ADAPTIVE CONTROL OF SACCADES

by
Ehsan Sedaghat-Nejad

A dissertation submitted to Johns Hopkins University in conformity with
the requirements for the degree of Doctor of Philosophy

Baltimore, Maryland
November 2021

Abstract

As we navigate through the world in our lifetime, our brains constantly adjust our movements to ensure their accuracy. How does the brain adaptively control our movements? In this dissertation, we looked at saccadic eye movements and performed behavioral and neural measurements to study different brain mechanisms for adaptive control of saccades.

In Chapter 2, we considered the effect of reward prediction error (RPE) as a strong modulator of dopamine on saccade vigor. We found that saccade vigor was affected in an orderly fashion by the magnitude and direction of the RPE event: the most vigorous saccades followed the largest positive-RPE, whereas the least vigorous saccades followed the largest negative-RPE. Thus, reward prediction error, and not reward per se, modulated the vigor of saccades.

In Chapter 3, we looked at corrective saccades during saccade adaptation and asked is there an implicit loss associated with the corrective movement that can modulate learning? We designed a novel paradigm that combined random dot motion discrimination with saccade adaptation to impose a cost on movement error. Our results demonstrated that when error cost was large, the brain learned more from error. Thus, during sensorimotor adaptation, the act of correcting for error carries an implicit cost for the brain which regulates the rates of learning.

Next, we focused on the potential neural mechanisms for adaptive control of saccades and the role of the principal cells of the cerebellar cortex, Purkinje cells (P-cells), in control of saccades. We introduced marmoset monkeys as a new animal model to study motor control and motor learning. In Chapter 4, we presented protocols to train marmoset monkeys to perform saccadic eye movement tasks and record from their cerebellum. In Chapter 5, we introduced an open-source software package, named *P-sort*, to analyze the cerebellar neurophysiological data.

Finally, in Chapter 6, we analyzed the P-cell data (n=149 cells) by organizing them into populations that shared the same preference for error as measured via their complex spike response. We found that the population activity of simple spikes produced a burst-pause pattern that started before saccade onset and ended with saccade termination. Next, we looked at synchronous activity among pairs of simultaneously recorded P-cells (n=42 pairs). Our results demonstrated that the synchronization index reached its peak probability after deceleration onset but before saccade end. Thus, the cerebellar cortex relies on spike synchronization within a population of P-cells, not individual firing rates, to predict when to stop a movement.

Primary Reader and Advisor: Dr. Reza Shadmehr

Secondary Reader: Dr. David S. Zee

Thesis Committee: Dr. Reza Shadmehr, Dr. David S. Zee, Dr. Jeremiah Y. Cohen, Dr. Abigail L. Person

Acknowledgements

Dedication

This dissertation is dedicated to my love, my best friend, and my sweet wife, Shima.

تقدیم به آنکه آفتاب مهرش

در آستانه قلبم پا برجاست

و هرگز غروب نخواهد کرد

Table of contents

List of figures

Chapter 1. Introduction

Saccades are one of the fastest movements generated by the human body (Leigh & Zee, 2015). Saccadic eye movements are so fast that the eyes are practically blind during the movement (Thiele et al., 2002), and there is no visual feedback to correct saccades as they take place (Xu-Wilson et al., 2011). Humans and other non-human primates generate around three saccades per second that are for the most part stereotypical. However, there are multiple factors that affect how the brain generates saccades. For example, when we make saccades toward familiar faces the vigor of movement (defined as the reciprocal of the latency) is higher (Xu-Wilson, Zee, et al., 2009), or when our eyes constantly land away from the target the brain adapts the saccades to compensate for the errors (Ethier et al., 2008b).

The study of saccadic eye movements has been an active area of research for decades and numerous laboratories have contributed to understanding the behavioral and neural aspects of saccades. Nevertheless, there are many more fascinating questions that can be asked and pursued to better understand how the brain controls these movements and what they reveal about the internal states of the brain (Zee, 2012). Eye movements are the window into the brain (van Gompel et al., 2007).

Here, we will start our journey by looking at the effect of reward and reward prediction error on the kinematics of eye movements (Chapter 2). In Chapter 3, we will explore how can we modulate error sensitivity during saccadic adaptation. Next, we will switch our focus from human psychophysics to electrophysiology in non-human primates. We will establish protocols to train marmoset monkeys to perform saccadic eye movement tasks (Chapter 4) and techniques to record from the cerebellum (Chapter 4) and analyze the cerebellar electrophysiological data (Chapter 5). Then, we will study the cerebellar cortex and the role of Purkinje cells in the control of eye movements (Chapter 6). At the end (Chapter 7), we will present the main takeaways and future directions.

In the rest of Chapter 1, we will establish some of the basic concepts about eye movements and then present the specific aims that we pursued.

1.1. Self-published content and references

Chapters 2-6 draws upon both published and unpublished work. Each chapter starts with an introduction on a topic, followed by methods, results, and then discussions. Certain passages and figures of Chapter 2 are derived from Sedaghat-Nejad, et al. (2019). Certain passages and figures of Chapter 3 are derived from Sedaghat-Nejad & Shadmehr (2021). Certain passages and figures of Chapter 4 are

derived from Sedaghat-Nejad, et al. (2019). Certain passages and figures of Chapter 5 are derived from Sedaghat-Nejad, et al. (2021). And certain passages and figures of Chapter 6 are derived from Sedaghat-Nejad, et al. (2021).

1.2.   Neural basis of saccadic eye movement

Every morning that I want to go to the lab on the 4th floor, I use the elevators in the lobby of our building. I push the up-arrow button and wait for one of the elevators to become ready to take me on. There are 5 elevators, each with a big indicator above them that lights up a couple of seconds before the elevator gets to the ground level. One of the indicators will light up, I will see that in my periphery, make a saccade to it to confirm that it is indeed on, then start moving toward that elevator. In this section, I will walk you through the brief moment starting from when the indicator lights up till the moment that my eyes land on it and can see it better.

When the indicator lights up in my field of view, its luminance excites the photoreceptors of my retina. At each moment, all the lights emitting and reflecting from the world in front of us, pass through the eye's lens and get projected on the retina. This is the first transformation that takes place in our visual system, the 3D world gets projected to the 2D retinal coordinate system. In other words, if we measure the activity of retinal cells, we can precisely estimate the location of a bright object with respect to the retinal coordinate system. The activity of retinal ganglion cells gets transferred through the retinal tract to the brain. It mainly projects to the lateral geniculate nucleus (LGN) in the thalamus with some projections to the superior colliculus. From LGN, the signal goes to the primary visual cortex (V1) where there is a retinotopic map of the visual world. From V1, the signal transformed to various cortical areas as part of two main streams: the ventral pathway and the dorsal pathway. The function and specification of these pathways are not of our interest in this section, please see Goodale & Milner (1992) for further reference on this topic. Eventually, all the cortical and subcortical pathways converge on the superior colliculus to initiate a saccadic eye movement. Did the elevator indicator light up? I should make a saccade to it and foveate on it to see it better.

The superior colliculus (SC) consists of multiple layers: superficial layer, intermediate layer, and deeper layers. The superficial layer mainly contains visual-related neurons that form a retinotopic map. The intermediate layer contains visuomotor-related neurons, and the deeper layers contain mainly motor-related neurons. The motor-related neurons in SC form a motor map, that is, their activation initiates an eye movement to a specific location that they represent in the retinotopic map. The motor-

related neurons can further get divided into two groups: rostral and caudal. The rostral neurons are active during fixation of the eye and small amplitude saccades (micro-saccades). The caudal neurons are responsible for initiating a saccadic eye movement. The activity of the motor-related neurons is often referred to as "all or none". This is due to the fact that when the activity in a given spot of SC reaches a threshold it will result in the bursting activity of that spot and suppression of nearby spots along with suppression of the SC on the other side of the brain (Leigh & Zee, 2015).

Motor-related neurons in SC project to downstream brainstem excitatory burst neurons (EBNs) and inhibitory burst neurons (IBNs). The burst neurons in turn project to eye muscle motoneurons that activate the eye muscles and move the eye as a saccade. The EBNs and IBNs work in tandem with each other. EBNs excite the ipsilateral motoneurons and drive the muscle of one eye, at the same time, the IBNs inhibit the contralateral motoneurons and squelch the activity of the muscle of the other eye. Together, they trigger a saccade in the ipsilateral direction (Fuchs et al., 1985; Sparks, 2002). When approaching the end of the saccade, the burst neurons also control the endpoint accuracy of the movement and prevent the eye from overshooting the target. Both IBNs (C A Scudder, 1988) and EBNs (Van Gisbergen et al., 1981) activate reciprocally with contraversive saccades to act as a brake and stop the eyes right on the target.

The process of generating a saccade can be thought of as a combination of two mechanisms. One that triggers the saccade and accelerates the eyes; and the other that chokes the movement to ensure that the fovea lands at the desired location. Most of the low-level neural circuitry involved in saccadic eye movements is understood. The ipsilateral burst neurons activate the agonist muscles and then the contralateral burst neurons activate the antagonist muscles (Leigh & Zee, 2015).

1.3. Saccadic adaptation

In section 1.2 we looked at some of the neural substrates that are involved in generating a saccadic eye movement. In this section, we go further and review some of the basic concepts of how the saccades get modulated. Let me provide you with another personal example. Both my parents wear glasses and during my childhood, I always wanted to wear glasses to be like them. Well, around age 14, I developed nearsightedness (myopic), and an optometrist prescribed me glasses. I never forget the first moment that I wore my glasses, the whole world was moving around, I became dizzy and had difficulty maintaining my balance. But gradually, over half an hour, everything started to stabilize. Why I got dizzy and how my brain adapted to my glasses is the topic of this section.

When I wore my glasses, I added two lenses in front of my eyes that refracted the entering light to my eyeballs. This refraction changed the way my brain used to see the world through my eyes. Now, there is a new dynamic (rotational magnification) in motion and my brain needs to adapt to it. This adaptation takes place gradually through plasticity in the brain. Finding the neural substrates of the saccade plasticity is an active area of research and there is a lot to be discovered in this field (Hopp & Fuchs, 2004; Shadmehr, 2020). In Chapter 3 we looked at a novel behavioral paradigm that modulates the rate of saccade adaptation, and in Chapter 6 we examined how the activity of cerebellar Purkinje cells contributes to saccadic eye movements.

As discussed in section 1.2 the SC projects to burst neurons, but it also sends a copy of that command to the oculomotor cerebellum. The oculomotor cerebellum then projects to burst neurons to form an indirect pathway from SC to burst neurons. One hypothesis is that the direct projection from SC to burst neurons provides the initial drive for saccades while the indirect pathway corrects the saccade endpoint accuracy (Soetedjo et al., 2019). Lesion of superior colliculus leads to an increase in saccade reaction time and slower velocities, but the accuracy of saccades recovers (Hanes et al., 2005). Lesions of the cerebellum leads to permanent saccade dysmetria (Quaia et al., 1999).

Data from cerebellar patients and animal lesion studies have suggested the cerebellum as one of the main loci for motor adaptation (Iwamoto & Yoshida, 2002; Ritchie, 1976; Zee et al., 1976). In the context of saccade adaptation, the cerebellum has strong connections to the brainstem burst neurons and is well suited for adaptation-related plasticity (Soetedjo et al., 2019). The cerebellar fastigial nucleus projects contralaterally to the burst neurons (Batton et al., 1977). In addition, during the saccade adaption, the activity of neurons in the fastigial nucleus gets modulated with saccade size (Charles A Scudder & McGee, 2003).

Cerebellar malfunction produces saccadic dysmetria; Results from Zee et al. (1976) showed that patients with cerebellar degeneration suffer from permanent saccadic dysmetria. In addition, the cerebellum is required for saccade adaptation. Results from Optican & Miles (1985) and Optican & Robinson (1980) demonstrated that when oculomotor muscles of one eye were weakened and the intact eye was patched, monkeys adapted to this perturbation by sending stronger motor commands to both eyes. This compensated for the effect of weaker muscles but made the intact eye hypermetric. When the patch was placed on the weak eye, the hypermetria of the intact eye waned gradually. However, total ablation of the cerebellum resulted in extensive hypermetria and the inability of the animal to adjust the saccade size. Other studies showed that when the cerebellar vermis was lesioned,

monkeys' saccade amplitude became variable (Barash et al., 1999; Takagi et al., 1998). These monkeys have an impaired ability to adjust their saccade amplitude in response to external perturbations.

1.4.   Specific Aims

The objective of this dissertation is threefold. First, we used human psychophysics to study the behavioral correlates of the dopaminergic system on the control of saccades. Second, we studied behavioral correlates of motor learning while human participants performed a novel eye movement task. Third, we used marmoset monkeys as a non-human animal model to study the neural correlates of motor control during saccades.

1.4.1.   Aim 1. The effect of reward prediction error on movement vigor.

Does dopamine release in response to a stimulus invigorate the ensuing movement and subsequently modulate learning? To test this hypothesis, we relied on the fact that reward prediction error (RPE) is a strong modulator of dopamine. In Chapter 2, we present how we designed an innovative task in which an RPE event occurred precisely before the onset of a stimulus-driven movement. We probabilistically produced a combination of large or small, negative- or positive-RPE events, and observed that saccade vigor carried a robust signature of the preceding RPE event: high-vigor saccades followed positive-RPE events, while low-vigor saccades followed negative-RPE events. This suggests that in humans, vigor is partly controlled through the release of dopamine in the moments before the onset of a movement.

1.4.2.   Aim 2. The effect of implicit loss associated with corrective movement on motor adaptation.

The nervous system often produces a reflexive corrective movement in response to movement error which consumes time. Is there an implicit loss associated with the corrective movement that can modulate learning? To test this hypothesis, in Chapter 3, we designed an innovative task that combines saccade adaptation with random dot motion discrimination. In this task, movement error results in corrective saccades, but those corrections take time away from the discrimination task. We altered the cost of error correction by using the coherence of images and investigated how the cost of error modulated motor adaptation. We found that increasing the cost increased how much the brain learned from error. Thus, the landscape of the loss associated with the act of correcting for error regulates the rates of sensorimotor learning.

1.4.3. Aim 3. The role of cerebellar Purkinje cells during control of movement.

Human psychophysics can help us understand the behavioral correlates of motor control and motor learning, however, to understand the involved mechanisms, we should investigate the neural substrates responsible for motor control and motor learning in the cerebellum. To pursue this aim, we used marmoset monkeys as a new promising non-human primate animal model to study motor control and motor learning.

1.4.4. Aim 3.1. Developing an animal protocol for behavioral training of marmosets and electrophysiological recording from the cerebellum.

Marmosets present the opportunity to investigate genetically based neurological diseases in primates, in particular, diseases that affect social behaviors, vocal communication, and eye movements. All of these behaviors depend on the integrity of the cerebellum. In Chapter 4, we present training methods that better motivate the subjects, allowing for improved performance. Next, we describe how we designed electrophysiological techniques that precisely target the subject's cerebellum, allowing for simultaneous isolation of multiple Purkinje cells.

1.4.5. Aim 3.2. Developing a software package for spike sorting the cerebellar electrophysiological data.

Recording neuronal activity from the cerebellum is a challenging venture. On the one hand, the principal cells of the cerebellum, Purkinje cells (P-cells), can be identified due to their unique electrophysiological properties. However, detection of P-cell complex spikes is difficult because these spikes are rare, and their waveforms vary from one spike to the next. In Chapter 5, we present the challenges that we faced during spike sorting the electrophysiological data from the cerebellar Purkinje cells. We explain how we developed a spike analysis software, called P-sort, that provides a set of tools to help detect simple and complex spikes, as well as quantify whether the two events are generated by a single P-cell.

1.4.6. Aim 3.3. Synchronous spiking of cerebellar Purkinje cells during eye movement.

The information that the brain transmits from one region to another is often viewed through the lens of firing rates. However, if the neurons could vary the timing of their spikes with respect to each other, then through synchronization they could highlight information that may be critical for control of behavior. In the cerebellum, the computations that are performed by the cerebellar cortex are conveyed

to the nuclei via inhibition. Yet, synchronous activity entrains nucleus neurons, making them fire. Does the cerebellar cortex rely on spike synchrony within populations of Purkinje cells to convey information to the nucleus? In Chapter 6, we present the results of Purkinje cells' activity of the marmoset cerebellum while they performed saccadic eye movements. We organized the cells into populations that shared a complex spike response to error. Before movement onset, P-cells transmitted information via a rate code: the simple spike firing rates predicted the direction and velocity of the impending saccade. However, during the saccade, the spikes became temporally aligned within the population, signaling when to stop the movement. Thus, the cerebellar cortex relies on spike synchronization within a population of P-cells, not individual firing rates, to predict when to stop a movement.

Chapter 2. Reward prediction error modulates saccade vigor

Movement vigor, defined as the reciprocal of the latency from availability of reward to its acquisition, changes with reward magnitude: movements exhibit shorter reaction time and increased velocity when they are directed toward more rewarding stimuli. This invigoration may be due to release of dopamine before movement onset, which has been shown to be modulated by events that signal reward prediction error (RPE). Here, we generated an RPE event in the milliseconds before movement onset and tested whether there was a relationship between RPE and vigor. Human subjects made saccades toward an image. During execution of the primary saccade, we probabilistically changed the position and content of that image, encouraging a secondary saccade. On some trials, the content of the secondary image was more valuable than the first image, resulting in a positive-RPE event that preceded the secondary saccade. On other trials, this content was less valuable (negative-RPE event). We found that reaction time of the secondary saccade was affected in an orderly fashion by the magnitude and direction of the preceding RPE event: the most vigorous saccades followed the largest positive-RPE, whereas the least vigorous saccades followed the largest negative-RPE. Presence of the secondary saccade indicated that the primary saccade had experienced a movement error, inducing trial-to-trial adaptation. However, this learning from movement error was not modulated by the RPE event. The data suggest that RPE events, which are thought to transiently alter the release of dopamine, modulate the vigor of the ensuing movement.

2.1. Introduction

We tend to move with shorter latency, and greater velocity, toward stimuli that we associate with greater value. For example, when the expected reward is large, saccades (Kawagoe et al., 1998; Milstein & Dorris, 2007; Xu-Wilson, Zee, et al., 2009) and reaching movements (Summerside et al., 2018) toward the reward site have shorter reaction times and higher peak velocities than when the expected reward is smaller. That is, the reciprocal of the time it takes to arrive at the reward site, which we can operationally define as vigor (Shadmehr et al., 2019), is modulated with reward magnitude. This link between expected reward and movement vigor may be partly due to the function of the basal ganglia (Kawagoe et al., 2004; Tachibana & Hikosaka, 2012) and release of dopamine (da Silva et al., 2018), raising the possibility that before every movement, the dopamine that is released in response to the stimulus partly controls the vigor of the ensuing movement.

Dopamine release appears to follow a simple rule. When the acquired reward is unexpectedly large, the neurons fire a burst, but if the same reward is expected, the neurons no longer respond (Bayer & Glimcher, 2005; Schultz et al., 1997). That is, dopamine neurons encode the difference between the predicted stimulus value and the acquired value, termed reward prediction error (RPE). This transient encoding of RPE provides an interesting prediction: if dopamine release in the

milliseconds before movement onset contributes to control of vigor, then movements that follow a positive-RPE (+RPE) event should exhibit high vigor, and those that follow a negative-RPE (-RPE) event should exhibit low vigor. That is, vigor modulation should depend on reward prediction error, not reward itself.

Unfortunately, the hypothesis that RPE (and not reward per se) drives vigor has been difficult to test because in a typical experiment, the RPE event occurs after a movement has been completed and the reward acquired, not before the onset of a movement. Here, we designed a task that overcame this limitation.

In our experiment, we relied on the idea that viewing of images carries some of the hallmarks of reward: when given the option of choosing from various image categories, people prefer face images, and are willing to spend a greater amount of effort in exchange for gazing at those images (Aharon et al., 2001; Yoon et al., 2018). Furthermore, viewing of face images activates the brain's reward system (O'Doherty et al., 2003). We used images as a proxy for reward, and then probabilistically controlled the image content to induce RPE events. We asked whether induction of an RPE event before a movement influenced vigor of that movement.

Subjects made saccades to view an image. However, upon initiation of the saccade we probabilistically altered the position and content of the image. The position change encouraged the subjects to follow their initial saccade with a secondary saccade. Our concern was vigor of this secondary saccade, i.e., its latency and velocity, which affected the total time it took from completion of the primary saccade to conclusion of the secondary saccade, thereby arriving at the reward site.

On some trials, value of image A (primary image) was higher than image B (secondary image), while in other trials value of B was higher than A. As a result, in some trials subjects expected to view a low valued image, but upon completion of their primary saccade, were presented with the opportunity to gaze at a high valued image. This resulted in conditions in which during the milliseconds before the onset of the secondary saccade (as A was replaced by B), there was a +RPE (B>A) or a -RPE (B<A) event. We asked whether the sign and magnitude of the RPE event altered vigor of the secondary saccade.

2.2. Methods

A total of n=55 healthy subjects (18-41 years of age, 23±7 mean±SD, 34 females) participated in this study. The procedures were approved by the Johns Hopkins School of Medicine Institutional Review Board. All subjects signed a written consent form.

2.2.1. Data collection procedure

Subjects viewed an LED monitor (27-inch, 2560x1440 pixels, light gray background, refresh rate 144 Hz) placed at a distance of 35 cm while we measured their eye position at 1000 Hz (Eyelink 1000+, SR Research). Each trial began with presentation of a fixation spot (a green dot, 0.3°×0.3°) that was randomly drawn near the center of the screen (the fixation spot was placed in a virtual box at -3° to +1° along the horizontal axis, and -1.5° to +1.5° along the vertical axis, where 0,0 refers to center of the screen). After a random fixation interval of 250-750 ms (uniform distribution), the fixation spot was erased, and a primary image was placed at 9° to the right along the horizontal axis. The size of this image was constant for each subject but varied between subjects: 1.5°×1.5° for some (n=20, 100 trials per block, 13 blocks of trials), 3.0°×3.0° for others (n=35, 145 trials per block, 13 blocks of trials). A green fixation dot always appeared at the center of every image. As the effect of RPE relied on within-subject analysis, data were combined in these two groups.

The removal of the central fixation dot and presentation of the primary image served as the go signal for the primary saccade. This saccade was detected in real-time via a speed threshold (20 °/s), or an eye position change of 2° from fixation, whichever happened first. Each session contained 13 blocks of trials. In block 1 the primary image remained unchanged following saccade onset. In blocks 2-13, after detecting saccade onset, the primary image was erased with probability of 50% and a new image was displayed at a distance of 3° from the original image. A green fixation dot also appeared at the center of the secondary image. As a result, following the completion of the primary saccade, subjects produced a secondary saccade to the secondary image. The location of the secondary image was random on each trial. The reason for this was to preclude accumulation of adaptation on the primary saccades that results from movement errors that subjects experience on each trial (Ethier et al., 2008b; Pekny et al., 2011). For some subjects (n=33), the secondary image was randomly located at either +3° or -3° along the horizontal axis with respect to the primary image. For other subjects (n=22), the secondary image was randomly located at each trial at either +3° or -3° along the vertical axis with respect to the primary image. As a result, the location of the secondary image was random along the horizontal or vertical axis. The size of the secondary image was always the same as the primary image.

Following completion of the secondary saccade, subjects were provided with 250 ms to view that image. At the end of this period the image was erased, and a center fixation dot appeared at a random location near the center of the screen, in the bounding box defined above. Each session

contained 13 blocks of trials, with 100 or 145 trials per block. Subjects were provided with a 30 second rest period between each block.

Images were chosen from two categories: face and noise images. The facial images were gathered from the Internet (500 total images) and were modified in a way that the center of the two eyes was located at the center of the image. The noise images were constructed by shuffling the pixels of each face image (500×500 pixels). This ensured that the luminance and color content of the two categories were identical.

2.2.2. Magnitude of the RPE event

In our experiment we presented a primary image (e.g., face), and then at random trials replaced it with another image (noise). We hypothesized that viewing each image was a rewarding event, and as a result a difference between the primary and secondary images would produce a reward prediction error prior to the execution of the secondary saccade. We estimated the magnitude of the RPE from probability of each image and its relative value.

An objective estimate of the value of a face image with respect to a noise image can be attained from the choices that people make when given the option of viewing these images. For the image types that we employed here, people on average choose the face image twice as often as the noise image (Yoon et al., 2018). This suggests that the relative value of face to noise is around 2.

On a given trial, the primary image was changed with a probability of 50%. Assuming a prior probability that target of the saccade is unlikely to change, and the observed likelihood that on 50% of trials the primary image changes, we can write the predicted value of the primary image as follows:

$$\widehat{F} = \alpha F + (1-\alpha)N \quad 0.5 < \alpha < 1$$
$$\widehat{N} = \alpha N + (1-\alpha)F$$

Equation 2.1

In Equation 2.1, F and N represent the subjective value of face and noise images, and $\widehat{F}$ and $\widehat{N}$ are the predicted value. Because the primary image can change, the first equation in the above expression implies that the predicted value for a primary face image is less than its subjective value (the face can become noise). The second equation implies that the predicted value for a primary noise image is greater than its subjective value (noise can become face).

Once the primary saccade concludes, the subject is presented with a secondary image. This is the image that they will actually have the opportunity to gaze at. We define RPE as the value of the

second image (reward we will receive) minus its predicted value (reward we had predicted). For example, if A is the primary image, and B is the secondary image, then $\hat{A}$ is the reward predicted, but B is the reward that will be received. That is, $RPE[A, B] = B - \hat{A}$.

There are four possible pairs of primary and secondary images. For each pair, we can compute the magnitude of the RPE event:

$$RPE[F, F] = F - \hat{F} = (1 - \alpha)(F - N)$$

$$RPE[F, N] = N - \hat{F} = \alpha(N - F)$$

$$RPE[N, N] = N - \hat{N} = (1 - \alpha)(N - F)$$

$$RPE[N, F] = F - \hat{N} = \alpha(F - N)$$

Equation 2.2

If we assume that the subjective value of a face image is roughly twice that of noise, $F \approx 2N$, we have:

$$RPE[N, F] = \frac{\alpha}{2}F$$

$$RPE[F, F] = \frac{1 - \alpha}{2}F$$

$$RPE[N, N] = \frac{\alpha - 1}{2}F$$

$$RPE[F, N] = -\frac{\alpha}{2}F$$

Equation 2.3

Equation 2.3 imply that a noise-face (NF) trial is a very positive RPE event (first equation in the above expression), a face-face (FF) trial is a mildly positive RPE event (second equation), a noise-noise (NN) trial is a mildly negative RPE event (third equation), and a face-noise (FN) trial is a highly negative RPE event (fourth equation).

If vigor is modulated by RPE, then the secondary saccades should exhibit their highest vigor in NF trials, and lowest vigor in FN trials. In comparison, FF trials should show smaller vigor as compared to NF trials, despite the fact that in both trials the secondary saccade is toward a face. Finally, NN trials should show a greater vigor than FN trials, despite the fact that in both trials the secondary saccade is toward noise.

2.2.3. Data analysis

Eye position data were filtered with a second-order Butterworth low-pass filter with cutoff frequency of 100 Hz. Eye velocity data in offline analysis were calculated as the first derivative of the filtered position

data. Saccades were identified with a speed magnitude threshold of 20 °/s, and minimum hold time of 10 ms at saccade end (i.e. velocity magnitude could not exceed the cutoff for a minimum 10 ms after endpoint). We measured reaction time of the secondary saccade via the time period between offset of the primary saccade and onset of the secondary saccade. Secondary saccades onset and offset were detected identically to the primary saccades, using 20 °/s threshold on velocity magnitude. The saccade duration was considered as the time between saccade onset and offset.

Statistical analyses were performed using SPSS and general linear models, with stimulus value (e.g., face or noise) serving as the within-subject factor. We reported results of tests of within-subject effects under the assumption of sphericity. We tested this assumption via Mauchly's test, which was confirmed in every case reported. We also performed two-sided t-tests on the between-subject effect of learning from movement errors.

2.3.   Results

To produce a +RPE event, the trial began with a noise image (Figure 2.1A). As subjects initiated their primary saccade, we probabilistically erased that image and replaced it with a face image at a new location (NF trials, Figure 2.1A, first column). The control condition for the +RPE event was a trial in which both the primary and secondary images were faces (FF control, Figure 2.1A, second column). Similarly, in order to produce a -RPE event, the trial began with presentation of a face image, which following saccade onset, was probabilistically replaced with a noise image (FN trials, Figure 2.1A, third column). The control condition for the -RPE event was a trial in which both the primary and secondary images were noise (NN control, Figure 2.1A, forth column). Therefore, the secondary saccade in both the control and +/-RPE trials was made toward the same image. Based on probability of the events, the trials produced four magnitudes of RPE: highly positive (NF trials), slightly positive (FF trials), slightly negative (NN trials), and highly negative (FN trials) (Equation 2.3).

2.3.1.   Effects of reward prediction error on vigor

Data from a representative subject are shown in Figure 2.1B. Position traces for single saccades are shown in the top panel of Figure 2.1B, and averaged velocity profiles are shown in the bottom panel. As expected, the primary saccade had a shorter reaction time and higher peak velocity when made toward a face image. During the primary saccade, on some trials the face image was changed to noise (-RPE event, FN trial). Similarly, on some trials the noise image was changed to face (+RPE event, NF trial). Because of the change in image location, at 100-150 ms following completion of the primary saccade the

subject generated a secondary saccade. We measured the reaction time of the secondary saccade with respect to end of the primary saccade. The reaction time and peak velocity of the secondary saccade were affected by not just the image at the destination (i.e., the secondary image), but more importantly, by the sign of the RPE event. Reaction time of the secondary saccade appeared shortest following the +RPE event, and longest following the -RPE event. Indeed, the properties of the secondary saccade appeared to follow a consistent pattern: shortest reaction time and highest velocity for the most positive RPE event (NF), longest reaction time and lowest velocity for the most negative RPE event (NF), and in between for the mildly positive (FF) and mildly negative (NN) RPE events.
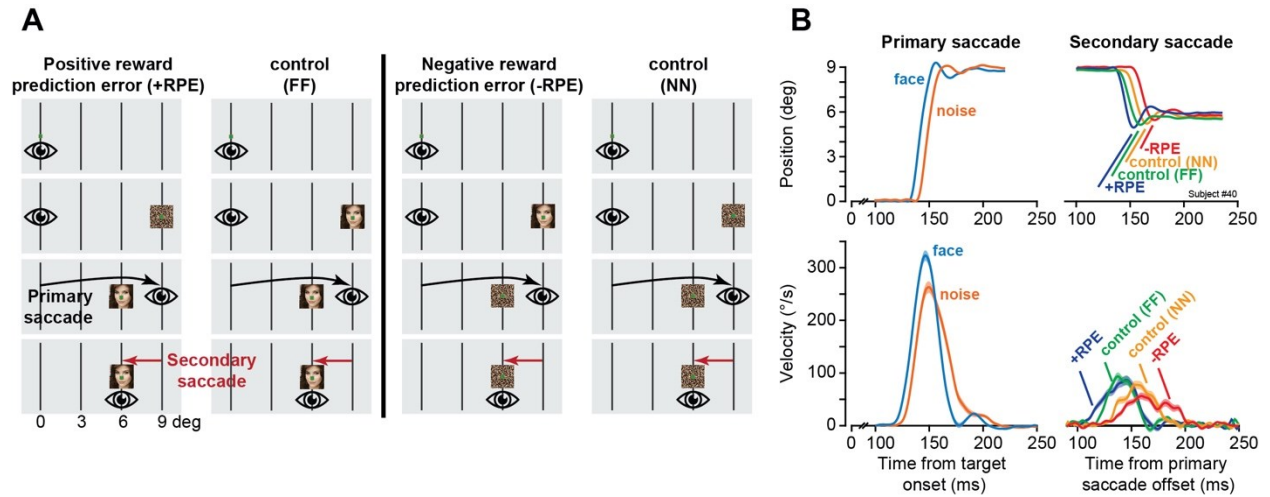
Figure 2.1. Experiment design and data from a representative subject.
A. Each trial began with a fixation dot near the center. Following a random fixation interval, we presented a primary image at 9° to the right along the horizontal axis. As the primary saccade took place, we erased the primary image and replaced it with a secondary image. In +RPE trials, a noise primary image was replaced with a face secondary image (1st column). A face-face trial (2nd column) served as control for the +RPE trial. In -RPE trials, a face primary image was replaced with a noise secondary image (3rd column). A noise-noise trial (4th column) served as control for the -RPE trial.
B. Saccade position and velocity traces for a representative subject. Position traces are for individual saccades. Velocity traces are within subject average of all saccades. Primary saccade exhibited a shorter reaction time and a higher velocity in response to a face image. Data for the secondary saccade are plotted with respect to termination of the primary saccade in the same trial. The secondary saccades exhibited shortest reaction times in +RPE trials, and longest reaction times in -RPE trials. Error bars are SEM.

These results were repeated in our population of subjects. The opportunity to view a face image strongly affected the vigor of the primary saccade (Figure 2.2). The distribution of reaction times shifted earlier (Figure 2.2A), from a mean of 150.2 to 140.95 ms, resulting in a within-subject reduction of 9.30±0.63 ms (mean±SEM) (within-subject change, $F_{(1,54)}$=217, $p<10^{-4}$). The face image also induced an increase in the velocity of the primary saccade (Figure 2.2B), particularly in the second half of that

movement. As a result, saccade peak velocity increased by 2.63±0.76 °/s (within-subject change, $F_{(1,54)}$=12.1, p=0.001). Maximum change occurred 10 ms after peak velocity with 6.01±0.84 °/s difference between two categories (within-subject change, $F_{(1,54)}$=51.5, $p<10^{-4}$).

We tried to minimize the changes in saccade amplitude that may arise from changes in image content by presenting a green dot at the center of every image. We observed a very small effect of image type on saccade amplitude: the primary saccades were 8.19±0.01° toward face images, and 8.10±0.01° toward noise images, a within-subject change of 0.09±0.015° ($F_{(1,54)}$=32, $p<10^{-4}$). However, this change was less than the resolution of our measurement instrument. Overall, the opportunity to view a face image produced a significant reduction in the total time it took for the eyes to arrive at the location of the primary target (Figure 2.2C, within-subject change, $F_{(1,54)}$=199, $p<10^{-4}$).
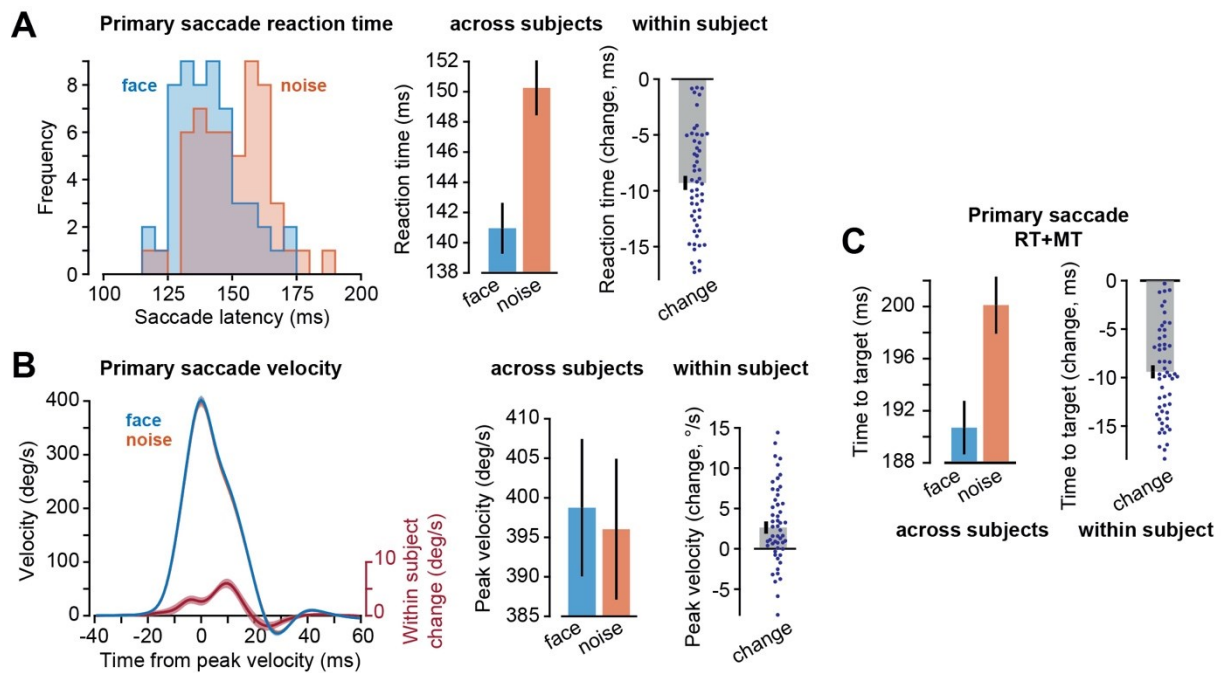
Figure 2.2. Primary saccades exhibited shorter reaction time and higher velocity in response to face images.
A. Distribution of mean reaction times and within-subject change in the mean reaction time (face minus noise). Each dot is a single subject. **B.** Saccade velocity traces, and within-subject change in velocity (face minus noise). The trace for noise is largely hidden behind the trace for face. **C.** Total time to target (reaction time plus movement duration) and within-subject change. Error bars and shaded traces are SEM.

As the primary saccade started, we displaced the primary image to a new location, encouraging the subjects to produce a secondary saccade. We found that the reaction times for the secondary saccade (Figure 2.3B) were shortest in the NF trials (+RPE event), and longest in the FN trials (-RPE

event). Indeed, there was an orderly increase in the reaction times in the precise pattern predicted by the RPE events (Figure 2.3C, RM ANOVA, $F_{(3,52)}$=34.7, p<10^{-4}). On average, the secondary saccade that followed the +RPE event had a reaction time that was 19.15±1.95 ms less than the -RPE event (148.6±3.1 ms in +RPE trial as compared to 167.75±3.3 ms in -RPE trial). Post-hoc pairwise comparisons indicated that reaction time of +RPE trials were 6±1.09 ms smaller than FF-control trials (p<10^{-4}), and reaction time of -RPE trials were 6.03±0.95 ms greater than NN-control trials (p<10^{-4}). Peak velocity of the secondary saccade appeared affected by the various events (Figure 2.3A&D, RM ANOVA, $F_{(3,52)}$=14.7, p<10^{-4}). However, post-hoc pairwise comparisons did not dissociate the +/-RPE events from their respective control trials.

Overall, the RPE events significantly affected the total time it took for the eyes to respond and acquire the secondary target (Figure 2.3E): the time to target, measured from completion of the primary saccade to arrival at the target, was smallest following the +RPE event (NF trials), and largest following the -RPE event (FN trials, RM ANOVA, $F_{(3,52)}$=34.9, p<10^{-4}). Post-hoc pairwise comparisons indicated that time to target of +RPE trials were 6.06±1.14 ms smaller than FF-control trials (p<10^{-4}), and time to target of -RPE trials were 6.15±0.93 ms larger than NN-control trials (p<10^{-4}). That is, the magnitude of the RPE event that preceded a saccade affected the vigor of that saccade.

Amplitude of the secondary saccade on average varied by less than 0.09° across the range of the various conditions: NF 3.02±0.04°, FF 3.08±0.04°, NN 3.04±0.05°, and FN 3.0±0.05°. Similarly, metrics of the primary saccade were not affected by whether the image was changed midflight. To check for this, we grouped NF and FN trials (image changed) and compared them to FF and NN trials (image unchanged). We found no difference between primary saccades in these trials (repeated measure ANOVA for the effect of stimulus on peak velocity, within-subject change, $F_{(1,54)}$=0.022, p=0.884; reaction time within-subject change, $F_{(1,54)}$=0.302, p=0.585; time to target within-subject change, $F_{(1,54)}$=0.919, p=0.342; and amplitude within-subject change, $F_{(1,54)}$=0.499, p=0.483).
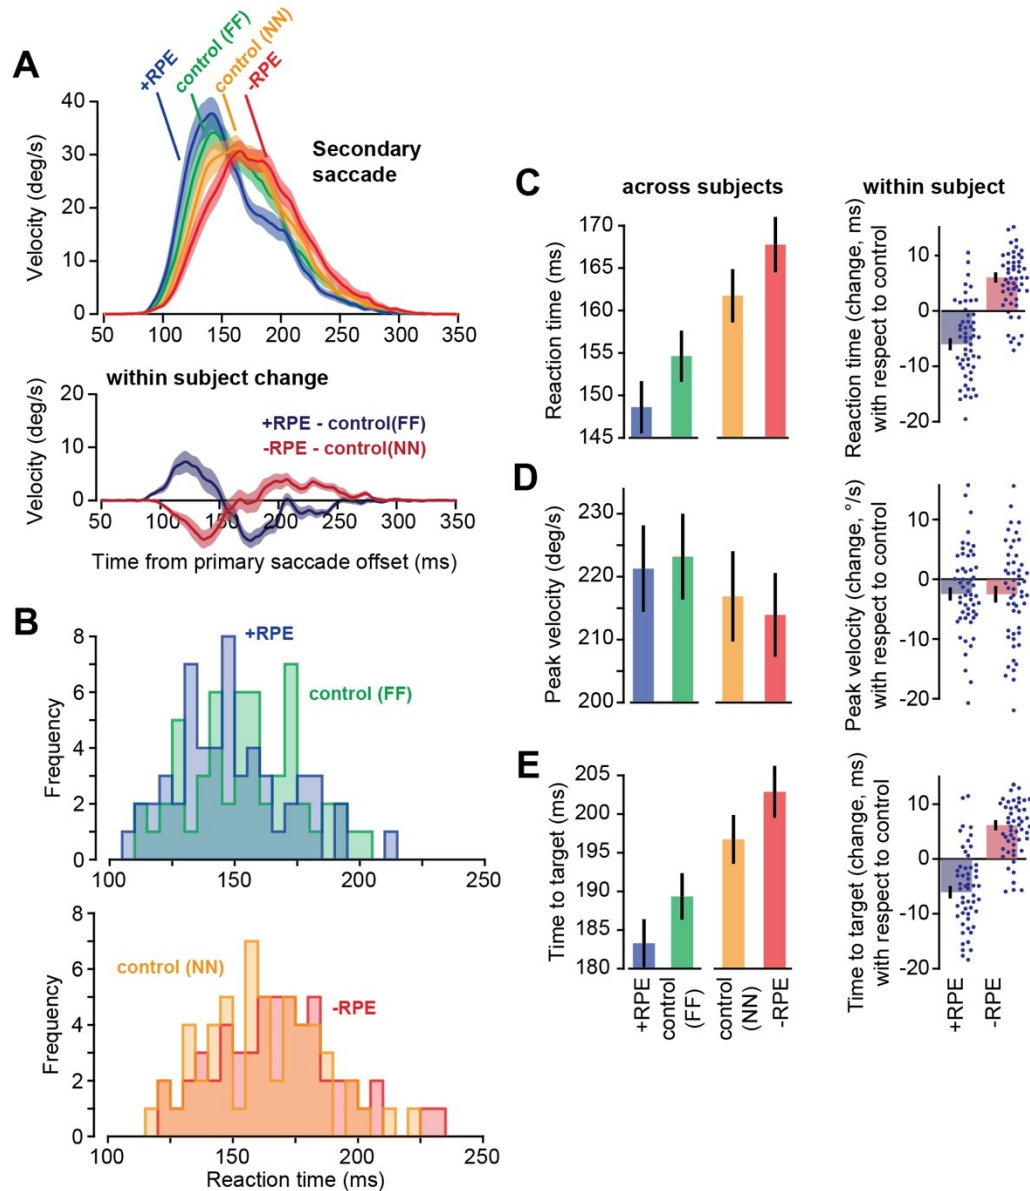
Figure 2.3. Secondary saccades were influenced by the preceding RPE event.
A. Saccade velocity and within-subject change in velocity. B. Distribution of mean reaction times across subjects. Reaction times are measured as the latency with respect to offset of the preceding primary saccade. C. Mean reaction times across subjects and the within-subject change in reaction times. The bars are +RPE condition with respect to control(FF), and -RPE condition with respect to control (NN). Dots are individual subjects. D. Peak velocity and within-subject change in peak velocity. E. Total time to target, measured from completion of the primary saccade to conclusion of the secondary saccade, i.e., the reaction time plus movement duration of the secondary saccade. The error bars and shaded traces are SEM. The dots show within-subject change with respect to control.

17

2.3.2. Effect of reward prediction error on learning

Presence of a secondary (or corrective) saccade indicates presence of a motor error: at the end of the primary saccade, the target was not on the fovea. The resulting movement error should induce plasticity in the cerebellum (Herzfeld et al., 2018), affecting the subsequent primary saccade. We wondered whether presence of the RPE event modulated learning from the motor error. In particular, would a +RPE event enhance learning?

In our experiment we displaced the primary image along four directions: positive and negative along the horizontal axis (H+ and H-), and positive and negative along the vertical axis (V+ and V-). In all cases the magnitude of the displacement was 3°. Each displacement resulted in a motor error, which in principle may have induced trial-to-trial learning. To measure this learning, we considered two consecutive trials in which the primary saccades were made to the same visual stimulus type, and then further divided the trials based on the direction of motor error. For example, suppose that on trial n the primary saccade was toward face, the subject experienced an H+ error on that trial, and that on trial n+1 the primary saccade as again toward face. In all such consecutive pairs of trials, we measured the change in the primary saccade made in trial n+1 with respect to trial n. This trial-to-trial change in the primary saccade is plotted in Figure 2.4A for each type of motor error. We found that across all error types, the largest change in the velocity profile was around 15 ms after the saccade peak velocity (Figure 2.4A). Following an H+ error, the tail of velocity trace (15 ms after peak velocity) increased by 6.62 ± 1.02 °/s (two-sided t-test, $t_{(32)}=6.49$, $p<10^{-4}$) along the horizontal direction. The trial-to-trial change in saccade amplitude showed 0.17 ± 0.018° (two-sided t-test, $t_{(32)}=9.10$, $p<10^{-4}$) increase following an H+ error. Similarly, following an H- error, the subsequent primary saccade exhibited a 4.96 ± 1.09 °/s (two-sided t-test, $t_{(32)}=4.55$, $p<10^{-4}$) decrease in the tail of velocity trace along the horizontal direction and 0.17 ± 0.016° (two-sided t-test, $t_{(32)}=10.75$, $p<10^{-4}$) reduction in amplitude.

Learning was also present following V+ and V- errors. Following a V+ error, there was 1.28 ± 0.42 °/s (two-sided t-test, $t_{(21)}=3.04$, p=0.006) change in velocity tail and 0.027 ± 0.012° (two-sided t-test, $t_{(21)}=2.35$, p=0.029) change in amplitude of vertical component of primary saccade in the subsequent trial. Following a V- error, there was 1.41 ± 0.42 °/s (two-sided t-test, $t_{(21)}=3.37$, p=0.003) change in velocity tail and 0.033 ± 0.0086° (two-sided t-test, $t_{(21)}=3.81$, p=0.001) change in vertical amplitude. These results demonstrated that experience of a motor error on a given trial induced learning, resulting in an error-dependent change in the subsequent primary saccade.

