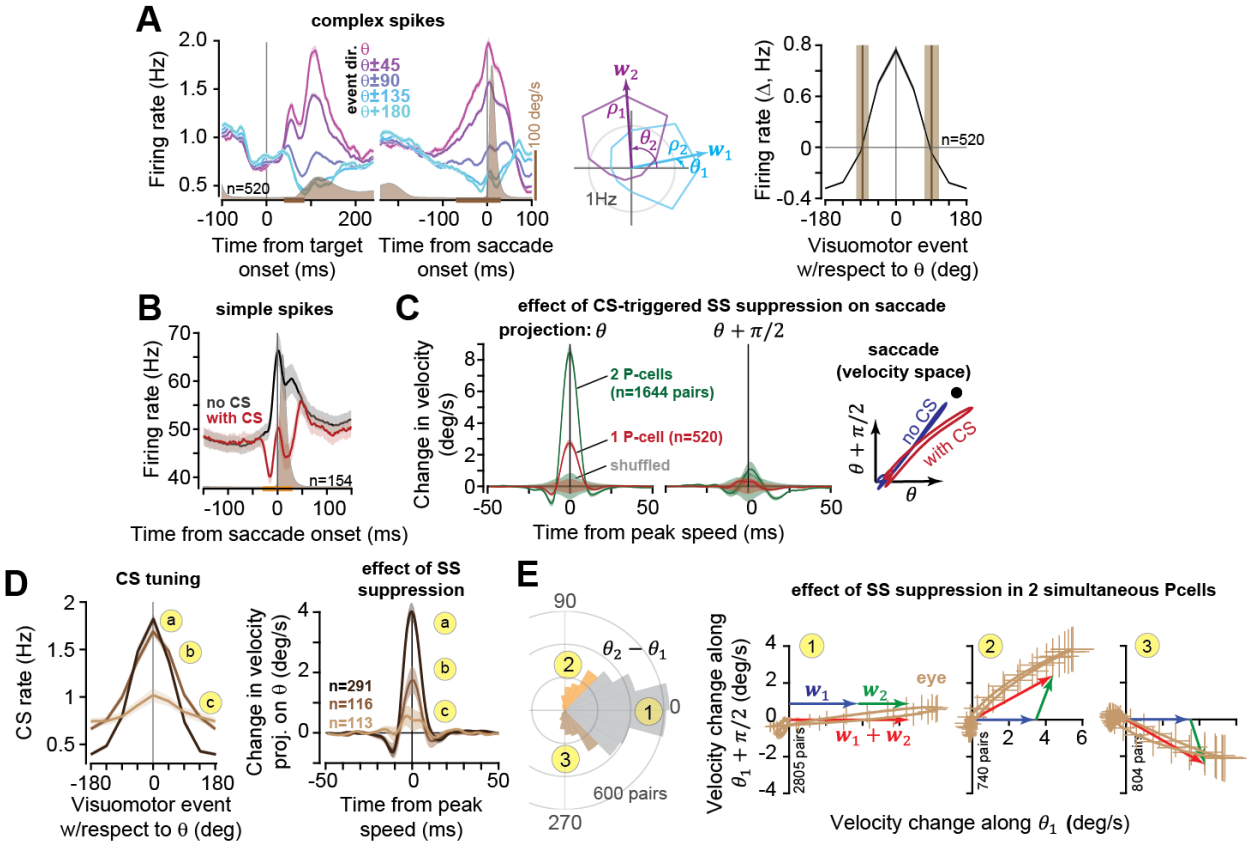

Figure 2. **CS-triggered suppression defined a potent vector in the behavioral space for each P-cell**. **A**. CS response aligned to the direction of visual and motor events. For each P-cell, the CS response defined a weight vector with amplitude ρ and direction θ. **B**. SS rates during saccades that did or did not experience a CS. **C.** Effects of the CS-triggered suppression in one P-cell, or two simultaneously recorded P-cells, on saccade trajectory. Here, pairs of P-cells had $|\Delta\theta_{ij}| < 22.5$ deg. Gray region shows 95% confidence interval for the null hypothesis. **D**. P-cells were divided into 3 groups based on the strength of their climbing fiber input ρ. The effect of CS-triggered SS suppression on saccade trajectory was largest in the group that had the largest ρ. **E**. Test of superposition. Pairs of simultaneously recorded P-cells were divided into three groups based on $\Delta\theta$. We then measured effects of simultaneous SS suppression on saccade trajectory. For example, in group 1, the two P-cells had vectors $\boldsymbol{w}_1$ and w_2 that were, on average, parallel. The sum of the two vectors is shown in red, along with the CS-triggered change in saccade trajectory (brown line and error bars). Error bars are SEM.

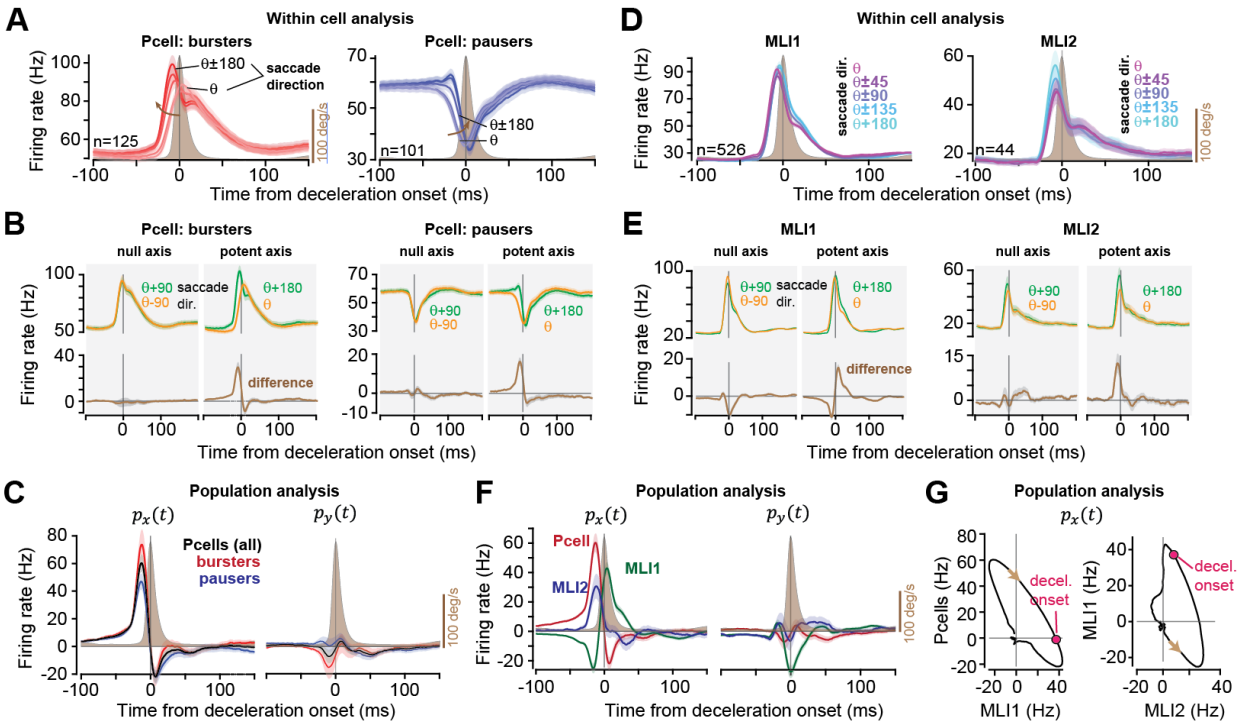

Figure 3. **Neurons within a clique had symmetric spike rates with respect to the potent axis, canceling the spikes that had downstream effects perpendicular to the direction of the movement**. **A**. Average P-cell response (red: bursters, blue: pausers) with respect to θ of each neuron. **B**. SS rates for saccades in direction $\theta + \frac{\pi}{2}$ are equal to the rates for saccades in direction $\theta - \frac{\pi}{2}$. In contrast, the rates for saccades parallel to the potent vector are not equal. **C.** Population response of all P-cells, i.e., the sum of rates, weighted by each neuron's 2D potent vector. **D**. Same as A for the pMLIs. **E**. Same as B for the pMLIs. **F**. Population response of the MLIs. **G**. Phase plots of P-cell potent response as a function of pMLI1 (left), and MLI1 as a function of pMLI2 (right). The negative slope indicates that in the population response, there is a negative correlation between the firing rates of the pMLI1s and the P-cells, and pMLI2s and pMLI1s. Error bars are SEM.