18
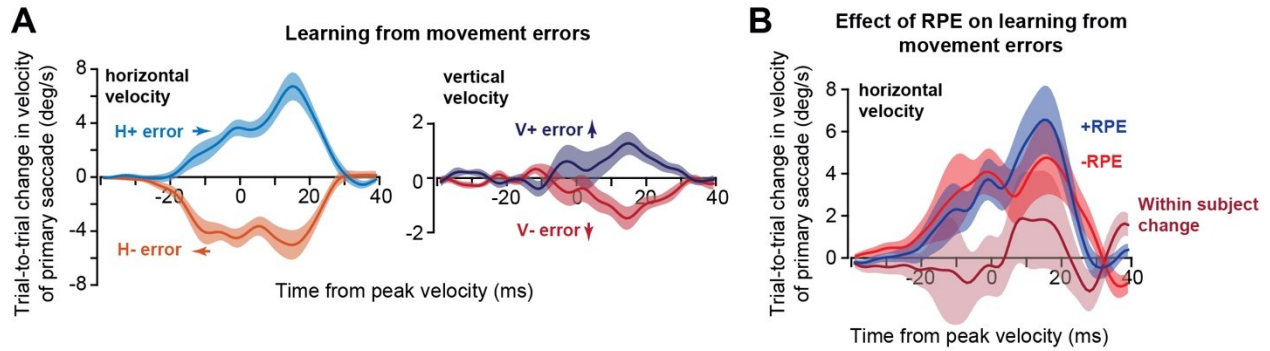
Figure 2.4. Learning from movement errors and the effect of RPE on learning.
A. Change in the velocity of the primary saccade from the trial in which the movement error was experienced to the next trial (the primary image type in both saccades were the same). Saccades are grouped based on the movement error experienced at the conclusion of the first primary saccade. Movement error is defined as the position of the secondary image with respect to the primary image. That is, H+ errors imply that the secondary image was further to the right along the horizontal axis than the primary image. The motor error along each axis produced learning along that axis. **B.** Trial-to-trial changes for horizontal motor errors only, grouped based on the RPE event that followed the first primary saccade. Data for H- was flipped and averaged with H+ trials. Within-subject change in learning was marginally larger following a +RPE event (as compared to a -RPE event), but this effect was not significant.

In some trials, a given movement error occurred in the context of a +RPE event, whereas in other trials the same movement error occurred in the context of a -RPE event. To measure the effect of the RPE events on learning, we focused on the horizontal error trials, as the vertical error trials produced substantially less learning. The trial-to-trial change in the primary saccade in the +RPE and -RPE trials are plotted in Figure 2.4B. To make this figure, we considered the H+ and H- pairs of trials together (with the positive axis now reflecting trial-to-trial change in velocity along direction of the movement error). We found that learning following a +RPE event (6.46 ± 1.59 °/s change in velocity tail, $t_{(32)}=4.07$, $p<10^{-4}$) was marginally stronger than following a -RPE event (4.64 ± 1.73 °/s, $t_{(32)}=2.68$, $p=0.012$), but the effect did not reach statistical significance (1.83 ± 2.32 °/s within-subject change, $F_{(1,32)}=0.620$, $p=0.23$).

2.4.   Discussion

It is possible that dopamine release in the milliseconds before onset of a movement serves to invigorate the ensuing movement. For example, in mice, self-initiated movements in an empty field that have higher than average acceleration tend to be preceded by higher than average activity in nigral dopamine cells (da Silva et al., 2018). Here, we attempted to indirectly test the link between dopamine release and vigor of self-generated movements in humans. Our approach was to probabilistically change the potential reward for completing a movement, thereby producing a combination of large or small,

negative or positive RPE events. Because dopamine release is affected by the magnitude and direction of the RPE event (Bayer & Glimcher, 2005; Schultz et al., 1997), we posited that vigor of the movement that followed the RPE event should exhibit a distinct pattern: highest vigor following a large +RPE event, lowest vigor following a large -RPE event.

Our innovation was a behavioral paradigm in which the RPE events occurred just before the onset of a movement. Subjects were presented with the opportunity to view face or noise images. Once they initiated their primary saccade, we probabilistically changed the location and content of the image. As a result, following completion of the primary saccade, the image value was higher or lower than expected. This resulted in a range of RPEs: highly negative (face change to noise, FN), slightly negative (noise not changed, NN), slightly positive (face not changed, FF), and highly positive (noise changed to face, NF). Notably, this RPE event took place just before onset of the secondary saccade, which had a latency of around 150 ms with respect to termination of the primary saccade. We found that reaction times were shortest following the large +RPE event, and longest following the large -RPE event. The time to target, defined as sum of reaction time and movement-duration, was shortest for the largest +RPE event, and longest for the largest -RPE event. That is, the magnitude and direction of the RPE event modulated vigor of the ensuing saccade.

Although there are currently no commonly accepted definitions of movement vigor, in the context of elementary, stimulus-driven movements such as saccades and reaching, one useful definition is the inverse of the time from stimulus onset to movement completion, conditioned on distance (Shadmehr et al., 2019). This definition is based on the empirical observation that both reaction time and movement duration are influenced by the subjective value of the reward at the destination (Haith et al., 2012; Kawagoe et al., 1998; Manohar et al., 2015; Milstein & Dorris, 2007; Reppert et al., 2015; Summerside et al., 2018; Xu-Wilson, Zee, et al., 2009).

2.4.1. Effects of visual stimulus properties on movement vigor

Because saccade reaction time and velocity depends on activity of neurons in the superior colliculus (Smalianchuk et al., 2018), and these neurons are influenced by luminance and other low-level properties of the visual stimulus (Marino et al., 2015), the differences in vigor may have arisen not from presence of an RPE event, but rather because of other variables associated with differences in properties of the secondary image. Therefore, we included control trials in which the saccade of interest was made to the same image as in RPE trials, but without the benefit of an RPE event. We found that in

+RPE trials, saccades had shorter reaction times as compared to image-matched control trials. Similarly, in -RPE trials, saccades had longer reaction times as compared to image-matched control trials. Furthermore, because the primary and secondary images were always presented at different locations with respect to the fovea, and often in opposite directions, we would expect little or no overlap between regions of collicular activity associated with the primary and secondary saccades. This dissociation between magnitude of primary and secondary saccades reduces the possibility of an interaction between the saccades at the level of colliculus. As a result, it seems likely to us that the vigor differences in the secondary saccades were not solely due to differences in the intrinsic response of the collicular neurons to information that they received from the retina. Rather, the link between RPE events in our experiments and modulation of vigor may have been because of reward-dependent regions that project to the colliculus, such as the basal ganglia and the frontal eye field.

Saccades toward a rewarding stimulus exhibit greater vigor partly because the opportunity for reward reduces the inhibition that the colliculus receives from the substantia nigra reticulata (SNr) (Yasuda et al., 2012). In addition, the opportunity for reward also increases the excitation that the colliculus receives from the frontal eye field (FEF) (Glaser et al., 2016; Heitz & Schall, 2012) It seems likely that dopamine plays an important role in controlling this drive. Just before onset of a spontaneous movement, there is diversity of responses among dopamine neurons: some show a transient increase, while others show a transient decrease (da Silva et al., 2018). Notably, for the dopamine cells that increase their activity, the amount of increase is positively corrected with the acceleration of the upcoming movement (da Silva et al., 2018).

In monkeys, dopamine neurons respond to presentation of a saccade target within 100 ms, and dissociate between reward and non-rewarding stimuli within 150 ms (Matsumoto & Hikosaka, 2007). For unconditioned, aversive stimuli, dopamine response can dissociate between various magnitudes in less than 100 ms (Matsumoto & Hikosaka, 2009). In our data, saccades that were affected by the RPE event were generated at extremely short reaction times: around 150 ms. Therefore, in principle, the time-range of dopamine response to reward is near the window of vigor modulation we observed in saccades. Because RPE events have a robust effect on dopamine release, it is possible that the RPE event driven change in saccade vigor is linked through dopamine dependent drive to the basal ganglia and FEF, affecting saccade-related discharge in the superior colliculus.

2.4.2. Reward dependent modulation of learning from movement error

If the primary saccade ends but the target is not on the fovea, the result is a sensory prediction error that produces unexpected activity on the superior colliculus (Kojima & Soetedjo, 2017b, 2018). This activity engages the inferior olive, resulting in modulation of complex spikes in the Purkinje cells of the cerebellum (Herzfeld et al., 2015; Kojima et al., 2010a; Soetedjo et al., 2008). The result is plasticity in the Purkinje cells (Herzfeld et al., 2018), altering the primary saccade on the subsequent trial. In our experiment we found robust evidence for this trial-to-trial learning: following a secondary saccade, the velocity trajectory of the next primary saccade was adjusted in the direction of the motor error.

Previous work has suggested that reward magnitude can modulate the rate of saccade adaptation. Kojima and Soetedjo (2017a) observed that when monkeys were rewarded for making saccades for one direction but not for another, they adapted their primary saccades more strongly in the direction of the rewarded trials. This reward-dependent effect developed slowly, accumulating over the course of about 400 trials.

Here, we did not see a significant effect of RPE on learning from movement errors. This may have been because we focused on adaptation in response to random errors, rather than consistent errors. Randomness of errors down-regulates learning from error (Herzfeld, Vaswani, et al., 2014), which makes it more difficult to measure any modulatory influence that reward may have had on learning from error.

2.4.3. Limitations

Although we observed robust effects of RPE on saccade latencies, we did not observe an effect on saccade velocities. One reason for this may have been that in our experiment, the saccades of interest (the secondary saccades) were only 3° in amplitude. While larger saccades produce higher velocities, they require presentation of stimuli farther from the fovea, which may make identification of that image more difficult. Regardless, the question of whether RPE events alter velocity of the ensuing movement can benefit from further exploration.

Our interpretations regarding RPE events relied on the assumption that the opportunity to gaze at an image served as proxy for reward acquisition. This assumption is based on the observation that for humans, gazing at images follows many of the behavioral characteristics associated with acquisition of primary rewards (i.e., food). For example, people make saccades that are faster toward images that they prefer (Xu-Wilson, Zee, et al., 2009), they gaze for a longer period of time at those images (Yoon et al.,

2018), and are willing to pay a greater effort cost in order to have the opportunity to view their preferred images (Yoon et al., 2018). Furthermore, viewing of images activates the reward system of the brain (O'Doherty et al., 2003). However, despite these observations, the question of whether gazing at an image engages the dopamine system remains to be explored.

In summary, whereas earlier work had demonstrated that movements are more vigorous toward more rewarding stimuli, here we found that the RPE event that takes place in the moments before onset of a movement, and not reward in itself, is necessary for modulation of movement vigor.

Chapter 3. The cost of correcting for error during sensorimotor adaptation

Learning from error is often a slow process. In machine learning, the learning rate depends on a loss function that specifies a cost for error. Here, we hypothesized that during motor learning, error carries an implicit cost for the brain because the act of correcting for error consumes time and energy. Thus, if this implicit cost could be increased, it may robustly alter how the brain learns from error. To vary the implicit cost of error we designed a task that combined saccade adaptation with motion discrimination: movement errors resulted in corrective saccades, but those corrections took time away from acquiring information in the discrimination task. We then modulated error cost using coherence of the discrimination task and found that when error cost was large, pupil diameter increased, and the brain learned more from error. However, when error cost was small, the pupil constricted, and the brain learned less from the same error. Thus, during sensorimotor adaptation, the act of correcting for error carries an implicit cost for the brain. Modulating this cost affects how much the brain learns from error.

3.1. Introduction

In machine learning, the error in the output of an artificial neural network is evaluated by a loss function that depends on the difference between the output and the desired one. This loss function is a mathematical description of the cost of error, which in turn is the principal driver of how much the network learns from error. In analogy to machine learning, during sensorimotor tasks, human learning also depends on a loss function (Körding & Wolpert, 2004) that tends to grow with error magnitude (Marko et al., 2012). This implies that in principle, altering the landscape of the sensorimotor loss function should affect the rate of learning. However, it has been difficult to find ways to modulate the sensorimotor loss function.

Previous approaches have considered inducements such as money (for humans) and food (for animals), thus associating an explicit cost to the movement error. For humans, associating error magnitude with monetary gains or losses can be effective in some cases (Galea et al., 2015; Nikooyan & Ahmed, 2015; Song & Smiley-Oyen, 2017), but not others (Quattrocchi et al., 2018; Spampinato et al., 2019). More recent results suggest that when there is an effect of reward on the rate of learning, it acts primarily through recruitment of the explicit, cognitive component of adaptation, not the implicit, unconscious component (Codol et al., 2018). On the other hand, in monkeys the presence of reward for one direction of saccade but not another (Kojima & Soetedjo, 2017a), or one target of visual pursuit but not another (Joshua & Lisberger, 2012), tends to increase the rate of learning for the rewarded direction or target.

Here, we began with the idea that when an error occurs during a movement, that error often engages a reflexive response that attempts to correct for the error, which in turn consumes time and energy. Thus, an implicit cost of the erroneous act is the penalty of time and energy paid during correction. For example, if a saccadic eye movement misses the target, the resulting error is followed by a corrective saccade. However, corrective movements carry a cost because they delay the acquisition of reward (Shadmehr et al., 2010). Thus, a natural loss function for movement error is the time that is expended in the act of producing the correction. If this time could be linked with a utility, then the landscape of the loss function may be altered, resulting in modulation of learning rates.

Here, we designed a paradigm that combined saccade adaptation with decision-making in a random dot motion discrimination task. Like traditional saccade adaptation tasks, subjects made a saccade toward a visual target and experienced an error that encouraged a corrective movement. However, unlike traditional tasks, the corrective saccade carried a cost: it consumed time needed to acquire information for the decision-making task. We varied this cost via coherence of the moving dots and found that increasing the error cost robustly increased how much the brain learned from error.

3.2. Methods

3.2.1. Subjects

A total of n=128 healthy subjects (18-54 years of age, 23±7 mean±SD, 66 females) participated in our study. The procedures were approved by the Johns Hopkins School of Medicine Institutional Review Board. All subjects signed a written consent form.

3.2.2. Data collection procedure

We considered three factors that could influence motor learning: cost of error, task difficulty, and reward. Our experiment design is summarized in Figure 3.1E, and Figure 3.9B.

In main experiments (1, 2, and 3), control-1, and control-2, participants sat in front of an LED monitor (27-inch, 2560x1440 pixels, light gray background, refresh rate 144 Hz) placed at a distance of 35 cm while we measured their eye position at 1000 Hz (Eyelink 1000+, SR Research). Each trial began with presentation of a fixation point (a green dot, 0.5°×0.5°) that was drawn near the center of the screen: the fixation point was placed randomly in a virtual box at -1° to +1° along the horizontal axis, and -1° to +1° along the vertical axis, where (0, 0) refers to center of the screen. After a random fixation

interval of 250-750 ms (uniform distribution), the fixation point was erased and a primary target was placed at 15° to the right or left along the horizontal axis.

In the main experiments (1, 2, and 3), removal of the central fixation and presentation of the primary target (a green dot, 0.5°×0.5°) served as the go signal for the primary saccade. After detecting primary saccade onset, the primary target was erased, and a random dot image was displayed. The image was a 3°×3° box with invisible borders containing a 0.5°×0.5° green dot at the center and 100 0.1°×0.1° white dots moving at 5 °/s either upwards or downwards with a predefined coherence.

The coherence was implemented via the portion of the dots that were moving either upwards or downwards (the rest of the dots were moving in a random direction). For example, if the coherence was 75% in upward direction, 75 points moved at 90° at 5 °/s speed and 25 points moved at 25 randomly assigned directions with 5 °/s speed. The locations of the 100 dots were assigned randomly at the beginning. When a dot hit one of the invisible borders, its position got reset to the opposite border while maintaining its angle and speed. The source code for generating the random dot image is available for download from the project's OSF repository (https://doi.org/10.17605/osf.io/H24J8, ESN_Moving_Dots.c). In addition, a MATLAB version of the code was included along with the c-code (ESN_Moving_Dots.m).

The location of random dot image was defined based on the trial type: during baseline trials the image was centered at primary target location. During perturbation trials the image was centered at 5° from the primary target toward the center of the screen. During error clamp trials the image was centered at the location of the primary saccade offset.

During adaptation trials, following completion of the primary saccade subjects produced a corrective saccade to place the random dot image on their fovea. This corrective movement carried a cost because it reduced the time available for the subject to view the image. This is because following detection of primary saccade offset, the image was available for only a limited time. The limited availability of the information, and the fact that corrective saccade's reaction time and execution time took away time from viewing this image, were key factors in our experiment design.

In the main experiments (1, 2, and 3) the image was present for only 300 ms after primary saccade offset. This was the only time available to view the image and decide on the direction of motion of the random dots. Following this 300 ms period, the image was erased, and two targets were displayed at 5° above and below the image. Subjects reported their perceived direction of motion by making an upward or downward saccade. Following this decision, they received feedback regarding their

decision accuracy via an auditory tone: a 1000 Hz (beep) 30 ms long sound for a correct decision, and a 500 Hz (boop) 30 ms long sound for an incorrect decision. At the end of this period the decision targets were removed, and the center fixation point appeared at a random location near the center of the screen, in the bounding box defined above.

3.2.3. Modulating cost of error

During the adaptation phase of the main experiments, the viewing period was set to be up to 300 ms from primary saccade offset, but in practice at the beginning of learning it was around 150 ms due to reaction time and duration of the corrective saccade. Thus, by adapting the primary saccade (reducing the size of the corrective saccade), subjects would have more time to view the image, increasing the accuracy of perceiving the direction of motion. To vary the cost of error, trials consisted of two types of stimuli. For targets on one side of the screen, coherence of the random dots was low (65%-75%), imposing a large cost on the error: the corrective saccade took precious time away from viewing the moving dots. For targets on the other side of the screen, coherence of the random dots was high (95%-100%). Here the error cost was small: the time consumed by the corrective saccade was relatively inconsequential for the ability to perceive motion of the dots.

In a control experiment (Control-1, described below), subjects received 300 ms to view the random dot image irrespective of the level of the adaptation. This served to remove the cost of error for the low and high coherence stimuli.

3.2.4. Main experiments

The logic of main experiments (1, 2, and 3) is illustrated in Figure 3.1E. N=60 subjects participated in main experiments (20 subjects in each). Each experiment started with 50 familiarization trials (no perturbations). During these trials the images appeared at the primary target location at various coherence levels to familiarize the participants with the saccadic task and motion discrimination paradigm. The collected data during the familiarization period was excluded from analysis.

After the familiarization block, the baseline block commenced. The baseline consisted of 100 trials and ended with a 30 sec set break. In this block, subjects experienced 50 low coherence trials on one side of the screen and 50 high coherence trials on the other side. The coherence side was counter-balanced between subjects. Since each subject experienced both type of stimuli (low vs high coherence), we used a within subject comparison for all statistical analysis.

Next, subjects experience 550 gain-down perturbation trials (trials 101-650), during which the random dots image was displayed 5° away from the primary target toward the center of the screen. The consistent experience of this perturbation gradually resulted in adaptation of the primary saccades. We asked how does cost of error modulated the rate of adaptation.

All experiments included an error-clamp period. In these trials, the perturbation was removed, and the image was centered at the end position of the primary saccade.

3.2.5. Control-1

This experiment (n=20 subjects) removed cost of error but maintained task difficulty and reward as factors that could influence the rate of learning. In contrast to main experiments in which the time to view the image was reduced because of the corrective saccade, in this control experiment (Figure 3.7B) the timer did not start until the eye landed around (4°×4°) the random dot image (3°×3°). This made it so that the time spent correcting for error did not compete with the time needed to view the random dot motion, thus equalizing the cost of error for the low and high coherence images.

3.2.6. Control-2

This experiment (Figure 3.8A, n=18 subjects) removed cost of error as well as task difficulty but maintained implicit reward as a factor that could influence the rate of learning. The primary target was always the image of a noise patch (3°×3°), presented at ±15° with respect to central fixation. A green dot always appeared at the center of every image. Upon initiation of the primary saccade, the primary target was erased and replaced by another image at 5° closer to central fixation. When the primary target was to one side, the replacement image had 50% probability of being a face, thus resulting in a positive reward prediction error (Sedaghat-Nejad, Herzfeld, & Shadmehr, 2019). For the other side the replacement image was always a noise patch, thus resulting in a zero reward prediction error. The facial images were gathered from the Internet (500 total images) and were modified in a way that the center of the two eyes was located at the center of the image. The noise images were constructed by shuffling the pixels of each face image (500×500 pixels). This ensured that the luminance and color content of the two categories were identical.

3.2.7. Experiment-4 and Control-3

Our main experiments (1, 2, and 3) introduced a high error cost for one direction of movement, and a low error cost for another direction. In a more general setting, one would wish to enhance learning rates

in all directions. To test whether this could be achieved, in Experiment-4 (n=16 subjects, Figure 3.9A) we tested the same subjects on two sessions, separated by one week.

We simplified the decision-making task: rather than judging the direction of motion of random dots, subjects had a limited amount of time (250 ms from primary saccade offset) to view a cue image. In some trials, this cue image provided them with information necessary for decision-making. In other trials the cue image was irrelevant for decision-making.

In Exp-4 and Control-3, participants sat in front of an LCD monitor (32-inch, 1920×1080 pixels, light gray background, refresh rate 60 Hz) placed at a distance of 40 cm while we measured their eye position at 1000 Hz (Eyelink 1000+, SR Research). Each trial began with presentation of a crosshair shaped fixation stimulus (1.5°×1.5°) around the center of the screen. During the fixation, a target (the image that represented the number "0", 1.5°×1.5°) appeared at ±15° for 200 ms, but the subjects were not allowed to saccade to it. Rather, they waited for the removal of the fixation stimulus (an additional 200-500 ms), and then made a saccade to the remembered location (memory-guided saccade). At the onset of this primary saccade the cue target (1.5°×1.5°) was placed 5° away from the previewed target. The cue target contained either a stimulus that carried a large cost, or no cost. For example, if the cue image contained 3 black dots, then the correct decision was a saccade to the target labeled "3" (1.5°×1.5°). If the cue image contained random noise, then the image was irrelevant for decision making—the correct decision was a saccade to any target.

In the main group (Figure 3.9B), during one session 90% of the cues on both sides were images that were important for decision-making, and thus carried a large error cost. During another session, 90% of the cues on both sides were irrelevant for decision-making, and thus had zero error cost. The sessions were counter-balanced across the participants.

In the control group (n=14 subjects), we sought to reproduce the results of the main experiments in this simplified decision-making paradigm. Thus, we tested the subjects during a single session, with movements to one side encountering images that had a high error cost, and movements to the other side encountering zero cost (Figure 3.9B, control group).

3.2.8. Data analysis

Eye position data were acquired using an EyeLink 1000+ system (SR Research) at 1000 Hz. During the online data acquisition, we used a second order Savitzky–Golay filter with seven datapoints to estimate the eye velocity. The onset of the primary saccades was detected when the eye velocity increased above

20 °/s, or the eye position left the area (4°×4°) around the start target (whichever happened first). The source code for online estimation of eye velocity is available for download from the project's OSF repository (https://doi.org/10.17605/osf.io/H24J8, ESN_Eye_Filt.c). Primary saccade offset was detected when the eye velocity fell below 75 °/s and the eye position was inside a rectangle (9°×4°) which contained the area around the primary and secondary target locations and the region in between.

Eye position data in offline analysis were filtered with a second-order Butterworth low-pass filter with cutoff frequency of 100 Hz. Eye velocity data were calculated as the derivative of the filtered position data. Saccades were identified with a speed magnitude threshold of 20 °/s, and minimum hold time of 10 ms at saccade end (i.e., velocity magnitude could not exceed the cutoff for a minimum 10 ms after endpoint). Corrective saccade onset and offset were detected identically to the primary saccades, using 20 °/s threshold on velocity magnitude. We measured change in primary saccade amplitude with respect to the average saccade amplitude in the first block (baseline block) in each condition.

Viewing period of the random dot image was measured based on the amount of time that the eye position was inside of an imaginary box of size 4°×4° centered on the image. In the case of the Control-1 experiment, the moment when the eye position entered this region started the 300 ms timer.

Decision accuracy was measured based on the number of correct decision responses divided by the total number of trials for each condition (high vs. low coherence).

Pupil area was measured by EyeLink 1000+ system (SR Reseach) and was reported in the system's arbitrary pixel coordinate system. We blanked this data during eyeblink events to account for divergence in eye tracking. To combine and compare the pupil data across participants and experiments, we measured the percentage change for each participant by dividing the pupil data by the average pupil area over the entire recording for that subject. To control for differences in visual stimulus properties, we computed the average normalized pupil area during 200 ms window of time when participants were fixating on the start target (0.5°×0.5° green dot). Next, we measured the change in normalized pupil area from one trial to the next to quantify how the conditions of each trial affect pupil dilation (Figure 3.10B).

Statistical analyses were performed using SPSS and Linear Mixed Models. We used stimulus type (2 levels, categorical) and bins of trials (multiple levels, hierarchical) and tested within-subject effect of those independent variables. To test between-subject effect of experiment type, we used Linear Mixed Models with bins of trials (multiple levels, hierarchical, within-subject variable) and experiment type (2 levels, categorical, between-subject variable) without repeated measurements as independent

variables. Statistical analysis were performed under the assumption of First-order Autoregressive (AR(1)) to model the covariance matrix. We used intercept, stimulus/experiment type, bins of trials, and trial by stimulus/experiment type interaction as fixed effects and the intercept as a random effect.

3.2.9. State-space model of learning

After the experience of a movement error, humans and other animals change their behavior on future trials. In the absence of error, adapted behavior decays over time. Here we used a state-space model (Albert & Shadmehr, 2018) to capture this process of error-based learning. Here, the internal state of an individual x, changes from trials n to n+1 due to learning and forgetting.

$$x^{(n+1)} = a x^{(n)} + b^{(n)} e^{(n)} + \varepsilon_x^{(n)}$$
Equation 3.1

Forgetting is controlled by the trial-to-trial retention a. The rate of learning is controlled by the error sensitivity b. Learning and forgetting are stochastic processes affected by internal state noise ε_x: a normal random variable with zero-mean and standard deviation of σ_x.

While we cannot directly measure the internal state of an individual, we can measure their movements. The internal state x leads to a movement y according to:

$$y^{(n)} = x^{(n)} + \varepsilon_y^{(n)}$$
Equation 3.2

The desired movement is affected by execution noise, represented by ε_y: a normal random variable with zero-mean and standard deviation of σ_y. To complete the state-space model described by Equation 3.1 and Equation 3.2, we must operationalize the value of an error, e. In sensorimotor adaptation, movement errors are determined both by motor output of the participant (y) and the size of the external perturbation (r):

$$e^{(n)} = r^{(n)} - y^{(n)}$$
Equation 3.3

We used Equation 3.1, Equation 3.2, and Equation 3.3 to estimate the trial-to-trial retention a and error sensitivity b during each experiment design.

We used permutation (10,000 iterations) of the population data to estimate the parameters of the single-state and two-states models. We formed population data by randomly sampling (with replacement) the subjects and then computed the average adaptation curve for the population. In each iteration, for each subject, we stacked all the data for a given stimulus type together and as a result formed two time series for each subject, one for each stimulus type. Then, since the number of trials

were balanced over subjects and stimuli, we combined the data for a given stimulus and computed the average population data for that stimulus. We fitted single-state and two-states models of the learning to the population data using Least Mean Square (LMS) optimization method.

After computing the distribution of each model parameter, we next computed the distribution of within-population difference between the large and small cost conditions. To test whether there was a significant effect of cost, we used the within-population difference and integrated the region from zero (no difference) to minus infinity, resulting in a p-value (Figure 3.5C and Figure 3.6D). It should be mentioned that the p-value computed here was a one-tailed p-value and should be compared to 0.025 for the two-tailed assumption.

3.3. Results

Subjects made center-out horizontal saccades to a visual target (a green dot, 0.5°×0.5°) at ±15°. At the conclusion of their primary saccade, they viewed an image that contained random dot motion (Figure 3.1A). Their objective was to detect the direction of motion of the dots, which was either upward or downward, and was reported by making a vertical saccade. After this vertical saccade, feedback was provided regarding decision accuracy.

In the baseline block, the random dot image was centered at the target (Figure 3.1B). However, during the adaptation block the image was centered 5° away from the target, resulting in a gain-down form of adaptation (Figure 3.1C). Thus, during adaptation the subjects made a saccade to the target, and then followed that with a corrective saccade to a location near the center of the image (Figure 3.2). Importantly, the movement error and the resulting corrective saccade carried a cost because the subject had only 300 ms from the end of their primary saccade to view the image. Thus, if the subject learned from movement error and adapted their primary saccade, the corrective saccade would consume less time, allowing them to view the random dots for a longer period, and therefore arrive at a more accurate decision.

In order to modulate the landscape of the loss function, we varied the coherence of the random dots. As the subjects viewed the random dots, their brain accumulated evidence for each possibility (upward or downward motion). Evidence accumulation (Figure 3.1D) is roughly the temporal integral of the instantaneous difference between the number of dots that move upward vs. downward. This means that evidence grows faster when the motion is more coherent (more of the dots move in a single direction). Because time is lost for correcting for error (reaction time plus movement duration of the

corrective saccades), for the low coherence image this loss will impose a large cost and produce a significant reduction in decision accuracy. In contrast, for the high coherence image, the same loss will have little or no effect on decision accuracy (Pilly & Seitz, 2009). Thus, by varying motion coherence, we varied the cost of error, which we hypothesized would change learning rates.

In summary, when the target was presented to one side of the screen, coherence of the image was low, making the corrective saccade costly because it took precious time away from viewing the image. However, when the target was presented to the other side of the screen, coherence was high, making the corrective saccade less costly. Critically, the probability of success (reward) was greater for the stimulus that had high coherence, and task difficulty (attention) was greater for the stimulus that had low coherence. Thus, we performed a series of experiments (Figure 3.1E) to disentangle the effects of cost of error, task difficulty, and reward.
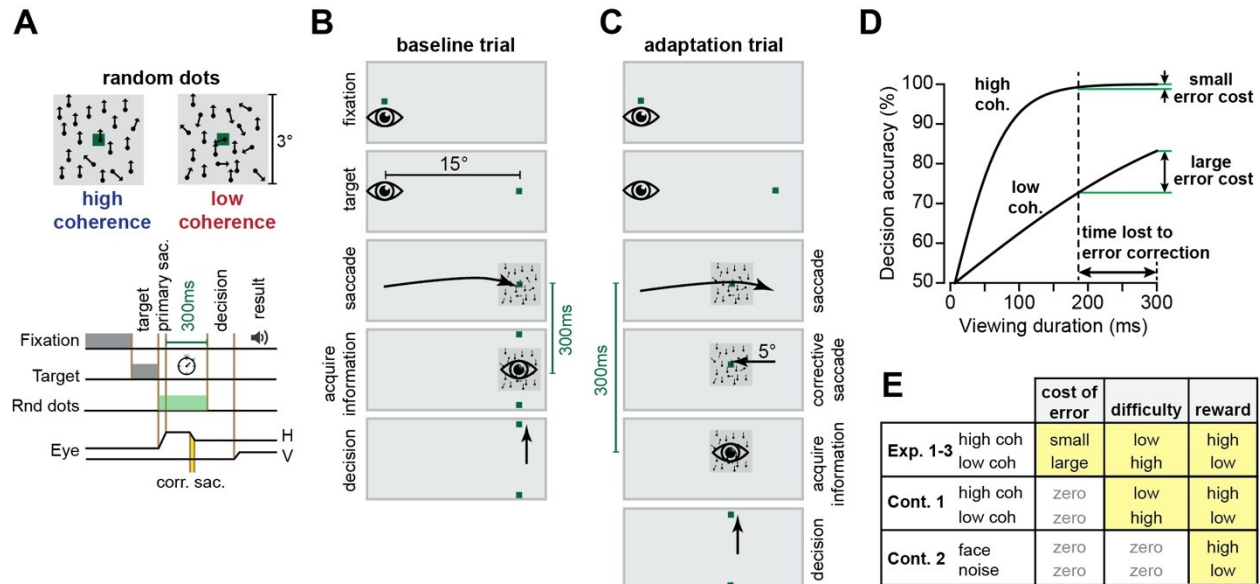
Figure 3.1. Experiment design.
A. Trials began with a random fixation interval (250-750 ms) after which a primary target was placed at ±15° along the horizontal axis. After conclusion of the primary saccade, subjects were presented with a random dot motion image with high or low coherence. The objective was to detect the direction of motion, which was either upward or downward, and was reported by making a vertical saccade, after which we provided feedback regarding decision accuracy. Subjects were limited to only 300 ms from the end of their primary saccade to view the image. Thus, the time consumed by the corrective saccade (reaction time and saccade duration) reduced the time available to place the image on the fovea. **B.** During the baseline block of trials the random dot image was presented at the primary target location. **C.** During adaptation trials the image was placed at a location 5° away from the primary target, producing a movement error that was followed by a corrective saccade (gain-down). **D.** Hypothetical landscape of the error cost. The period spent correcting for error carried a cost that depended on stimulus coherence. **E.** We sought to explore the effects of three factors on learning: cost of error, task difficulty, and reward.
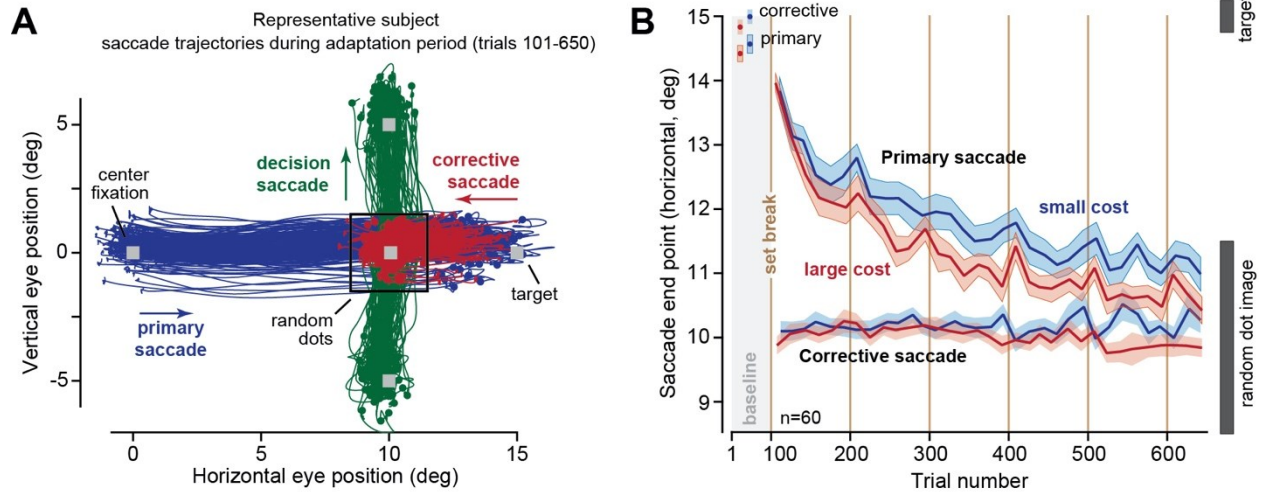
Figure 3.2. Saccade trajectories and end positions for a representative subject and for the main groups. **A.** Data for a representative subject during the adaptation trials for the target on the right. Primary saccade (blue) is followed by a corrective saccade (red), and then a decision saccade (green). **B.** End position of the primary and corrective saccades for all subjects in the main experiments (n=60). In the baseline condition, the image was centered at the target, located at 15° along the horizontal axis. Primary saccades were slightly hypometric, which were then followed by a corrective saccade that brought gaze to center of the image. In the adaptation trials (101-650), the image was centered 5° away from the target. Following the primary saccade, a corrective saccade brought gaze to near the center of the image. This corrective movement carried a large cost for some images, and a small cost for other images. The rate of adaptation in the primary saccade was greater when the cost of error was large. Bin size in B is 8 trials. Error bars are SEM.

3.3.1. Cost of error increased both the rate and the asymptote of adaptation

In Experiment-1 (n=20 subjects, Figure 3.3A), in adaptation trials one target was always associated with large error cost (low image coherence), and another target that was always associated with small error cost (high image coherence). In baseline trials, as well as during adaptation, the probability of a correct decision was much higher for the small cost target (Figure 3.3F, Linear Mixed Models, trials 401-650, within-subject effect of cost, $F_{(1,72.4)}=6.14$, p=0.016). Yet, the subjects learned more from errors that carried a large cost (Figure 3.3B), as indicated by the fact that adaptation rate was faster for the low coherence image (Linear Mixed Models on amplitude change, trials 51-400, within-subject effect of trial, $F_{(1,279.3)}=840.9$, $p<10^{-4}$, and trial by cost interaction, $F_{(1,114.8)}=4.87$, p=0.029).

Following a block of adaptation, we imposed a block of error-clamp trials that eliminated movement error. As expected, without errors to sustain adaptation, saccade amplitude returned toward baseline (Figure 3.3B, Linear Mixed Models, trials 651-800, within-subject effect of trial, $F_{(1,169.1)}=22.7$, $p<10^{-4}$, no trial by cost interaction, $F_{(1,78.9)}=0.022$, p=0.882). Following the error-clamp block, further training brought performance toward a plateau (Figure 3.3B, trials 951-1250). However, adaptation

remained higher for the side with the larger error cost (Linear Mixed Models, trials 951-1250, within-subject effect of cost, $F_{(1,130.8)}$=5.47, p=0.021).
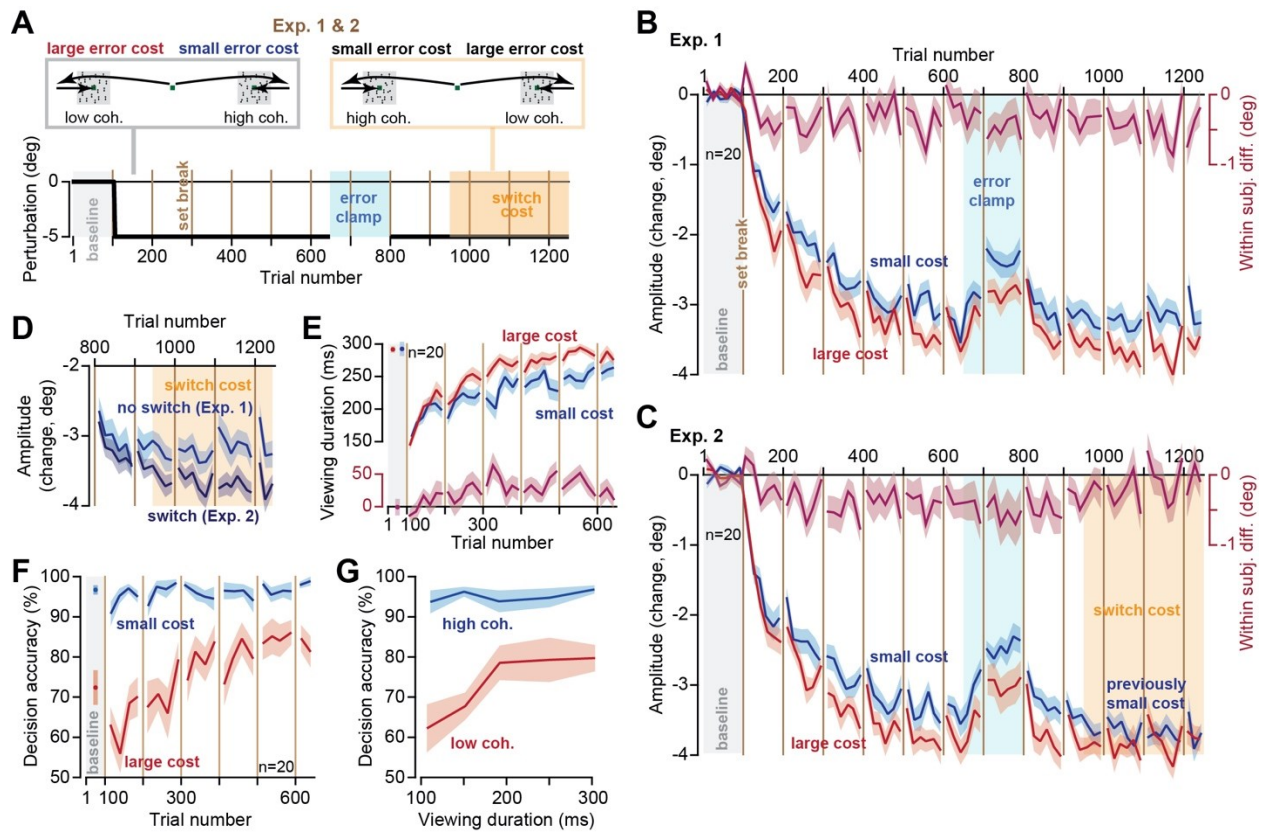
Figure 3.3. Cost of error modulated the rate and the asymptote of learning.
A. Subjects experienced 550 adaptation trials followed by 150 error clamp trials. In Experiment-1, this was followed by another 450 adaptations trials (trials 801-1250) which established the asymptote of performance. In Experiment-2, after an initial 150 adaptation trials (trials 801-950), at trial 951 without warning the side that previously displayed large cost stimuli (low coherence) switched to displaying small cost stimuli (high coherence). **B & C.** Amplitude of primary saccades with respect to baseline. Large error cost increased both the rate and the asymptote of adaptation. **D.** Increasing the error cost rescued adaptation levels. In Experiment-1, the side that had small cost retained that cost. In Experiment-2, the side that had small cost was switched to a large cost. **E.** The time spent viewing the random dot image increased for both the small and the large cost stimuli, but the increase was greater for the large cost stimulus. Faster adaptation provided more viewing time to perform motion discrimination. Data from Experiment-1. **F.** Decision accuracy for small and large cost stimuli over the course of the experiment. Reward rate was higher for the high coherence (small error cost) image. Data from Experiment-1. **G.** A change in viewing period produced large changes in decision accuracy for the low coherence image, but it had little or no consequence for the high coherence stimulus. Compare this error landscape to the hypothetical one in Figure 3.1D. Data from Experiment-1. Bin size in B-E is 8 trials, in F is 12 trials, and in G is 50 ms. Error bars are SEM.

Thus, the rate of adaptation, as well as the asymptote of adaptation, was greater toward the target that carried a large cost of error, not the target that carried a greater probability of success.

3.3.2. Increasing the cost of error rescued low adaptation

If error cost is a causal mechanism that modulates learning from error, then a change in error cost should produce a change in adaptation. Because Experiment-1 had established that the asymptote of adaptation was greater for the large error cost stimulus, we checked whether we could rescue adaptation by increasing its cost.

In Experiment-2, subjects (n=20) began with stimuli that were identical to Experiment-1: large error cost to one side, small error cost to the other. However, at trial 951 (Figure 3.3A, switch cost), without warning the side that previously displayed large cost images (low coherence) switched to displaying small cost images (high coherence). Similarly, the side that previously displayed small cost images switched to displaying large cost images.

During the initial phase of the experiment (trials 101-650), adaptation rate was faster toward the side that carried a large error cost (Figure 3.3C, Linear Mixed Models, trial by cost interaction, $F_{(1,105.7)}$=8.00, p=0.006), thus confirming the findings of Experiment-1. However, the switch from small to large cost produced convergence of saccade amplitudes for the two sides (Figure 3.3C, Linear Mixed Models, trials 801-1250, trial by cost interaction, $F_{(1,234.9)}$=9.05, p=0.003). We compared saccade amplitude in Experiment-1 when there was no switch in cost (during trials 801-1250), with the condition in which the cost switched (Experiment-2). The switch in cost rescued a zero slope learning curve to one that exhibited further learning (Figure 3.3D, Linear Mixed Models, trial (within-subject) by switch (between-subject) interaction, $F_{(1,148.0)}$=7.14, p=0.008).

It is noteworthy that adaptation rate was greater for the large cost stimulus, despite the fact that the stimulus on the opposite side was more rewarding (Figure 3.3F, Linear Mixed Models on probability of success, trials 101-650, Experiment-1, within-subject effect of cost, $F_{(1,196.7)}$=165.7, $p<10^{-4}$, as well as a trial by cost interaction, $F_{(1,195.1)}$=27.2, $p<10^{-4}$). The consequences of greater reward for the small cost stimulus was readily present in the reaction time of the primary saccades: as in many previous experiments (Manohar et al., 2015; Milstein & Dorris, 2007; Sedaghat-Nejad, Herzfeld, & Shadmehr, 2019; Shadmehr & Ahmed, 2020; Takikawa et al., 2002; Yoon et al., 2018, 2020), saccades toward the more rewarding stimulus exhibited a shorter reaction time (Figure 3.4, Linear Mixed Models, trials 101-650, Experiment-1, within-subject effect of trial by cost interaction, $F_{(1,339.6)}$=5.32, p=0.022). That is,

greater reward was associated with greater vigor (earlier reaction time), but not better adaptation. Rather, adaptation rate was higher for the stimulus that carried a greater cost.
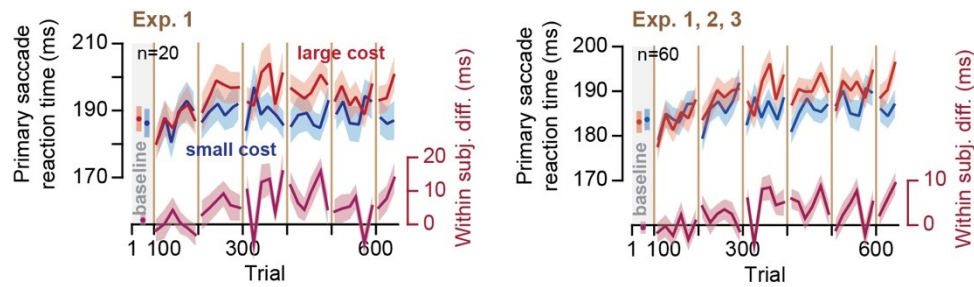
Figure 3.4. Reaction times of primary saccades
  Reaction times of primary saccades were faster toward the stimuli that promised greater reward, not the stimuli that carried a greater error cost. Note that the reaction times for the two stimuli start at approximately the same value but separate as the experiment continues. Bin size is 8 trials. Error bars are SEM.

We assumed that the time spent correcting for error carried a cost that depended on stimulus coherence (Figure 3.1D). To check the validity of our assumption, we quantified the relationship between decision accuracy (Figure 3.3F) and viewing time (Figure 3.3E) for each stimulus. For the high coherence stimulus, a change in the viewing period produced little or no change in decision accuracy (Figure 3.3G). On the other hand, for the low coherence stimulus a change in viewing period produced large changes in decision accuracy (Figure 3.3G, interaction of viewing period by decision accuracy, $F_{(1,69.7)}$=5.87, p=0.018). This confirmed that the time spent correcting for the movement error carried little or no cost for the high coherence stimulus (small cost), whereas the same expenditure was quite costly for the low coherence stimulus (large cost).