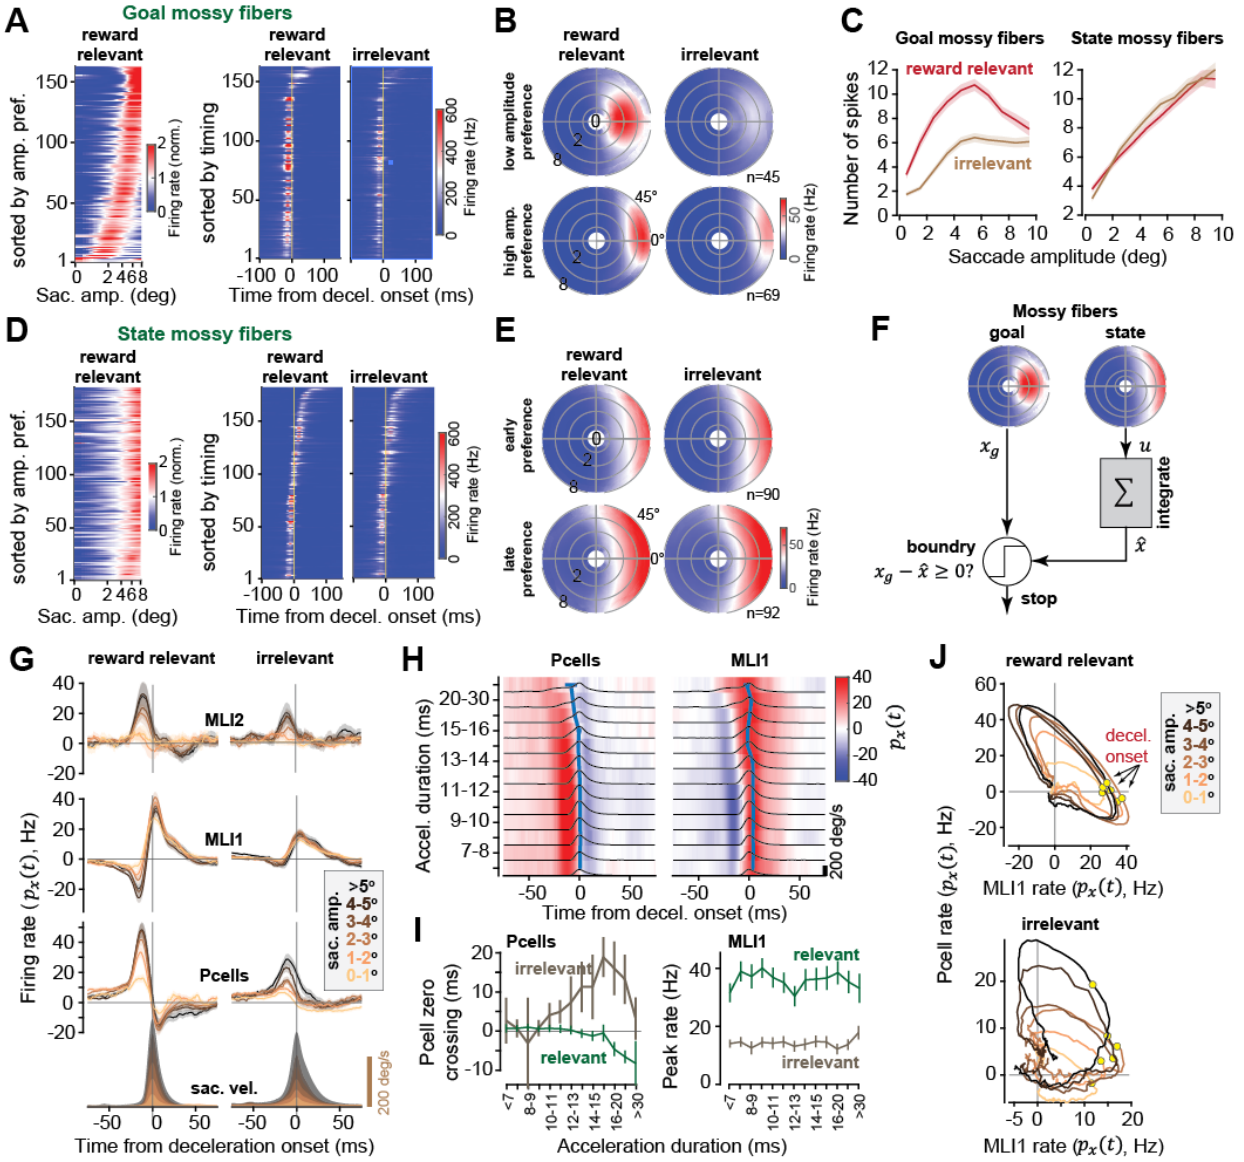

Figure 4. **P-cells signal when the saccade should begin decelerating, but only if the mossy fibers provide goal location information**. Two groups of mossy fibers carried information regarding the sensory goal location (goal mossy fibers A-C) and motor commands (State mossy fibers D-F). **A**. left, Normalized responses of goal mossy fibers to different saccade amplitudes sorted by the preferred direction, right, average responses of the same mossy fibers aligned to saccade deceleration onset for reward relevant vs reward irrelevant saccades. **B**. Polar plots presenting the movement fields of the goal mossy fibers divided into low and high amplitude preference groups for reward relevant and reward irrelevant saccades. **C.** Number of spikes as a function of saccade amplitude for goal mossy fibers for the two types of movement. **D-F**. same as A-C for state mossy fibers. **G**. A feasible model for estimation of the deceleration time based on the mossy fiber input. An integration of motor commands in time would estimate the displacement and a comparison between the displacement and goal mossy fiber input can signal when to stop the saccade. **H**. Clique pMLIs and P-cell responses on potent axis for reward relevant and reward irrelevant saccades binned by peak velocity. **I**. P-cell and pMLI1 potent responses for reward relevant saccades binned by acceleration duration. **J**. P-cell zero-crossing and pMLI1 peak rate for different acceleration durations. Error bars are SEM.

1. M. M. Churchland, K. V. Shenoy, Preparatory activity and the expansive null-space. *Nat. Rev. Neurosci.* **25**, 213–236 (2024).

2. M. T. Kaufman, M. M. Churchland, S. I. Ryu, K. V. Shenoy, Cortical activity in the null space: permitting preparation without movement. *Nat Neurosci* **17**, 440–448 (2014).