In summary, adaptation rate was greater toward the stimulus that carried a greater error cost, not the stimulus that was more rewarding. When the error cost increased (switch cost), so did the asymptote of performance, suggesting a causal relationship between the cost of error and adaptation.

3.3.3. Cost of error increased learning from error, not retention

The fact that cost of error modulated the rate of adaptation as well as the asymptote of performance raised the question of whether this cost affected sensitivity to error, trial-to-trial retention, or both. To consider these possibilities, in Experiment-3 we implemented a spontaneous recovery paradigm and then analyzed the results using a state-space model of adaptation (Albert & Shadmehr, 2018).

In Experiment-3 (Figure 3.5A), subjects (n=20) began with stimuli that were identical to Experiment-1 and 2: large error cost to one side, small error cost to the other. We again observed that

adaptation was faster toward the stimulus with large error cost (Figure 3.5B, Linear Mixed Models, trials 51-400, trial by cost interaction, $F_{(1, 105.3)}$=4.64, p=0.034). After the initial adaptation period we reversed the direction of movement errors (trials 651 to 750) to induce "extinction", resulting in a sharp change in saccade amplitude toward baseline. Following error reversal subjects experienced a long period of error clamp trials (trials 751-1200). As expected (Ethier et al., 2008a), during the error clamp period saccade amplitude exhibited spontaneous recovery toward the adapted state (Figure 3.5B, trials 751-850, Linear Mixed Models, within-subject effect of trial, $F_{(1,151.6)}$=6.02, p=0.015).
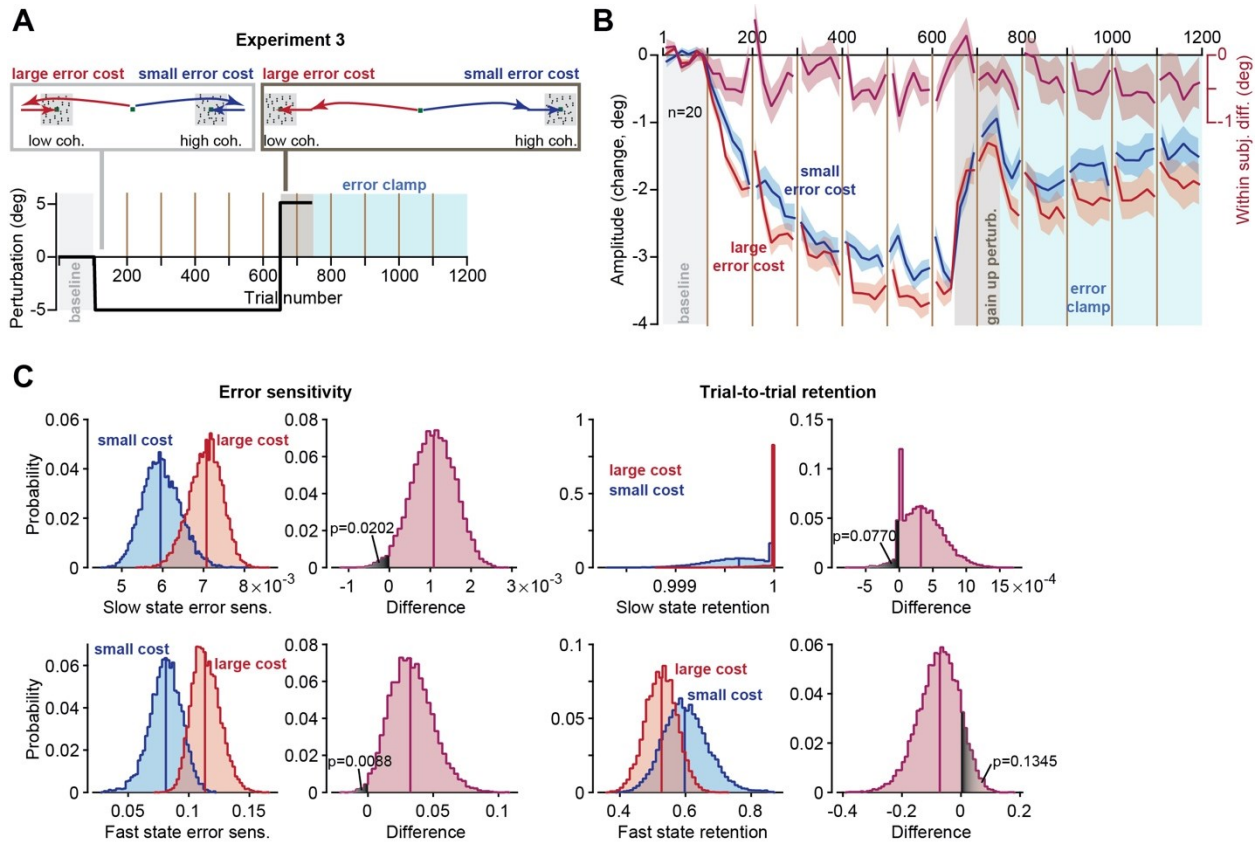
Figure 3.5. Cost of error affected sensitivity to error but not retention.
A. Experiment design for the spontaneous recovery paradigm. **B.** Amplitude of primary saccades with respect to baseline. **C.** A two-state model was fit to the data and error sensitivity and retention were estimated for the fast and slow states using a permutation procedure. The plots show the resulting distribution of parameter values and the within-population differences in parameter values due to change in cost. Bin size in B is 8 trials. Error bars are SEM.

We next applied a state space model and estimated error sensitivity and trial-to-trial retention. When the cost of error was large, error sensitivity was elevated for both the slow (Figure 3.5C, permutation n=10,000, p=0.0202, one-tailed) and the fast state (Figure 3.5C, p=0.0088, one-tailed). In

contrast, cost of error did not appear to affect trial-to-trial retention (Figure 3.5C, slow state: permutation n=10,000, p=0.0770, one-tailed; fast state: p=0.1345, one-tailed).

To check the robustness of this result, we reconsidered the data during the adaptation period in main experiments (Figure 3.6A), with the caveat being that because this data set did not contain a spontaneous recovery period, we did not have sufficient power to consider a two-state model, and thus fitted a single-state set of equations. We again found that error sensitivity was larger for the large cost target (Figure 3.6D, permutation n=10,000, p=0.0167, one-tailed), with no significant effect on the trial-to-trial retention (Figure 3.6D, permutation n=10,000, p=0.4671, one-tailed).
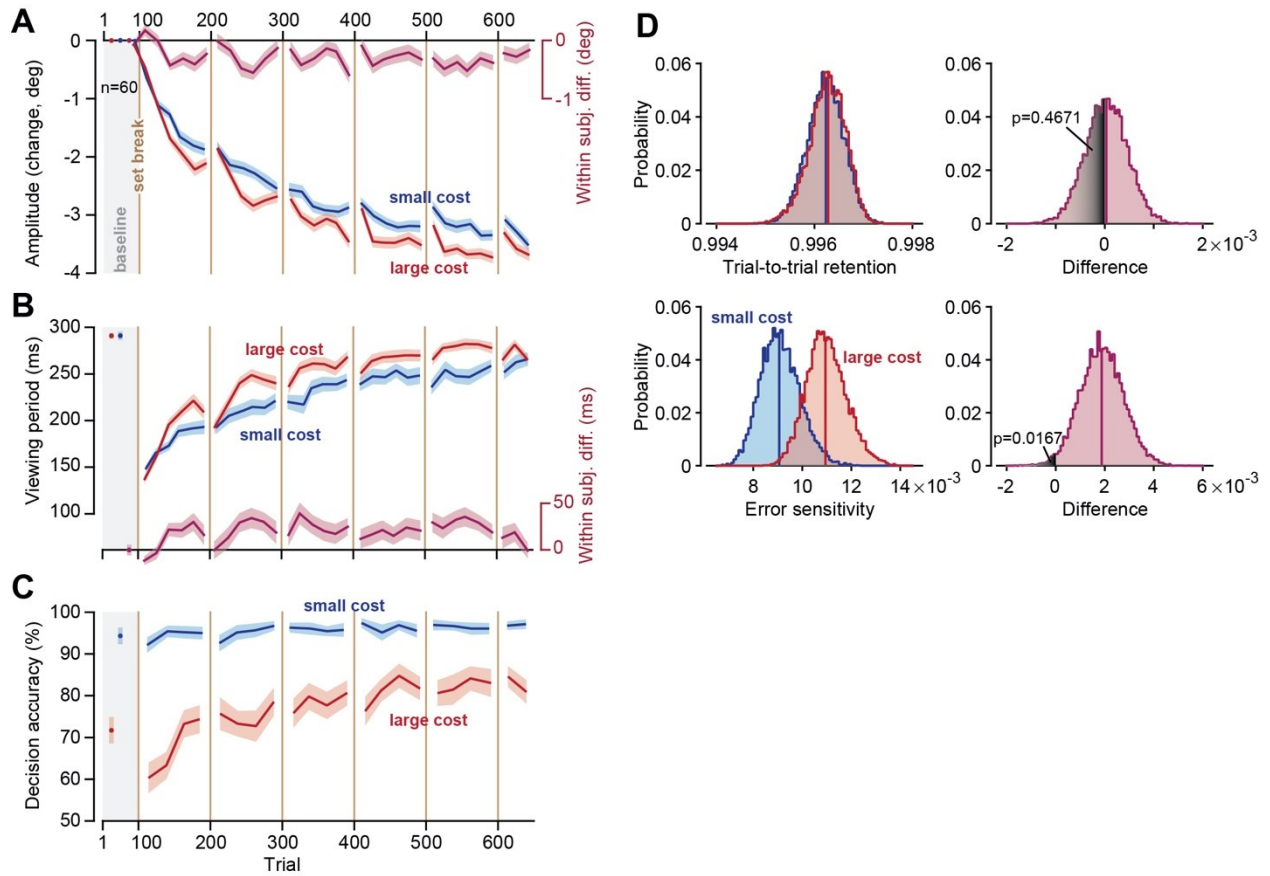
Figure 3.6. The faster adaptation of the primary saccade provided more viewing time to perform motion discrimination.

A. Change in primary saccade amplitude from baseline. **B.** Image viewing period over the course of adaptation. Faster adaptation provided more viewing time to perform motion discrimination. **C.** Decision accuracy over the course of the adaptation. The impact of the increased viewing period on decision accuracy was much greater for large cost stimulus as compared to the small cost stimulus. **D.** A single-state model was fit to the data and error sensitivity and retention were estimated using a permutation procedure. The plots show the resulting distribution of parameter values and the within-population differences in parameter values due to change in cost. Data from main experiments (1, 2, and 3). Bin size in A & B is 8 trials and in C is 12 trials. Error bars are SEM.

In summary, cost of error increased sensitivity to error, not retention.

3.3.4. Faster adaptation provided more time for decision-making

We had assumed that adaptation would afford subjects more time to view the image, and thus help them make more accurate decisions. To check for this, we combined the data for the three experiments (n=60 subjects). As expected, large error cost coincided with a faster rate and a greater extent of adaptation (Figure 3.6A, Linear Mixed Models, trials 51-400, within-subject effect of trial, $F_{(1,805.3)}$=2375.8, $p<10^{-4}$, trial by cost interaction, $F_{(1,325.9)}$=11.7, p=0.001, and trials 401-650, within-subject effect of cost, $F_{(1,314.4)}$=7.31, p=0.007). This increased rate of adaptation for the large cost stimulus provided more time to view the moving dots (Figure 3.6B, Linear Mixed Models, trials 101-400, within-subject effect of trial, $F_{(1,777.6)}$=474.6, $p<10^{-4}$, trial by cost interaction, $F_{(1,315.4)}$=8.40, p=0.008, and trials 401-650, within-subject effect of cost, $F_{(1,286.6)}$=5.41, p=0.021). Finally, as saccades adapted and the viewing period increased, so did decision accuracy (Figure 3.6C). The impact of increased viewing period on decision accuracy was much greater when time was more valuable (i.e., low coherence stimulus, Figure 3.6C, trials 101-400, trial by cost interaction, $F_{(1,218.8)}$=15.7, $p<10^{-4}$).

3.3.5. Control experiments: eliminating the error cost equalized rates of adaptation

There are potential confounds in our interpretation. First, the task was harder for the low coherence stimulus, making it possible that learning rate was not driven by the cost of error, but rather the greater attention or cognitive load required for the more difficult task. Second, the probability of reward (success) was lower for the low coherence stimulus. It is conceivable that the increased reward impaired adaptation rates. Thus, we performed a series of additional experiments (Figure 3.1E).

To test for the effect of task difficulty, we performed an experiment (Control-1) in which the cost of error was equal for the two stimuli, but task difficulty was greater for one of them. Subjects (n=20) made a primary saccade to targets at ±15°, and again were presented with an image that was centered 5° away (Figure 3.7A). However, unlike the main experiments, here the subjects were provided with 300 ms to view the random dot image regardless of their primary saccade. That is, the time allowed to view the image did not start until eye position was within 2° of the center of the image (Figure 3.7B). With this subtle change we removed the cost associated with movement error: now the time expended during error correction did not affect the period available to view the image.

As before, decision accuracy was greater for the side that contained the high coherence image, confirming that on one side the task remained more difficult than the other (Figure 3.7D, Linear Mixed Models, trials 401-650, within-subject effect of coherence, $F_{(1,81.7)}=30.5$, $p<10^{-4}$). Indeed, decision accuracy at both start and end of adaptation was similar among Control-1 and the main experiments groups (Linear Mixed Models, no between-subject effect of experiment type, $F_{(2,294.0)}=2.26$, $p=0.106$, or trial by experiment type interaction, $F_{(2,294.0)}=2.04$, $p=0.133$). However, while saccade amplitude exhibited adaptation (Figure 3.7C, Linear Mixed Models, trials 51-400, within-subject effect of trial, $F_{(1,280.4)}=808.5$, $p<10^{-4}$), we found no significant within-subject difference between the low and high coherence stimuli (Figure 3.7C, Linear Mixed Models, trials 401-650, no within-subject effect of coherence, $F_{(1,94.5)}=0.650$, $p=0.422$, trials 51-400, no trial by coherence interaction, $F_{(1,116.3)}=2.55$, $p=0.113$). Furthermore, we found no significant within-subject difference in the asymptotic learning between the two types of stimuli (Figure 3.7C, Linear Mixed Models, trials 951-1250, no within-subject effect of coherence, $F_{(1,108.0)}=0.189$, $p=0.665$).
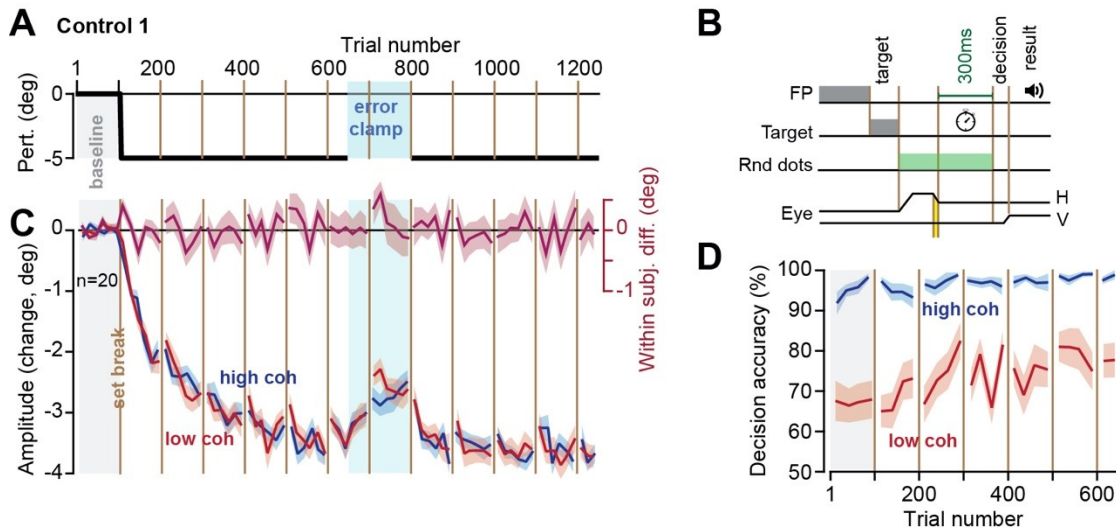
Figure 3.7. Control-1 experiment: eliminating the error cost equalized rates of adaptation.
A & B. Design was similar to main experiments, with one crucial difference; the time available for viewing the random dot image did not start until the eye landed around the image. Thus, the subjects had 300 ms to view the random dots regardless of their primary saccade amplitude. **C.** Primary saccade amplitude with respect to baseline. There were no significant within-subject differences in the rate or asymptote of adaptation. **D.** Decision accuracy (and thus success rate) remained high for the high coherence image. Bin size in C is 8 trials, in D is 12 trials. Error bars are SEM.

To check for the influence of reward independent of both the cost of error and task difficulty, we performed a second control experiment. In Control-2 (Figure 3.8), the primary target was always a noisy image. However, when this image appeared to one side, there was a 50% probability that

following the onset of the primary saccade it would be replaced with the image of a face, thus producing a positive reward prediction error (15). When it appeared on the other side, it remained the same noisy image. We found that both the primary saccade (Figure 3.8B, lower plot, Linear Mixed Models, trials 1-650, within-subject effect of reward prediction error, $F_{(1,308.9)}=4.36$, p=0.038) and the corrective saccade (Figure 3.8C, Linear Mixed Models, trials 101-650, within-subject effect of reward prediction error, $F_{(1,266.9)}=18.1$, $p<10^{-4}$) had shorter reaction times for the stimulus that was paired with positive reward prediction errors. However, the increased reward had no significant effect on adaptation rates (Figure 3.8B, upper plot, Linear Mixed Models, trials 51-400, no within-subject effect of trial by reward prediction error interaction, $F_{(1,85.6)}=1.13$, p=0.291, trials 401-650, no effect of reward prediction error, $F_{(1,106.9)}=0.866$, p=0.354).
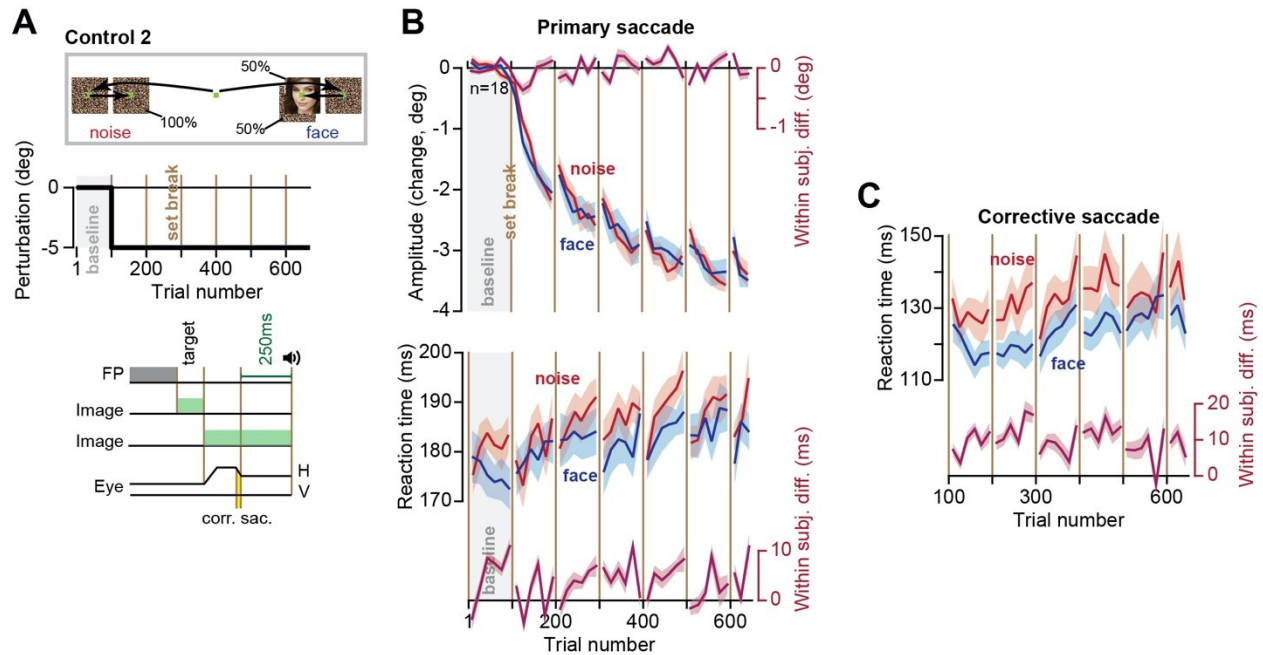
Figure 3.8. Control-2: learning rates were unaffected by increased reward rate.
   We modulated the reward rate via presentation of noise and face images, eliminating both the cost of error and task difficulty. **A.** Experiment design. The primary target was always a noisy image. On 50% of the trials to one side, this image was replaced by image of a face. **B.** Primary saccades had a shorter reaction time toward the side that had a higher probability of a face image, but the rate of adaptation was not affected. **C.** Corrective saccades had a lower reaction time toward the face image. Bin size is 8 trials. Error bars are SEM.

Another limitation of our main experiments was that while one direction of movement experienced a high error cost, another direction experienced a low cost. Can error cost be used to increase learning for all directions? To answer this question, we performed a multi-day experiment (Experiment-4, n=16 subjects, Figure 3.9A). During one session all targets were paired with a high error

cost, but on a different session (a week apart) this cost was removed. We found that the increased cost of error enhanced the learning rate (Figure 3.9C, main group, Linear Mixed Models, trials 41-320, within-subject effect of trial by cost interaction, $F_{(1,117.3)}$=5.90, p=0.017, trials 241-320, within-subject effect of cost, $F_{(1,44.8)}$=4.86, p=0.033). In addition, during a control experiment (Figure 3.9C, control group, n=14 subjects, Linear Mixed Models, trials 41-320, within-subject effect of trial by cost interaction, $F_{(1,137.5)}$=5.75, p=0.018, trials 241-320, within subject effect of cost, $F_{(1,52.4)}$=5.24, p=0.026) in which one direction of movement was paired with cost of error while the other direction was not, we again confirmed the findings of our main experiments.
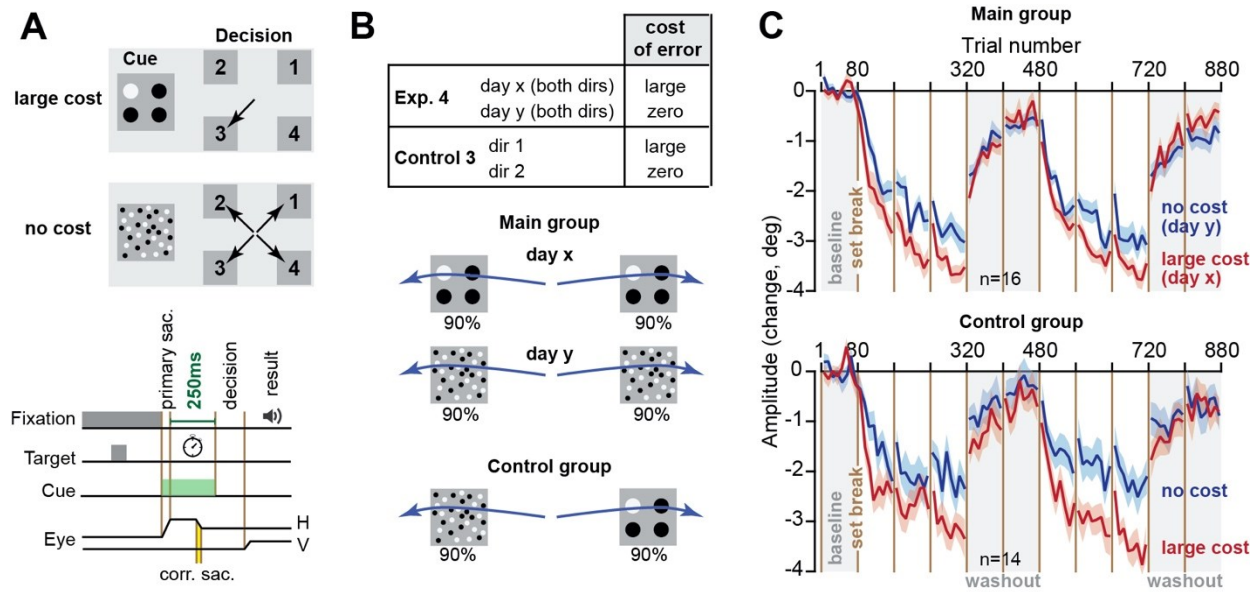
Figure 3.9. Increasing the error cost for all movement directions increased the learning rates.
  A. Experiment design. Subjects had a limited time to view the cue image. If the cue image contained large circles, then the correct decision was determined by the number of large black circles. If the cue image contained noise, then any decision was correct. **B.** In the main group, during one session (day x, high-cost day) there was 90% probability that the cue on both directions would have a large cost. During another session (day y, low-cost day, a week apart) there was only 10% probability that the cues would have a large cost. In the control group, cues in one direction had a 90% probability of having a high cost, while in the other direction this probability was 10%. **C.** Primary saccade amplitude with respect to baseline. Large cost consistently increased the learning rates. Bin size in C is 10 trials for main group and 5 trials for Control group. Error bars are SEM.

In summary, eliminating the cost of error while maintaining task difficulty eliminated the effects on the learning rates. Increasing the reward rate while keeping task difficulty and cost of error constant produced no changes in learning rates. Thus, among the three variables that we considered (cost of error, task difficulty, and reward rate), cost of error alone was a robust factor in modulating the learning rates.

3.3.6. Pupil dilation coincided with increased cost of error

What might be the neural mechanism that links cost of error with adaptation? To approach this question, we considered pupil dilation as a proxy for activation of the brainstem neuromodulatory system (Vazey et al., 2018). We measured pupil diameter as subjects fixated the center target and found that at the onset of each block the pupil was dilated, but then progressively constricted as the trials continued (Figure 3.10A). Following a set break, the pupil once again dilated. These patterns were present in the main experiments, as well as in Control-1 (Figure 3.10A, Linear Mixed Models, trials 101-600, within-subject effect of trial, main experiments: $F_{(1,1728.9)}=476.1$, $p<10^{-4}$, Control-1: $F_{(1,568.8)}=244.8$, $p<10^{-4}$). If we view pupil diameter as a proxy for arousal, then it appears that there was a general decline in arousal within each block of trials, followed by sharp recovery due to set breaks.

Next, we asked how the conditions of each trial affected pupil diameter. For each subject and each trial, we compared pupil size at center fixation (trial onset) to the fixation at the onset of the next trial before the target was displayed. This within trial response served as our proxy for how the neuromodulatory system reacted to the events that had transpired during that trial.

We found that in the baseline block, the trials that were more difficult (low coherence) produced pupil dilation, whereas trials that were easy (high coherence) produced pupil constriction (Figure 3.10B). The difference in the pupil response to the stimulus content of each trial was present in the baseline block of both the main group, and the control group (Figure 3.10B, Linear Mixed Models, trials 1-100, within-subject effect of coherence, main experiments: $F_{(1,330.5)}=55.8$, $p<10^{-4}$, and Control-1 experiment: $F_{(1,106.2)}=5.19$, $p=0.025$). Thus, as has been noted before (Kahneman & Beatty, 1966), when the trial included a more difficult decision-making process, requiring a greater cognitive load, it produced greater pupil dilation.

In the main experiments, as the adaptation blocks began the pupil continued to dilate in trials that were difficult and had large cost (Figure 3.10B, Linear Mixed Models, trials 101-650, within-subject effect of cost, $F_{(1,1256.6)}=111.1$, $p<10^{-4}$). In the Control-1 experiment the trials were still more difficult for the low coherence stimulus, but the error cost was equalized between the two stimuli. Interestingly, in the Control-1 experiment the pupil response to trial difficulty appeared to dissipate (Figure 3.10B, Control-1, Linear Mixed Models, trials 101-650, no within-subject effect of coherence, $F_{(1,507.8)}=1.22$, $p=0.269$). We were concerned that this difference in the two groups may have been because of the larger group size in the main experiments. However, the statistical pattern was also present in each of the main experiments (Linear Mixed Models, trials 101-650, within-subject effect of cost, Exp-1:

$F_{(1,424.5)}$=91.2, p<10^{-4}, Exp-2: $F_{(1,404.0)}$=19.1, p<10^{-4}, Exp-3: $F_{(1,435.7)}$=21.1, p<10^{-4}). In addition, comparing the within-subject difference in within trial pupil dilation between Experiment-1 and Control-1 showed an effect of experiment type (Linear Mixed Models, trials 101-650, between-subject effect of experiment type, $F_{(1,1316)}$=20.4, p<10^{-4}).
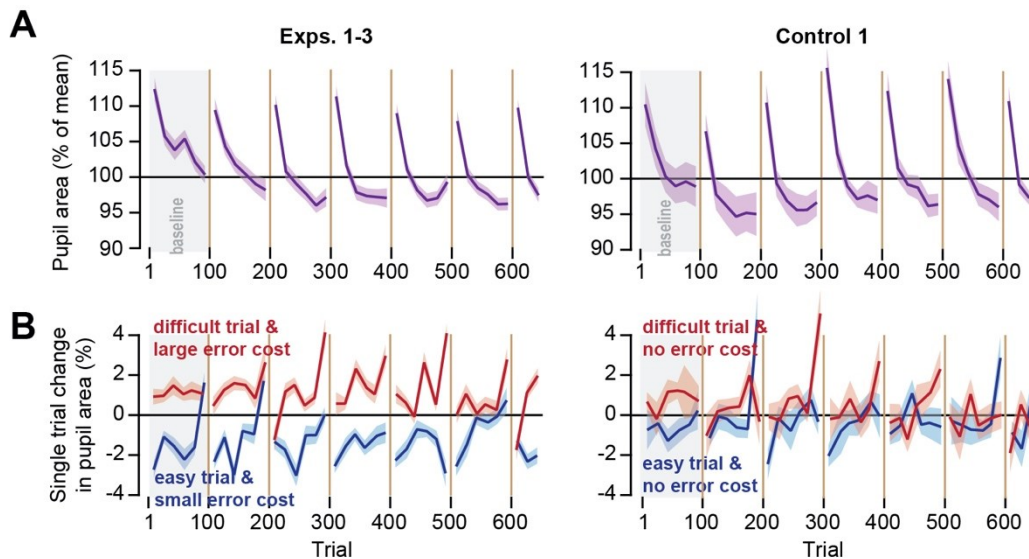
Figure 3.10. Pupil diameter changed in response to trial conditions.
A. Pupil area at the beginning of each trial, normalized to the overall mean. Pupil was most dilated when the block of trials started. It gradually constricted as the trials continued within the block, then recovered following a set break. **B.** Within trial change in pupil diameter as a function of trial conditions. The plots show the change in pupil diameter from fixation at start of each trial to onset of fixation at the start of the subsequent trial. Pupil area increased in trial that had increased difficulty, but error cost of that trial modulated this increase. Bin size is 8 trials. Error bars are SEM.

In summary, the pupil progressively constricted during each block of trials, suggesting a waning of attention, but then dilated following each set break, suggesting partial recovery. The pupil also responded to the conditions of each trial: in more difficult trials (low coherence) the pupil dilated. However, cost of error modulated this response: the within trial change in pupil diameter was greater when the trial was both difficult and incurred a large error cost.

3.4. Discussion

When movements produce an unexpected outcome, the nervous system often produces a reflexive response that corrects for error. This corrective movement consumes time and energy, providing an implicit loss that may affect learning. Here, we used saccade adaptation as a model of sensorimotor learning and explored whether imposing a cost on the time spent correcting for error could modulate learning from error.

To impose a cost on error, we combined decision-making with adaptation. In the resulting paradigm, it was costly to make a movement that ended with a large error because both visual acuity and the ability to detect differential motion drops rapidly with distance from the fovea (McKee & Nakayama, 1984). Thus, the erroneous saccade was followed by a corrective movement that placed the image on the fovea but took time away from the limited period available to acquire information and make an accurate decision. By modulating the value of time, we modulated the cost of error and found that an increased error cost robustly increased the rates of sensorimotor adaptation.

However, imposing a cost of error also changed other aspects of the task: its difficulty, and its reward rate. For example, errors that were more costly occurred in trials in which decision-making was more difficult, and thus the reward rate was lower. To find the effects of error cost independently of task difficulty and reward rate, we performed two control experiments. In Control-1, we maintained the disparity in difficulty and the rate of reward but removed the error cost. This eliminated the effects on the learning rates. In Control-2, we maintained the disparity in the rate of reward but removed task difficulty and error cost. This also eliminated the effects on the learning rates. Thus, error cost stood out as the critical factor that modulated learning rates.

Earlier studies have considered the effects of explicit inducements on retention of motor memories (Abe et al., 2011; Galea et al., 2015; Wächter et al., 2009). For example, Abe et al. (2011) showed that monetary reward was effective in improving retention of motor memories during a force-tracking task. Galea et al. (2015) also found that during a visuomotor rotation task presence of reward was associated with greater retention. In contrast, Steel et al. (2016) reported that during serial reaction time and force tracking tasks, neither reward nor punishment benefitted retention. We implemented a spontaneous recovery paradigm and assessed the effects of error cost on both error sensitivity and trial-to-trial retention. We found that while error cost increased error sensitivity, it did not have a significant effect on retention.

Meermeier et al. (2016) examined effects of implicit reward on saccade adaptation by using neutral (noise) or highly engaging images. They found that adaptation rates were not affected when the time allowed to view the images was unrestricted. Our Control-2 experiment confirmed this finding. However, Meermeier et al. (2016) observed that when they limited the time that the subjects had to view the highly engaging images, the result was an increase in the learning rates. If we view their results in the framework of cost of error, we see that while increasing reward had no effect on adaptation

rates, imposing a cost on the corrective saccade (limiting the time available to view the valuable image) increased the rate of adaptation.

There are of course other factors that influence how much the brain learns from error (Herzfeld, Pastor, et al., 2014; Leow et al., 2020; Marko et al., 2012; Wei & Körding, 2009). For example, Marko et al. (2012) and Hanajima et al. (2015) noted that error sensitivity was relatively high for small errors and low for large errors. Herzfeld et al. (2014) and Leow et al. (2016) showed that in environments where errors were likely to be consistent, subjects increased their error sensitivity. Albert et al. (2021) observed that large variability in the trial-by-trial sequence of errors tended to suppress learning from error. Conscious of these potential pitfalls, we kept the perturbation size consistent over the course of the experiment, and also controlled the statistics of the error that the subjects experienced as they made saccades. Despite this, the rate of learning was greater toward the stimulus that carried a greater error cost.

We found that increasing the cost of error rescued low adaptation, suggesting a potentially causal relationship between the cost of error and adaptation rates. Previous studies (Albert et al., 2021; Kim et al., 2019) have also found that modulating error sensitivity affected the asymptote of performance during motor learning.

In our saccade task, the pupil progressively constricted as the trials wore on within a block of trials, suggesting a decline in arousal (Mathôt, 2018), but then dilated following the set break at the start of the next block, suggesting a partial recovery. The resulting saw-tooth pattern in pupil diameter was reminiscent of behavioral changes during adaptation in many other experiments: rapid adaptation that follows set breaks, and gradual adaptation that ensues with progression of trials (Chen-Harris et al., 2008; Coltman et al., 2019; Ethier et al., 2008a; Xu-Wilson, Chen-Harris, et al., 2009).

Within each trial, during the baseline block the pupil dilated in response to stimuli that required greater mental effort (low coherence stimuli) and constricted in response to stimuli that required smaller effort (high coherence stimuli). During the adaptation block, when the stimuli carried an error cost the pupil continued to dilate in response to the high cost, greater mental effort stimuli. However, when the error cost was equalized in the Control-1 experiment, the dilation in response to stimuli that required greater mental effort waned. These results raise the possibility that pupil dilation is not only a correlate of attention and mental effort invested in the task, but also a correlate of learning from error.

The potential link between pupil diameter and learning from error is noteworthy because it may highlight one pathway with which the brain modulates learning. Changes in pupil size are due to a band

of muscles that surround the pupil, which in turn are controlled by motoneurons that reside in the Edinger-Westphal nucleus in the brainstem. Neurons in the intermediate layers of the superior colliculus project to this nucleus (May et al., 2016). As a result, weak micro-stimulation of the intermediate layers of the superior colliculus can produce a transient increase in pupil diameter that reaches its peak at around 300-500 ms (Joshi et al., 2016; C.-A. Wang et al., 2012). Notably, some superior colliculus neurons project to the contralateral inferior olive (Harting, 1977), which provide climbing fibers that carry error information to Purkinje cells of the cerebellum. For example, the climbing fiber carries information regarding the visual error following conclusion of a saccadic eye movement (Herzfeld et al., 2015; Sedaghat-Nejad, Herzfeld, Hage, et al., 2019; Shadmehr, 2020; Soetedjo et al., 2008), which in turn guides plasticity in Purkinje cells and affects trial-to-trial learning (Herzfeld et al., 2018). Notably, the amount that the cerebellum learns from error may be related to the state of the superior colliculus: in trials in which collicular neurons respond more strongly to the visual error, there is greater trial-to-trial learning (Kojima & Soetedjo, 2017b, 2018). Thus, on the one hand the superior colliculus contains the neural machinery to control pupil size, and on the other hand, it provides information to the cerebellum regarding saccade related visual errors.

Improving how we learn from our erroneous movements is a critical factor in applied settings such as rehabilitation (Quattrocchi et al., 2018). Previous work has generally focused on reward as a variable that may modulate learning. Here, we took advantage of the fact that movements that contain an error are often followed by corrective actions. By imposing a cost on this corrective action, we found a way to help the brain learn faster.

Chapter 4. Behavioral training of marmosets and electrophysiological recording from the cerebellum

The common marmoset (Callithrix Jacchus) is a promising new model for study of neurophysiological basis of behavior in primates. Like other primates, it relies on saccadic eye movements to monitor and explore its environment. Previous reports have demonstrated some success in training marmosets to produce goal-directed actions in the laboratory. However, the number of trials per session has been relatively small, thus limiting the utility of marmosets as a model for behavioral and neurophysiological studies. Here, we report the results of a series of new behavioral training and neurophysiological protocols aimed at increasing the number of trials per session while recording from the cerebellum. To improve the training efficacy, we designed a precisely calibrated food regulation regime that motivated the subjects to perform saccade tasks, resulting in about a thousand reward-driven trials on a daily basis. We then developed a multi-channel recording system that used imaging to target a desired region of the cerebellum, allowing for simultaneous isolation of multiple Purkinje cells in the vermis. In this report, we describe (1) the design and surgical implantation of a CT guided, subject specific head-post, (2) the design of a CT and MRI guided alignment tool for trajectory guidance of electrodes mounted on an absolute encoder microdrive, (3) development of a protocol for behavioral training of subjects, and (4) simultaneous recordings from pairs of Purkinje cells during a saccade task.

4.1. Introduction

The common marmoset (Callithrix jacchus) has gained attention for neurophysiological investigation in recent years because of its potential for transgenic manipulation (Kishi et al., 2014; Miller et al., 2016; Sasaki et al., 2009). Marmosets are small, New World primates with no known lethal zoonotic diseases that are transmittable to humans (Wakabayashi et al., 2018). Behaviorally, they share important attributes with us: they are social and live in family units, they are vocal and rely on species-specific sound production for their communication, and they are visual and use saccadic eye movements to explore their environment. Notably, for a primate they have a particularly short gestation period (~5 months) and regularly give birth to twins or triplets. Thus, they have a high breeding efficiency with potential for germline transmission of genetically modified models.

Indeed, marmosets research is benefiting from transgenic (Sasaki et al., 2009), gene-editing (Kishi et al., 2014), and optogenetic tools that target specific neurons (MacDougall et al., 2016). Thus, marmosets have the potential to become a model system for investigating cognitive and social behaviors (Mustoe et al., 2015; Takahashi et al., 2017), auditory (X. Wang, 2018) and visual perception (Solomon & Rosa, 2014), as well as neural control of vocalization (Eliades & Miller, 2017; Eliades & Wang, 2013; Roy et al., 2011) and eye movements (Johnston et al., 2018; Mitchell et al., 2014).

Many aspects of social behavior, vocalization, and eye movements depend on the integrity of the cerebellum. For example, children who suffer from autism spectrum disorder (ASD), a developmental disorder that leads to impairments in social and communication skills, exhibit anatomic abnormalities in their cerebellum, including reduced number of Purkinje cells (Whitney et al., 2008), particularly along the vermis (Courchesne et al., 2001; Hashimoto et al., 1995; Murakami et al., 1989; Scott et al., 2009). Notably, in children with ASD, we found that damage to the cerebellum is prevalent in lobule VI and parts of lobule VIII (Marko et al., 2015), regions that in the macaque are critical for control of saccadic eye movements (Barash et al., 1999; Takagi et al., 1998). Control of saccades in healthy subjects shows exquisite sensitivity to decision related variables such as reward prediction error (Sedaghat-Nejad, Herzfeld, & Shadmehr, 2019), and history of effort expenditure (Yoon et al., 2018). Thus, marmosets may present a particularly good opportunity to investigate the role of the cerebellum in eye movements, social communication, and neurological diseases such as ASD.

However, despite these attractive features, there is concern that marmosets may be difficult to train for studies that investigate the neural basis of goal-directed behavior in a laboratory setting. This concern is due to the fact that in the current literature, the number of trials that these animals can perform within a session appears to be quite limited. For example, one recent study that trained marmosets to perform saccades found that on average, they performed 80-100 trials per session (Johnston et al., 2018), while another study reported 300-800 trials per session (Mitchell et al., 2014).

These numbers are on the lower bound of what is needed to identify task-related neurons in the cerebellum. For example, identification of an eye-related Purkinje cell requires 200-300 trials to measure complex spike tuning, and an additional 500-1000 trials to measure the simple spike response to saccadic eye movements (Soetedjo et al., 2008; Soetedjo & Fuchs, 2006). As a result, the relatively low numbers of trials per session reported earlier makes it unclear whether marmosets can serve as a model for neurophysiological studies of operant conditioned behavior.

Over the past four years, our team has been building a new marmoset laboratory, aiming to develop behavioral and neurophysiological protocols that allow for electrophysiological recording from the cerebellum of head-fixed animals. From a behavioral perspective, our goal was to ask whether the subjects could be motivated to produce a sufficiently large number of rewarded trials on a daily basis in the head-fixed configuration. From an electrophysiological perspective, our aim was to produce precise targeting of a region of the cerebellum, thus providing the means to simultaneously isolate multiple Purkinje cells using high-density electrodes.

To solve the training problem, we designed a carefully calibrated food regulation regime that motivated the subjects to perform a saccade task. To solve the electrophysiological problem, we developed a high-precision, acute multi-channel recording system that resulted in simultaneous recording from multiple Purkinje cells within the cerebellar vermis.

In this report, we provide an account of our experience, both in terms of procedures that failed and those that succeeded, going from initial behavioral training through neurophysiological recording. We begin with the design of a CT guided subject-specific titanium printed head-post, one that may be the first of its kind in the marmoset. We then describe the training and the MRI guided electrode alignment procedures. We conclude with examples of simultaneously isolated pairs of Purkinje cells, a difficult accomplishment only once before reported in the awake behaving primate (Medina & Lisberger, 2007).

4.2. Methods

4.2.1. Subjects

The procedures were carried out on three marmosets: two female (subject B, 420g, 11 years old; subject M, 350g, 4 years old) and one male (subject R, 400g, 4 years old). All three subjects were born and raised in a colony that Dr. Xiaoqin Wang has maintained at the Johns Hopkins School of Medicine since 1996. The experimental procedures were evaluated and approved by Johns Hopkins University Animal Care and Use Committee in compliance with the guidelines of the United States National Institutes of Health.

4.2.2. Design of head-post and recording chamber

Recording from Purkinje cells of the cerebellum imposes particular demands on stability: one needs to not only isolate a neuron, but also establish identity of that neuron via presence of both simple and complex spikes, which requires movable electrodes and longer periods of recordings. In order to stabilize the head during behavioral sessions, the current technique is to build an acrylic based thick "helmet" atop the skull, which then serves as a base that holds two or more metal posts (Lu et al., 2001a). The glue-inserted metal posts are attached to bars that are fixed to a primate chair, thereby holding the head in a fixed position. A more recent method is a halo-like device that holds the skull inside a ring (Johnston et al., 2018). Both approaches have been used for recordings from the cerebral cortex. However, reaching the cerebellum requires open access to the posterior parts of the skull,

making the halo approach less desirable. Furthermore, we aimed to produce an approach that reduced disruption of the underlying skin and muscles, resulting in a design that allowed for skin coverage of the skull and integrity of the temporalis muscles. This led us to develop a titanium based head-post and chamber design that eliminated use of the acrylic based helmet, and maximized the open area in the posterior part of the skull (Figure 4.1).

To prepare for head-fixation, we designed a head-post based on the specific geometry of the surface of each subject's skull and manufactured it using 3D printed titanium. We began with a preoperative CT (Figure 4.1A) and used it to build a 3D model of the skull (Figure 4.1B). This was done using 3D Slicer, an open-source image analysis and visualization software (Fedorov et al., 2012). The 3D skull model was then imported into a CAD environment (SolidWorks and Autodesk Fusion 360), which informed the design of an X-shaped titanium head-post (Figure 4.1B&C), a titanium base recording chamber (Figure 4.1D), and a plastic protective chamber cap. The head-post and base chamber designs were optimized to remain light weight and within the anatomical constraints of the frontal eye ridge and the lateral ridges on the skull, allowing us to minimize damage to the temporalis muscle attachment. The head-post and base chamber were 3D printed with laser melted grade 5 titanium (6AI-4V, Sculpteo.com), producing a lightweight (3 g), biocompatible, corrosion resistant structure. The protective cap that covered the chamber was 3D printed with PLA plastic filament (0.5 g).
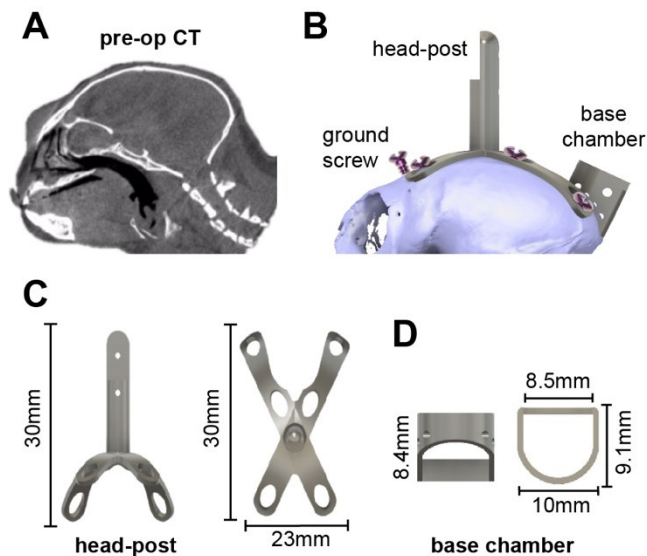
Figure 4.1. Design of the head-post and base chamber.
   A. Preoperative CT image that we used to build a geometric model of the surface of the skull. **B.** Model of the head-post and chamber fitted to the CT-based model of the animal's skull. The screws, head-post, and chamber are titanium. **C.** The head-post (2.39 g), designed to precisely fit the curved geometry of the subject. **D.** Base chamber (0.48g).

4.2.3. Surgery

Implanting the head-post and chamber was conducted by a surgical team led by lab members and aided by veterinary technicians and veterinarians. The animal was anesthetized with Alfaxalone (12 mg/Kg), then intubated with #2.0 uncuffed endotracheal (ET) tube, and administered isoflurane (1-1.25%, adjusted as needed) to maintain the anesthetic plane, as well as atropine and dexamethasone via intramuscular (IM) injections. The subject was then head-fixed by inserting bilateral ear bars into the osseous ear canal, which were fastened to a stereotaxic frame, while a bite bar was slid into place with a forehead clamp attachment to ensure that the head was positioned appropriately and securely with respect to the chest. 2% Lidocaine and 1:100000 epinephrine was then injected into the subcutaneous space beneath the scalp, and the operative region was treated with alternating cycles of diluted chlorhexidine and alcohol and at the very end sprayed with betadine. Next, we used a surgical marker to draw an outline for a midline incision on the scalp, extending 2mm posterior to the supraorbital fat pad and extending to the occipital ridge. We then made the incision using a scalpel and the skin was separated from the underlying fascia as well as the fascia from the underlying muscle. From here, the surface of the skull was prepared by removing residual tissue using curettes and 3% hydrogen peroxide solution, while the head-post was positioned for implantation. The temporalis muscle was pushed backed from its bony insertions on the spots that intersected with the head-post.

To attach the head-post, phosphoric acid etchant was applied for 15-20 seconds and then rinsed. Next, a thin layer of Optibond Solo Plus (Kerr Corp.) was applied to the surface of the skull, covering the entire area to be occupied by the head-post, cranial ground screw, and base chamber. Optibond Solo Plus was also applied to head-post X-shaped legs and lower half of the base chamber. After application, the Optibond was cured using a UV gun. The screw locations were then marked and drilled using drill bit #54 (1.4 mm diameter), with a stopper limiting the hole depth to 1.1mm. NX3 dual-cure dental cement (Kerr Corp.) was applied to the center of the bottom of the head-post, placed on the skull, and then cured using UV light.

Titanium screws (self-tapping, 2.0x4 mm, 0.04 g, or 2.0x5 mm, 0.05 g) were then used to affix the legs of the head-post onto the skull. A cranial ground screw (titanium, self-tapping 2.0x8 mm, 0.08 g) was screwed into the prepared hole (Figure 4.1B). We confirmed that screw depth into the skull did not exceed 1.1 mm. Next, NX3 cement was applied inside head-post screw holes, and around the ground screw to ensure structural integrity of the implanted screws.

NX3 cement was also applied to the bottom edge of the base chamber and the base chamber was placed between the rear legs of the head-post (Figure 4.1B) so that it was approximately located above the visual cortex and directed straight down. NX3 cement underwent a UV light curing. A final layer of NX3 cement was applied over the whole structure and was smoothed using Parafilm tape, and then cured with UV light.

The total weight of the NX3 cement, the installed apparatus, and the screws was less than 8.5 g. With the implantation complete, the temporalis muscle and fascia were replaced to their original positions. Finally, the two skin flaps were re-attached together via sutures using a 0/4 single filament thread. As a result, the surgical procedure produced an animal in which the skin covered almost the entire skull, except for small regions dedicated to the frontal ground screw, the main bar of the head-post, and the base chamber.

At this point anesthesia was discontinued and the ET tube was removed. The animal was monitored as consciousness and motor function were regained. A second dose of Dexamethasone (0.25 mg/Kg) was administered via IM injection 12 hours after the first dose. Recovery was monitored with daily postoperative check-ups for 13 days, including inspection of the surgery site for infection, checking for integrity of the sutures, recording of vital signs, feeding, and the administering of drugs.

4.2.4. Establishing a chamber-based coordinate system for the cerebellum

Because of the small size of the skull and our chamber design, we found it difficult to attach a micromanipulator to the head-post. Given that the micromanipulator was external to the head-post, we were concerned that we may not be able to accurately target and reproducibly approach a desired location in the cerebellum. This led us to design an electrode guide system that we integrated into the chamber design (Figure 4.2 and Figure 4.3). Our problem was that during surgery, the base chamber was placed by hand on the skull, and therefore its location with respect to the cerebellum was unknown. Furthermore, post-surgical CT or MRI imaging of the implanted animal was possible but problematic because the titanium pieces generated substantial artifacts, making it difficult to precisely localize the base chamber.

To solve this problem, we designed a CT visible reference ruler (Figure 4.2) that we inserted into the base chamber. The reference ruler consisted of 7 layers of a 1.5mm marked grid system, 3D printed with PLA plastic filament (Figure 4.2B). With the reference ruler inserted into the chamber, we performed a post-operative CT scan (Figure 4.2C). This produced a clear image of the ruler, but with

artifacts from the titanium. We then co-registered pre-operative CT (Figure 4.1A) and post-operative CT (Figure 4.2C) using bony areas and co-registered pre-operative CT (Figure 4.1A) and the pre-operative MRI (Figure 4.2A). Fixing pre-operative CT between 2 registrations produced a co-registered image of all three images. The result was a full representation of the subject's head, including the reference ruler, the skull, and the brain (Figure 4.2D).
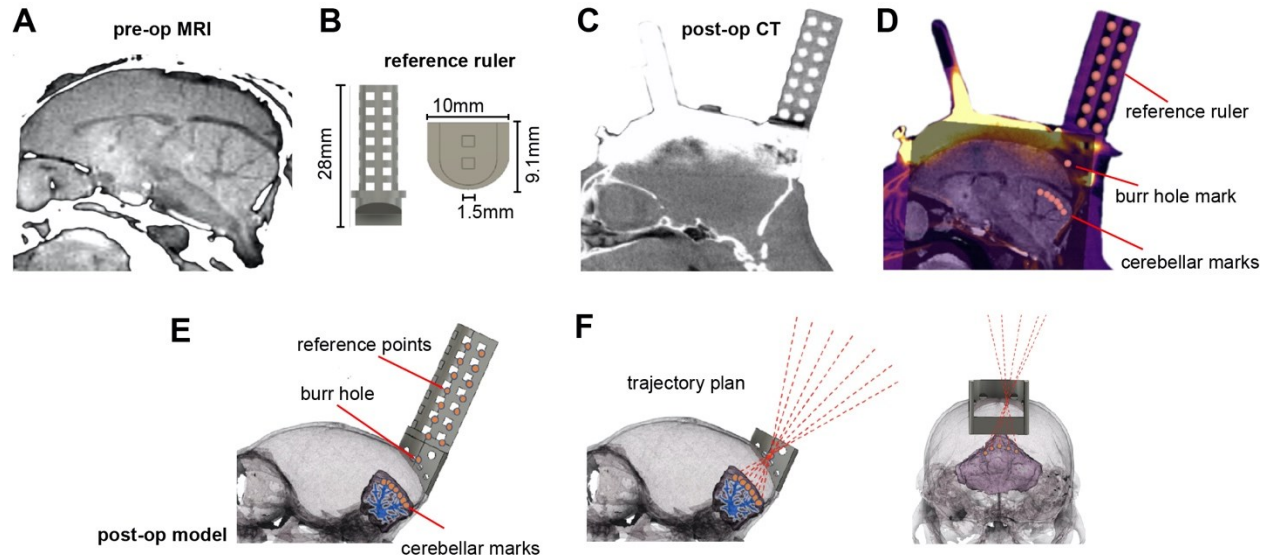
Figure 4.2. Development of a chamber-based geometric model
  A. Pre-operative MRI image used for identifying the desired regions of interest in the cerebellum. **B.** Reference axis ruler (1.104g) that was inserted into the base chamber before post-op CT imaging. **C.** Post-operative CT image after the surgical installation of the head-post, base chamber, and reference ruler. While the reference ruler is clearly visible, the titanium head-post and chamber have produced significant artifacts. **D.** Co-registered pre-operative CT, post-operative CT, and pre-operative MRI. Markers identify points on the reference axis and points in the cerebellum. A mark identifies the burr hole location. **E.** A 3D model that has co-registered the skull, cerebellum, chamber, and reference ruler. **F.** Using the 3D model, we drew trajectories that began at points of interest in the cerebellum, converged on a single 1.5mm diameter burr hole on the skull, and then diverged beyond the chamber as cylinders within the guidance tool. These trajectories represented desired electrode paths.

We used the ruler, chamber, and skull geometry to plan electrode trajectories in order to arrive at 20 distinct points of interest in the cerebellum. In this design, all 20 trajectories traveled through a single 1.5 mm diameter burr hole. To calculate each trajectory, we placed 21 points at various positions on the reference ruler (red marks, Figure 4.2D), 1 point on the surface of the skull for a craniotomy burr hole (burr hole point, Figure 4.2D), and 20 points at various locations along lobules V, VI, and VII of the vermis (cerebellar marks, Figure 4.2D). The point representing the burr hole was placed at 1.1 mm lateral to the midline in order to avoid the superior sagittal sinus.

The various points were then imported into a CAD environment, along with a model of the brain that we generated with 3DSlicer from the MRI images (Figure 4.2E). We aligned the imported points to their respective targets, guided by the co-registered image. The process relied on the alignment of all reference axis ruler points from the co-registered image (Figure 4.2D) to the reference ruler within the model environment (Figure 4.2E). This automatically aligned the burr hole and points on the cerebellum to the skull and brain models, creating an accurate representation of the distances and geometry of the desired cerebellar points, the desired burr hole, and base chamber.

Using the coordinates of the various points we created a set of 20 trajectories that originated from the desired recording locations in the cerebellum, converged through the single burr hole, and traveled out past the base chamber (Figure 4.2F). Each line in 3D space represented a trajectory that an electrode would travel to reach a given cerebellar destination.

4.2.5. Design of an electrode guidance tool

Based on the electrode trajectories that we had defined in the post-op model (Figure 4.2F), we designed a tool that would be attached to the base chamber and provide alignment of the electrode to the specific trajectory that traveled through the burr hole and arrived at the desired cerebellar point of interest (Figure 4.3A). This alignment tool featured 20 cylindrical tracks, each a cylinder of 1/32" in diameter. All cylinders converged onto the single burr hole, from which point the trajectories diverged to arrive at the various cerebellar destinations. The 20 cylinders (guide cylinders, Figure 4.3B) were configured in a 4X5 grid which corresponded to the four targeted cerebellar lobules (VIIa, VIa-c, VIe, and Vd), spanning from 0 to 2.5mm along the midline-lateral axis. The cylinders in the alignment tool provided a physical representation of the electrode trajectories above the burr hole. The tool was 3D printed using an Objet260 printer with VeroClear plastic filament.

4.2.6. Alignment of the electrode

To advance the electrode, we chose 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). The microdrive held the electrode and advanced it along a single axis defined by the chosen cylinder of the electrode alignment tool. Thus, the next problem was to position the microdrive in 3D space so that its single direction of motion was precisely aligned with the desired cylinder in the electrode alignment tool.

To align the axis of the microdrive with the desired cylinder, we began by attaching the microdrive to a mechanical stereotaxic micromanipulator (SM-11-HT, Narishige, Japan). With the alignment tool installed in the base chamber, we inserted a 1/32" rod (outer diameter) inside the desired cylinder of the alignment tool (Figure 4.3C), and then attached a 16G tube (0.033 inch inner diameter) to the microdrive using the same electrode holder that would also hold the electrode. The tube was fixed to the microdrive, serving as a model for the axis of travel of the electrode.
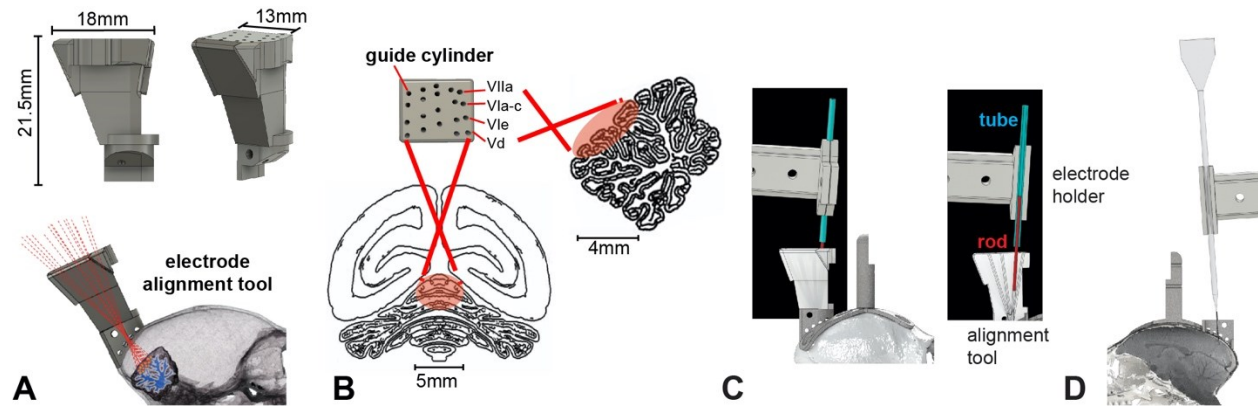
Figure 4.3. Design of the electrode trajectory guidance tool.
A. The alignment tool was designed to allow precise placement of the electrode along each of the 20 desired trajectories that converged onto various cerebellar destinations. The tool consisted of a 3D printed block (2.03g) that had a single cylinder for each of the desired electrode trajectories. The figure shows a model of the tool with 20 guide cylinders that align with the recording regions of interest and converge through the single craniotomy burr hole. **B.** Guide cylinder trajectories with posterior and sagittal views of the recording range on the cerebellum. The labels identify the cerebellar lobule that the trajectory is aimed toward. **C.** Alignment of the electrode with the desired trajectory. We installed the alignment tool in the base chamber and then placed a rod in the cylinder corresponding to the desired trajectory. The image on the right shows a section view of the alignment tool. We then placed a tube in the electrode holder (the same holder that would later hold the electrode), and attached the electrode holder to the microdrive. The microdrive was then maneuvered by a micromanipulator so that the axis of motion of the microdrive holding the tube was aligned with the rod. **D.** An image of the electrode and the electrode holder aligned to advance along the desired trajectory.

We then used the stereotaxic micromanipulator to maneuver the microdrive, thereby positioning the tube held by the electrode holder so that it precisely traveled along the axis defined by the rod. Alignment of the tube (representing the electrode) with the rod (representing the guide) was confirmed under a microscope by advancing the microdrive and demonstrating that this motion inserted the rod into the tube. Once this alignment was confirmed, the stereotactic manipulator that held the microdrive was locked into place. We then removed the tube from the microdrive and replaced it with

the recording electrode (Figure 4.3D). All further motion of the electrode was now under the control of the microdrive, limited to a single axis of travel that was along the desired trajectory.

4.2.7. Craniotomy

Our design required that the 20 trajectories specified by the alignment tool converge on a single point on the skull and then travel within the brain to reach the cerebellum. The next step was to localize this point on the skull and drill a 1.5 mm diameter burr hole. This was achieved by using the alignment tool to mark the skull and localize the craniotomy. During the procedure the subject remained awake, following a protocol described earlier (Lu et al., 2001a).

The subject was head-fixed while the general craniotomy area, the base chamber, and the alignment tool were sterilized using alcohol, rinsed with saline, and then suctioned dry. The alignment tool was installed on the base chamber and a 22G blunt needle tip was dipped into ink and then inserted into one of the cylinders (the choice was arbitrary as all cylinders converged to the same burr hole location). The ink tipped needle was advanced until contact was made with the bone and a mark was left on the surface.

Using the procedures for aligning an electrode, we aligned a sterilized drill bit so that its trajectory of travel precisely matched the trajectory specified by a cylinder on the alignment tool. We attached a DC motor to the bit, and then attached the motor to the stereotaxic micromanipulator. At this point the alignment tool was removed, exposing the marked skull. From our co-registered pre-op and post-op CT we estimated that the total thickness of the skull and dental cement was between 900-1100 μm. Under a microscope, a miniature electric drill mounted on the micromanipulator was advanced to a near touch of the dental cement. The skull and dental cement thickness estimate were used to guide the advancement of the drill bit through the thin layer of dental cement and about 80-90% of the thickness of the bone (before reaching the dura). Drilling concluded with observation of wet bone chips and/or liquid. The remaining bone covering the dura surface was then carefully removed using hand-held fine instruments under a microscope. Over the past two decades the Wang Laboratory (Lu et al., 2001a) has safely used this approach, and we were able to confirm that the animals never exhibited signs of discomfort or distress during this procedure.

4.2.8. Sealing the burr hole with a transparent silicone gel

Our design included a single burr hole, thus requiring us to continue using a single entry for many weeks during daily electrode penetrations. To accomplish this, we experimented with a silicone gel to seal the exposed dura at the burr hole.

Silicone gels can help preserve the integrity of the intra-cranial space after a craniotomy and in prolonging the life and functionality of burr holes. We wanted to use a gel that was non-toxic, transparent, elastic, easy to apply, and had resealing capabilities, allowing multiple penetrations by electrodes over an extended period of time.

We experimented with a commercially available soft silicone gel (DOWSIL 3–4680, Dow Corning, DuraGel, Cambridge Neurotech). This gel is a polydimethylsiloxane (PDMS) based silicone and is primarily used for hydrophobic encapsulation of electronic microchips. However, due to its biocompatible and antibacterial properties, it has recently been used as a dura substitute/cover for chronic recording in mice (Jackson & Muthuswamy, 2008; Jiang et al., 2017). These previous reports had tested biocompatibility, cytotoxicity, and sealing capability of the gel in mice, suggesting that it is a safe product, effective in sealing the dura from air contact, thereby minimizing possibility of infection. The published reports suggested that the gel had a number of attractive features: it was biocompatible, reduced possibility of inflammation, reduced cerebro-spinal fluid leakage, reduced humidity loss from the craniotomy, and remained soft for a period of days to weeks, thereby allowing microelectrode penetration with minimal force.

Immediately after drilling the burr hole and cleaning the remaining bone chips we covered it with the silicone gel (DuraGel, Cambridge Neurotech). Once the gel was cured (about 30 minutes), its mechanical properties allowed the electrodes to penetrate it for weeks. Moreover, the viscosity of the gel allowed us to place it inside of the burr hole, thereby covering the dura completely, sealing it from air. The gel is transparent, which allowed us to monitor the dura beneath the gel. We replaced the gel at two-week intervals.

4.2.9. Data acquisition

We recorded from the cerebellum using three types of electrodes: quartz insulated 4 fiber (tetrode) or 7 fiber (heptode) metal core (platinum/tungsten 95/05) electrodes (Thomas Recording), and 64 contact high density silicon probes (Cambridge Neurotech). None of these electrodes could penetrate the marmoset dura. Therefore, we performed a micro-duratomy using a 30G needle, which was installed on

the stereotaxic micromanipulator frame and advanced to the surface of dura until a puncture was made. Once the puncture was made, the various electrodes could travel through the dura.

We connected each electrode to a 32 or 64 channel headstage amplifier and digitizer (Intan Technologies, USA), and then connected the headstage to a communication system (RHD2000, Intan Technologies, USA), as shown in Figure 4.4 In addition, we connected to the RHD2000 system digital outputs from our custom behavioral software and analog output from an optical sensor mounted on the TV screen.

We used OpenEphys (Siegle et al., 2017), an open-source extracellular electrophysiology data acquisition software, for interfacing with the RHD2000 system and recording of signals. The signals were post processed and analyzed using MATLAB (Mathworks, USA) and Python.
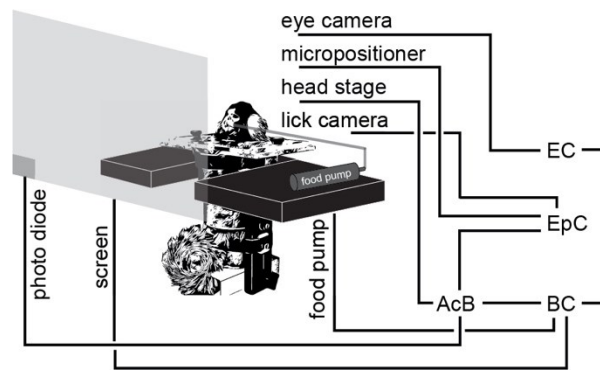
Figure 4.4. Behavioral and electrophysiological recording system.
AcB: acquisition board. BC: behavioral computer. EpC: electrophysiological computer. EC: eye recording computer.

4.2.10. Behavioral training

In our earlier work in the macaque we had discovered that a key step in decoding activity of Purkinje cells in the oculomotor vermis was to organize them into groups in which all the cells within a group shared the same preference for visual error (Herzfeld et al., 2015, 2018). The preference for error was signaled via the cell's complex spike tuning, as measured when the animal made a saccade but at saccade end, the target was not on the fovea, resulting in a sensory prediction error which we represented as a vector. When the Purkinje cells were organized based on their error-dependent complex spike tuning, each population predicted in real-time the motion of the eyes during a saccade as a gain-field. Thus, we trained the marmosets in a task in which we could identify the complex spike tuning of each Purkinje cell with respect to visual error.

Immediately following surgical recovery, the animal was placed on a food regulated diet. For 5 days a week, this diet consisted of 15 g of lab diet powder and 10 g of apple/mango sauce mixed in 30 g of water. This resulted in net 40 mL of deliverable food, which we administered via a syringe pump during the task. During the weekends, the food consisted of 30 g of a solid lab diet for the first day and 25 g for the second day.

Weight was monitored daily to ensure the health of the animals: weight was maintained within 85-100% of the average weight before the food regulation regime. If the weight fell below 85%, food regulation was stopped and the animal was fed in the colony until weight recovered to at least 90%. The animal was progressively acclimated to being held by a handler, entering its carrier independently, and sitting in its task chair. The experiment room was maintained at 74-84°F.

Saccade training proceeded at a frequency of 5 days/week. Visual targets were presented on a TV screen (Curved MSI 32" 144 Hz - model AG32CQ) while binocular eye movements were tracked using an EyeLink-1000 eye tracking system (SR Research, USA). We began with fixation training, followed by saccade training to a primary target. As the subject became accustomed to these paradigms, we introduced a secondary target, thus encouraging a corrective saccade (Figure 4.5A).

Following fixation training, saccade training initiated with presentation of a primary target (0.5°×0.5° square) at a random location and distance of 3°. Reward was provided if the primary saccade was within 1.5° radius from the center of the primary target. In the fully trained subject the trial began with fixation of a center target for 200 ms, after which a primary target (0.5°×0.5° square) appeared at a random location at a distance of 5°-6° (Figure 4.5A). As the animal made a saccade to this primary target, that target was erased and a secondary target was presented at a distance of 2°-2.5°. The subject was rewarded if following the primary saccade it made a corrective saccade to the secondary target, landed within 1.5° radius of the target center, and maintained fixation for 200 ms.

Correct trials produced a distinct auditory tone, and engagement of the food pump at a rate of 0.020 mL/trial. We found that in the trained subjects, this low rate encouraged them to complete a few consecutive trials before stopping to lick the food tube, thus further increasing the number of correct trials per session.
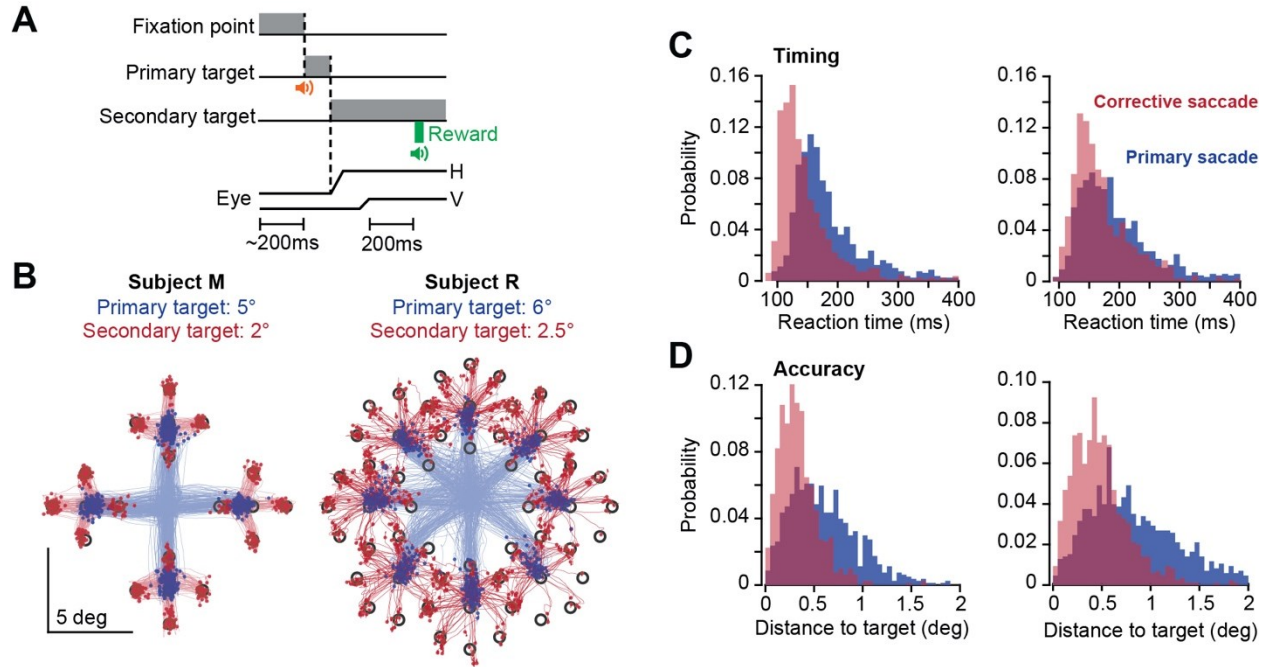
Figure 4.5. Task design and behavioral results.
   A. A trial began with 200 ms of fixation, followed by presentation of a primary target at 5° or 6°. During the primary saccade the target was erased, and a secondary target was presented at 2° or 2.5° displacement with respect to the primary target. Reward was presented following 200 ms fixation of the secondary target. **B.** Saccade trajectories for subjects M and R during one session. Primary saccades are plotted in blue and corrective saccades are plotted in red. Targets were 0.5°×0.5° square. **C.** Reaction time distribution for primary and corrective saccades. D. Distance to target at conclusion of primary and corrective saccades.

4.2.11. Data analysis

To detect simple spikes, we began with a high-pass filtered version of the signal (300 Hz), and then subtracted at each time point the mean of the signal across all contacts, a method known as common average referencing. We then used a threshold-based technique to hand-sort the data. Here, we consider only the largest spike recorded by each contact. To detect complex spikes, we relied on frequency-domain analysis that tentatively identified these spikes via their power spectrum properties.

To confirm that the complex and simple spikes originated from the same cell, we compared the conditional probability Pr(S(t)|C(0)) with Pr(S(t)|S(0)). These probabilities describe spike-triggered histograms. For example, Pr(S(t)|C(0)) is the probability that a simple spike occurred at time t, given that a complex spike was generated at time zero. Pr(S(t)|S(0)) is the probability that a simple spike occurred at time t, given that a simple spike was generated at time zero. Simple spikes that originate from a single cell produce a refractory period. Thus, Pr(S(t)|S(0)) should exhibit a low probability period of roughly 10 ms in duration centered at time zero. On the other hand, a complex spike suppresses production of

future simple spikes, but not those that occurred before. As a result, Pr(S(t)|C(0)) should be asymmetric, with a long period of low simple spike probability following time point zero.

Multi-contact electrodes allow for analysis of simultaneously recorded neurons. However, spiking activity in one neuron can easily influence the data recorded by two nearby contacts, thus giving an illusion that the two contacts are picking up two distinct neurons. To guard against this, after we sorted the data in each contact, we waveform triggered the data recorded by contact A by the spikes recorded on contact B. This identified the waveform of the neuron recorded by contact B on the spike recorded on contact A. We compared this cross-contact triggered waveform with the within contact triggered waveform generated by the spikes recorded by contact A. The cross-contact triggered waveform should produce a different cluster of spikes in A than the main neuron isolated by A. If there were spikes that occurred within 1 ms of each other on contacts A and B, we used these coincident-spike events to trigger the waveform in A. The spikes in A that were identified to be coincident with B should look approximately the same as the non-coincident spikes in A.

To quantify coordination between activities of two Purkinje cells, we computed conditional probabilities. For contacts 1 and 2, we computed Pr(S1(t)|C2(0)), testing whether complex spikes on contact 2 produced any changes in the simple spikes on contact 1. If the data were sorted properly, this conditional probability should be essentially flat. To visualize coordination between the neurons, we next plotted the spike-triggered waveform of voltages recorded by contact 1, triggered by the simple spikes on contact 2. Finally, we computed Pr(S1(t)|S2(0)). This measure quantified whether the occurrence of a simple spike on contact 2 altered the probability of simple spikes on contact 1.

4.3.   Results

The methods that we described above allowed us to train subjects to produce around a thousand trials per session during cerebellar recordings. However, we arrived at these methods following a learning process that involved a number of unsuccessful attempts. Here, we describe our experience in terms of both types of results.

4.3.1.   Behavioral results

For subject B, we experimented with saccade training without the head-post using a loose fitting head-restraining system (days -416 to 0, training without head-post, Figure 4.6A). Following a number of attempts we abandoned this approach because we were unable to produce robust calibration.

We began behavioral training of subject B after head-post surgery using the traditional method (Johnston et al., 2018; Mitchell et al., 2014) of feeding the animals in the colony and supplementing that food with reward during task performance in the head-fixed condition (training without food regulation, Figure 4.6A). We found that this was insufficient to motivate this subject.

We next eliminated home-cage feeding on experiment days, and instead provided task-based food rewards during the head-fixation period (food regulation, Figure 4.6). We ensured that regardless of performance, the subjects always received their required allotment of food (at least 25 ml of food per day) and weight was maintained at around 90% of the pre-food regulation regime. However, food was delivered only while the animal was head-fixed. This produced improved performance in subject B (days 846-917, Figure 4.6A). We then fully tested this approach in subjects M and R (Figure 4.6B and Figure 4.6C).
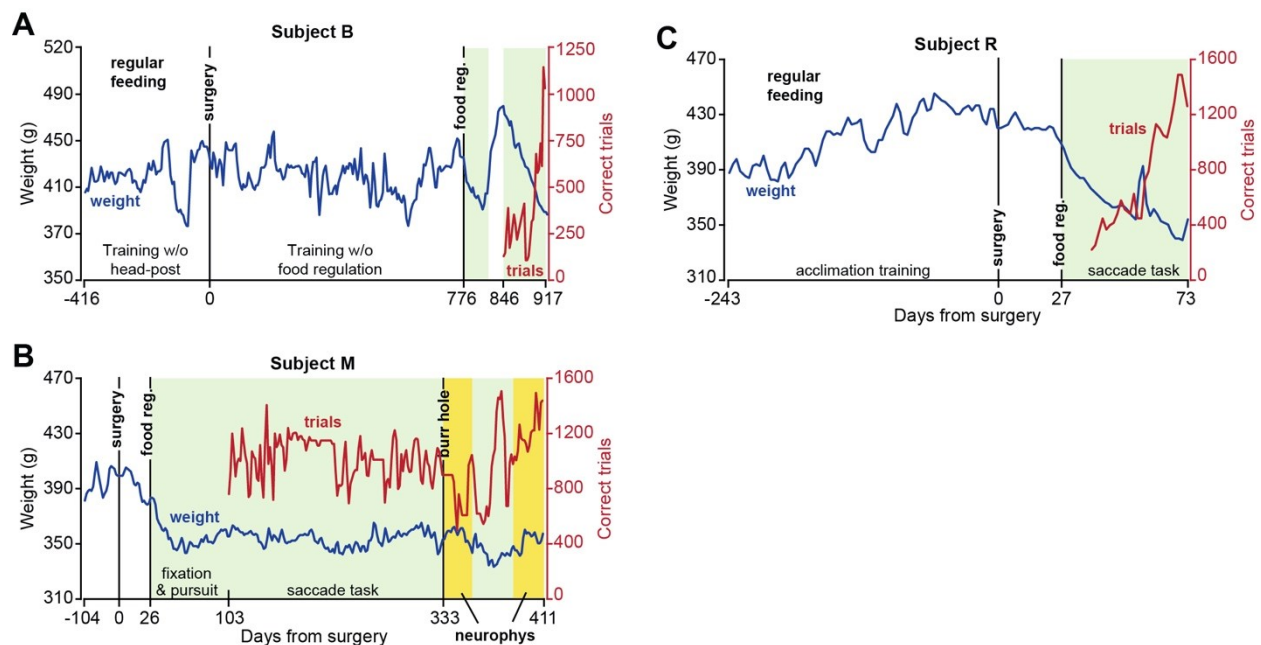
Figure 4.6. Weight patterns and number of correct trials following food regulation.
A. Record of weight and correct trials in subject **B.** Training began without a head-post, but was unsuccessful because of an inability to produce reliable calibration of eye measurements. Following surgery, a long period of training without food regulation ensued, but was ultimately abandoned because of the unwillingness of the animal to work for more than a few hundred trials. Following food regulation, performance in the task dramatically improved. B. Weight and correct trial data in subject M. **C.** Weight and correct trial data in subject R. Correct trials data represent running average of bin size 2. Green color indicates periods of food regulation. Yellow color indicates periods of food regulation and neurophysiological recording.

Following recovery from surgery, subject M was food regulated and fed only in the head-fixed condition. After 30 sessions of fixation training, subject M was trained for 20 sessions on trials with

primary saccades only (gradually increasing from 3° to 5°). Following this, subject M began training on the main task (Figure 4.5A), which included both primary and corrective saccades. The correct number of trials shown in Figure 4.6B begins with the first day of training on the primary saccade and continues as the subject transitioned to the main task. We found that within three months after start of fixation training subject M was able to perform around 1000 correct trials in the main task, and then maintained this performance for over a year. During neurophysiological experiments, subject M produced 600-1500 correct trials per session (neurophysiology trials, running average data presented in Figure 4.6B).

We experimented with a different approach to saccade training in subject R. Instead of starting with fixation training, we began with pursuit training for 3 days. This subject's good performance allowed us to subsequently train on a saccade task for 15 sessions during which only the primary target was presented, and amplitude gradually increased from 3° to 5°. The subject was then trained on the main task (Figure 4.5A), during which the primary target was placed at 6° and the secondary target was at 2.5°. The correct number of trials in Figure 4.6C begins with the first day of training on the primary saccade and continues as the subject graduated to the main task. This subject was able to produce 1000-1500 correct trials in the main task within two months after food regulation (Figure 4.6C).

Notably, during a year of food-regulated training in subject M, and two months of food-regulated training in subject R, both remained healthy, as suggested by their stable weight and lack of complications.

Our aim was to train the subjects in a task in which we could measure contributions of Purkinje cells to control of saccades. Each trial began with 200 ms fixation of a central location and was followed by a primary saccade to a peripheral target (0.5°×0.5° square) at 5° or 6° displacement, corrective saccade to a secondary target at 2° or 2.5° displacement, and 200 ms of fixation at the secondary target (Figure 4.5A). Saccades for a representative session are plotted in Figure 4.5B, and the distribution of reaction times are plotted in Figure 4.5C. Primary saccade reaction times for subjects M and R were 188±65 ms (mean±SD) and 197±69 ms. Corrective saccade reaction times for subjects M and R were 146±50 ms and 169±50 ms. To quantify saccade accuracy, we measured the distance between saccade endpoint and target center and found that for the primary saccade, this distance was 0.59±0.35° (mean±SD) for subject M, and 0.85±0.50° for subject R (Figure 4.5D). For the corrective saccade, distance to target was 0.34±0.22° for subject M and 0.49±0.32° for subject R. The reward region was 1.5° radius around the target. As a result, across the two subjects 81% (1011 out of 1251) and 88% (1464

out of 1672) of the trials that initiated with a primary saccade to the target concluded correctly and were rewarded.

An important component of our training was that the subjects received a small amount of food per correct trial: 0.02 mL per trial. At completion of a successful trial, the computer generated a beep, and food delivery was signaled by the sound of the pump. However, the low reward rate encouraged the subjects to withhold licking until the food accumulated in the transparent tube that was placed in front of them. The effect was to teach the subjects to perform the task in blocks of 3-4 uninterrupted consecutive trials, thus increasing the total number of successful trials per session.

4.3.2. Simultaneous recordings from pairs of Purkinje cells

We explored lobules V, VI and VII of the cerebellum using tetrodes, heptodes, and high density silicon arrays. We thought that the cranial ground screw would be required to provide a reference for our electrophysiological recordings. However, with experience with subject M we learned that the ground screw was unnecessary: the head-post by itself acted as a reliable reference for the electrical signals during neurophysiological recordings. This is likely because the titanium screws established electrical continuity between the head-post and the bone. Thus, in subject R we eliminated the ground screw from the surgical procedures.

As the electrode approached the tentorium, there was a noticeable increase in background activity, indicative of proximity to Purkinje cells and their high baseline discharge rate. When the electrode passed through the tentorium, there was a distinct "pop", which was then followed by a gradual increase in spiking activity.

Identification of a Purkinje cell required presence of both simple and complex spikes. Examples of putative complex and simple spikes, recorded by a silicon array, a tetrode, and a heptode are presented in Figure 4.7A. Complex spikes have a stereotypical positive wave that has slow dynamics, resulting in greater power in the lower frequencies (<300 Hz range) than the waveform of simple spikes (Warnaar et al., 2015). Furthermore, complex spikes are relatively rare events, presenting an inter-spike interval that is roughly two orders of magnitude greater than simple spikes.

A useful real-time tool for identification of complex spikes is a low-pass filter (300 Hz), which under good recording conditions can provide a tentative label for the complex spikes. Figure 4.7B shows an example of recording data from a heptode that simultaneously isolated three neurons: two putative Purkinje cells, and a third unidentified neuron. In this figure, the voltages for each of the three

simultaneously recorded neurons are represented as black traces, and the low-pass filtered version of the same signal is shown in red. The complex spikes are tentatively identified by the sharp increase in the low-frequency power of the signal.

An important feature of a Purkinje cell is that production of a complex spike suppresses the generation of simple spikes for a period of 10 ms or longer. Thus, the shared cellular origin of complex and simple spikes can be verified by using the complex spikes to trigger the voltage waveform. The upper trace of Figure 4.7C provides an example of this, demonstrating that the simple spikes occur with roughly equal probability before the complex spike, but then cease entirely after the complex spike.
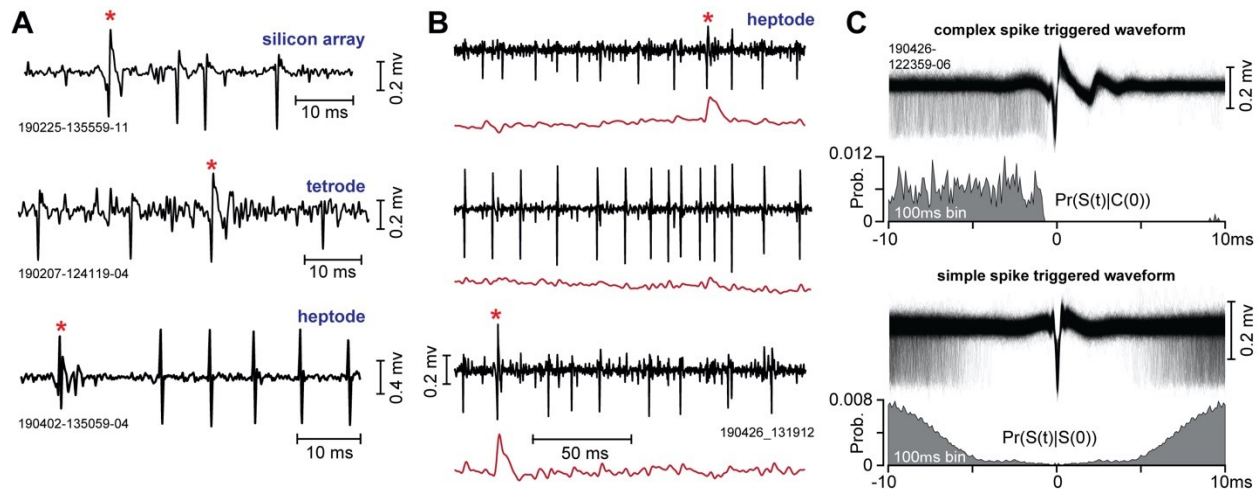
Figure 4.7. Recordings from the marmoset cerebellum.
A. Complex and simple spikes recorded by a silicon array, a tetrode, and a heptode. The complex spike is noted by an asterisk. **B.** Example of multiple neurons simultaneously recorded by a heptode. The black traces are the high-pass filtered voltages and the red trace is the low-pass filtered version of the same signal (300 Hz cutoff). The top and bottom traces are putative Purkinje cells while the middle trace is an unidentified third neuron. **C.** Spike-triggered waveforms and histograms. In the upper trace, the voltages recorded by a single contact on a heptode were triggered by the complex spikes on the same contact. The resulting complex spike-triggered waveform exhibits a uniform pattern of simple spikes before the complex spike, and then simple spike suppression. This pattern is quantified as $Pr(S(t)|C(0))$, which is the probability of a simple spike at time t, given that a complex spike occurred at time 0. In the lower trace, the voltages were triggered by the simple spikes. The resulting simple spike-triggered waveform exhibits a symmetric period of spike suppression, reflecting the Purkinje cell's refractory period. This pattern is quantified in $Pr(S(t)|S(0))$.

We quantified this pattern via the condition probability $Pr(S(t)|C(0))$, which measured the probability of a simple spike at time t, given that a complex spike occurred at time 0. This measure illustrated that the simple spikes occurred with equal frequency before the complex spike but were then completely suppressed for around 10 ms following the complex spike.

A second feature of a Purkinje cell is that the simple spikes have a short refractory period. This feature can be verified by using the simple spikes to trigger the voltage waveform. The lower trace of Figure 4.7C provides an example of this. We quantified this pattern via the conditional probability $Pr(S(t)|S(0))$, which measured the probability of a simple spike at time t, given that a simple spike occurred at time 0. This measure illustrated that the simple spikes exhibited a 5 ms refractory period.

Two examples of simultaneously recorded neurons from lobule VII during the saccade task are shown in Figure 4.8. In Figure 4.8A, the heptode isolated three neurons, two of which were Purkinje cells (high-pass filtered signals are shown in black, the low-pass filtered signal is shown in red). In this trial, the animal was presented with a target at 5° on the horizontal axis and made a saccade at 220 ms latency (bottom row, Figure 4.8A). At saccade onset the target was erased and redrawn at a displacement of 2° along the vertical axis. Thus, at the conclusion of the vertical saccade the target was not on the fovea, generating a visual error, and encouraging the animal to make a corrective vertical saccade. This is an example of cross-axis adaptation (Deubel, 1987; Xu-Wilson, Chen-Harris, et al., 2009). The saccade landed about 0.5° short of the target, and was followed by a corrective saccade. The Purkinje cell on contact 6 (2nd row, Figure 4.8A) produced a complex spike before the primary saccade. A second Purkinje cell, recorded on contact 5 (top row, Figure 4.8A), did not produce a complex spike before the onset of the primary saccade, but rather produced a complex spike following experience of a visual error (before onset of the corrective saccade). A third neuron was isolated by contact 7, but this neuron was not a Purkinje cell as it lacked complex spikes.

The waveforms for the complex and simple spikes in this recording session for each contact are shown in the left column of Figure 4.8B. The waveforms illustrate the slower dynamics of the complex spikes, thus providing the basis for why a low pass filter is generally useful in their identification.