3. E. A. Kirk, K. T. Hope, S. J. Sober, B. A. Sauerbrei, An output-null signature of inertial load in motor cortex. *Nat Commun* **15**, 7309 (2024).

4. H. Gao, C. Solages, C. Lena, Tetrode recordings in the cerebellar cortex. *J.Physiol Paris* **106**, 128–136 (2012).

5. K.-S. Han, C. H. Chen, M. M. Khan, C. Guo, W. G. Regehr, Climbing fiber synapses rapidly and transiently inhibit neighboring Purkinje cells via ephaptic coupling. *Nature Neuroscience* **23**, 1399–1409 (2020).

6. J. V. Walsh, J. C. Houk, E. Mugnaini, Identification of unitary potentials in turtle cerebellum and correlations with structures in granular layer. *J.Neurophsyiol.* **37**, 30–47 (1974).

7. P. L. van Kan, A. R. Gibson, J. C. Houk, Movement-related inputs to intermediate cerebellum of the monkey. *Journal of Neurophysiology* **69**, 74–94 (1993).

8. S. Tomatsu, T. Ishikawa, Y. Tsunoda, J. Lee, D. S. Hoffman, S. Kakei, Information processing in the hemisphere of the cerebellar cortex for control of wrist movement. *J.Neurophysiol.* **115**, 255–270 (2016).

9. G. Szapiro, B. Barbour, Multiple climbing fibers signal to molecular layer interneurons exclusively via glutamate spillover. *Nat.Neurosci.* **10**, 735–742 (2007).

10. S. Malhotra, G. Banumurthy, R. L. Pennock, J. H. Vaden, I. Sugihara, L. Overstreet-Wadiche, J. I. Wadiche, Climbing Fiber-Mediated Spillover Transmission to Interneurons Is Regulated by EAAT4. *J. Neurosci.* **41**, 8126–8133 (2021).

11. E. P. Lackey, L. Moreira, A. Norton, M. E. Hemelt, T. Osorno, T. M. Nguyen, E. Z. Macosko, W.-C. A. Lee, C. A. Hull, W. G. Regehr, Specialized connectivity of molecular layer interneuron subtypes leads to disinhibition and synchronous inhibition of cerebellar Purkinje cells. *Neuron* **112**, 2333-2348.e6 (2024).

12. D. F. English, S. McKenzie, T. Evans, K. Kim, E. Yoon, G. Buzsáki, Pyramidal Cell-Interneuron Circuit Architecture and Dynamics in Hippocampal Networks. *Neuron* **96**, 505-520.e7 (2017).