The middle column of Figure 4.8B displays the probability densities of the inter-spike intervals for each type of spike (note that the x-axis for the complex spikes is 5 seconds, whereas for the simple spikes is 50 ms). The number of simple spikes in a given period of time was roughly two orders of magnitude larger than the number of complex spikes.
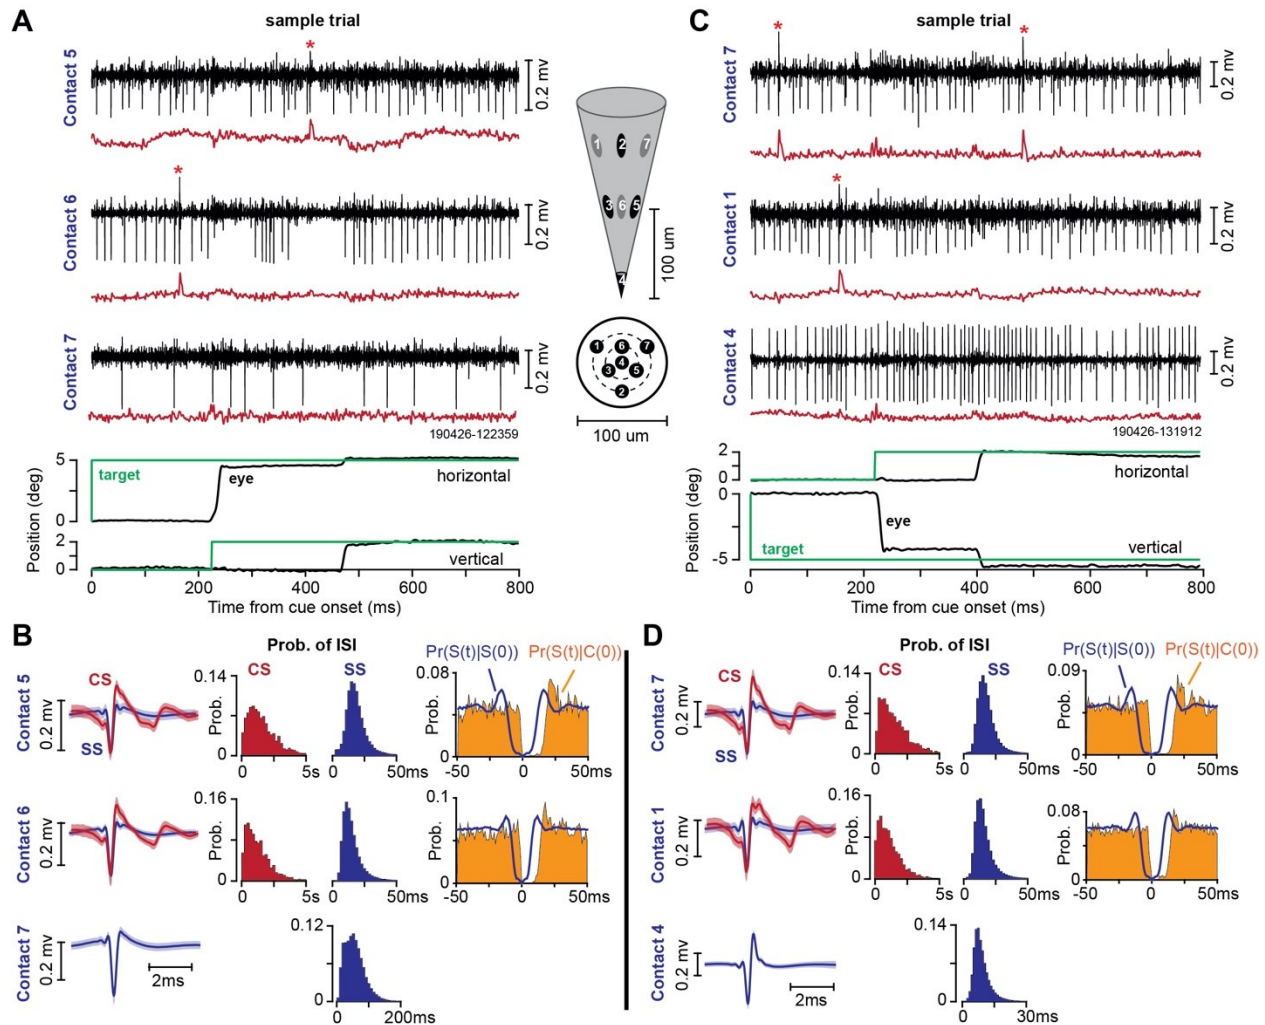
Figure 4.8. Example of heptode recordings during two sessions in the saccade task.
A. Sample trial, illustrating activity of a Purkinje cell on contact 5, another Purkinje cell on contact 6, and an unidentified neuron on contact 7. The black traces are high-pass filtered voltages and the red traces are low-passed filtered signal (300 Hz cutoff). The middle figure shows the geometry of the heptode and the location of each contact. **B.** Spike waveforms and patterns of spike timing recorded from the session in part A. CS is complex spike and SS is simple spike. ISI is inter-spike interval. Pr(S(t)|S(0)) is the probability of a simple spike at time t, given that a simple spike occurred at time zero. Pr(S(t)|C(0)) is the probability of a simple spike at time t, given that a complex spike occurred at time zero. **C.** Sample trial recorded in a different session, illustrating activity of a Purkinje cell on contact 7, another Purkinje cell on contact 1, and an unidentified neuron on contact 4. **D.** Spike waveforms and patterns of spike timing recorded from the session in part B. Bin size is 1 ms for conditional probabilities, 200 ms for probability of complex spike ISI, and 2 ms for probability of simple spike ISI. Error bars on the spike waveforms are standard deviation.

The right column of Figure 4.8B illustrates the within-cell timing properties of the spikes via conditional probabilities. Pr(S(t)|C(0)) demonstrates that following a complex spike at time 0, there was a silent period of ~10 ms during which simple spikes were absent. Thus, production of a complex spike

69

briefly suppressed production of simple spikes, demonstrating that the complex and simple spikes on this contact originated from the same Purkinje cell. Pr(S(t)|S(0)) demonstrates a 5 ms refractory period, which is critical for demonstrating that the simple spikes were correctly attributed to a single Purkinje cell, and not misattributed because of cross-contaminated spikes from a neighboring contact. Finally, a comparison of Pr(S(t)|C(0)) with Pr(S(t)|S(0)) demonstrates that the complex-spike induced suppression of simple spikes was longer than the typical simple spike refractory period, another indication that the source of complex and simple spikes was a single Purkinje cells.

Another example of a recording session is shown in Figure 4.8C. In this trial, the target was presented at 5° along the vertical axis, and the animal made a saccade at a latency of 220 ms. At saccade onset, the primary target was erased and a new target at 2° to the right of the primary target was displayed. A Purkinje cell was isolated on contact 1, while a second Purkinje cell was isolated on contact 7. Contact 4 isolated a third neuron, but this was not a Purkinje cell.

Figure 4.8D illustrates the within contact timing properties of the complex and simple spikes. Like the data in Figure 4.8B, the conditional probabilities demonstrate that simple spikes exhibited a 5 ms refractory period, and following a complex spike, there was suppression of simple spikes for a period of 10 ms.

In our experience, the head-post and the electrode holder provided good recording stability. Over the course of 22 sessions in which single or pairs of Purkinje cells were isolated, we were able to maintain single unit isolation during an average of 39.1±3.8 minutes (mean±SEM), producing 516±62 correct trials.

Activity of a Purkinje cell over the course of a representative session (700 trials, 40 minutes) is illustrated in Figure 4.9. In this session, subject M performed the task shown in Figure 4.5B. The trial began with a primary target at 5° displacement. Upon saccade initiation, the primary target was erased and replaced with a secondary target at 2° displacement (Figure 4.9A). A corrective saccade followed termination of the primary saccade at a latency of 120-150 ms. The rasters in Figure 4.9A present the timing of simple and complex spikes during each trial via thin blue and thick red lines, respectively. This cell produced a burst of spikes around primary saccade onset, exhibiting elevated activity that continued long after the saccade had ended (Figure 4.9B, left panel). There was a second burst of simple spikes around the time of the corrective saccade (Figure 4.9B, right panel), which also exhibited a duration that was much longer than the saccade. This disparity between burst duration of simple spikes and duration

70

of saccade is a common feature of saccade related Purkinje cells (Herzfeld et al., 2015; Thier et al., 2000).
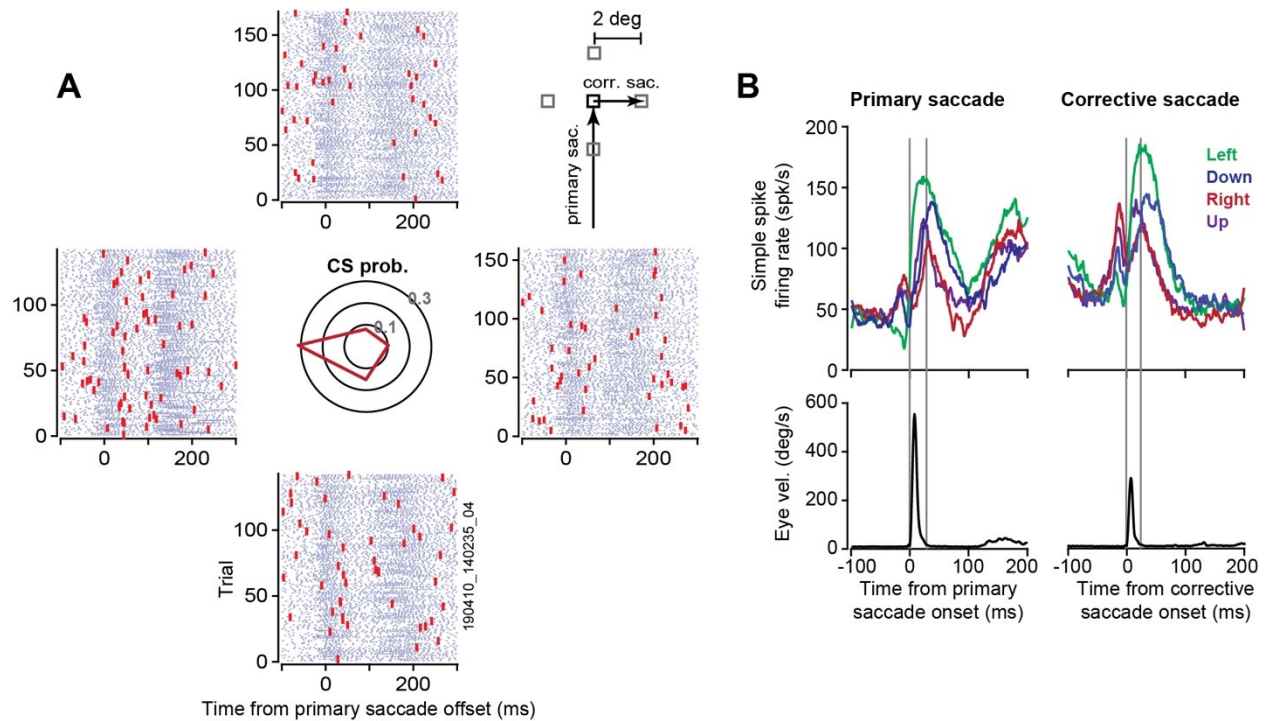
Figure 4.9. Example of data recorded from a Purkinje cell (subject M) during the behavioral task. **A.** The primary target (black square) was presented at 5° displacement. Upon initiation of saccade, the target was replaced with a secondary target (gray square) at 2° displacement in one of four directions, inducing a visual error and encouraging a corrective saccade. Rasters are aligned to primary saccade offset, and grouped based on the direction of the visual error vector. For example, the rasters on the right refer to trials in which the corrective saccade was to the right. The central plot shows probability of complex spikes during the 200 ms period following offset of the primary saccade. This cell had a preference for visual errors toward -180°. **B.** Simple spike activity and eye velocity aligned to onset of primary saccade and corrective saccade. In the left subplot the simple spikes are grouped based on direction of the primary saccade. In the right subplot the spikes are organized based on direction of the corrective saccade. The gray vertical lines denote onset and offset of the saccade. The cell exhibited a burst of activity around saccade onset, but burst duration was much longer than saccade duration.

Upon termination of the primary saccade the target was not on the fovea, resulting in a sensory prediction error. This visual event likely engaged neurons in the superior colliculus (Kojima & Soetedjo, 2017b) and the inferior olive, leading to occasional complex spikes. We measured probability of complex spikes during the 200 ms period following completion of the primary saccade (central panel of Figure 4.9A). This Purkinje cell was tuned to visual errors along -180°, suggesting that it received visual error information from the right superior colliculus via the left inferior olive.

Herzfeld et al. earlier work (Herzfeld et al., 2015, 2018) found that when Purkinje cells are organized into populations that share a common preference for error (i.e., complex spike tuning is similar), their combined simple spike activity produces a pattern of spikes that no longer has a duration disparity with respect to the saccade. That is, while the simple spike activity of individual Purkinje cells does not have an obvious relationship to eye motion during a saccade, as a population the Purkinje cells can precisely predict motion of the eye in real-time.

4.3.3. Spike timing properties of simultaneously recorded Purkinje cells

The ability to record simultaneously from multiple Purkinje cells allows one to ask whether nearby cells coordinate their activity, a function that may be fundamental in driving neurons in the cerebellar nucleus (Person & Raman, 2011). However, with multiple contacts there is a danger that some of the spikes that are recorded on contact A and attributed to Purkinje cell 1 are in fact generated by nearby Purkinje cell 2, which in turn is also being recorded by contact B. This cross-contact influence (cell 2 affecting recordings in both contacts A and B) would result in spurious coordination. One way to guard against this misattribution is to compare the spike-triggered waveforms.

Figure 4.10A (left column) illustrates the waveform of simple spikes recorded by contact 6, during the same session as that shown in Figure 4.8A. The middle column illustrates the waveform recorded by contact 6 but triggered by the simple spikes on contact 5. The data suggest that on average, the voltages recorded by contact 6 reflect in only a minor way the spikes being recorded by contact 5. Occasionally, the two P-cells produced a simple spike within 1 ms of each other. The waveform on contact 6 as triggered by only these synchronous events (right column of Figure 4.10A) looks very similar to the waveform of the usual spikes (left column of Figure 4.10A). Thus, these features suggest that any coordinated activity present in the data of these two contacts would be because two independent cells fired simple spikes at approximately the same time.

We visualized the coordination between two Purkinje cells by triggering the voltages recorded by one contact via the simple spikes generated by a different Purkinje cell on a separate contact (Han et al., 2018). The resulting spike-triggered histogram of waveforms recorded by contact 6, triggered by spikes on contact 5, is shown in Figure 4.10B. The term $Pr(S6(t)|S5(0))$ measures the probability of simple spikes on contact 6 at time t, given that a simple spike occurred on contact 5 at time 0. The conditional probability exhibits a peak at around 0 ms, suggesting a measure of coordination. This coordination could have arisen because both Purkinje cells were saccade related, or driven by

presentation of a visual stimulus. It is also possible that these two cells received similar inputs (Heck et al., 2007), or that activity in one affected activity in the other (Han et al., 2018).
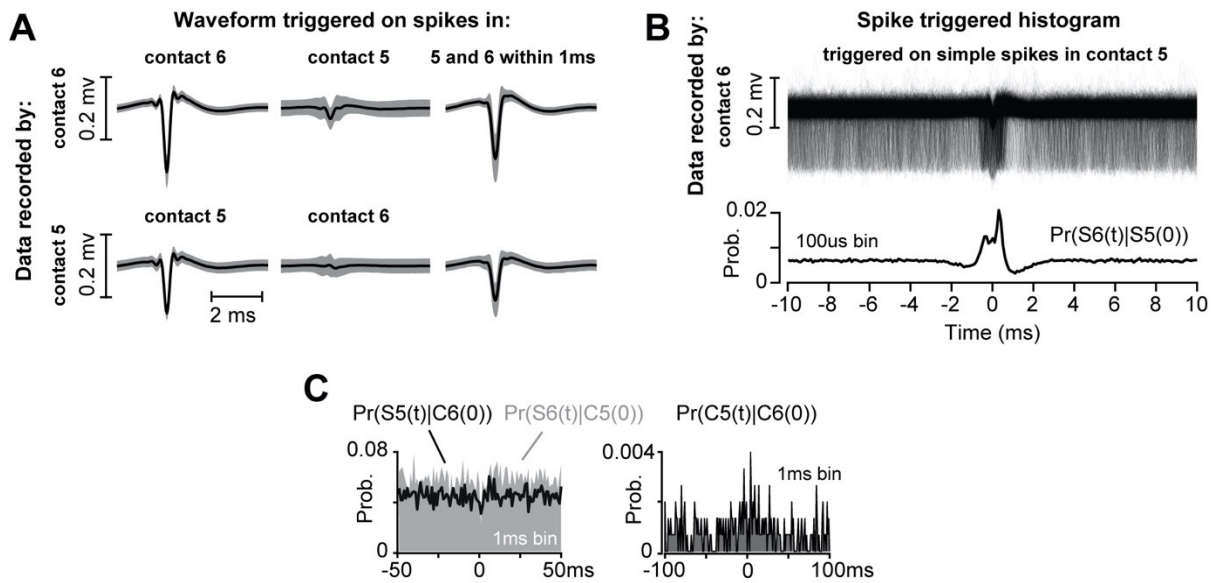
Figure 4.10. Spike timing property of a pair of simultaneously recorded Purkinje cells during the saccade task.

A. A Purkinje cell was isolated by contact 5, and another one by contact 6. Upper row: data recorded by contact 6 were triggered by simple spike events on contact 6 (left column), and simple spike events on contact 5 (middle column). The right column is data recorded by contact 6 and triggered if there was a simple spike event in both contacts 6 and 5 within 1 ms of each other. Lower row shows the waveform for contact 5. On average, spikes in one contact did not produce a significant voltage change in another contact, and the shapes of synchronous (within 1 ms) and non-synchronous spikes were nearly identical. **B.** To measure coordination among the Purkinje cells, we triggered the waveform on contact 6 by the spikes on contact 5. The pattern of coordination is quantified via the conditional probability $Pr(S6(t)|S5(0))$. **C.** Left plot shows the probability of simple spikes at time t, given that a complex spike occurred at time zero in another channel, quantified by $Pr(S5(t)|C6(0))$ and $Pr(S6(t)|C5(0))$. There is no suppression of simple spikes following a complex spike. The plot on the right shows the probability of co-occurrence of the complex spikes, $Pr(C5(t)|C6(0))$. The two cells did not produce complex spikes that were strongly coordinated. Prob., probability.

The left plot of Figure 4.10C illustrates the spike timing properties of these pairs of Purkinje cells. The conditional probability $Pr(S5(t)|C6(0))$ quantifies the effect of complex spikes on contact 6 at time 0 on the simple spikes on contact 5 at time t. It demonstrates that the complex spike on contact 6 did not suppress the simple spikes on contact 5. Similarly, the conditional probability $Pr(S6(t)|C5(0))$ demonstrates that the complex spike on contact 5 did not suppress the simple spikes on contact 6. This provides further evidence that the two nearby contacts picked up two distinct Purkinje cells.

The conditional probability Pr(C5(t)|C6(0)), displayed in the right plot of Figure 4.10C, quantifies the relationship between complex spikes on contacts 5 and 6. This plot illustrates that the complex spikes in these two Purkinje cells were not well coordinated (little or no co-occurrence at 0 ms latency).

In summary, these data from a pair of simultaneously recorded Purkinje cells, cells that were likely less than 50 µm apart, suggested that the cells were not served by the same inferior olive neuron. However, the simple spikes exhibited a measure of coordination, suggesting that they may share common inputs, or that the electrical activity in one cell influenced the activity in the other (Han et al., 2018).

4.3.4. Silicone gel coating of the burr hole

We sealed the burr hole with a silicone-based polymer (DuraGel) that we hoped would provide long-term protection to the dura, while allowing for daily penetrations. Our experience with this product has been positive. Figure 4.11 provides images of the burr hole after drilling and following application of the gel. The gel remained in place for two weeks as we performed daily recordings. At two-week intervals we removed the gel, examined the dura under a microscope, and took samples and performed tissue culture. In every case the results suggested a healthy dura without any evidence of infection. Thus, DuraGel appears compatible with the marmoset dura, providing long-term protection while allowing for daily electrode penetration.
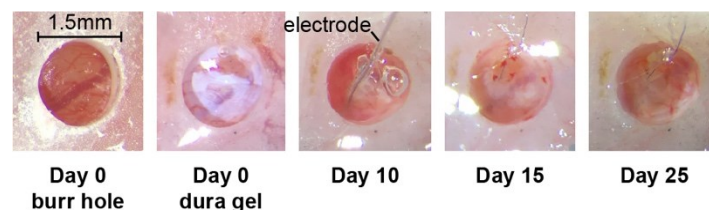
Figure 4.11. Images from the burr hole before (day 0) and after application of the silicone gel.

4.4. Discussion

While the marmoset presents a number of attractive features as a primate model for the study of neural control of social behaviors, verbal communication, and saccadic eye movements, there is concern regarding whether the animals can be motivated to produce a sufficient number of trials. Previous reports on head-restrained marmosets have demonstrated operant conditioned behaviors in auditory tasks (Osmanski et al., 2013; Osmanski & Wang, 2011; Remington et al., 2012) as well as eye movement tasks (Johnston et al., 2018; Mitchell et al., 2014). However, performance has been a concern: around

100 trials per session in the auditory task (Osmanski et al., 2013; Osmanski & Wang, 2011; Remington et al., 2012), 80-100 trials per session in the saccade task (Johnston et al., 2018), and 300-800 trials per session in the eye movement task (Mitchell et al., 2014). Here, we aimed to improve on these results, and found that through a behavioral training protocol, it was possible to motivate the animals to consistently perform around 1000 trials per session.

In addition, in order to investigate the neural basis of motor learning, we aimed to build tools that allowed for neurophysiological investigation of the cerebellum. We began with a CT+MRI based geometric model of the subject's skull and used that data to compute the curvature of the base of a 3D printed titanium head-post and chamber. Using the imaging data, we designed an alignment tool that had cylinders that guided an electrode through the burr hole via an absolute encoder microdrive. This resulted in stable recordings from multiple Purkinje cells during the saccade task.

4.4.1. Behavioral training

In previous reports, marmosets typically underwent a moderate food regulation regime and then were presented with the opportunity to acquire a form of liquid reward such as diluted apple sauce. In our approach, the subjects were gradually trained to receive their main source of nutrition as a consequence of rewarded behavior. Although our experience is limited to only three subjects, the results are encouraging: we were able to maintain the subjects in excellent health, while motivating them to perform the task.

In previous reports, the amount of reward delivered per trial was relatively large: 0.1-0.2 mL per trial (Osmanski et al., 2013; Osmanski & Wang, 2011; Remington et al., 2012), 0.07 mL per trial (Johnston et al., 2018), or 0.05-0.06 mL per trial (Mitchell et al., 2014). Here, we trained the subjects to expect a lower reward rate, 0.02 mL per trial. The sound of the successful trial and the engagement of the food pump indicated food delivery, but the low reward rate encouraged the subjects to wait three or more correct trials before harvesting the accumulated food. This low reward rate may have been an important factor in motivating the subjects to consistently produce a relatively large number of trials per session.

We were able to substantially reduce the training period in subject R by eliminating fixation training and instead starting with pursuit, finding that after only a few days it was possible to begin training on the saccade task, and arrive at around a thousand correct trials within 1.5 months after start of food regulation.

4.4.2. Head-post design

Our head-post design departed significantly from epoxy-based helmet design that was employed in the past two decades in the study of the marmoset auditory cortex (Eliades & Wang, 2008; Lu et al., 2001b; Roy & Wang, 2012; X. Wang et al., 2005). Our design was made possible because of recent advances in 3D printing of titanium, allowing the use of imaging to build a subject-specific head-post.

The X-shaped design of our head-post was in contrast to the halo design recently demonstrated by the Everling group (Johnston et al., 2018) in the marmoset. Each design has its own advantages. The halo design maximizes the open surface of the skull, thus allowing for access to many regions of the brain, which is ideal for simultaneous recording from multiple cortical areas. In contrast, our X-design limits the access to only the posterior regions but has the advantage that the animal's skin can cover the entire head-post base, leaving only the chamber and the head-post pole exposed. Thus, our design allowed the skull to remain largely covered by skin.

4.4.3. Simultaneous recording from multiple Purkinje cells in the marmoset cerebellum

Although the marmoset is roughly the same weight as a rat, its cerebellum is about twice the size (Fujita et al., 2010). Similar to the rat, stripes of different aldolase C expression intensities separate the marmoset cerebellum into longitudinal compartments. However, olivo-cortical and cortico-nuclear projections indicate that the marmoset cerebellum has several compartments that are not present in rodents (compartments in the flocculus, nodulus, and the most lateral hemispheres) (Fujita et al., 2010). In comparison to the macaque, marmoset cerebellum has fewer folia, but a relatively larger vermis. Because the vermis is likely the region critical for saccadic eye movements, as well as control of the tongue during vocalization, the marmoset vermis may be particularly well suited for study of internal models in motor control.

We found that with modern electrodes it was possible to isolate multiple Purkinje cells from the vermis during goal-directed saccades, something that to our knowledge has only once before been reported in the primate (Medina & Lisberger, 2007). We have experimented with tetrodes, heptodes, and high density silicon arrays, finding that each has its own advantages. Tetrodes and heptodes have a 3D geometry that give them the potential to record from multiple Purkinje cells within the same folium, whereas silicon arrays provide the possibility to record from Purkinje cells across folia. In our hands, the 3D geometry of the heptode has provided a robust ability to routinely isolate and stably maintain

recording from pairs of Purkinje cells. These cells often share important functional properties such as complex spike tuning and simple spike coordination (Figure 4.10).

We also found that the computational tools used for identifying Purkinje cells in other animals worked well in the marmoset. Like other animals, the complex spikes in the marmoset Purkinje cells had a waveform that exhibited slow dynamics, making it possible to use frequency-based filtering techniques to tentatively distinguish complex spikes from simple spikes. Furthermore, following a complex spike, there was a 10-15 ms period of simple spike suppression, helping to confirm that a single Purkinje cell produced both spikes.

Finally, we were able to simultaneously record from pairs of Purkinje cells and observe sub-millisecond coordination in their timing. Populations of Purkinje cells can coordinate their spiking, and this coordination may play a critical role in regulating activity of neurons in the deep cerebellar nucleus (Person & Raman, 2011). Perhaps the earliest demonstration of simple spike coordination was by Bell and Grimm (Bell & Grimm, 1969), who made microelectrode recordings in anesthetized cats, finding that Purkinje cells that were located <70 µm apart often fired nearly simultaneously. Later work found further evidence in support of this idea (C I De Zeeuw, Koekkoek, et al., 1997; Ebner & Bloedel, 1981; Heck et al., 2007; Shin & De Schutter, 2006; Wise et al., 2010). For example, Heck et al. (2007) demonstrated that in rats, P-cells within a few hundred microns fired synchronously during the period of reach to grasp of a food pellet, but not during the period after grasp completion.

Here we observed some evidence for sub-millisecond coordination in simple spikes of nearby Purkinje cells during performance of a saccade task. A recent study in mice (Han et al., 2018) also found that firing of many neighboring Purkinje cells exhibited sub-millisecond coordination. That study noted that when chemical synapses were blocked in brain slices, some synchrony persisted. They suggested that sub-millisecond coordination of neighboring Purkinje cells was not via shared synaptic inputs, but via electrical activity that in one Purkinje cell resulted in opening of sodium channels in the neighboring Purkinje cell. The question of whether Purkinje cells in the primate produce coordinated activity during goal-directed movements remains to be explored.

4.4.4. The need to record from Purkinje cells during sensorimotor learning

In order to move accurately, the brain relies on the cerebellum to build internal models and predict sensory consequences of motor commands (Donchin et al., 2012; Golla et al., 2008; Izawa et al., 2012; Maschke et al., 2004; Rabe et al., 2009; Roth et al., 2013; Smith & Shadmehr, 2005). However, it has

been difficult to decipher how the cerebellum learns internal models because the neural encoding of movements has been difficult to decode in the simple spikes of Purkinje cells (Helmchen & Büttner, 1995; Hewitt et al., 2011; Roitman et al., 2005, 2009; Thier et al., 2000). As learning produces a change in the neural coding, this poor understanding has made it difficult to relate the rich behavioral changes that have been observed during motor learning, such as multiple timescales of memory (Kording et al., 2007; Smith et al., 2006), spontaneous recovery (Ethier et al., 2008a), and savings (Pekny et al., 2011), with the neural mechanisms of learning in the cerebellum. Our long-term goal in building this marmoset lab is to search for the neural basis of learning internal models. However, the first step is to understand the neural code with which Purkinje cells make predictions.

Recently, using data in the macaque we proposed that the Purkinje cells may be organized in micro-clusters (Herzfeld et al., 2015, 2018), wherein each micro-cluster is composed of Purkinje cells that share a common preference for error. The preference for error is expressed through the complex spike tuning of the Purkinje cell (Soetedjo et al., 2008; Soetedjo & Fuchs, 2006). We found that if Purkinje cells were organized in this way, the language of internal models as expressed by Purkinje cells became decipherable: simple spikes as a population predicted the real-time motion of the eye during a saccade (Herzfeld et al., 2015). Thus, this theory suggested a potential solution to the encoding problem: organize the simple spikes of a population of Purkinje cells based on the functional properties of each cell's complex spikes.

If the neural representation of internal model relies on a population coding of Purkinje cells, the challenge is to simultaneously record simple and complex spikes from multiple Purkinje cells in the awake, behaving animals. However, this has proven to be difficult. Array electrodes have simultaneously isolated multiple Purkinje cells in rodents (Blenkinsop & Lang, 2011; Sugihara et al., 2007; Tang et al., 2016, 2019; Welsh et al., 1995) and decerebrated cats (Ebner & Bloedel, 1981), but this has often been in the anesthetized animal (however, see Han et al. (2018)). In the awake behaving subject, calcium imaging has been used to track complex spikes in multiple Purkinje cells (Kostadinov et al., 2019; Najafi et al., 2014), but this technique does not have the temporal resolution yet to track simple spikes. In primates, simultaneous recordings from Purkinje cells are quite rare. Until our work, the report by Medina and Lisberger (Medina & Lisberger, 2007) is the only example we are aware of in which pairs of Purkinje cells were recorded in the primate during goal-directed behavior.

4.4.5. Limitations

Although the X-shaped head-post design had the advantage that the skin would cover nearly the entire skull, we found that within about a year following surgery the skin in subjects B and M gradually receded to the boundary of the dental cement. One of our challenges is to find ways to slow and eventually prevent this process of skin contraction, which might include experimentation with different forms of cement or its elimination.

Our experience with array type electrodes that have contacts along a single surface suggests that despite their high density, these electrodes can provide isolation of multiple Purkinje cells along neighboring folia, but rarely within a single folium. The goal of simultaneously recording from many Purkinje cells within a single folium will likely require employment of multiple electrodes.

Isolation of multiple Purkinje cells introduces the computational problem of spike attribution, as one contact can record complex spikes while the other records only simple spikes. Here, we showed that conditional probabilities that consider the effect of complex spikes on simple spikes can help with the attribution problem, but the task becomes much harder when many contacts are involved.

In summary, in order to better understand the neural basis of internal models, we have built a new laboratory that focuses on marmosets. Here, we reported on our attempts to improve marmoset behavioral training and electrophysiological recording methods. We found that through behavioral shaping, it was possible to motivate the animals to perform around 1000 trials in a head-fixed saccade task. We also found that with the aid of imaging, it was possible to build a system for simultaneous isolation of multiple Purkinje cells in the cerebellum.

Chapter 5. P-sort: an open-source software for cerebellar neurophysiology

Analysis of electrophysiological data from Purkinje cells (P-cells) of the cerebellum presents unique challenges to spike sorting. Complex spikes have waveforms that vary significantly from one event to the next, raising the problem of misidentification. Even when complex spikes are detected correctly, the simple spikes may belong to a different P-cell, raising the danger of misattribution. To address these identification and attribution problems, we wrote an open-source, semi-automated software called P-sort, and then tested it by analyzing data from P-cells recorded in three species: marmosets, macaques, and mice. Like other sorting software, P-sort relies on nonlinear dimensionality reduction to cluster spikes. However, it also uses the statistical relationship between simple and complex spikes to merge disparate clusters and split a single cluster. In comparison with expert manual curation, occasionally P-sort identified significantly more complex spikes, as well as prevented misattribution of clusters. Three existing automatic sorters performed less well, particularly for identification of complex spikes. To improve development of analysis tools for the cerebellum, we provide labeled data for 313 recording sessions, as well as statistical characteristics of waveforms and firing patterns of P-cells in three species.

5.1. Introduction

Recording neuronal activity from the cerebellum presents both opportunities and challenges. The principal cells of the cerebellum, Purkinje cells (P-cells), can be identified based on their unique electrophysiological properties. Among cells in the cerebellum only P-cells produce simple and complex spikes (Thach, 1967). This makes it possible to use statistical methods to identify a P-cell: generation of a complex spike should be followed by a pause in the simple spikes (Eccles et al., 1966; Sato et al., 1992). However, detection of complex spikes is difficult because these spikes are not only rare, their waveforms change dramatically within 25 microns of the P-cell soma. Whereas in the dendritic field of a P-cell the complex spike is a slow, negative waveform, as the electrode approaches the soma the slow waveform becomes positive, develops a sharp initial negative spike, and is then followed by a variable number of spikelets that alter the waveform from one spike to the next (Gao et al., 2012; Han et al., 2020; Monsivais et al., 2005). Thus, it is common to detect the simple spikes but not the complex spikes, or alternatively, detect the complex spikes but later realize that they are not followed with simple spike suppression and therefore do not belong to the same P-cell. To address these issues, we developed a spike analysis software that aids detection of simple and complex spikes, as well as quantifies whether the two events are generated by a single P-cell.

Unlike simple spikes, complex spikes tend to have their greatest power in the low-frequency range (30-800 Hz). As a result, a typical complex spike can produce a "broad spike" in the low-pass

filtered representation of the data (local field potential, LFP). Indeed, two recent developments in complex spike detection are novel algorithms that depend partly on the LFP waveform (Markanday et al., 2020; Zur & Joshua, 2019). Once the simple and complex spikes are labeled, the final step is to determine whether the simple spikes have been suppressed after a complex spike. If so, then one may conclude that the two kinds of spikes were generated by a single P-cell. However, in some data sets complex spikes do not have an LFP signature. Moreover, even if the complex spikes are detected, some or all of the simple spikes may belong to a different P-cell, or even a non P-cell.

As a result, the problem is twofold: in the identification step, we need to label the simple and complex spikes, whereas in the attribution step, we need to determine which group of complex spikes was generated by the P-cell that produced a particular cluster of simple spikes. To consider these challenges, we formed a collaboration that included laboratories that focused on marmosets, macaques, and mice. Our software was developed using a database of over 300 P-cells recorded in three species.

The diversity of species and recording electrodes helped us identify some of the critical issues that are present in cerebellar electrophysiology. The presence of experts from the various laboratories provided a diversity of opinions, helping us verify the algorithms, as well as highlight their limitations. Here we report the results of this effort.

P-sort is an open-source, Python-based software that runs on Windows, MacOS, and Linux platforms. To cluster waveforms and identify simple and complex spikes, P-sort uses both a linear dimensionality reduction algorithm and a recently developed nonlinear algorithm called UMAP (Uniform Manifold Approximation and Projection) (McInnes et al., 2018). Importantly, it quantifies the probabilistic interaction between complex and simple spikes, providing an objective measure that can split a single cluster, or merge two different clusters, despite similarities or differences in their waveforms. Thus, P-sort helps the user go beyond waveforms to improve labeling of spikes.

However, a limitation of P-sort is that it requires user interaction. To encourage development of automated algorithms for the cerebellum, and to allow others to sort the same datasets and compare their results with P-sort, here we provide a large database of labeled spikes from all three species. P-sort's source code is available at https://github.com/esedaghatnejad/psort. The labeled data are available at https://doi.org/10.17605/osf.io/gjdm4. Instructional videos that introduce the user to P-sort are available at https://www.youtube.com/channel/UCSifaCZt1HmnZJ3440AOf8g/videos.

5.2. Methods

5.2.1. Subjects

Neurophysiological data were collected from two marmosets (Callithrix Jacchus, male and female, 350-370 g, subjects M and R, 4 years old), six rhesus macaques (Macaca mulatta; males; 5.0-7.4 kg, subjects B, F, K, P, S, and W), and ten mice (Mus musculus; male C57BL/6J mice at least 12 weeks of age, subjects N082, N083, N086, N089, T029, T052, T083, T101, T122, T124).

The marmosets were born and raised in a colony that Prof. Xiaoqin Wang has maintained at the Johns Hopkins School of Medicine since 1996. The procedures on the marmosets were evaluated and approved by Johns Hopkins University Animal Care and Use Committee in compliance with the guidelines of the United States National Institutes of Health.

The procedures on the macaques were performed in accordance with the Guide for the Care and Use of Laboratory Animals (2010) and exceeded the minimal requirements recommended by the Institute of Laboratory Animal Resources and the Association for Assessment and Accreditation of Laboratory Animal Care International. The procedures were approved by the local Animal Care and Use Committee at the University of Washington.

The procedures on the mice were approved by the Baylor College of Medicine Institutional Animal Care and Use Committee based on the guidelines of the US National Institutes of Health. The experimental mice were singly housed on a reverse light/dark cycle (8:00 lights-off to 20:00 lights-on).

5.2.2. Marmoset data acquisition

Following recovery from head-post implantation surgery, the new world monkeys were trained to make saccades to visual targets and rewarded with a mixture of applesauce and lab diet (Sedaghat-Nejad, Herzfeld, Hage, et al., 2019). Using MRI and CT imaging data for each animal, we designed an alignment system that defined trajectories from the burr hole to various locations in the cerebellar vermis, including points in lobule VI and VII. 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.

We recorded from the cerebellum using three types of electrodes: quartz insulated 4 fiber (tetrode) or 7 fiber (heptode) metal core (platinum/tungsten 95/05) electrodes (Thomas Recording), and 64 contact high density silicon probes (Cambridge Neurotech). We connected each electrode to a 32 or

64 channel head stage amplifier and digitizer (Intan Technologies, USA), and then connected the head stage to a communication system (RHD2000, Intan Technologies, USA). Data were sampled at 30 kHz and band-pass filtered (2.5 Hz to 7.6 kHz). We computed a common average reference signal (median of all simultaneously recorded channels, computed at 30 kHz) and subtracted this signal from all channels. We used OpenEphys (Siegle et al., 2017), an open-source extracellular electrophysiology data acquisition software, for interfacing with the RHD2000 system and recording of signals.

5.2.3. Macaque data acquisition

The data were collected during previous studies (Kojima et al., 2010a, 2010b; Soetedjo et al., 2008). Following recovery from surgery, the monkeys were trained to make saccades to visual targets and rewarded with applesauce. A recording chamber was implanted on the midline of the cerebellum (14.5mm posterior of the interaural axis and directed straight down), providing access to the oculomotor vermis (lobule VI and VII). Single-unit activity was recorded with homemade tungsten electrodes with an iron-particle coating (100 kΩ impedance at 1 kHz). Neurophysiology data was sampled at 50 kHz by a Power 1401 digitizer (Cambridge Electronic Design, Cambridge, UK) and subsequently band-pass filtered (30 Hz to 10 kHz). Data were displayed in real-time on a computer monitor running the data acquisition and analysis software Spike2 (Cambridge Electronic Design, UK) and saved for offline analysis (Soetedjo & Fuchs, 2006).

5.2.4. Mice data acquisition

Single-unit extracellular recording was performed as previously described (Heiney et al., 2018). In brief, a 2-3 mm diameter craniotomy was opened over the right side of the cerebellum (6.5 mm posterior and 2.0 mm lateral from bregma) to access lobule V and the eyeblink microzone, and the dura was protected by a layer of Kwik-Sil (WPI). A custom 3D printed recording chamber and interlocking lid (NeuroNexus) was secured over the craniotomy with dental acrylic to provide additional protection. After 5 days of recovery, the mouse was fixed in place on a treadmill via a previously implanted headplate, and Purkinje cell simple spikes and complex spikes were isolated using a tetrode (Thomas Recording, AN000968) acutely driven into the cerebellar cortex with microdrives mounted on a stereotactic frame (Narishige MMO-220A and SMM-100). The voltage signal was acquired at a 24,414 Hz sampling rate, and band-pass filtered between 100 Hz to 10 kHz (AP channel) and between 2 Hz to 300 Hz (LFP channel) using an integrated Tucker-Davis Technologies and MATLAB system (TDT RZ5, medusa, RPVdsEx) running custom

code (https://github.com/blinklab/neuroblinks). The data include Purkinje cells from a previously collected dataset (Achilly et al., 2021; Ohmae & Medina, 2015).

5.2.5. P-sort main window

To allow P-sort to run on Windows, MacOS, and Linux, the code was written using Python-based libraries (Behnel et al., 2011; Harris et al., 2020; Pedregosa et al., 2011; Virtanen et al., 2020). The GUI was written using PyQt5 (The Qt Company and Riverbank Computing Ltd.) and PyQtGraph to provide a fast and intuitive interaction for the user. To facilitate further development of P-sort by the user community, we used object-oriented coding. P-sort's source code is available for download here https://github.com/esedaghatnejad/psort and instructional videos are provided here https://www.youtube.com/channel/UCSifaCZt1HmnZJ3440AOf8g/videos.

A process starts by loading the data and dividing it into one or more periods of time (called slots). The slot framework helps the user to account for potential drift and fluctuation in spike quality and shape over time. After sorting one slot, the parameters and spike waveform templates will be copied to the next slot to facilitate the sorting, but the user can further change the parameters independently in each slot.

The sorting process starts by filtering the signal into two channels, Action Potential (AP) and Local Field Potential (LFP). The default is a 4th order Butterworth filter with the 50-5000 and 10-200 Hz range for AP and LFP channels, respectively. However, these parameters can be modified using the GUI to better fit the specific conditions of the data. The default assumption is that simple spikes generate negative peaks in AP channel and complex spikes generate positive peaks in the LFP channel. However, this assumption can be changed via the GUI. Once the respective peaks are detected, the next question is what should be the threshold to reject or accept a peak as being a potential spike. P-sort computes the histogram of the peaks and fits a Gaussian Mixture Model (GMM) with two basis functions to the histogram for each channel. The lower bound Gaussian is considered the noise and the upper bound Gaussian is the signal of interest. GMM assigns query data points to the multivariate normal components that maximize the component posterior probability, given the data:

$$p(\theta|x) = \sum_{i=1}^{K} \tilde{\phi}_i \mathcal{N}(\tilde{\mu}_i, \tilde{\Sigma}_i)$$ Equation 5.1

where $K = 2$, and i^{th} basis function is characterized by a normal distribution with weight $\tilde{\phi}$, mean $\tilde{\mu}$, and covariance matrix $\tilde{\Sigma}$.

The lower bound Gaussian is considered the noise and the upper bound Gaussian is the signal of interest. Based on this assumption, the intersection of the two fitted Gaussians is used as the default threshold to prune the detected peaks. However, the user is provided with a GUI to manually change the SS and CS thresholds, either by using the interactive plots, or directly by assigning their values.

The next question is how to relate a peak in the LFP channel, which may potentially be a complex spike, to its waveform in the AP channel. In a typical recording, a CS waveform consists of an initial sharp negative peak and a broad positive bump. The peak in LFP happens due to the broad positive bump but its timing is variable. Thus, using the LFP peak to align the CS waveform is unreliable. P-sort provides three different methods to find the onset of the complex spikes: SS index, SS template, and CS template. Initially, it uses the detected sharp negative peak (SS index) to align CS waveforms. This provides a set of putative CS spike waveforms to calculate a CS template (average of the waveforms). However, due to variability in CS waveforms, not all CSs express the initial sharp negative peak. To address this problem, after forming a CS template, the user can redetect CS and use CS template to align them which is not sensitive to the lack of initial sharp peak.

For the template alignment, P-sort uses the LFP peak time point as the starting point. However, as mentioned above, this time point is variable from one spike to the next. To account for variability in LPF peak time point, P-sort moves the CS template signal (3.5 ms in duration) along a 5 ms time window around the LFP peak (4 ms before and 1 ms after the LFP peak) on the AP signal and selects the point of time which results in maximum correlation between the two signals. Correlation between two signals is computed as:

$$C_{TS}(t)|_{t=-4ms}^{t=+1ms} = \sum_{t'=0}^{t'=+3.5ms} T(t').S(t+t') \qquad \text{Equation 5.2}$$

where t is time point relative to LFP peak, T is the template signal, and S is the AP channel signal. Optimum time (t^*) is when the $C_{TS}(t)$ is maximum.

For the recordings in which the LFP peak is later than 4 ms after the sharp negative peak, this default value should be adjusted using the Preferences interface. Alignment of the simple spikes relies on the timing of the sharp peak value of the waveform.

P-sort ensures that a candidate spike is labeled as either a simple spike, or a complex spike, but not both. Moreover, due to the refractory period in a spiking neuron, two arbitrary spikes cannot happen closer than 0.5 ms with respect to each other. Based on these constraints, P-sort addresses

potential conflicts between CS-CS, CS-SS, and SS-SS candid events. The default values for each scenario can be changed using the Preferences interface.

After resolving the potential conflicts, P-sort provides sets of potential simple and complex spikes. For each set of spikes, P-sort represents spike waveforms, instant-firing-rate distribution, peak distribution, conditional probabilities, and feature scatter plots. Numerous features can be used for clustering of these data, including UMAP, principal components, timing of the spikes, relative time with respect to next or previous spike, similarity to templates, peak value, and instantaneous firing rate. Using the interactive plots, the user can select subset of the spikes based on the waveform plot or feature scatter plot and further prune the data or even change their label from simple to complex spikes and vice versa. As these clusters are manipulated, P-sort provides immediate feedback on their statistical features, thus allowing the user to determine whether the simple and complex spikes are likely to have been generated by a single P-cell, and whether the latest manipulations of the clusters improved the probabilities.

Overall, P-sort's main window aims to provide a balance between the ability to visualize each spike waveform, and the ability to cluster the spikes and visualize their interactions. From this main window P-sort branches into two additional windows: the Cluster Module, and the Dissect Module.

5.2.6. Cluster Module

The Cluster Module gives the user the ability to use various algorithms to form groups of spike waveforms, assign labels to each waveform cluster (or part of a cluster), and immediately assess the statistical interactions between the labeled groups. Cluster Module includes manual and automatic labelling tools to label datasets based on their features or waveforms. This is the key feature of P-sort that allows the user to go beyond spike waveforms and rely on statistical interactions between spikes to merge seemingly different clusters or split a single cluster into two.

P-sort relies on various algorithms such as UMAP, principal component analysis, etc., to extract features from spike waveforms and reduce their dimensionality. UMAP uses a manifold learning approach for dimensionality reduction. The algorithm performs the reduction in two steps. First a weighted neighbor graph is formed based on the local distances of the samples in high dimensional space. The weights represent the probability of the two samples being connected in high dimensional space. Then an optimization step finds a low dimensional layout of data that minimizes the cross-

entropy between the two representations. For a complete description of the UMAP algorithm, please see (McInnes et al., 2018).

One of the important aspects of UMAP optimization is the trade-off between preserving the global structure or forming local clusters. In contrast, principal component analysis preserves the pairwise distances amongst all samples. UMAP uses two hyper-parameters to tune the aforementioned trade-off, number of nearest neighbors (n_neighbors) and minimum distance (min_dist). During the first step of forming the high dimensional graph, n_neighbors was used to determine the local distance based on nth nearest neighbor of each data point. Hence, smaller n_neighbors values lead to a more local view of the data and loss of the global structure. On the other hand, min_dist determines the density of the points in the low dimensional space. This hyper-parameter affects the density of the clusters with higher values leading to denser representations and lower values resulting in more spread points. We used default values of min_dist=0.1 and n_neighbors=15 for our application. For further information regarding the comparison of UMAP algoritm with other dimensionality reduction techniques please visit https://umap-learn.readthedocs.io/en/latest/performance.html#.

Once these features are computed, P-sort provides various algorithms to automatically cluster and label the spikes. The statistical interaction between the labeled spike groups is the key information that helps the user confirm the identity of the simple and complex spikes, and determine whether the two belong to a single P-cell.

The clustering algorithms used by P-sort include: (1) Gaussian Mixture Models (GMM), which requires the user to specify the number of clusters and their initial centroids, (2) Iso-split algorithm (Chung et al., 2017; Magland & Barnett, 2015), which is automatic and determines the number of clusters based on bimodality tests, and (3) HDBSCAN algorithm (Campello et al., 2013), which is also automatic and requires no user inputs. We implemented a post-processing layer for HDBSCAN's outputs and restricted the number of clusters to less than 10. We did this by setting extra clusters with least number of members as noise (assigned as label -1). In addition to automatic clustering algorithms, an outlier detection method was implemented based on Local Outlier Factor density (Pedregosa et al., 2011) which receives the quantile threshold as input. All these algorithms use the selected elements of the feature scatter plot by default; however, the user can perform multi-dimensional clustering by selecting the multiple features from the feature list in the GUI.

To use the cluster module, the user will select the spikes of interest either from the feature scatter plot or the waveform plot and assign a label to the selection. Cluster Module provides four

interactive subpanels for representing (1) color-coded feature scatter plot of spikes, (2) spike waveforms of each label, (3) peak histogram of each label including the firing rates information, and (4) cross and auto correlograms of chosen labels. Once the user has labeled a group of putative simple spikes, the module provides the probability $Pr(S(t)|S(0))$, thus allowing the user to visualize the refractory period of the chosen spikes. When the user labels a group of putative complex spikes, the module provides the probability $Pr(S(t)|C(0))$, thus allowing the user to visualize the suppression period. groups of spikes and assigned SS-1, SS-2, and CS-1 labels to the clusters. Then, the user can address the attribution problem of potential simple spikes with the candidate complex spike by checking the correlogram plots.

A unique challenge in cerebellar neurophysiology is finding the simple and complex spike clusters that belong to a single P-cell. It is possible that on certain recordings, one or more neurons will contribute to the signals that are recorded by a single contact. For example, it is possible that the large amplitude simple spikes are not produced by the P-cell that has produced the complex spikes in the recording. Rather, the smaller amplitude simple spikes should be paired with the complex spikes.

5.2.7. Dissect Module

P-sort dissect module is designed to provide more tools for reviewing individual spikes. In some scenarios, looking at the individual spikes and their surrounding events provides a better insight than the average features. Dissect module provide a tool set to move between spike events and look at each one over time. This module also provides the tool to manually overwrite a spike or change its alignment.

5.2.8. Comparison with expert curation

50 sessions (34 in the macaque and 16 in mouse) were sorted using Spike2 (Cambridge Electronic Design, UK) and P-sort by different experts. For Spike2 sorted data, we ensured that a given spike was not labeled as both a CS and a SS and removed the overlapping labels. Next, we compared P-sort data with Spike2 data by finding shared complex/simple spikes that happened in a 0.5/5.0 ms window of each other. We used the window of time to account for the variability in alignments of Spike2 data due to lack of template matching. If a given spike was not shared between the two datasets, it was labeled as an exclusive spike. In order to compare the number of exclusive spikes over datasets, we normalized the number of the exclusive spikes by the total number of the spikes in each dataset.

5.2.9. Comparison with automatic sorters

We quantified performance of various automatic sorting algorithms on three different cerebellar data, ranging from easy to medium to hard. The hard data set was named the "P-sort challenge". Each dataset included around 15 minutes of recording from the marmoset oculomotor vermis and contained some of the difficulties that are present in cerebellar neurophysiology, including diverse patterns of spikes from neighboring cells. Thus, sorting of the data required isolating different spike types, as well as addressing the CS-SS attribution problems. The same 0.5/5.0 ms window was used to detect shared and exclusive spikes between automatic sorters and P-sort data.

For automatic sorters, we asked the authors of some of the algorithms to sort the data. P. Yger sorted the data using SpyKING CIRCUS (Yger et al., 2018), and A. Markanday sorted the data using a neural network (Markanday et al., 2020). In addition, we used Mountainsort4 (Chung et al., 2017) and KiloSort2 (Pachitariu et al., 2016) to sort the same data. We post-processed the outputs by merging and associating simple spikes, complex spikes, and multi-unit activity (MUA) based on rates and cross correlograms to arrive at the best match. For Mountainsort4, we used the default parameters and then the output units were manually merged and labelled as SS1, SS2, CS1, CS2, and MUA. Similarly, for Kilosort2, we used the default parameters and manually merged and labeled the outputs using the Phy2-interface correlograms and rates.

We quantified performance of each algorithm by computing a spike "agreement score" (Buccino et al., 2020). This score measured the percentage of spikes that were agreed upon between pairs of algorithms. This involved computing the number of matched spikes (using a 1 ms time window), divided by the total number of spikes reported by each sorter, minus the number of matches. Using the same score we quantified reproducibility of P-sort by having five experts sort the same data sets that were sorted by the automated algorithms.

5.3. Results

To illustrate the variety of challenges that we face in cerebellar neurophysiology, consider the data shown in Figure 5.1. Here, the LFP channel (10-200 Hz) is plotted in red, and the AP channel (50-5000 Hz) is plotted in black. Occasionally, one is lucky enough to isolate a P-cell that exhibits easily identifiable complex and simple spikes (Figure 5.1A). In this example, LFP shows a large positive peak for the complex spike. To confirm that the complex and simple spikes originate from the same cell, we compute the conditional probabilities Pr(S(t)|C(0)), and Pr(S(t)|S(0)), over a domain of ±50 ms. The term

Pr(S(t)|C(0)) indicates the probability that a simple spike occurred at time t, given that a complex spike was generated at time zero. The term Pr(S(t)|S(0)) is the probability of a simple spike at time t, given that another simple spike was generated at time zero. Spikes that originate from a single cell produce a refractory period. Thus, Pr(S(t)|S(0)) exhibits a near zero probability period centered at time zero. On the other hand, a complex spike is followed by a pause in the simple spikes. As a result, Pr(S(t)|C(0)) is asymmetric, with a long period of near zero simple spike probability following the time point zero. The presence of simple spike suppression following a complex spike, as shown in Figure 5.1A (right panel), confirms that these two groups of spikes are generated by the same P-cell.

Figure 5.1B presents a more challenging example. Here, the complex spikes do not have an LFP signature. However, analysis of the AP channel using UMAP (McInnes et al., 2018) identifies potential complex spikes. The identified events are genuine complex spikes, as evidenced by their spikelets (more examples from the same cell are shown in Figure 5.2A), and the fact that the simple spikes pause after a complex spike (Figure 5.1B, right panel).

If the electrode is near the soma, the complex spike waveform can exhibit spikelets that are similar to simple spikes (Burroughs et al., 2017; Davie et al., 2008; Khaliq & Raman, 2005; Monsivais et al., 2005; Schmolesky et al., 2002; Warnaar et al., 2015). For example, the complex spike in Figure 5.1B exhibits large spikelets. These spikelets are difficult to dissociate from ordinary simple spikes, as evidenced by the fact that Pr(S(t)|C(0)) shows a small non-zero probability between 0 and +10 ms (right panel of Figure 5.1B), during the period in which we would expect a near complete suppression.

Another example of the diversity of complex spike waveforms is shown in Figure 5.1C. In this case, the complex spike exhibits a negative LFP peak, which may indicate that the electrode is in the dendritic field of the P-cell (Han et al., 2020). Nevertheless, once the complex spikes are correctly identified, the simple spikes show a period of suppression.

While detection of complex spikes may be challenging, a more crucial problem is attribution: sometimes the prominent groups of simple and complex spikes are generated by different P-cells. An example of this problem is shown in Figure 5.1D. Here, the LFP signal easily provides for the detection of complex spikes. However, there are two groups of simple spikes, SS1 and SS2. The spikes labeled SS2 are the larger amplitude events. Notably, the complex spikes do not produce suppression of SS2 spikes. Rather, the smaller amplitude events SS1 are the appropriate simple spikes that pair with the complex spikes.
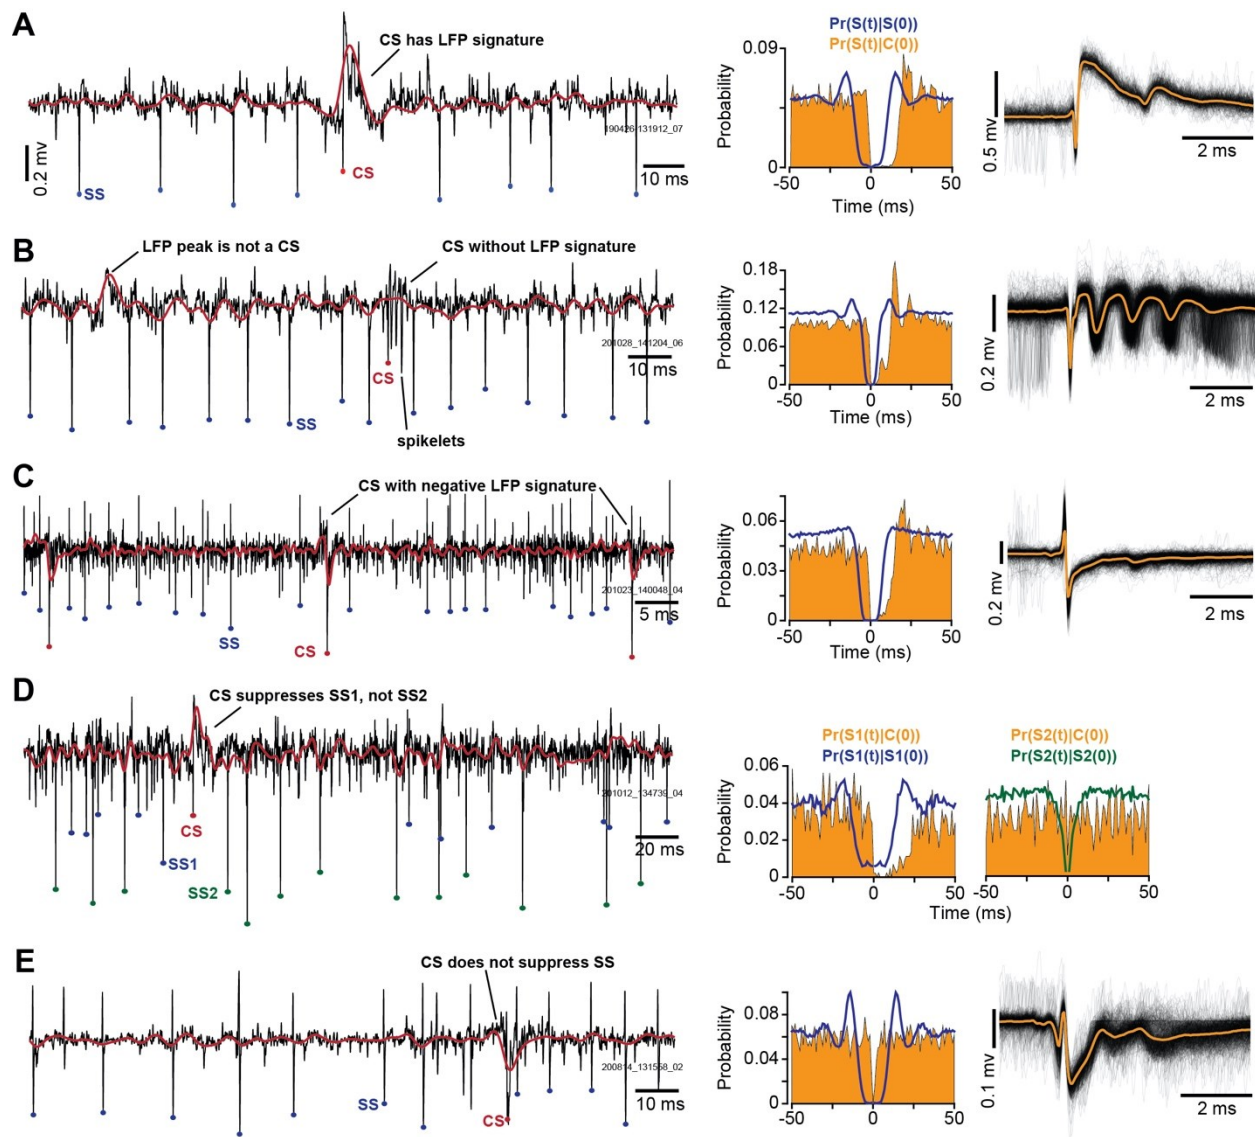
Figure 5.1. Examples of challenges in cerebellar spike sorting.

In the left column, LFP channel (10-200 Hz) is plotted in red, and the AP channel (50-5000 Hz) is plotted in black. The middle column displays the conditional probability of a simple spike at time t, given that a complex spike occurred at time zero, labeled as Pr(S(t)|C(0)). Note the asymmetric suppression following a complex spike. The conditional probability Pr(S(t)|S(0)) indicates the probability of a simple spike at time t, given that another simple spike occurred at time zero. The right column includes individual complex spike traces, as well as the average trace. **A.** In this recording, complex spikes have a positive LFP peak. **B.** Here, complex spikes cannot be identified from their LFP waveform as they lack an LFP signature. **C.** In this recording, complex spikes have a negative LFP peak. **D.** Here, complex spikes coincided with the suppression of one group of simple spikes (SS1), but not a second group (SS2). The probability pattern (right column) suggests that SS2 is likely not a P-cell. **E.** In this recording, complex spikes do not coincide with suppression of the simple spikes. Thus, the two groups of spikes are not generated by the same P-cell. Bin size is 1 ms for the probability plots.
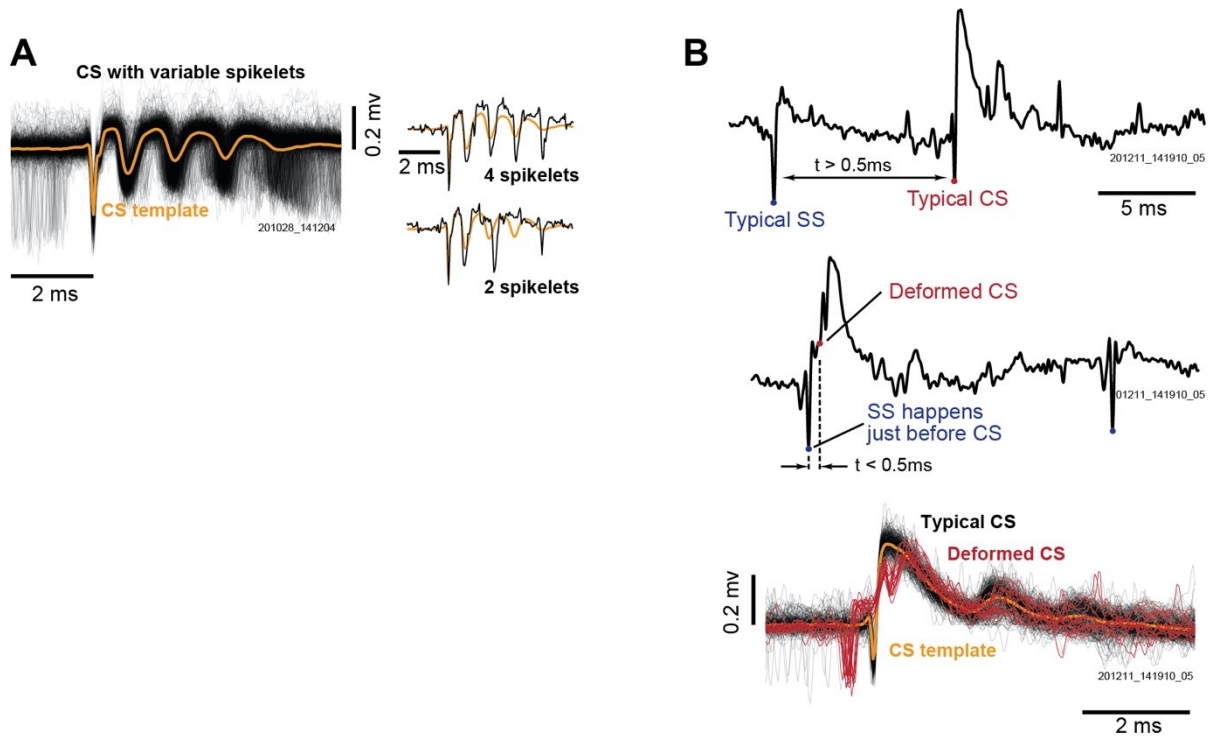
Figure 5.2. Additional challenges in identification of complex spikes.
A. Complex spike waveforms vary because of the timing and number of spikelets. **B.** Complex spike waveforms vary because of the temporal proximity of a simple spike. Here, a simple spike before onset of a complex spike distorts the complex spike waveform. Top and middle rows: simple spike occurs at 3 or 0.2 ms before a complex spike. Bottom row: red traces are complex spikes with deformed waveforms.

A different form of the attribution problem is shown in Figure 5.1E. In this case, the simple spikes are easily identified, and the complex spikes have a negative LFP. Remarkably, despite the excellent isolation, the simple spikes are not suppressed following the complex spikes. Rather, the two groups of spikes in this recording are generated by two different P-cells.

Identification of complex spikes suffers from additional problems. There are variable number of spikelets in the complex spike waveform (Burroughs et al., 2017; Davie et al., 2008; Ito & Simpson, 1971; Monsivais et al., 2005; Yang & Lisberger, 2014b), and thus template matching may have difficulty labeling all complex spikes. Examples of the variable spikelets are shown in Figure 5.2A. Paradoxically, the when the electrode is near the P-cell's soma, the signal strength is high but there is also a large impact of waveform variations caused by spikelets.

Finally, the complex spike waveform can be distorted because of the proximity of a simple spike, as shown in Figure 5.2B (Markanday et al., 2020; Servais et al., 2004). This is because simple and complex spikes are driven by different inputs to the P-cell. Simple spikes are produced spontaneously, as

well as due to excitatory inputs from granule cells. On the other hand, complex spikes are produced because of the input from the climbing fiber. As a result, a simple spike can be generated up to a fraction of a millisecond before a complex spike (middle plot of Figure 5.2B). Here, the complex spike that follows the simple spike at short latency lacks the sharp component that initiated more typical complex spikes. As a result, the complex spike waveform is distorted by the proximity of the simple spike (lower plot, Figure 5.2B).

In summary, identification of complex spikes is difficult because they can lack an LFP signature, their waveforms can be distorted because of nearby simple spikes, or their waveforms can incorporate variable number of spikelets. After the simple and complex spikes are identified, one still faces the problem of finding the simple spike cluster that belongs to the P-cell that generated the complex spikes. We wrote the open-source tool P-sort to help with these identification and attribution problems.

5.3.1. Clustering waveforms

The diversity of complex spike waveforms suggests that it may be difficult to find a single mathematical technique that could generalize to all situations. For example, in some cases it is possible to identify the complex spikes from the LFP channel (Figure 5.1A), whereas in other cases it is necessary to search the AP channel (Figure 5.1B). In the case where two types of simple spikes are present, often the larger amplitude simple spikes is suppressed after the complex spikes (Figure 5.1C), but occasionally the smaller amplitude spikes are the correct choice (Figure 5.1D). P-sort provides tools for clustering as well as hypothesis testing. The tools include traditional dimensionality reduction methods such as principal component analysis (PCA), as well as novel algorithms such as UMAP. As the user identifies putative groups of simple and complex spikes, the software provides statistical feedback regarding the probability that the spikes are from the same P-cell.

P-sort works best with the raw, broad-spectrum recording such as the data generated by Open Ephys, an open-source data acquisition software (Siegle et al., 2017). However, P-sort can also import data in which the LFP and AP channels are acquired separately. Upon loading the data, the user is provided with a GUI to specify how the data should be chunked into "slots". Each slot is a region of data that is analyzed in turn, but once a slot is analyzed, other subsequent slots inherit features such as spike templates and threshold. In case of broad-spectrum data, the user can specify the filter properties for the LFP and AP channels (Figure 5.3, part 1). P-sort automatically selects a threshold for each channel

(Figure 5.3, part 2, see Methods) and displays the statistics of the resulting simple and complex spikes (Figure 5.3, parts 3 and 4).

In the example shown in Figure 5.3, P-sort automatically selected thresholds for LFP and AP channels and identified putative simple and complex spikes. The UMAP space indicated a single simple spike cluster (Figure 5.3, part 3), which was confirmed by Pr(S(t)|S(0)), illustrating that the simple spikes had around 10 ms suppression period (centered at zero). However, the UMAP space was not homogeneous for the complex spikes (Figure 5.3, part 4), meaning that there was variability in the waveforms. Regardless, Pr(S(t)|C(0)) exhibited a suppression period. Thus, there was statistical confirmation that the simple spikes were well isolated, and were suppressed after the complex spikes.
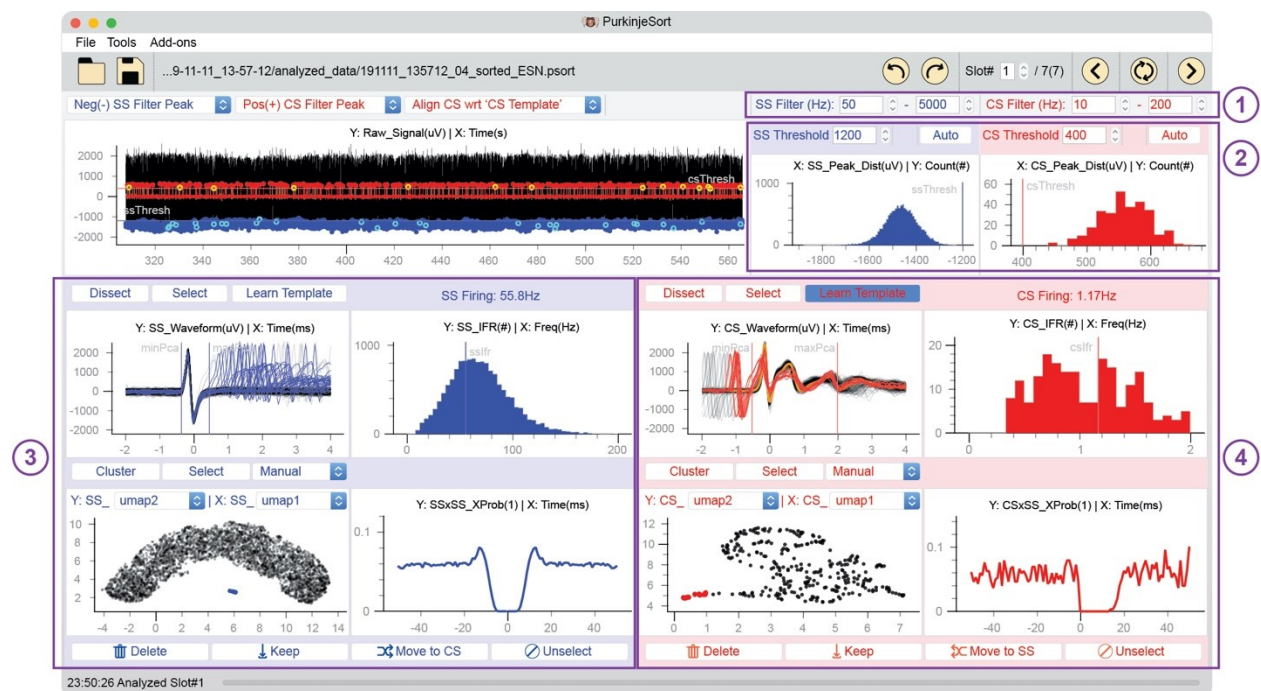
Figure 5.3. Main window of P-sort.
1. Specification of filter properties for the AP and LFP channels. **2.** Threshold selection and the distribution of voltage peaks for events detected in the LFP and AP channels. **3.** Identification of simple spikes. The windows include spike waveforms aligned to peak voltage, distribution of instantaneous firing frequency, conditional probability Pr(S(t)|S(0)) demonstrating a suppression period, and clustering using UMAP. **4.** Identification of complex spikes. The windows include spike waveforms aligned to a learned template (yellow), distribution of instantaneous firing frequency, conditional probability Pr(S(t)|C(0)) demonstrating suppression of simple spikes, and clustering using UMAP. The red traces indicate complex spikes that have been distorted because of nearby simple spikes.

Of course, in most cases the data are not as easily sorted as the case shown in Figure 5.3. Another frequent case is one in which the complex spikes do not have an LFP signature (Figure 5.4A),

and thus one must search the AP channel. In the recording shown in Figure 5.4C, the simple and complex spike waveforms happen to be quite similar. As a result, in the PCA space the data present a single cluster (Figure 5.4B, left subplot). Thus, if we were to rely on PCA alone, we might conclude that only simple spikes are present.
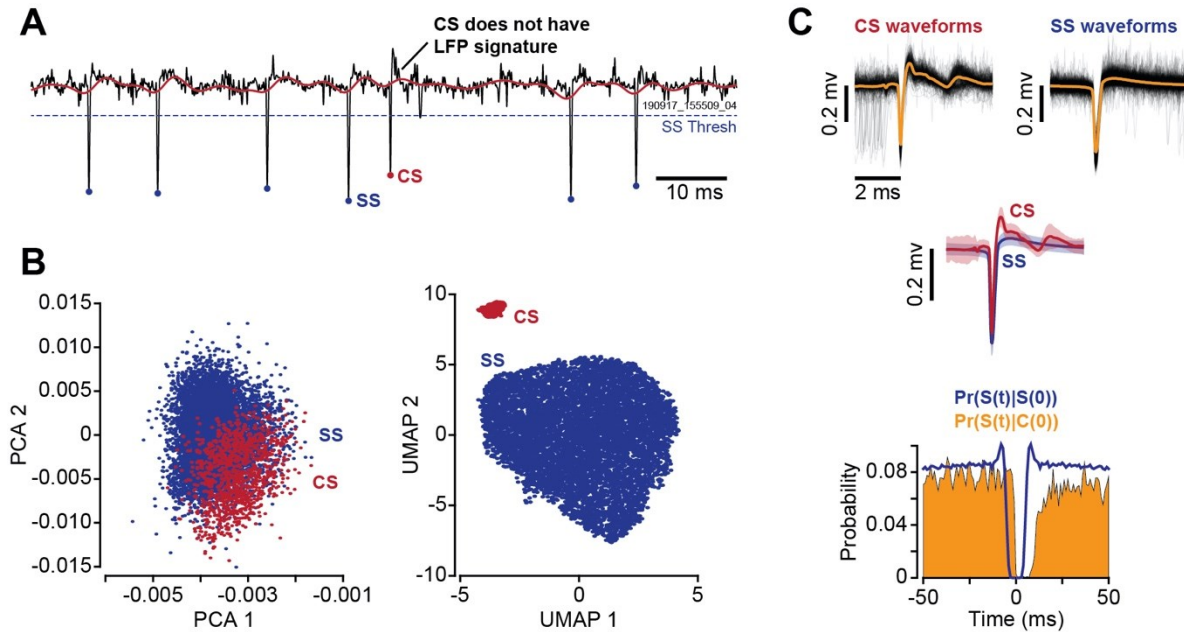
Figure 5.4. UMAP dissociates simple and complex spikes.
A. In this recording, complex spikes do not exhibit an LFP signature (red trace). **B.** Clustering of the spikes in PCA space does not produce a clear separation. However, the two groups of spikes separate in the UMAP space. The complex spikes identified by UMAP are shown in red in the PCA space. **C.** The waveforms and average traces for the complex spike and simple spike clusters as identified by UMAP. The conditional probabilities demonstrate that the complex spike events coincide with suppression of the simple spikes, suggesting that the two groups of spikes are likely generated by the same P-cell. Error bars are standard deviation.

P-sort relies on a dimensionality reduction algorithm called UMAP (McInnes et al., 2018). UMAP is a nonlinear technique that, in our experience, is particularly powerful for clustering waveforms and identifying complex spikes, as also shown by the work of Markanday et al. (2020). Unlike PCA which tries to preserve the pairwise distances amongst all samples, i.e., global structure, manifold learning algorithms like tSNE and UMAP rely on local distances. UMAP first constructs a high-dimensional weighted neighbor graph and then aims at finding its low dimensional layout with similar structure. The trade-off between global and local structure in UMAP space can be tuned using the n_neighbors hyper-parameter by constraining the local radius of connectedness in high-dimensional graph based on the nth neighbor of each point (default value is n_neighbors=15). This leads to better dissociation of different clusters with local similarities, especially in unbalanced cluster sizes scenarios. For example,

discrimination of simple and complex spike clusters faces the unbalanced cluster size issue because of the intrinsic difference in rates of these two kinds of spikes.

Indeed, in the case of the data in Figure 5.4, projecting the waveforms onto the UMAP space unmasks two clusters (Figure 5.4B, right subplot). Using the graphical user interface (GUI) we select the smaller group of spikes and tentatively label them as complex spikes. Immediately, P-sort updates the probability Pr(S(t)|C(0)) window, as shown in Figure 5.4C, illustrating that these putative complex spikes were followed by a pause in the simple spikes.

5.3.2. Finding onset of complex spikes

One of the issues in identifying complex spikes is that the waveform can be significantly distorted by a preceding simple spike. Indeed, a P-cell can produce a complex spike at a fraction of a millisecond following a simple spike, as shown in Figure 5.2B. The simple spike proximity distorts the complex spike waveform. If the distortion is small, template matching can identify the onset of the complex spikes (Figure 5.3, parts 4). However, as we will see below, template matching can sometimes produce an incorrect alignment of complex spikes. P-sort can solve the problem by first aligning the complex spikes using a simple spike template, and then re-aligning using a complex spike template.

An example of this two-step process is shown in Figure 5.5A. Initially, the putative complex spikes are aligned by P-sort based on a simple spike based template. This alignment results in two groups of complex spikes (shown by black and red traces in Figure 5.5A). Next, the user selects the correctly aligned complex spikes, and defines a new template based on their mean waveform (Figure 5.5A, black traces). Finally, the user re-aligns all complex spikes using this template (Figure 5.5B). This improves the complex spike alignment, but the problem of aligning the deformed complex spikes persists. Indeed, the UMAP space indicates presence of two groups of complex spikes (red and black dots in Figure 5.5B). However, because both complex spike groups suppress the simple spikes, they are likely a single cluster that need to be merged. To correct the error, P-sort provides the Dissect Module. The Dissect Module is a semi-manual platform to re-label the onset of the distorted complex spikes (Figure 5.5C), and correctly identify the simple spikes that precede them. Following this re-labeling, the complex spikes are correctly aligned (Figure 5.5D).

A. Align CS based on SS features

500 uv

190925_152228_02

1 ms

Delete selected spikes
and then form
CS template

B. Align CS based on CS template

Deformed CS which
was mis-aligned

Deformed CS

UMAP 2

UMAP 1

C. Alignment of deformed CS using Psort's dissect module

400 uv

CS Thresh

SS Thresh

2 ms

Change CS alignment to compensate for deformation
Label as simple spike

D. Final CS alignment

Deformed CS
CS template

Figure 5.5. Correcting for misalignment of complex spikes.
A. P-sort initially aligns complex spikes based on the sodium/potassium peak which resembles simple spikes. This results in misalignment of some complex spikes because these complex spikes do not express sodium/potassium peak. The user can delete the misaligned spikes (red) and compute a complex spike template (yellow). **B.** Alignment of the spikes to the complex spike template does not solve the problem for all complex spikes: some of the deformed complex spikes remain misaligned. **C.** P-sort provides a Dissect Module with which the user can change the spike alignment to compensate for the deformation. **D.** Final complex spike alignment. Note the deformation in the complex spikes caused by the proximity of the simple spikes (red traces).

5.3.3. Identifying complex spike spikelets

A source of variability in complex spike waveforms is the stochastic nature of spikelets (Warnaar et al., 2015). In our experience, the waveforms of spikelets are not easily dissociable from simple spikes. For example, the complex spike waveforms in Figure 5.6A suggest a well-isolated P-cell, but the simple spikes that follow a complex spike have a probability distribution $Pr(S(t)|C(0))$ that exhibits a sharp peak in the putative suppression period (Figure 5.6C, top row, red arrow). This raises the possibility that spikelets are being mis-labeled as simple spikes.

A heuristic that can help identify some of the spikelets is their temporal precision with respect to CS onset. To check for the presence of this regularity, P-sort provides a tool that produces a raster plot in which the simple spikes are aligned with respect to the onset of each complex spike (Figure 5.6B). To produce this plot, in the simple spike section of the main window (or in the cluster module) the user selects the drop-down menu, selects t_prev_CS for the x-axis, and selects spike peak for the y-axis. This will present the simple spikes in a format where the x-axis is the time of each simple spike with respect

to the onset of the immediately preceding complex spike, and the y-axis is the voltage of that simple spike at the peak of its waveform.

If spikelets are present, their timing is likely to be highly regular, producing a pattern of spike alignment in the 4-10 ms post CS onset (Figure 5.6B). In this example the spikelets can be dissociated from simple spikes both because of their timing regularity and their smaller peak voltages. The user can select the putative spikelets and view their waveforms in the main window. In our experience the waveforms tend to have a somewhat smaller amplitude but slightly wider shape than regular simple spikes. However, the main heuristic is the strong temporal alignment of spikelets following the onset of the complex spike.
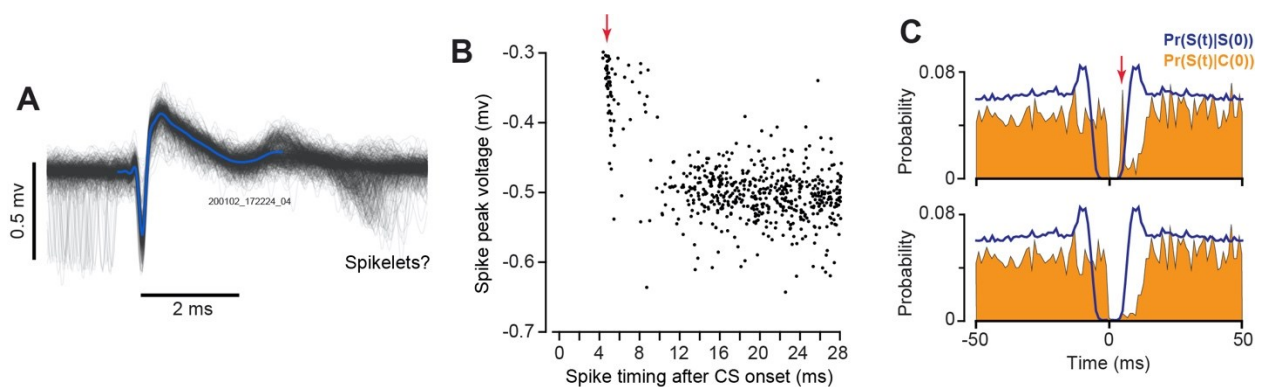
Figure 5.6. Identification of spikelets in the complex spike waveform.
A. Example of complex spike waveforms and potential spikelets. **B.** To help identify whether the spikelets are part of the complex spike waveform or are regular simple spikes, P-sort allows the user to visualize the temporal regularity of the spikes. In this graph, the x-axis is the timing of each simple spike with respect to the immediately preceding complex spike, and the y-axis is the peak voltage of that simple spike. There appears to be two clusters of spikes. The cluster indicated by the red arrow has higher temporal precision and smaller peak voltage. **C.** Before spikelet identification, the probability Pr(S(t)|C(0)) shows a sharp peak at around 5 ms following onset of the complex spike (red arrow, top subfigure). Removal of these spikelets produces the probability in the bottom row.

5.3.4. Merging clusters or splitting them using statistical interactions between simple and complex spikes

A unique challenge in cerebellar neurophysiology is finding the simple and complex spike clusters that belong to a single P-cell. P-sort provides cluster labeling and statistical tools to help with this problem. However, the main strength of P-sort is to go beyond waveform clustering and use the statistical relationship between complex and simple spikes to merge clusters, or even split a single cluster.

To illustrate this, let us begin with a particularly challenging example (Figure 5.7A). Here, visual inspection of the waveforms suggests presence of multiple neurons. Indeed, UMAP clustering produces

numerous groups of simple (Figure 5.7B) and complex spikes (Figure 5.7C). The task is to find the clusters that are signal and not noise, and more importantly, determine which simple spike cluster(s) can be attributed to which complex spike cluster(s).

The reason for the numerous clusters in the simple spike UMAP space (Figure 5.7B) is because the electrode is picking up signals from multiple neurons, and sometimes spikes from one neuron can distort the spike from another neuron. For example, the cluster labeled SS1-1 in Figure 5.7B is due to simple spikes from neuron 1 (Figure 5.7D, SS1-1). The nearby cluster SS1-2 in Figure 5.7B is due to simple spikes from neuron 1 that are in close temporal proximity with a spike from another neuron (Figure 5.7D, SS1-2). There are also clusters associated with spikes from neuron 2 that occur in isolation (Figure 5.7D, SS2-1), or in close temporal proximity with a spike from another neuron (Figure 5.7D, SS2-2). Sometimes, SS1 and SS2 co-occur, resulting in a larger spike, as labeled by cluster SS1&2-1. Finally, there are spikelets in the complex spike waveform that can be mis-labeled as simple spikes. The spikelets are noted in Figure 5.7D.

There are four clusters in the complex spike UMAP space (Figure 5.7C). One large cluster is associated with complex spikes labeled CS1 (Figure 5.7E, CS1-1), while the other large cluster is associated with CS2 (Figure 5.7E, CS2-1). Near each of these large clusters there are smaller clusters, reflecting the variability in the complex spike waveform. A source of variability in the complex spike waveform is presence of simple spikes from multiple P-cells. Thus, the complex spike will coincide with suppression of one group of simple spikes, but not all groups. As a result, CS1 can be distorted by arrival of simple spike labeled SS2, and CS2 can be distorted by arrival of simple spikes labeled SS1 (Figure 5.7E, CS1-2 and CS2-1).

P-sort provides tools to label the clusters, examine their statistical properties, and determine whether clusters should be merged or not. For example, both the conditional probabilities and the statistics of the firing rates suggest that SS1 and SS2 are two different simple spikes, as shown in Figure 5.7F, first row. Furthermore, the complex spike cluster CS1 coincided with the suppression of simple spike cluster SS1, but not SS2 (Figure 5.7F, second row). Similarly, the complex spike cluster CS2 coincided with the suppression of simple spike cluster SS2, but not SS1 (Figure 5.7F, third row). As a result, in this recording we have two distinct P-cells.
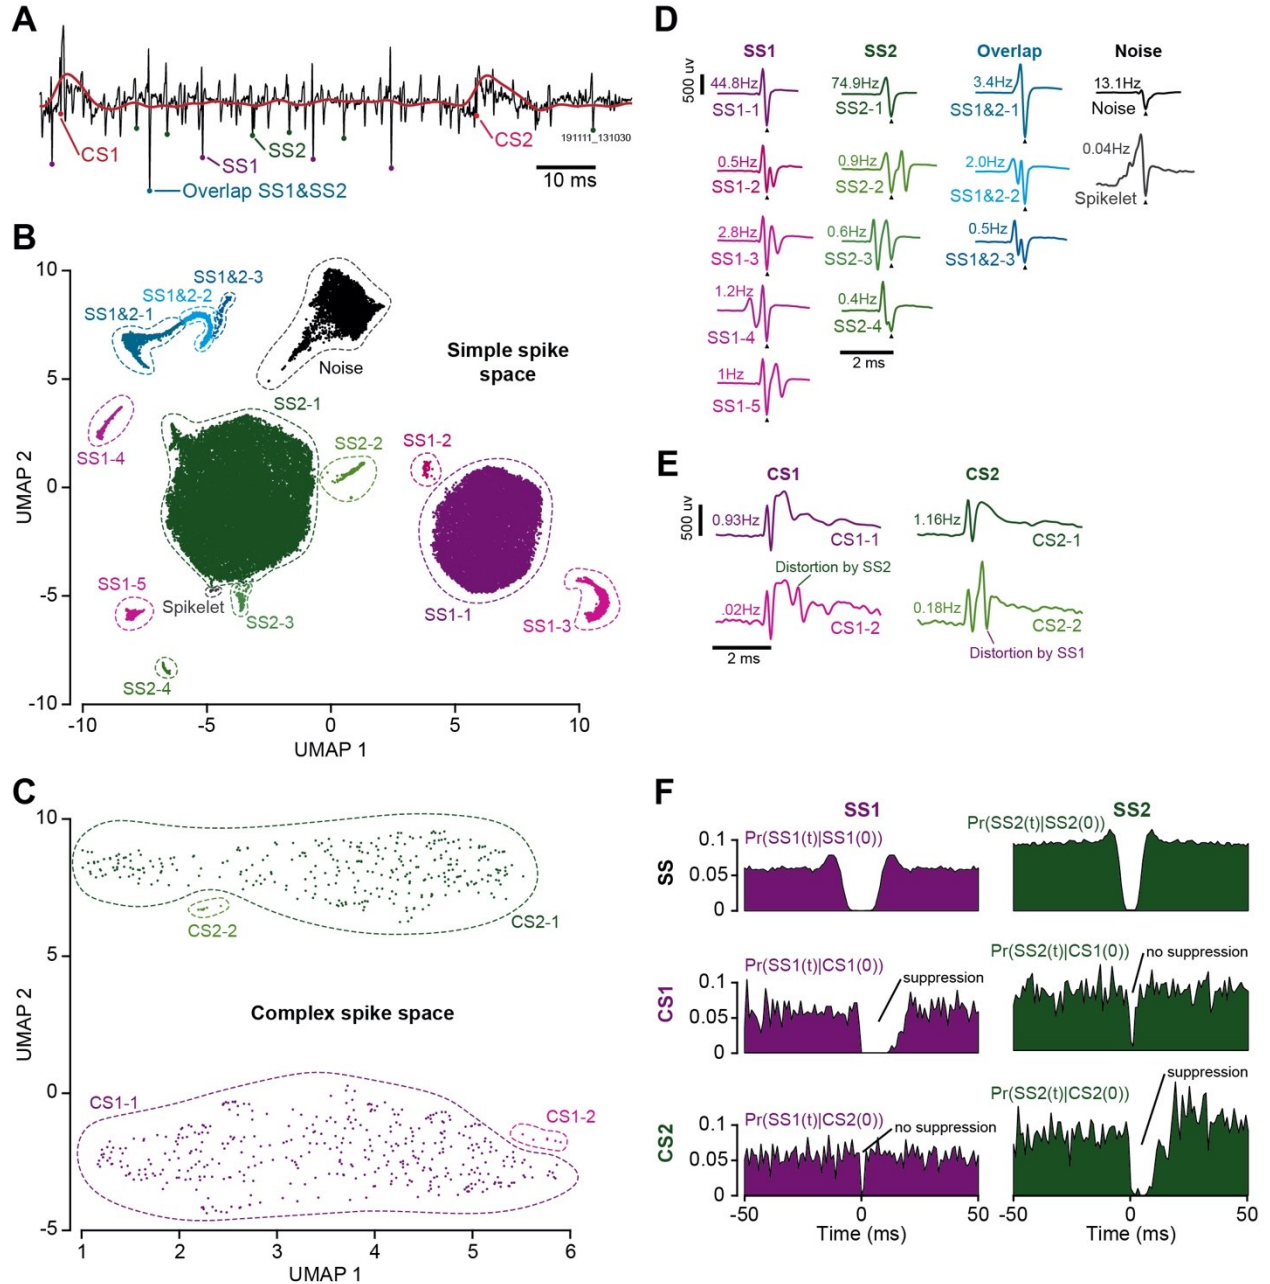
Figure 5.7. Finding clusters of simple and complex spikes that are likely generated by the same P-cell. **A.** This recording includes at least four groups of spikes: two that appear to be complex spikes, and two that are simple spikes. **B.** UMAP clustering of the simple spike space. The two major clusters are SS1 and SS2. Their waveforms are shown by SS1-1 and SS2-1 in part D. The smaller clusters are distorted spikes that are due to the temporal proximity of these major spikes, as well as other spikes, as shown in part D. A smaller cluster of spikes are labeled as spikelets of complex spikes. **C.** UMAP clustering of the complex spike space. The two major clusters are CS1 and CS2. Their waveforms are labeled as CS1-1 and CS2-1 in part E. Their waveforms can be distorted by simple spikes, as shown by CS1-2 and CS2-2. **D.** Waveforms of various clusters labeled in the simple spike space. **E.** Waveforms of the four clusters labeled in the complex spike space. **F.** Suppression period of SS1 and SS2 is quantified by the probability Pr(S1(t)|S1(0)) and Pr(S2(t)|S2(0)). The probability

Pr(S1(t)|C1(0)) quantifies the suppression following CS1. Thus, CS1 coincides with suppression of SS1 but not SS2. CS2 coincides with suppression of SS2 but not SS1. Bin size is 1 ms.

In a second example, let us show that the statistical interactions between simple and complex spikes can provide evidence suggesting that a single cluster of spikes may in fact be composed of two different cells. This data set is shown in Figure 5.8. In this recording, the complex spike waveforms produce a single cluster in the UMAP space (Figure 5.8A). However, only a part of this complex spike cluster is followed by suppression of the simple spikes. The sub-cluster that preceded the suppression of the simple spikes is labeled as CS1, and its waveform is shown in Figure 5.8B. Notably, the sub-cluster CS2 has a waveform that is similar to CS1, but CS2 does not induce suppression of the simple spikes, as illustrated in Figure 5.8C.
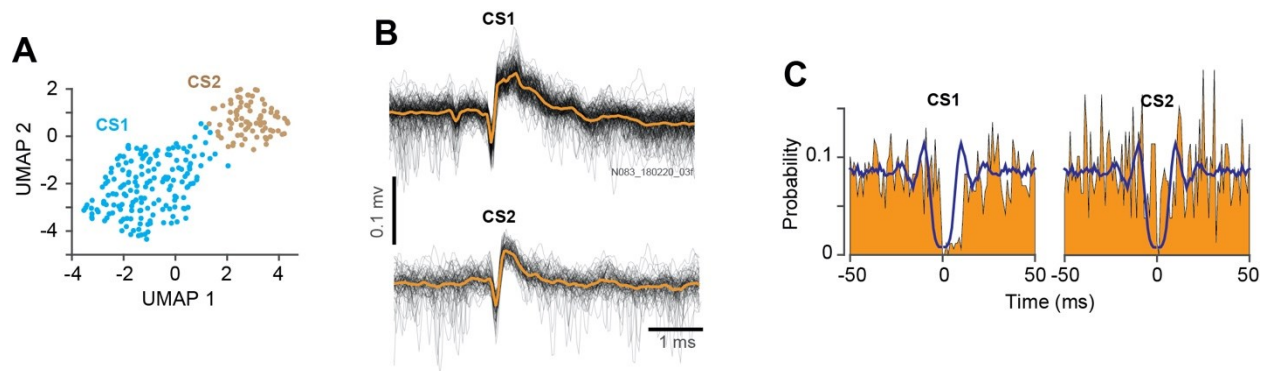
Figure 5.8. Splitting a complex spike cluster based on its statistical properties with simple spikes.
A. In this dataset, the UMAP space indicates a single complex spike cluster. However, statistical considerations raise the possibility of two subgroups, CS1 and CS2. **B.** The waveforms for the CS1 and CS2 subgroups are similar. **C.** The yellow regions are the conditional probabilities of simple spike suppression following a CS1 at time zero, or CS2 at time zero. Note that CS1 coincides with suppression of the simple spikes, but CS2 does not. Thus, while it is difficult to split the complex spike cluster into two groups based on their waveforms, they exhibit very different statistical properties.

This example highlights the possibility that on occasion, the complex spike waveforms may suggest presence of a single cluster, but a consideration of the statistical interactions can reveal that the simple spikes have been suppressed by only a sub-group within that cluster. This interaction between simple and complex spikes is a crucial feature that allows the user to go beyond waveforms and identify and attribute spikes of a P-cell.

In summary, P-sort provides clustering tools to identify simple and complex spikes. It relies on template matching to find onset of spikes, but also provides tools to correct instances where template matching can fail. A critical feature of P-sort is to provide tools for labeling and pairing of simple and

complex spikes, thus allowing the user to visualize the likelihood that specific groups of spikes are generated by a single P-cell. This can lead to merging of nearby complex spike clusters because both groups are followed by suppression of the simple spikes, or splitting of a single cluster because only a sub-group is followed by suppression.