13. J. Platkiewicz, E. Stark, A. Amarasingham, Spike-Centered Jitter Can Mistake Temporal Structure. *Neural Computation* **29**, 783–803 (2017).

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

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

16. D. J. Herzfeld, Y. Kojima, R. Soetedjo, R. Shadmehr, Encoding of error and learning to correct that error by the Purkinje cells of the cerebellum. *Nat.Neurosci.* **21**, 736–743 (2018).

17. E. Sedaghat-Nejad, J. S. Pi, P. Hage, M. A. Fakharian, R. Shadmehr, Synchronous spiking of cerebellar Purkinje cells during control of movements. *Proceedings of the National Academy of Sciences* **119**, e2118954119 (2022).

18. J. S. Pi, M. A. Fakharian, P. Hage, E. Sedaghat-Nejad, S. Z. Muller, R. Shadmehr, The olivary input to the cerebellum dissociates sensory events from movement plans. *Proceedings of the National Academy of Sciences* **121**, e2318849121 (2024).

19. M. L. Streng, L. S. Popa, T. J. Ebner, Climbing Fibers Control Purkinje Cell Representations of Behavior. *J.Neurosci.* **37**, 1997–2009 (2017).

20. S. Z. Muller, J. S. Pi, P. Hage, M. A. Fakharian, E. Sedaghat-Nejad, R. Shadmehr, Complex spikes perturb movements and reveal the sensorimotor map of Purkinje cells. *Current Biology* **33**, 4869-4879.e3 (2023).

21. P. Hage, M. A. Fakharian, A. M. Shoup, J. S. Pi, E. Sedaghat-Nejad, S. P. Orozco, I. K. Jang, V. Looi, T. Arginteanu, R. Shadmehr, Control of tongue movements by the Purkinje cells of the cerebellum. bioRxiv [Preprint] (2024). https://doi.org/10.1101/2024.07.25.604757.

22. J. F. Medina, S. G. Lisberger, Links from complex spikes to local plasticity and motor learning in the cerebellum of awake-behaving monkeys. *Nat.Neurosci.* **11**, 1185–1192 (2008).

23. S. Hao, X. Zhu, Z. Huang, Q. Yang, H. Liu, Y. Wu, Y. Zhan, Y. Dong, C. Li, H. Wang, E. Haasdijk, Z. Wu, S. Li, H. Yan, L. Zhu, S. Guo, Z. Wang, A. Ye, Y. Lin, L. Cui, X. Tan, H. Liu, M. Wang, J. Chen, Y. Zhong, W. Du, G. Wang, T. Lai, M. Cao, T. Yang, Y. Xu, L. Li, Q. Yu, Z. Zhuang, Y. Xia, Y. Lei, Y. An, M. Cheng, Y. Zhao, L. Han, Y. Yuan, X. Song, Y. Song, L. Gu, C. Liu, X. Lin, R. Wang, Z. Wang, Y. Wang, S. Li, H. Li, J. Song, M. Chen, W. Zhou, N. Yuan, S. Sun, S. Wang, Y. Chen, M. Zheng, J. Fang, R. Zhang, S. Zhang, Q. Chai, J. Liu, W. Wei, J. He, H. Zhou, Y. Sun, Z. Liu, C. Liu, J. Yao, Z. Liang, X. Xu, M. Poo, C. Li, C. I. De Zeeuw, Z. Shen, Z. Liu, L. Liu, S. Liu, Y. Sun, C. Liu, Cross-species single-cell spatial transcriptomic atlases of the cerebellar cortex. *Science* **385**, eado3927 (2024).

24. L. M. Optican, D. A. Robinson, Cerebellar-dependent adaptive control of primate saccadic system. *J.Neurophysiol.* **44**, 1058–1076 (1980).

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

26. C. Liu, F. Q. Ye, C. C.-C. Yen, J. D. Newman, D. Glen, D. A. Leopold, A. C. Silva, A digital 3D atlas of the marmoset brain based on multi-modal MRI. *NeuroImage* **169**, 106–116 (2018).

27. A. M. Shoup, N. Porwal, M. A. Fakharian, P. Hage, S. P. Orozco, R. Shadmehr, Rejuvenating silicon probes for acute neurophysiology. *Journal of Neurophysiology* **132**, 308–315 (2024).