5.3.5. Comparison of P-sort with manual curation (Spike2)

A common method currently employed for sorting of cerebellar data is via manual curation by an expert user via Spike2. To validate P-sort results, we analyzed data from mice and macaques, and then compared them with those generated by the experts in the two laboratories.

Figure 5.9A presents results from an example data set from the macaque cerebellum. Here, the expert and P-sort agreed on roughly 99% of the simple spikes, but around 92% of the complex spikes (the numbers are the average of percent spikes identified by one method, say P-sort, and confirmed by another method, i.e., the expert). The median difference between the two methods in determining the timing of the spikes was 0.06±0.01 ms (median±MAD, median absolute deviation) for simple spikes, and -0.12±0.05 ms for complex spikes (Wilcoxon Signed-Ranks Test, SS: Z=0.0077, p=0.99, CS: Z=-0.052, p=0.96). P-sort labeled 74 (0.29%) simple spikes that were not identified by the expert. These are labeled as P-sort exclusive simple spikes. The expert labeled 382 (1.48%) simple spikes that were not identified by P-sort. These are labeled as expert exclusive simple spikes. P-sort's rate of simple spikes was 45.5 Hz, while that of the expert was 46.0 Hz.

For complex spikes, there were 111 (16.7%) events picked by P-sort that were not picked by the expert, producing a complex spikes rate of 1.19 Hz for P-sort vs. 1.0 Hz for the expert. Thus, P-sort identified 19% more complex spikes than the expert. The waveforms suggest that the P-sort exclusive complex spikes are likely valid. However, the 8 complex spikes picked by the expert and not P-sort are also likely valid, as indicated by their waveforms. The reason why P-sort missed these complex spikes is because in some cases, a simple spike was in temporal proximity and distorted the complex spike waveform. Thus, in this data set there was general agreement between the two methods, although P-sort identified 16.7% more complex spikes than the expert.

Across 34 recording session in the macaque data set, a median of 2.53% of the complex spikes were detected exclusively by the expert, and 2.79% of the complex spikes were detected exclusively by P-sort. For simple spikes, 0.12% (median) of the spikes were detected only by the expert, as compared to 0.35% for P-sort. The two methods converged in their estimates of simple and complex spike firing

rates (Figure 5.9C, firing rate of P-sort spikes minus expert, Wilcoxon Signed-Ranks Test, SS: Z=-0.23, p=0.82, CS: Z=-0.31, p=0.76). The resulting conditional probabilities of the data sorted by P-sort and the expert were indistinguishable. Some of the complex spikes missed by P-sort were due to significant distortions that were present in the waveform. In general, agreement between P-sort and expert were higher for better isolated recording data.
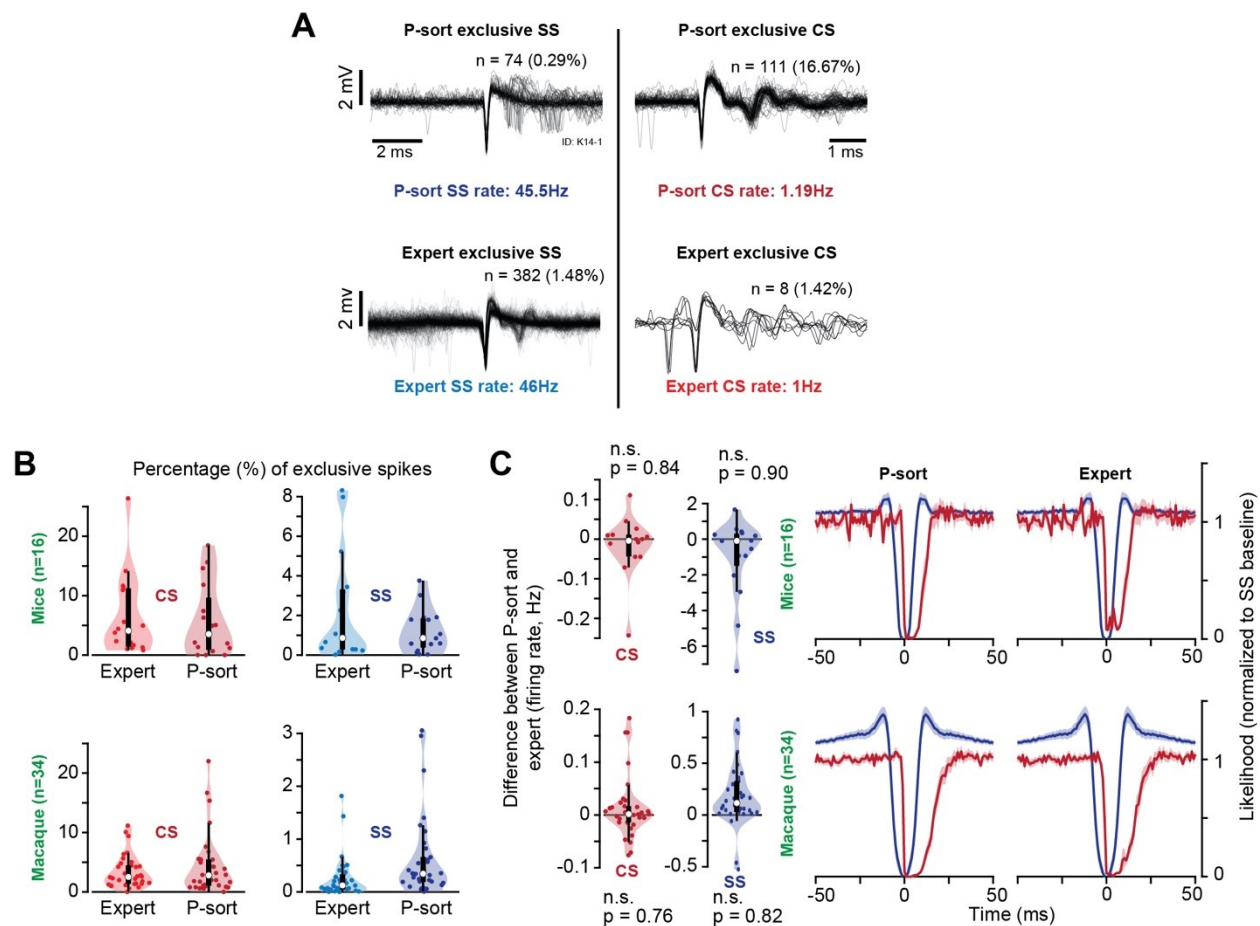
Figure 5.9. Comparison of P-sort with expert curation on mice and macaque data sets.
A. Data from a macaque recording session. P-sort picked out 74 simple spikes that were not identified by the expert (0.29% of total), labeled as P-sort exclusive. Expert picked 382 simple spikes that were not identified by P-sort (1.48% of total), labeled as expert exclusive. Complex spikes that were exclusive to P-sort and the expert are also plotted. **B.** Summary statistics on the mice (n=16 sessions) and macaque (n=34 sessions) data sets. Percentage of exclusive simple and complex spikes are plotted for the expert and P-sort. The central mark indicates median of the distribution, and the bottom and top edges of the box indicate the 25th and 75th percentiles. The thin line indicates the range of the data excluding the outliers. **C.** Difference between P-sort and expert in terms of firing rate. Right columns show the likelihoods, normalized to the baseline simple spike probability in each session, averaged over all recording sessions for each species (bin size is 1 ms). Error bars are SEM.

In the mice data set (16 sessions), a median of 4.10% of the complex spikes were detected exclusively by the expert, and 3.53% of the complex spikes were detected exclusively by P-sort, as shown in Figure 5.9B. For simple spikes, the agreement between the two methods was better: on average 0.86% (median) of the simple spikes were detected only by the expert, as compared to 0.857% for P-sort. Overall, there were no significant differences between P-sort and the expert in terms of rates of simple and complex spikes (Figure 5.9C, firing rate of P-sort spikes minus expert, Wilcoxon Signed-Ranks Test, SS: Z=0.25, p=0.90, CS: Z=0.21, p=0.84).

Importantly, in one recording session P-sort highlighted the possibility that the expert paired the wrong sub-group of complex spikes with the simple spikes. In this data set (Figure 5.8), the expert labeled a single complex spike cluster, which is of course reasonable because of the similarity of the waveforms. However, using P-sort we were able to split this cluster into two, labeling them as CS1 and CS2. Only the CS1 sub-group were the complex spikes that were generated by the P-cell that also produced the simple spikes.

Overall, a comparison of P-sort with expert manual curation suggested a general agreement: the rates of simple and complex spikes were generally similar. However, for a few sessions P-sort was able to identify more complex spikes (Figure 5.9A), and correctly label spikelets. In one instance, P-sort prevented pairing of the wrong subgroup of complex spikes with the simple spikes (Figure 5.8).

5.3.6. Reproducibility of P-sort results, and comparisons with automated algorithms

A major limitation of P-sort is that it is relies on the user to select tools and explore their efficiency in identifying and attributing spikes. This may produce substantial differences in the results that are produced by different users. In comparison, automated software can identify spikes with little or no user intervention. To quantify reproducibility of P-sort results, we compared the sorts produced by 5 experienced users on three data sets. To compare P-sort results with automated algorithms, we sorted the same three data sets with several automated algorithms.

In these comparisons we sorted data sets of various difficulty, ranging from easy to hard. In the easy data set (Figure 5.10A), complex spike waveforms were readily dissociable from simple spikes. In the medium difficulty data set (Figure 5.10B), complex spikes waveforms were more similar to simple spikes. Finally, in the more difficult data set (Figure 5.7), two different pairs of simple and complex spikes were present on a single electrode. Because the automated algorithms did not attribute pairs of

simple and complex spikes, we manually performed this step following conclusion of the automated

spike sorting.
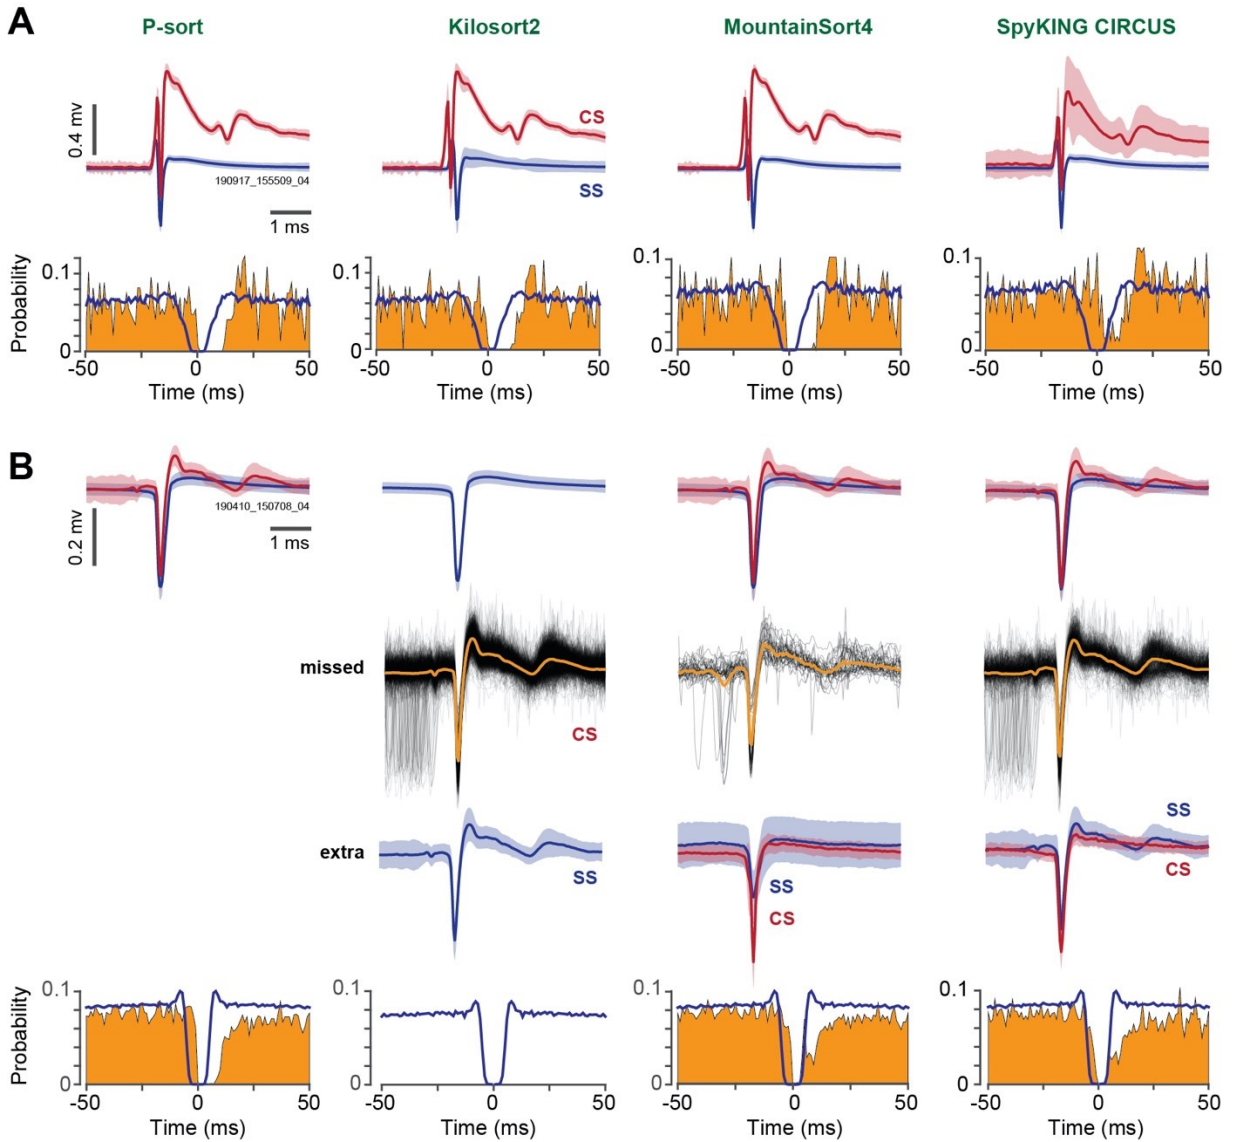
Figure 5.10. Comparison of P-sort with automated spike sorting algorithms on two data sets.
A. Easy data set. The simple and complex spike waveforms are illustrated in the first row. The conditional probability for simple spikes Pr(S(t)|S(0)) is plotted in blue. The conditional probability for simple spike suppression following a complex spike Pr(S(t)|C(0)) is plotted in yellow. All algorithms identified the simple and complex spikes. **B.** More difficult data set. First row shows the spikes identified by each algorithm. Kilosort did not identify the complex spikes. Second row shows the complex spikes missed by each algorithm, with respect to P-sort. Third row shows the spikes that were identified by each algorithm but not P-sort. Fourth row is the conditional probabilities for the labeled spikes. Error bars are standard deviation.

The lack of ground truth made it difficult to validate our results. Our recourse was to rely on the

waveform of each spike type, and the statistical interaction between them. In addition, we compared

the P-sort labeled complex spikes to the results of a neural network trained specifically to find them (Markanday et al., 2020).

We quantified performances of P-sort users and algorithms by computing a spike "agreement score" (Buccino et al., 2020). This score measured the percentage of spikes that were agreed upon between pairs of users, pairs of algorithms, and pairs of users and algorithms. Briefly, we computed the number of matched spikes (using a 1 ms time window), divided by the total number of spikes reported by each sorter, minus the number of matches.

In the easy data set, the agreement score between P-sort users was 1.0±0.0 (mean±SEM) for the simple spikes, and 0.992±0.003 for the complex spikes, demonstrating near unanimous agreement (Figure 5.11B, Psort-Psort score). The agreement score among the automated algorithms was 0.99±0.004 for the simple spikes, but only 0.750±0.066 for the complex spikes. The agreement between P-sort and the automated algorithms was 0.99±0.002 for the simple spikes, but 0.87±0.028 for the complex spikes.

Kilosort2 (Pachitariu et al., 2016) found 10452 simple spikes and 146 complex spikes (first row of Figure 5.10A). MountainSort4 (Chung et al., 2017) found 10460 simple spikes and 127 complex spikes. These results compared well with 10461 simple spikes and 147 complex spikes identified by P-sort. SpyKING CIRCUS (Yger et al., 2018) found 10488 simple spikes and 198 complex spikes in the easy data set (Figure 5.10A). Some (43/198) of the complex spikes found by SpyKING were likely invalid, as suggested by the relatively poor suppression period of simple spikes that followed.

In the medium difficult data set the agreement score between P-sort users was 0.985±0.004 for the simple spikes, and 0.945±0.01 for the complex spikes (Figure 5.11B). The agreement score among the automated algorithms was 0.874±0.05 for the simple spikes, but only 0.127±0.127 for the complex spikes. As a result, the agreement between P-sort and the automated algorithms was 0.927±0.198 for the simple spikes, and 0.383±0.081 for the complex spikes.

Kilosort2 was unable to identify any complex spikes (first row of Figure 5.10B), resulting in missed spikes shown in the second row of Figure 5.10B. It also mis-labeled many complex spikes as simple spikes (third row of Figure 5.10B). To examine the labeling differences more closely, we plotted the waveforms for spikes that were labeled by P-sort but not the automated algorithms ("missed", Figure 5.10B), as well as the waveforms for spikes that were labeled by the automated algorithms but not P-sort ("extra", Figure 5.10B). MountainSort4 found 70328 simple spikes and 996 complex spikes. For simple spikes the labeling was essentially identical with P-sort, but for complex spikes the labeling

differed by 178 events (17.2%). There were 36 P-sort identified complex spikes that were not labeled by MountainSort4 (second row of Figure 5.10B), and some simple spikes that were labeled as complex spikes (third row of Figure 5.10B). The labeling produced by P-sort resulted in a clean suppression period (Figure 5.10B, bottom row), suggesting that in the medium difficulty data set, MountainSort4 mislabeled or missed around 15% of the complex spikes. The ability of MountainSort4 to find the complex spikes, but not Kilosort2, may be because MountainSort4 uses a PCA branching algorithm that is better in discrimination of local differences between waveforms in comparison to PCA. SpyKING found 70381 simple spike and 796 complex spikes, missing 493 complex spikes (second row of Figure 5.10B) that were found by P-sort, and mislabeled some complex spikes as simple spikes (third row of Figure 5.10B).

In the medium difficulty data set, the complex spikes that were labeled by P-sort but missed by the automated algorithms had waveforms that were consisted with complex spikes, suggesting that they were erroneously missed by the automated algorithms. In contrast, the "extra" simple spikes found by Kilosort (as compared to P-sort) were likely to be complex spikes (Figure 5.10B, "extra"). The extra simple spikes found by MountainSort were likely from a different cell because they had a smaller amplitude, and the extra complex spikes found by MountainSort were likely mislabeled simple spikes. Finally, the complex spikes that were missed by SpyKING were likely valid, as suggested by their waveforms (Figure 5.10B, "missed"). The simple spikes that were labeled by SpyKING but not P-sort were likely mislabeled complex spikes, and the complex spikes that were labeled by SpyKING but not P-sort were likely mislabeled simple spikes.

We also tested the algorithms on a difficult data set that had simple and complex spikes from two different P-cells on a single contact, labeled as SS1, SS2, CS1, and CS2 in Figure 5.7. Sorting this type of data can benefit significantly from statistical interactions between spike clusters. In this difficult data set the agreement score between P-sort users was 0.989±0.0175 for SS1, 0.970±0.004 for CS1, 0.92±0.009 for SS2, and 0.965±0.003 for CS2 (Figure 5.11B). The agreement score among the automated algorithms was much lower, ranging from a low of 0.361 for CS2, to a high of 0.765 for SS1. The agreement between P-sort and the automated algorithms was 0.853±0.012 for SS1, 0.745±0.06 for CS1, 0.696±0.032 for SS2, and 0.612±0.082 for CS2.
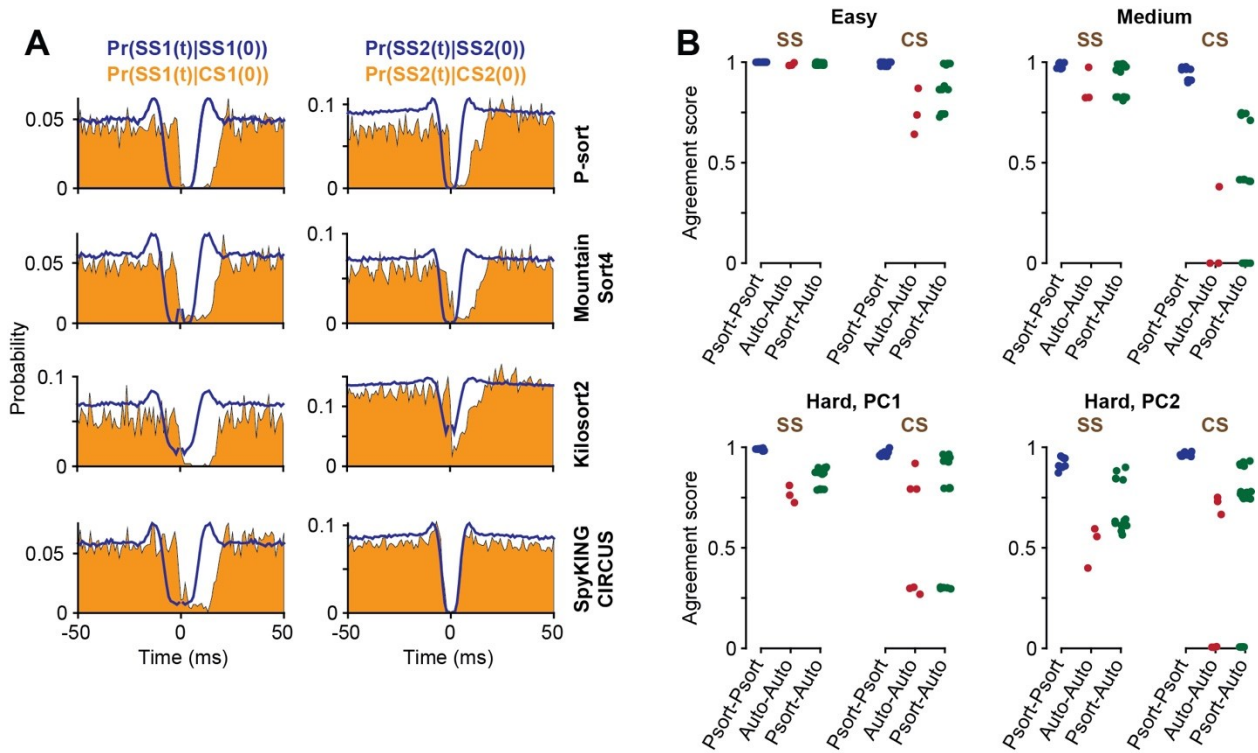
Figure 5.11. Performances of the automated algorithms in comparison with each other, and with P-sort. **A.** Performance on the difficult data set (data in Figure 5.6). The plots show conditional probabilities for the simple and complex spikes identified by the automated algorithms and P-sort. Left column is the SS1 and CS1 relationships. Right column is the SS2 and CS2 relationship. For simple spikes, the suppression period is particularly poor for Kilosort2 for both SS1 and SS2. For complex spikes, SpyKING CIRCUS and Kilosort2 produce little or no suppression of simple spikes SS2. Bin size is 1 ms. **B.** Spike agreement scores for 5 P-sort users and 3 automated algorithms, measured on the easy, medium, and difficult data sets. The spike agreement score was computed separately for simple and complex spikes (SS and CS), for each pair of P-sort users (Psort-Psort column, blue dots), for each pair of automated algorithms (red dots), and finally each pair of algorithm-Psort users (green dots).

MountainSort4 found 48478 simple spikes SS1, and 51471 simple spikes SS2 (Figure 5.12A). This agreed with 98.6% of SS1 spikes labeled by P-sort, but only 67.4% of SS2 spikes. MountainSort4 found 1156 complex spikes CS1 and 892 complex spikes CS2. This agreed with 88.9% of CS1 and 85.3% of CS2 spikes labeled by P-sort. The complex spikes that were exclusively labeled by MountainSort4 or P-sort are shown on the right panel of Figure 5.12A. Some MountainSort4 CS1 spikes were labeled as CS2 by P-sort. Similarly, some MountainSort4 CS2 spikes were labeled as CS1 by P-sort. The conditional probabilities (Figure 5.11A) provide a method to compare these results. For SS1 spikes, the probability Pr(S(t)|S(0)) for P-sort (blue lines, Figure 5.11A, first row) exhibited a cleaner suppression period than MountainSort4. For the complex spikes, the probability Pr(S(t)|C(0)) for P-sort for SS1 by CS1, and SS2 by CS2, both showed a cleaner suppression for P-sort. Thus, the main disagreements in the two approaches

were regarding the smaller amplitude simple spikes SS2, and memberships of complex spikes in CS1 and CS2.
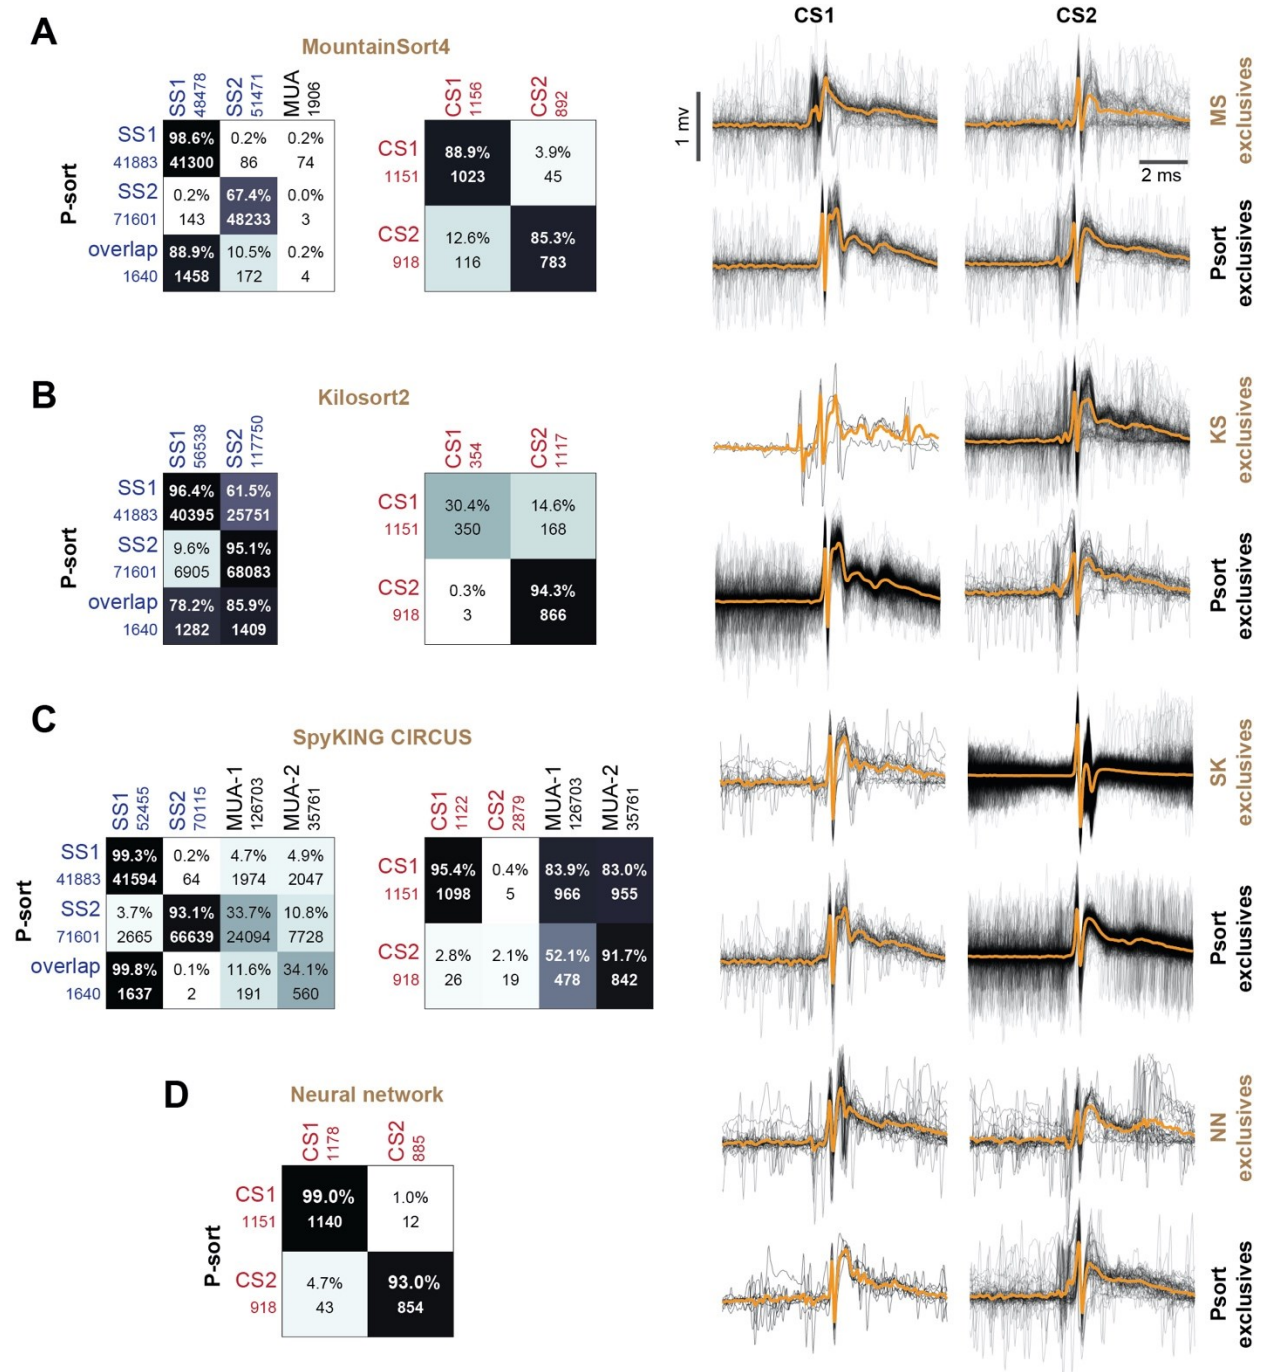
Figure 5.12. Comparison of P-sort with automated spike sorting algorithms on a difficult data set
The data are from a marmoset recording that contained multiple clusters of simple and complex spikes (the same data were presented in Figure 5.7). **A.** Comparison to MountainSort4. Simple spike numbers are shown in the left subfigure, complex spike numbers are shown on the right subfigure. The numbers indicate the number of spikes found by P-sort or MountainSort, and the percentage

indicates the agreement (spikes found by both). The right column shows complex spikes CS1 and CS2 that were labeled exclusively by each algorithm. **B.** Kilosort2 missed roughly 70% of the CS1 complex spikes labeled by P-sort. **C.** SpyKING CIRCUS disagreed entirely with P-sort regarding complex spike CS2 cluster. **D.** Comparison of P-sort complex spike identification with a neural network trained to identify complex spikes (Markanday et al., 2020).

Kilosort2 found 56538 simple spikes SS1, and 117750 simple spikes SS2 (Figure 5.12B). The smaller magnitude spikes SS2 formed a much larger group in Kilosort2 as compared to both MountainSort4 and P-sort. For complex spikes, Kilosort2 found only 354 CS1, missing roughly 70% of the CS1 complex spikes labeled by P-sort. In contrast, it found 1117 CS2 events, agreeing with 94.3% of CS2 events found by P-sort. Many of the CS2 complex spikes labeled by Kilosort2 were labeled as CS1 by P-sort (right column, Figure 5.12B). In contrast, Kilosort2 missed many CS1 complex spikes labeled by P-sort. The conditional probabilities (Figure 5.11A, third row) suggest a poor suppression period for both SS1 and SS2 simple spikes labeled by Kilosort2. The Kilosort2 CS1 complex spikes demonstrate an excellent suppression period, suggesting they were real. However, Kilosort2 missed a larger number of complex spikes that were labeled by P-sort and exhibited good simple spike suppression (Figure 5.11A, first row).

SpyKING CIRCUS produced SS labels that agreed somewhat better with P-sort than other automated software. For simple spikes, 99.3% of the SS1 labels and 93.1% of SS2 labels in P-sort were also labeled by SpyKING CIRCUS. An important difference, however, were the coincidence simple spike events, i.e., SS1 and SS2 spikes that occurred simultaneously. P-sort labeled 1640 events as coincidence of SS1 and SS2, but SpyKING CIRCUS labeled all of these as SS1 events. Another notable disagreement was among the complex spikes. While SpyKING CIRCUS found nearly the same CS1 events as P-sort (95.4% agreement), the two approaches disagreed entirely regarding the CS2 events (2.1% agreement). As the traces in the right column of Figure 5.12C illustrate, the CS2 complex spikes labeled by SpyKING are likely to be simple spikes. This conjecture appears to be confirmed by a lack of suppression in the probability traces in Figure 5.11A, right column, fourth row.

Because the main difference between P-sort and the automated algorithms was regarding identification of complex spikes, we thought to further verify P-sort's performance specifically on complex spikes. We did this by comparing P-sort with a recently developed neural network (Markanday et al., 2020) that was specifically trained on complex spike waveforms. Here we found near unanimous agreement between the two approaches (Figure 5.12D). The neural network labeled 1178 CS1 and 885 CS2 events. This corresponded to 99.0% of the CS1 and 93.0% of the CS2 events labeled by P-sort. The

few disagreements in the labeled events are shown on the right column of Figure 5.12D. In almost all cases, the complex spikes were preceded by a temporally adjacent simple spike, thus producing waveform distortions, making the labeling process particularly challenging.

In summary, we compared results of P-sort with automated algorithms in three data sets and found that the main difference was in the identification of complex spikes. For example, in the medium difficultly data set in which complex and simple spikes had similar waveforms, Kilosort2 missed the complex spikes. Conditional probabilities suggested that in both the easy and the difficult data sets, performance of MountainSort4 was close to P-sort, though it also missed around 15% of the complex spikes. In the medium difficulty data set, performances of MountainSort4 and SpyKING were close to P-sort. In all three data sets, P-sort results appeared highly reproducible, resulting in a lower bound of 92% spike agreement among the users in the difficult data set. As compared to the automated algorithms, P-sort produced cleaner simple spike suppression periods, as well as more robust complex spike periods of simple spike suppression. Finally, there was near unanimous agreement between P-sort and a neural network trained specifically to identify complex spikes.

5.3.7. A database for testing and development of algorithms

Automated algorithms can analyze tens or hundreds of simultaneously recorded electrodes, making them essential for high-density probes. Unfortunately, current automated algorithms may not perform ideally in the cerebellum, as illustrated by the data here, and documented elsewhere (Hall et al., 2021). While we have successfully used P-sort on data sets that include 4-7 electrodes, we have found it very time consuming to use P-sort when the data set is 64 electrodes or more. Thus, there remains a clear need to develop automated algorithms for the cerebellum.

Development of these algorithms relies on training and validation of on large data sets. Unfortunately, to our knowledge none are currently available. Thus, we recorded and analyzed a database of over 300 P-cells from the marmosets, macaques, and mice cerebella. We then used P-sort to label and attribute the simple and complex spikes in each recording. Recordings in the primates were from the vermis, lobules VI and VII. Recordings from the mice were from the eye blink region of lobule V.

We found that on average, simple spike rate in the marmoset (Figure 5.13A) was somewhat higher than in the macaque (63.8±1.29 Hz vs. 55.9±2.45 Hz, Mean±SEM, independent samples t-test, $t_{(285)}$=2.75, p=0.006). In contrast, complex spike rate in the marmoset was somewhat lower than in the

macaque (0.88±0.013 Hz vs. 0.98±0.027 Hz, $t_{(285)}$=-3.67, p=0.0003). For the simple spikes, the conditional probability Pr(S(t)|S(0)) was very similar in the two primate species (Figure 5.13C), suggesting that the simple spike suppression periods are similar.

The complex spikes in the macaque were followed by a somewhat longer period of simple spike suppression (Figure 5.13C). To measure the duration of simple spike suppression that followed a complex spike, for each P-cell we computed the time T when the probability Pr(S(t)|C(0)) increased beyond a threshold of 63% of the baseline before onset of the complex spike (Figure 5.13D). In the marmoset, P-cells had an average of 14.8±0.26 ms suppression period, significantly less than the suppression duration of 20.4±0.89 ms we quantified in the macaque ($t_{(285)}$=-8.02, p=3×10^{-14}).

These differences are difficult to interpret because spike rates and suppression durations depend on the precise location of the P-cell. For example, P-cells located in zebrin negative regions display higher frequency simple spikes than those located in neighboring zebrin positive stripes (Zhou et al., 2014). Furthermore, the simple spike suppression following a complex spike is longer in zebrin positive zones. In both marmoset and macaque, a fraction of a millimeter difference in the recording site along the medial-lateral direction can change the zebrin band characteristics, particularly in lobule VII of the vermis (Fujita et al., 2010). Thus, the differences in rates and suppression duration between species may be due to small differences in sites of recordings.

In the mouse data the complex spike rate was somewhat higher than the marmoset, 1.42±0.07 Hz ($t_{(254)}$=12.03, p=1×10^{-26}), and also somewhat higher than the macaque ($t_{(81)}$=7.13, p=4×10^{-10}). Furthermore, the complex spikes in the mouse were followed by a shorter suppression period (12.96±0.80 ms) than in the macaque ($t_{(81)}$=-5.21, p=1×10^{-6}) and in the marmoset ($t_{(254)}$=-2.25, p=0.026).

Finally, in all three species we observed a consistent pattern in the distribution of inter-spike intervals (ISI). To compute the ISI distribution, we first measured the average ISI for each type of spike for each P-cell and then normalized the distribution by setting the average ISI to one. The resulting simple and complex spike normalized ISIs were different from each other, but nearly identical in the three species (Figure 5.13E). These patterns may be useful priors that can aid identification of these spikes in an automated software.
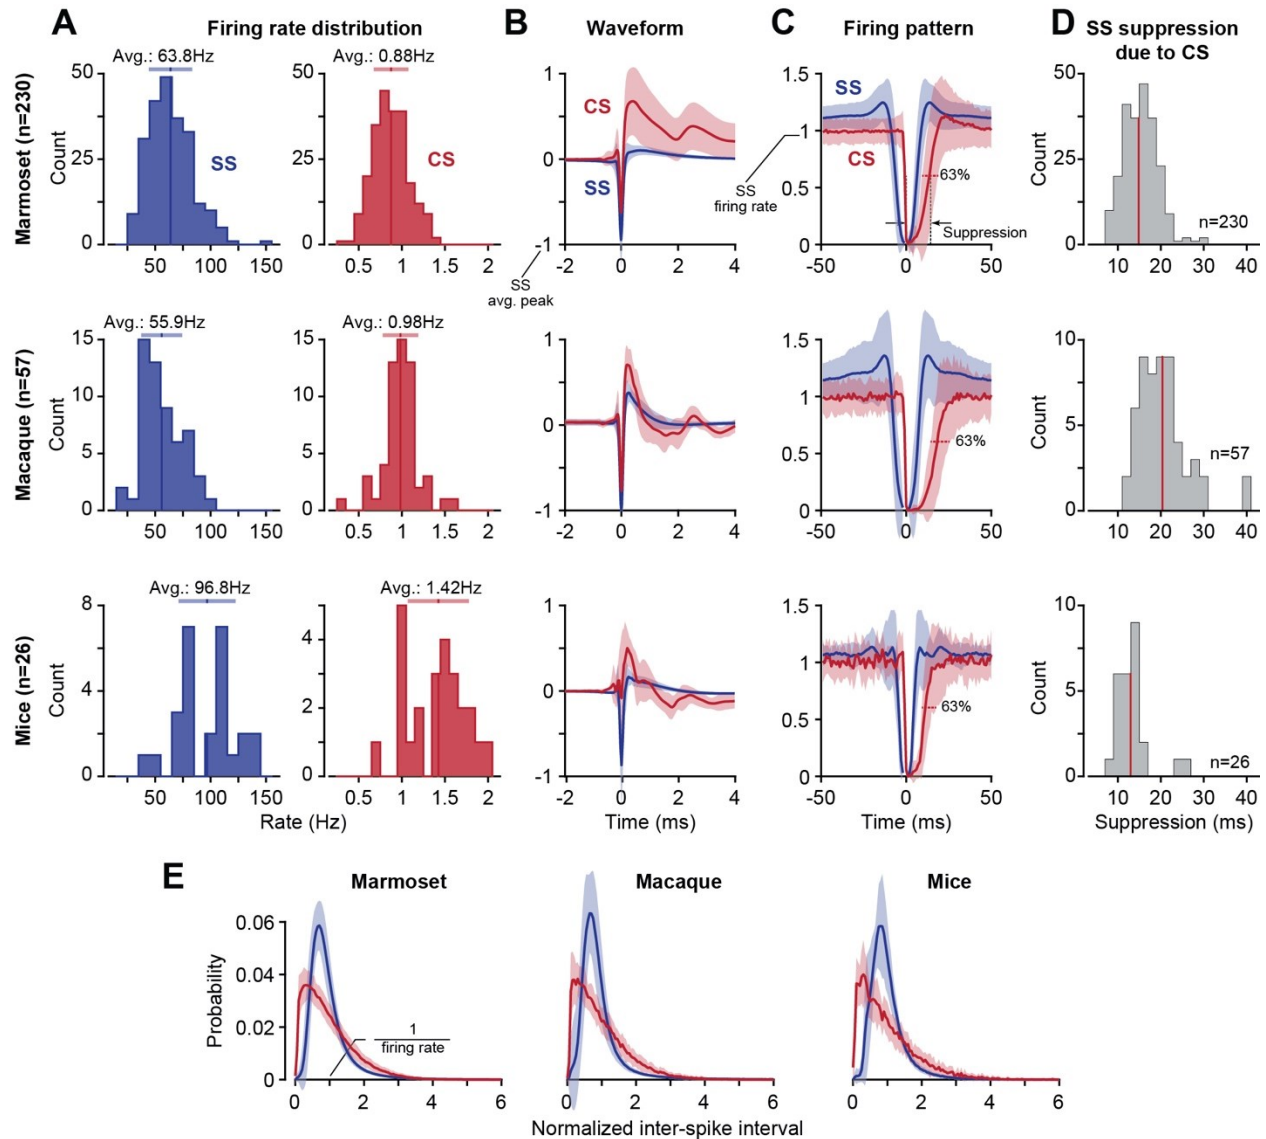
Figure 5.13. Statistical properties of simple and complex spikes in three species.
A. Distribution of average firing rates. **B.** Waveform of simple and complex spikes. Simple and complex spikes of each P-cell were both normalized by setting to -1 the negative peak of the simple spike waveform. Error bars are standard deviation. **C.** Suppression period of simple spikes (blue, SS|SS) and the suppression coincided with complex spikes (red, SS|CS). SS|SS indicates the rate of simple spikes at time t when another simple spike occurs at time zero. SS|CS indicates the rate of simple spikes at time t when a complex spike occurs at time zero. Simple and complex spike rates for each P-cell were normalized with respect to average simple spike firing rate. Error bars are standard deviation. **D.** Suppression period of simple spikes following arrival of a complex spike. Suppression period for each P-cell was defined as the duration of time after a complex spike that was required before the simple spike rate recovered 63% of its pre-complex spike value. The red line indicates mean. **E.** Inter-spike interval distribution for simple (blue) and complex spikes (red). ISI data for each spike type in each cell was normalized so that the average ISI, defined as the inverse of the average firing rate, was equal to one. Error bars are standard deviation.

In summary, we labeled simple and complex spikes in over 300 recordings in three species and quantified their statistical properties. The raw data as well as the P-sort labels are available at https://doi.org/10.17605/osf.io/gjdm4.

5.4. Discussion

Spike sorting in the cerebellum can be a joy, something akin to a treasure hunt: finding the correct clusters of waveforms produces a satisfying statistical pattern in which the complex spikes are followed with suppression of the simple spikes. However, identifying spikes that belong to a P-cell can be difficult both for those who prefer manual curation and those who employ automatic algorithms. While most complex spikes leave an LFP signature because of their broad spike (Zur & Joshua, 2019), some complex spikes can lack this characteristic (Figure 5.1B), possibly because of the location of the electrode with respect to the soma (Han et al., 2020). Moreover, complex spike waveforms can differ from one event to the next because of their temporal proximity to simple spikes (Markanday et al., 2020), because of spikelets (Burroughs et al., 2017; Davie et al., 2008; Ito & Simpson, 1971; Monsivais et al., 2005), or because of intrusion of spikes from neighboring neurons. Even after the spikes are identified based on their waveforms, the simple spikes may belong to one P-cell, while the complex spikes may belong to another (Figure 5.1E).

To help with the identification and attribution problems, we organized a set of clustering and labeling tools in a GUI based cross-platform analysis software called P-sort. Like other sorting software, P-sort clusters spikes based on their waveform properties. However, P-sort emphasizes the statistical relationship between complex and simple spikes: there is a suppression of simple spikes following a complex spike. Although there may be rare cases in which this property is not present, usually this information can be used to justify merging of disparate clusters, for example when spike waveforms are distorted (Figure 5.7), or justify splitting of a single cluster, for example when spikes from two different cells have similar waveforms (Figure 5.8). Thus, P-sort provides the user with the means to go beyond waveform-based clustering by providing statistical information regarding how spikes interact with each other.

Our development of P-sort was aided by a diverse collection of data from marmosets, macaques, and mice cerebellums. To validate the results, we compared P-sort's results with manual curation performed by the experts in the mice and macaque laboratories. In addition, we compared P-sort with automatic algorithms MountainSort4 (Chung et al., 2017), Kilosort2 (Pachitariu et al., 2016),

and SpyKING CIRCUS (Yger et al., 2018), as well as a neural network trained to identify complex spikes (Markanday et al., 2020). On the easy data sets, the performances of these automated algorithms were generally good, but on medium and high difficulty data sets, they occasionally missed large groups of complex spikes, or mis-labeled them as simple spikes. A comparison of P-sort selected complex spikes with a neural network trained to identify complex spikes produced near unanimous agreement.

P-sort follows the traditional approach in which features of the spike waveforms are identified and then clustered. It relies on a recently developed algorithm called UMAP (McInnes et al., 2018), a nonlinear dimensionality reduction technique that, in our experience, is particularly powerful for clustering waveforms. UMAP is related to another non-linear dimensionality reduction algorithms called t-SNE (t-distributed stochastic neighborhood embedding) (van der Maaten & Hinton, 2008). Unlike t-SNE, UMAP returns an invertible transform onto which new data can be projected without having to re-compute the map. This has the unique advantage of allowing for cross validation, which can be employed by semi-supervised learning methods in which the expert labels a subset of the data and leaves it to UMAP to make predictions on the unlabeled data set. Indeed, UMAP was recently adopted by Markanday et al. (2020) to cluster complex spikes from the cerebellum, and by Lee et al. (2021) to cluster spikes from the cerebral cortex.