765 28. P. Hage, I. K. Jang, V. Looi, M. A. Fakharian, S. P. Orozco, J. S. Pi, E. Sedaghat-Nejad, R. Shadmehr,
766 Effort cost of harvest affects decisions and movement vigor of marmosets during foraging. *eLife* **12**,
767 RP87238 (2023).

768 29. M. E. Bellet, J. Bellet, H. Nienborg, Z. M. Hafed, P. Berens, Human-level saccade detection
769 performance using deep neural networks. *Journal of Neurophysiology* **121**, 646–661 (2019).

770 30. R. J. Leigh, D. S. Zee, *The Neurology of Eye Movements* (Oxford University Press, 2015).

771 31. J. H. Siegle, A. C. Lopez, Y. A. Patel, K. Abramov, S. Ohayon, J. Voigts, Open Ephys: an open-source,
772 plugin-based platform for multichannel electrophysiology. *J.Neural Eng* **14**, 045003 (2017).

773 32. M. Pachitariu, N. Steinmetz, S. Kadir, M. Carandini, H. K. D, Kilosort: realtime spike-sorting for
774 extracellular electrophysiology with hundreds of channels. *bioRxiv*, 061481 (2016).

775 33. E. Sedaghat-Nejad, M. A. Fakharian, J. Pi, P. Hage, Y. Kojima, R. Soetedjo, S. Ohmae, J. F. Medina, R.
776 Shadmehr, P-sort: an open-source software for cerebellar neurophysiology. *Journal of*
777 *Neurophysiology* **126**, 1055–1075 (2021).

778 34. M. Kase, D. C. Miller, H. Noda, Discharges of Purkinje cells and mossy fibres in the cerebellar vermis
779 of the monkey during saccadic eye movements and fixation. *J.Physiol.* **300**, 539–555 (1980).

780 35. K. Ohtsuka, H. Noda, Burst discharges of mossy fibers in the oculomotor vermis of macaque
781 monkeys during saccadic eye movements. *Neurosci.Res.* **15**, 102–114 (1992).

782 36. M. Prsa, S. Dash, N. Catz, P. W. Dicke, P. Thier, Characteristics of responses of Golgi cells and mossy
783 fibers to eye saccades and saccadic adaptation recorded from the posterior vermis of the
784 cerebellum. *J.Neurosci.* **29**, 250–262 (2009).

785 37. G. Van Dijck, M. M. V. Van Hulle, S. A. Heiney, P. M. Blazquez, H. Meng, D. E. Angelaki, A. Arenz, T.
786 W. Margrie, A. Mostofi, S. Edgley, F. Bengtsson, C.-F. Ekerot, H. Jörntell, J. W. Dalley, T. Holtzman,
787 Probabilistic Identification of Cerebellar Cortical Neurones across Species. *PLOS ONE* **8**, e57669
788 (2013).

789 38. A. Markanday, S. Hong, J. Inoue, E. De Schutter, P. Thier, Multidimensional cerebellar computations
790 for flexible kinematic control of movements. *Nat Commun* **14**, 2548 (2023).

791 39. X. Jia, J. H. Siegle, C. Bennett, S. D. Gale, D. J. Denman, C. Koch, S. R. Olsen, High-density
792 extracellular probes reveal dendritic backpropagation and facilitate neuron classification. *Journal of*
793 *Neurophysiology* **121**, 1831–1847 (2019).

794 40. C. D. Thielert, P. Thier, Patterns of projections from the pontine nuclei and the nucleus reticularis
795 tegmenti pontis to the posterior vermis in the rhesus monkey: a study using retrograde tracers.
796 *J.Comp Neurol.* **337**, 113–126 (1993).

797 41. J. K. Harting, Descending pathways from the superior collicullus: an autoradiographic analysis in the
798 rhesus monkey (Macaca mulatta). *J.Comp Neurol.* **173**, 583–612 (1977).

42. W. F. Crandall, E. L. Keller, Visual and oculomotor signals in nucleus reticularis tegmenti pontis in alert monkey. *J.Neurophysiol.* **54**, 1326–1345 (1985).

43. D. L. Sparks, The brainstem control of saccadic eye movements. *Nat.Rev.Neurosci.* **3**, 952–964 (2002).

44. U. von Luxburg, A tutorial on spectral clustering. *Stat Comput* **17**, 395–416 (2007).

45. P. Berens, CircStat: A MATLAB Toolbox for Circular Statistics. *Journal of Statistical Software* **31**, 1–21 (2009).