However, UMAP has certain disadvantages. In UMAP space, cluster sizes and distances between them do not contain information regarding the waveform structures. This issue can be resolved by viewing clusters in the PCA space, which is also provided in P-sort. Moreover, UMAP relies on a stochastic optimization process that produces non-deterministic outcomes of different runs. To help with this problem, P-sort supports a GPU implementation of UMAP that provides results for a typical 5-minute recording session in around 1-2 seconds. This rapid response makes it possible for the user to evaluate the same data set multiple times, perhaps with different waveform window sizes. In addition, in most cases UMAP can reproduce the same number of clusters with similar topological relationships. Thus, although UMAP is not deterministic, it can extract clusters that are reproducible. Regardless, there remains a possibility of overfitting in noisy recording scenarios. To help with this, the Cluster Module depicts waveforms of each cluster, as well as their statistical relationship to other spike clusters.

Once the clusters are identified, P-sort provides automated tools to find cluster boundaries. One way to improve this step is to use Louvain clustering, i.e., finding community of spikes that are highly inter-connected in the UMAP space. This approach was recently demonstrated by Lee et al. (2021) in sorting of cortical spikes.

Finding cluster boundaries, however, is not the only method to sorting, as illustrated by SpyKING CIRCUS (Yger et al., 2018). That work uses a template to define the centroid of each cluster, not their precise borders. In our tests, this approach produced good results on the large amplitude simple and complex spikes, but poor results on the smaller amplitude spikes, as illustrated by the conditional probabilities (Figure 5.11).

The development of high-density electrodes highlights the need for automatic sorters. These sorters have addressed many issues including cluster matching between different channels, as well as multi-unit sorting. Moreover, their software features GUI-based visualization and manual curation toolboxes, thus allowing the user to post-process the results. However, the automated approaches currently lack GUIs for identification of dependent spikes, i.e., simple and complex spikes. In contrast, P-sort was designed to efficiently illustrate and interactively handle the simultaneous sorting and attribution of complex and simple spikes. The addition of features like multi-unit sorting opens the way for automatic sorters to be used as an initial starting point, followed by clustering and attribution by P-sort.

The rapidly evolving silicon probe technology makes it essential that we encourage development of automated sorters for the cerebellum. Thus, we used P-sort to label spikes in recordings made from over 300 P-cells in various species and provide this labeled database to help software developers test and improve their algorithms. P-sort software is available at https://github.com/esedaghatnejad/psort. The labeled neurophysiological data are available at https://doi.org/10.17605/osf.io/gjdm4.

We hope that the availability of this data set can help the cerebellar community by facilitating the development of automatic algorithms, which rely on the availability of large data sets for training and validation. Furthermore, because the labeled data reflect the results of P-sort, investigators can try their current sorters and compare their results, thus uncovering potential limitations of P-sort. Finally, the labeled data allows the novice P-sort user to compare their sorting results with the expert user, thus serving as a teaching tool.

Chapter 6. Synchronous spiking of cerebellar Purkinje cells during control of movements

The information that the brain transmits from one region to another is often viewed through the lens of firing rates. However, if the output neurons could vary the timing of their spikes with respect to each other, then through synchronization they could highlight information that may be critical for control of behavior. In the cerebellum, the computations that are performed by the cerebellar cortex are conveyed to the nuclei via inhibition. Yet, synchronous activity entrains nucleus neurons, making them fire. Does the cerebellar cortex rely on spike synchrony within populations of Purkinje cells (P-cells) to convey information to the nucleus? We recorded from multiple P-cells while marmosets performed saccadic eye movements and organized them into populations that shared a complex spike response to error. Before movement onset, P-cells transmitted information via a rate code: the simple spike firing rates predicted the direction and velocity of the impending saccade. However, during the saccade, the spikes became temporally aligned within the population, signaling when to stop the movement. Thus, the cerebellar cortex relies on spike synchronization within a population of P-cells, not individual firing rates, to convey to the nucleus when to stop a movement.

6.1.   Introduction

To understand how neurons in a region of the brain respond to sensory information or participate in control of movements, we typically search for correlates of the sensory and motor variables in the patterns of spikes. These patterns are usually quantified via the average firing rates of neurons. However, there can be additional information in the timing of each spike, as exemplified by the independent rate and temporal codes in the hippocampus (Huxter et al., 2003), the thalamus (Y et al., 1998), and the somatosensory cortex (Lankarany et al., 2019). A central question is whether neurons use spike timing to transmit functionally relevant information from one region of the brain to another.

A special form of temporal coding is synchronization of spikes among a group of neurons. For example, synchronization among glutamatergic thalamic neurons increases the efficiency of driving post-synaptic neurons in the somatosensory cortex (Bruno & Sakmann, 2006). However, unlike the thalamus, the sole output from the cerebellar cortex is via GABAergic Purkinje cells (P-cells). As a result, asynchronous activity of P-cells inhibits the cerebellar nucleus neurons. Indeed, previous analysis of spike timing in single P-cells did not find evidence that timing of spikes affected ongoing movements (Payne et al., 2019). Yet, there are specialized mechanisms in the cerebellar cortex that promote synchronization of nearby P-cells (Han et al., 2018), raising the question of whether the cerebellum relies on synchronization to transfer information from its cortex to its nuclei.

In principle, when a population of P-cells synchronizes their spikes, they can drive cerebellar output in a way that is not possible via asynchronous spiking (Gauck & Jaeger, 2000). For example, when P-cells are synchronously stimulated (in slice, and anesthetized preparations), they entrain the nucleus cells, transforming their inhibitory inputs to the nucleus into production of spikes (Özcan et al., 2020; Person & Raman, 2011). This raises the possibility that analogous to the thalamic input to the cerebral cortex, P-cells may rely on synchronization to convey information to the nucleus, possibly affecting a specific part of the ongoing movement (Heck et al., 2007).

Here, we focused on saccadic eye movements because they are so brief as to preclude the possibility of sensory feedback, requiring the brain to rely entirely on its internal predictions (Chen-Harris et al., 2008; Keller et al., 1996; Quaia et al., 2000). These predictions depend critically on the cerebellum (Catz et al., 2005; Xu-Wilson, Chen-Harris, et al., 2009). For example, firing rates of populations of P-cells, but not individual cells, predict the direction and velocity of the ongoing saccade (Herzfeld et al., 2015, 2018). However, to check for synchrony we needed to simultaneously record from multiple P-cells during saccades, something that to our knowledge had not been accomplished in any primate species.

6.2. Methods

Neurophysiological data were collected from two marmosets (Callithrix Jacchus, male and female, 350-370 g, subjects M and R, 6 years old. The marmosets were born and raised in a colony that Prof. Xiaoqin Wang has maintained at the Johns Hopkins School of Medicine since 1996. The procedures on the marmosets were evaluated and approved by the Johns Hopkins University Animal Care and Use Committee in compliance with the guidelines of the United States National Institutes of Health.

6.2.1. 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 (Sedaghat-Nejad, Herzfeld, Hage, et al., 2019). Visual targets were presented on an LCD screen (Curved MSI 32" 144 Hz - model AG32CQ) while binocular eye movements were tracked using an EyeLink-1000 eye tracking system (SR Research, USA). Timing of target presentations on the video screen was measured using a photo diode.

We used the MRI and CT imaging data for each animal and designed an alignment system that defined trajectories from the burr hole to various locations in the cerebellar vermis, including points in

lobule VI and VII. 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.

We recorded from the cerebellum using four types of electrodes: quartz insulated 4 fiber (tetrode) or 7 fiber (heptode) metal core (platinum/tungsten 95/05) electrodes (Thomas Recording), and 64 channel checkerboard or linear high density silicon probes (M1 and M2 probes, Cambridge Neurotech). We connected each electrode to a 32 or 64 channel head stage amplifier and digitizer (RHD2132 and RHD2164, Intan Technologies, USA), and then connected the head stage to a communication system (RHD2000 Evaluation Board, Intan Technologies, USA). Data were sampled at 30 kHz and band-pass filtered (2.5 Hz to 7.6 kHz). We used OpenEphys (Siegle et al., 2017), an open-source extracellular electrophysiology data acquisition software, for interfacing with the RHD2000 system and recording of signals.

6.2.2. Behavioral protocol

Each trial began with fixation of a center target for 200 ms, after which a primary target (0.5°×0.5° square) appeared at one of 8 randomly selected directions at a distance of 6.5°. Onset of the primary target coincided with presentation of a distinct tone. As the animal made a saccade to this primary target, that target was erased, and a secondary target was presented at a distance of 2°-2.5°, 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 1.5° radius of the target center, and maintained fixation for at least 200 ms. Onset of reward coincided with presentation of another distinct tone. Following an additional 150-250 ms period (uniform random distribution), the secondary target was erased, and the center target was displayed.

6.2.3. Data analysis

All saccades, regardless of whether they were instructed by presentation of a visual target or not, were identified in the behavioral data using a velocity threshold. Saccades to primary, secondary, and central targets were labeled as targeted saccades, while all other saccades were labeled as task irrelevant.

Electrophysiological data were sorted into spikes using P-sort (Sedaghat-Nejad, Fakharian, et al., 2021), a newly developed open-source software that we developed specifically for identification of simple and complex spikes. P-sort provides tool to help confirm that the complex and simple spikes

originate from the same P-cell. Briefly, we compared the conditional probability Pr(S(t)|C(0)) with Pr(S(t)|S(0)). For example, Pr(S(t)|C(0)) is the probability that a simple spike occurred at time t, given that a complex spike was generated at time zero. Pr(S(t)|S(0)) is the probability that a simple spike occurred at time t, given that a simple spike was generated at time zero. Simple spikes that originate from a single P-cell produce a refractory period. Thus, Pr(S(t)|S(0)) exhibits a low probability period of roughly 5 ms in duration after time zero. On the other hand, a complex spike coincides with suppression of future simple spikes, but not those that occurred before. As a result, Pr(S(t)|C(0)) is asymmetric, with a long period of low simple spike probability (around 15 ms) following time point zero.

Simple and complex spike baseline firing rates were computed by dividing the total number of spikes by the duration of the entire recording. Simple and complex spike instantaneous firing rate were calculated from peri-event time histograms with 1 ms bin size. Events of interest included: visual events (target onset), saccade onset, deceleration onset, and saccade offset. We used a Savitzky–Golay filter (2nd order, 31 datapoints) to smooth the traces for visualization purposes.

Complex spike tuning was computed by measuring the CS probability following target onset. We counted the number of complex spikes after target onset up to saccade onset or a fixed 200 ms window, whichever happens first. This approach ensured that the complex spikes during saccades or after saccade offset did not get included in the measurements. Dividing the spike count by the number of events resulted in the CS probability in each direction.

Suppression duration of simple spikes following a complex spike was computed by measuring the period during which the simple spikes recovered 63% of their pre-complex spike firing rate.

To compute population response during saccades, we began by computing the change in simple spike firing rate of each P-cell with respect to its baseline. Next, we labeled each saccade by measuring its direction with respect to the CS-on of the recorded P-cell. Finally, we summed the activities in all P-cells (i.e., changes with respect to baseline) for saccades in direction CS-on, CS+45, etc., using a bin size of ±22.5°.

6.2.4. Analysis of the simultaneously recorded P-cells

Multi-channel electrodes allowed for analysis of simultaneously recorded neurons. However, spiking activity in one neuron can easily influence the data recorded by two nearby channels, thus giving an illusion that the two channels are picking up two distinct neurons. To guard against this, after we sorted the data in each channel, we waveform triggered the data recorded by channel A by the spikes recorded

on channel B. This identified the waveform of the neuron recorded by channel B on the spike recorded on channel A. We compared this cross-channel triggered waveform with the within channel triggered waveform generated by the spikes recorded by channel A. The cross-channel triggered waveform must produce a different cluster of spikes in A than the main neuron isolated by A. If there were spikes that occurred within 1 ms of each other on channels A and B, we used these coincident-spike events to trigger the waveform in A. The spikes in A that were identified to be coincident with B should look approximately the same as the non-coincident spikes in A. Examples of this approach are provided in Sedaghat-Nejad et al. (2019).

To quantify coordination between activities of two P-cells, we computed joint probabilities, corrected for chance (Han et al., 2018; Wong et al., 1993). We computed Pr(S2(t),S1(0)) / (Pr(S2).Pr(S1)), which is equal to Pr(S2(t)|S1(0)) / (Pr(S2)). This quantified whether the occurrence of a simple spike on channel 1 at time zero altered the probability of simple spikes on channel 2 at time t, corrected for probabilities expected from their average firing rates. Because channel labels 1 or 2 are interchangeable, we considered the average of the two cases as the corrected conditional probability for a pair of P-cells. We implemented a similar analysis to quantify the coordination between complex spikes in two cells, or complex spikes in one cell and simple spikes in another cell.

To compute the probability of synchronization of simple spikes during saccades, we began by computing the joint probability of spiking at time t, given that a saccade took place at time zero in a particular direction, Pr(S1(t),S2(t)|sac(0)). To correct for the fact that firing rates changed during the saccade, we divided the joint probability by the independent probabilities of spike production in each cell, measured when a saccade took place in the given direction at time zero. Thus, the synchronization index was defined for each saccade direction as:

$$SI = \frac{\Pr\left(S1(t), S2(t)|sac(0)\right)}{\Pr\left(S1(t)\big|sac(0)\right)\Pr\left(S2(t)|sac(0)\right)}$$

Equation 6.1

6.2.5. Modeling

To check whether the increased synchronization index during saccades was an artifact of the change in firing rates of P-cells, we performed a simulation of spiking neurons that burst and paused like cells in our population but had independent probabilities of spike timing. We pre-defined the firing rate pattern for 11 hypothetical neurons all with a 60 spk/s baseline firing rate and then add a 130 ms duration modulation. The result produced a population response firing pattern that mimics the P-cells in our

dataset. Using a Bernoulli process and pre-defined firing rates, we simulated the spiking activity for the 11 hypothetical neurons for 1000 trials. Next, we used the methods described above to compute the estimated firing rate of each cell and the population response. Finally, we computed the joint probability as well as the synchronization index between 55 (2 choose 11) pairs of cells. Our results confirmed that while the joint probability was modulated according to the change in firing rates in the population of cells, the synchronization index stayed at chance level.

6.3.   Results

We focused on marmosets, a primate that like macaques and humans relies on saccadic eye movements to explore its visual scene, but is a fraction of the size of macaques, thus making it possible to record from the cerebellum using short, multi-channel probes. We used MRI and CT-aligned maps of each animal's cerebellum 18 to guide electrodes and record from P-cells in lobule VI and VII of the vermis (Figure 6.1C). Because there were no previous electrophysiological data from the marmoset cerebellum, we searched for saccade related activity and found that P-cells in the posterior lobule VI and anterior lobule VII produced simple spikes that were modulated during saccades. Thus, we focused on these regions and recorded from n=149 well-isolated P-cells (Figure 6.2 provides characteristics of the entire data set). Crucially, our data included n=42 pairs of simultaneously isolated P-cells that were recorded from separate channels.

We trained the animals to fixate a central target and make a saccade to a primary target that appeared at random in one of 8 directions (Figure 6.1A&B). At the onset of the primary saccade the target was erased and replaced with a secondary target, also at a random location. Following a random period of fixation and delivery of reward, the secondary target was erased, the central target was displayed, and the trial re-started. Whereas production of saccades accompanied modulation of simple spikes (SS), the random nature of the visual stimuli produced sensory prediction errors, promoting modulation of complex spikes (CS) (Herzfeld et al., 2015, 2018; Soetedjo et al., 2008; Soetedjo & Fuchs, 2006).
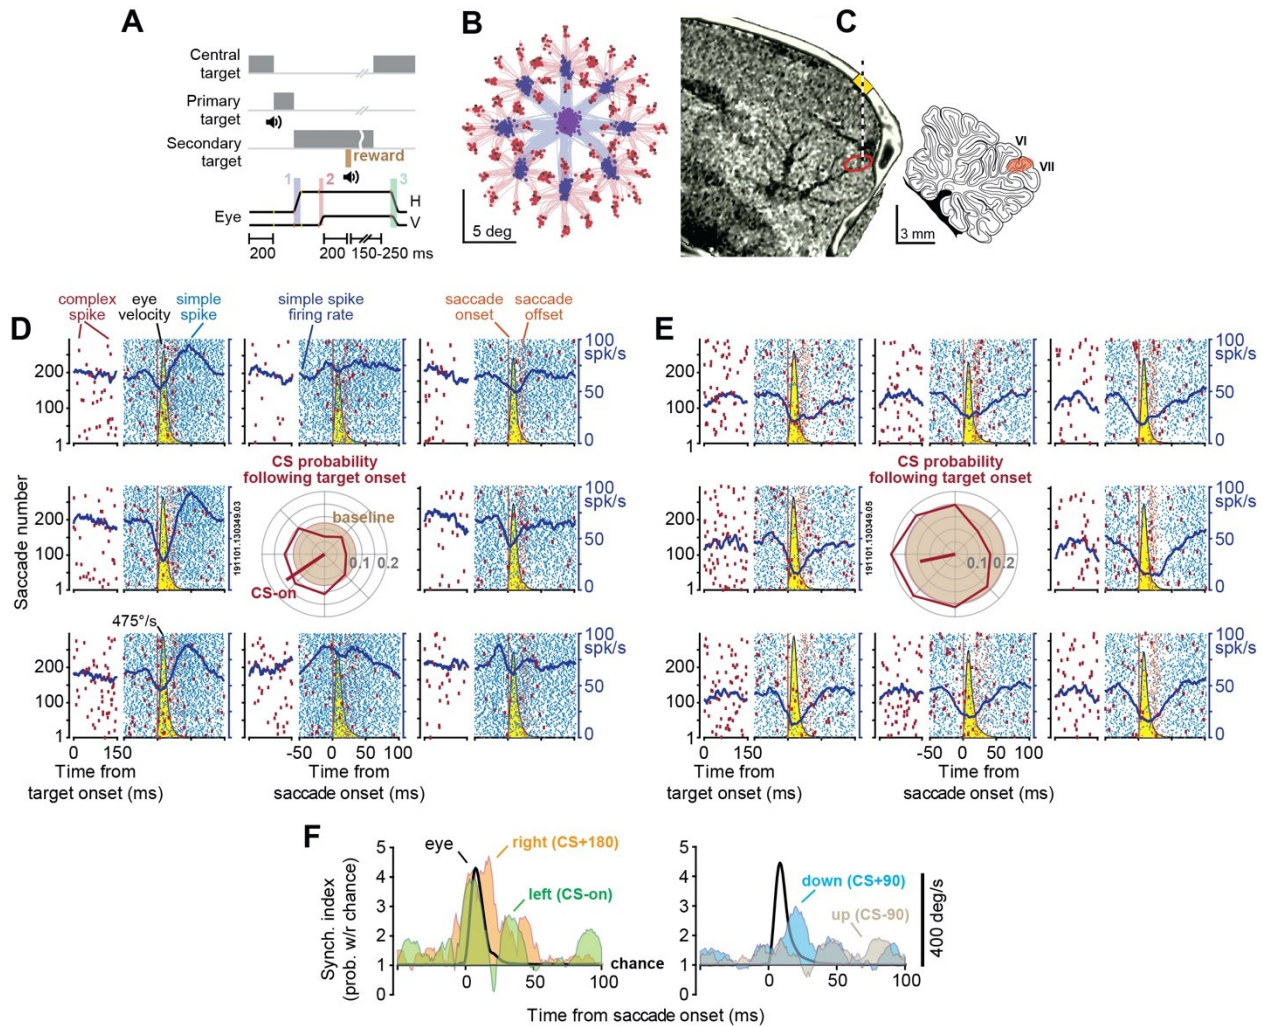
Figure 6.1. P-cells synchronized their simple spikes during saccades.
A. Experimental paradigm. Marmosets were trained to make saccades to visual targets that appeared randomly at one of 8 directions. Onset of the primary saccade (labeled 1 in the lowest trace) resulted in the replacement of the primary target with a secondary target, also at a random direction. Following the secondary saccade (labeled 2), and a fixation period, reward was delivered, and the center target was displayed, resulting in a centripetal saccade (labeled 3). **B.** Eye position for the primary (blue) and secondary (red) saccades in a typical experiment. **C.** We used the MRI and CT images of each animal to guide the electrodes to lobule VI or VII of the cerebellar vermis. **D & E.** Simple (blue) and complex spikes (red) in two simultaneously recorded P-cells during saccades to various directions. Eye velocity is shown via the yellow curve. The complex spikes are also aligned to the onset of the visual target. Both cells exhibited a reduction in simple spikes during saccades, with a modulation pattern that lasted much longer than the saccade. CS probability during the 200 ms following target onset is quantified via the center plot. Baseline CS probability is shown by the brown circle at center. The target direction that produces the highest CS probability (CS-on) is estimated by the red line at center. **F.** Synchronization index during saccades to various directions. This index quantified the probability of synchronization with respect to chance at 1 ms time bins. Eye velocity is indicated by the black curve. Probability of synchronization is greatest for saccades in direction CS+180, reaching a peak at around saccade deceleration.
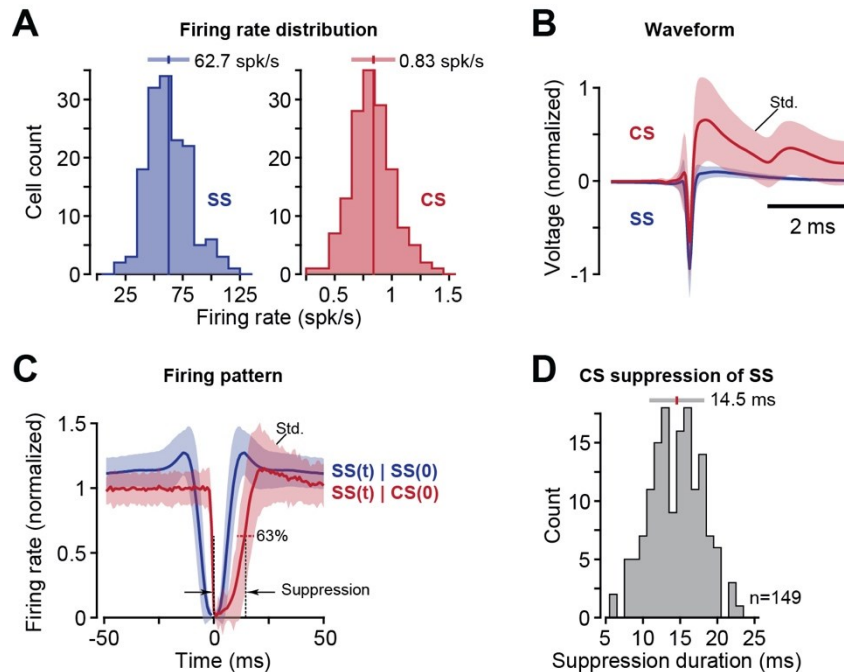
Figure 6.2. Properties of saccade-related P-cells in the marmoset vermis lobule VIa and VIIa-c.
A. Distribution of P-cell (n=149) average firing rates for simple spikes (blue) and complex spikes (right). The bar at the top indicates mean and standard deviation. **B.** Waveforms for the simple and complex spikes. The waveforms were normalized by setting the cell's mean voltage to 0 and the maximum negative going simple spike deflection to -1. Error bars are standard deviation. **C.** Within-cell interactions between simple and complex spikes. The blue curve shows the firing rate of simple spikes at time t, given that the cell produced a simple spike at time zero, labeled as SS(t)|SS(0). The red curve shows the firing rate of simple spikes at time t, given that the cell produced a complex spike at time zero, labeled as SS(t)|CS(0). Simple spike rates for each P-cell were normalized with respect to average simple spike firing rate as computed over the entire recording session. Error bars are standard deviation. **D.** Suppression period of simple spikes following production of a complex spike. Suppression period for each P-cell was defined as the duration of time after a complex spike that was required before the simple spike firing rate recovered 63% of its pre-complex spike value. The bar at the top indicates mean and standard deviation.

Data from a pair of simultaneously recorded P-cells are shown in Figure 6.1D&E. Despite their proximity (50 μm), one cell tended to pause its SS activity with saccades, while the neighboring cell tended to pause then burst. Indeed, in both cells the SSs remained modulated long after the saccade ended. However, when the target appeared to the left of the fovea, both P-cells responded with an increased probability of CS (center subplot of Figure 6.1D&E), and when the target appeared to the right, both decreased their CS probability. Production of a CS in one P-cell was followed by 10-20 ms suppression of SS in that cell but not the neighboring cell (Figure 6.3B&D). Yet, the SSs shared a degree of temporal coordination: the probability of observing a SS in P-cell 2 in a 1 ms window of time increased by 39.3% if P-cell 1 happened to generate a SS during the same period (Figure 6.3C).

To quantify SS coordination during saccades, we measured the probability of synchronization with respect to chance (Han et al., 2018; Wong et al., 1993). The synchronization index quantified the probability that both cells fired a spike during a 1 ms interval of time, corrected for the independent probabilities of spiking in each cell (all probabilities were conditioned on a saccade to a specific direction at time zero). Thus, the index determined whether there was greater synchrony than expected, where chance was quantified from the saccade related changes in the average firing rates of each neuron.
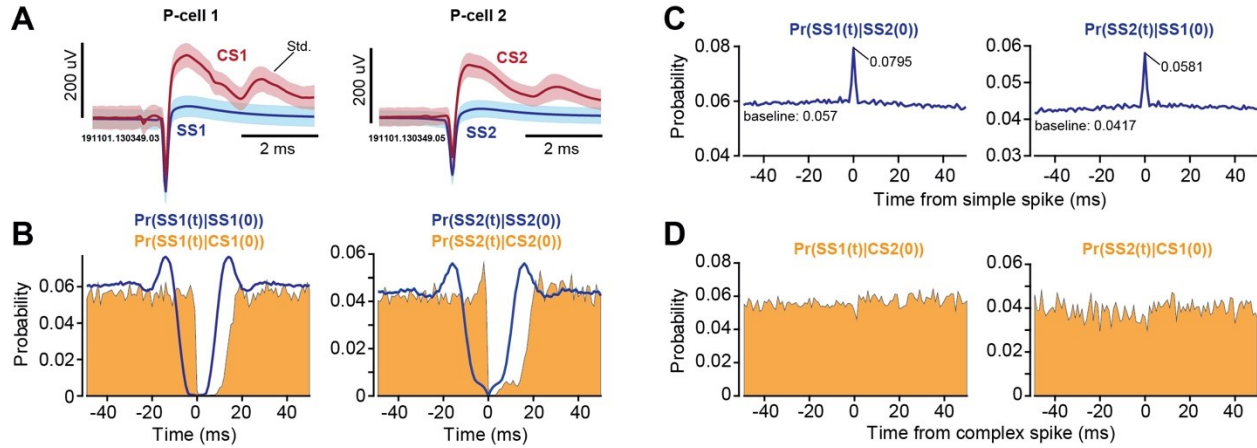
Figure 6.3. Spike timing properties of a sample pair of simultaneously recorded P-cells.
These are the same cells as in Figure 6.1. **A.** Simple and complex spike waveforms. Error bars are standard deviation. **B.** The curve Pr(S1(t)|S1(0)) quantifies the probability of a simple spike in P-cell 1 at time t, given that P-cell 1 produced a simple spike at time zero. This quantifies the simple spike refractory period. The curve Pr(S1(t)|C1(0)) quantifies the probability of production of a simple spike in P-cell 1 at time t, given that P-cell 1 produced a complex spike at time zero. This indicates the complex spikes induced suppression of simple spikes. **C.** The curve Pr(S1(t)|S2(0)) quantifies the probability of production of a simple spike in P-cell 1 at time t, given that a simple spike was produced by P-cells 2 at time zero. **D.** The curve Pr(S1(t)|C2(0)) quantifies the probability of production of a simple spike in P-cell 1 at time t, given that a complex spike was produced by P-cells 2 at time zero. Bin size is 1 ms in the probability plots.

Remarkably, while both cells reduced their firing rates during saccades, their spikes became more synchronized (Figure 6.1F). Moreover, the probability of synchronization depended on the direction of the saccade: it was greatest when the saccade was toward the direction for which complex spikes were least likely (CS+180). In these two P-cells, the probability of SS synchronization reached a maximum around the time when the saccade decelerated and came to a stop.

To analyze the data in our population, we began by measuring the CS response of each P-cell to the various visual events (primary, corrective, or central target). For each event, we estimated the target direction that produced the largest CS probability (CS-on, Figure 6.4A). We found that the direction of CS-on remained consistent across the various targets (Figure 6.4C, within cell comparison of direction of

CS-on in response to visual event type, primary vs. secondary target 0.1±5.4°(SEM), $t_{(148)}$=0.03, p=0.98; primary vs. central target -3.7±5.2°, $t_{(148)}$=-0.71, p=0.48, secondary vs. central target -3.8±5.3°, $t_{(148)}$=-0.72, p=0.47). Thus, we combined the CS response for all three visual events and used the results to define the CS-on direction of each P-cell (Figure 6.4A, all targets).

The distribution of CS-on directions across the P-cells varied widely (Figure 6.4D). However, the CS-on direction was not random. Rather, it varied with the location of the cell in the vermis: P-cells in the left vermis tended to have right-ward CS-on, while P-cells on the right had left-ward CS-on (Figure 6.4F). When target directions were represented with respect to CS-on, the result was a unimodal tuning function that described the CS response of P-cells following presentation of a visual target (Figure 6.4E). The probability of CS increased by 43.5±2.6% (mean±SEM) above baseline when the stimulus was presented in direction CS-on but decreased by 34.6±1.7% below baseline when it was presented in direction CS+180.

Because each target instructed a saccade, we wondered whether the CS response was due to the sudden onset of the stimulus or associated with the movement that followed. Although our experiment was not designed to specifically answer this question, we made an interesting observation. Saccades that were made in response to visual targets were preceded with large changes in CS firing rates (Figure 6.4B, saccades to targets). However, saccades that were not instructed by a target, but were in the same direction and amplitude, were preceded by significantly smaller modulation of CS firing rates (Figure 6.4B, "other saccades", average CS firing rate 50 ms before saccade onset in direction CS-on, paired t-test, $t_{(296)}$=6.4, p=5×10^{-10}, direction CS+180, $t_{(296)}$=-2.8, p=0.005). Thus, the complex spikes were modulated primarily in response to sensory events that instructed movements, but not when similar movements were made spontaneously.
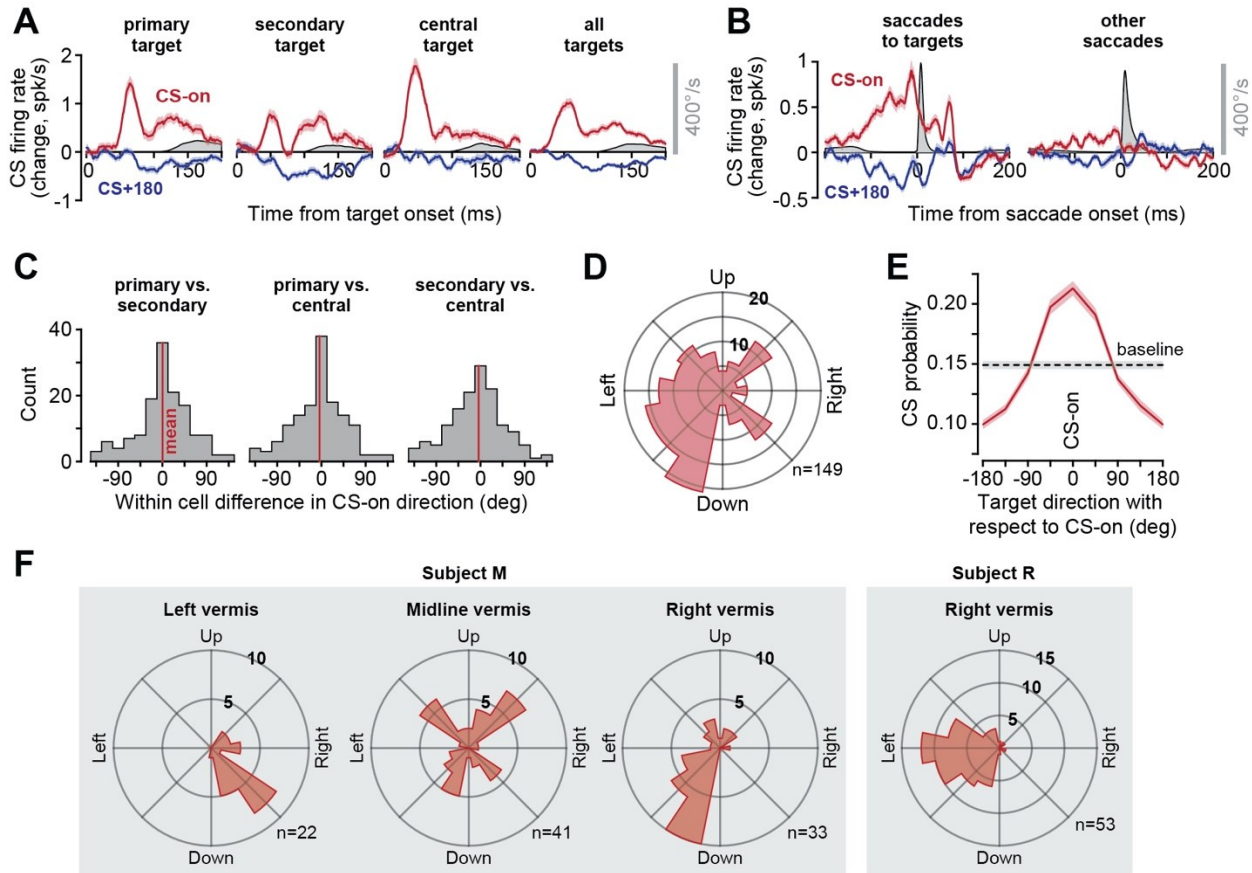
Figure 6.4. Complex spikes exhibited tuning with respect to the direction of target.
A. CS response aligned to target onset. For each type of target, the direction of stimulus that produced the greatest probability of CS was labeled as CS-on. Eye velocity is shown in gray. **B.** CS response aligned to saccade onset. Modulation of CS response was present before saccades that were visually instructed. The response was muted before "other saccades". **C.** Within cell difference between CS-on directions as computed following the onset of the primary target, the secondary target, and the central target. We found no systematic differences in the estimate of CS-on between various types of targets, and thus combined the response for all targets to compute the CS-on of each P-cell. **D.** Distribution of CS-on across the population of P-cells. **E.** CS tuning function. **F.** Distribution of directions of CS-on in various regions of the vermis in two animals. Error bars are SEM.

The simple spikes exhibited a variety of patterns during saccades: some P-cells increased their activity, some decreased their activity, while others produced more complicated patterns (Figure 6.6A). The activity patterns did not separate the cells into clusters, but rather formed a continuum (Figure 6.6B). For the sake of labeling, we divided the P-cells into two groups: pausers and bursters. 48% of our cells were bursters, while 52% were pausers (Figure 6.5C). To quantify how well their activities were modulated during saccades, for each P-cell we measured the change in SS rates aligned to saccade onset

and computed a z-score (Figure 6.5A). Indeed, the P-cell SS rates were modulated strongly during the movements (Figure 6.5B, z-score 7.5±0.3, mean±SEM).
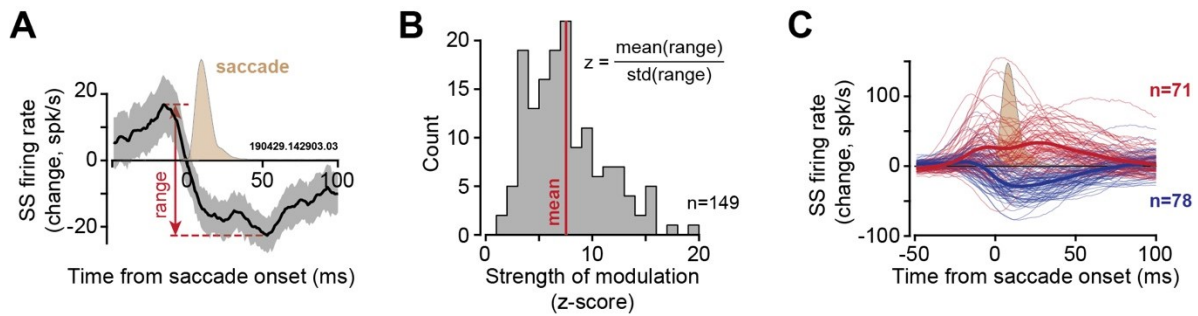
Figure 6.5. Modulation of simple spikes in individual P-cells during saccades.
A. Data from a single P-cell, quantifying the change in simple spike rates, aligned to saccade onset. Range was defined as the maximum change in firing rate in the pre- to post-saccade period. The light brown curve shows the average saccade velocity (peak value is 475 °/s). Error bars are standard deviation computed via bootstrapping. **B.** Strength of saccade-related modulation of each P-cell was defined via a z-score. The z-score computed the mean of the range of the firing rates divided by the standard deviation of the range (7.5±0.3, mean±SEM). Strongly modulated P-cells were those that had a z-score greater than 3, composing 96% (142 out of 149) of the cells in the population. **C.** Change in simple spike rates with respect to baseline for the bursters (red) and pausers (blue). Baseline firing rate is defined as the average firing rate as measured during the entire recording session.

We next organized the P-cells based on a computational model that incorporated an important anatomical feature of the cerebellum: P-cells that have similar CS tuning not only receive similar olivary inputs, but they also are likely to be part of a single olivo-cerebellar module (Apps & Garwicz, 2000; C I De Zeeuw, Van Alphen, et al., 1997; Ruigrok & Voogd, 2000; Shadmehr, 2020; Sugihara, 2011; Sugihara & Shinoda, 2007). This anatomical organization implied that to estimate activity of a population of P-cells that belonged to an olivo-cerebellar module, we needed to compute saccade direction with respect to the CS-on of each P-cell. By using this coordinate transformation, we estimated the population SS response in a hypothetical olivo-cerebellar module when saccades were in direction CS-on, CS+90, etc. (Figure 6.6C).

Unsurprisingly, activities of the bursters and pausers were modulated long after the saccade ended (Figure 6.6C, top row). However, when the activities across all cells were organized into a population and summed, the response exhibited a clear pattern: there was a burst that preceded saccades in all directions. Notably, for direction CS+180 the burst was followed by a pause that ended near saccade termination (Figure 6.6C, bottom row). This burst-pause pattern was somewhat weaker for saccades in direction CS±90, and the pause was missing entirely for saccades in direction CS-on.
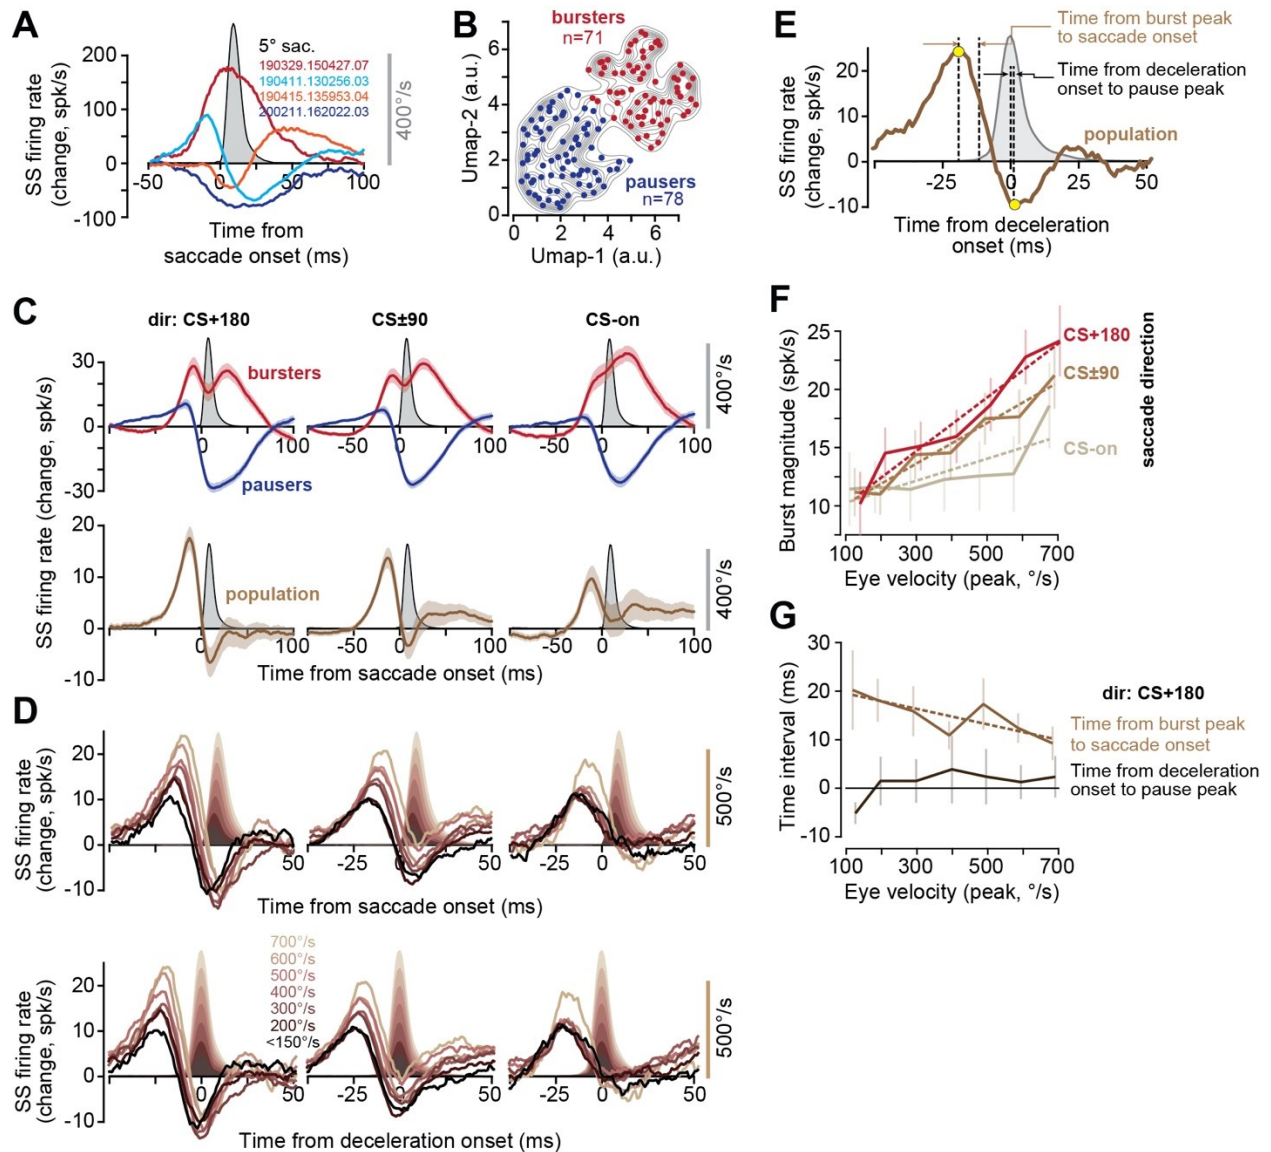
Figure 6.6. Population response of simple spikes encoded saccade direction, peak velocity, and the onset of deceleration.

A. Average change in the firing rates of four representative P-cells with respect to baseline, during saccades (data collapsed across all directions). **B.** Clustering of saccade-aligned change in firing rates for all P-cells, using the algorithm UMAP. Separating the data into two clusters produces bursters (red) and pausers (blue). **C.** Activities of the bursters and pausers during saccade in various directions. The population response is the sum of firing rates in all P-cells. **D.** Population response aligned to saccade onset and deceleration onset. The burst tends to grow with saccade velocity and shifts forward in time, but the pause remains invariant with respect to the onset of deceleration. **E.** Quantification of the population response with respect to saccade kinematics. **F.** Magnitude of the burst before saccade onset as a function of saccade peak velocity in various directions. **G.** Timing of the burst with respect to saccade onset decreased with increased velocity, while the timing of the pause with respect to deceleration onset remained invariant.

The burst increased with the velocity of the impending saccade (Figure 6.6D). However, the rate of increase as a function of velocity was direction dependent, showing the greatest gain for saccades in direction CS+180 (Figure 6.6F, dir CS+180, $F_{(1,5)}$=88.3, p=0.0002). This pattern is called a gain-field, confirming earlier findings in the macaque cerebellum 16. Thus, before saccade onset the P-cells appeared to inhibit the nucleus with a magnitude that depended on the velocity and direction of the forthcoming saccade.

In direction CS+180, the magnitude of the burst increased with saccade velocity, but its timing shifted forward: the period from the peak of the burst to the onset of the saccade (Figure 6.6E) became smaller as saccade velocity increased (Figure 6.6G, r^2=0.65, $F_{(1,5)}$=9.4, p=0.027). As the saccade started toward direction CS+180, the activity changed from a burst to a pause (Figure 6.6D). However, unlike the burst that preceded the saccade, the pause magnitude and timing remained invariant with respect to saccade velocity (Figure 6.6G, time of deceleration onset to pause peak as a function of velocity, $F_{(1,5)}$=2.4, p=0.18, Figure 6.6D, rate of pause as a function of velocity, $F_{(1,5)}$=1.1, p=0.34). Critically, despite a 7-fold change in velocity, the timing of the maximum pause was unchanged with respect to saccade deceleration onset (Figure 6.6G). Thus, with increased saccade velocity the burst magnitude increased, and its timing shifted forward. However, regardless of saccade velocity, the pause that followed the burst was time-locked to the onset of saccade deceleration.

This invariant relationship between the timing of the pause in firing rates and the onset of saccade deceleration (in direction CS+180) raised the possibility that the P-cells were signaling when the nucleus cells should fire, presumably stopping the saccade. However, entraining the nucleus neurons would be more efficient if the P-cells not only reduced their firing rates (thus disinhibiting the nucleus), but also synchronized their spikes (Gauck & Jaeger, 2000; Person & Raman, 2011). To test this hypothesis, we computed the probability of synchronized firing in our population of simultaneously recorded P-cells.

Production of a SS in a P-cell was associated with 31% increase in the probability (with respect to chance) that there would be a simultaneous (1 ms window) SS in another P-cell (Figure 6.7A, top row). Similarly, production of a CS in a P-cell increased the probability of observing a CS in another P-cell at ±5 ms latency by 227% with respect to chance (Figure 6.7A, bottom row). Finally, production of a CS in one P-cell reduced the probability of SS in another P-cell at 1 ms latency by 27% (Figure 6.7A, middle row). All of these observations are consistent with earlier findings in mice, demonstrating that nearby P-cells not only share a degree of spike synchrony 6, but that a CS in one cell can briefly suppresses SS in

another cell (Han et al., 2020). Furthermore, the simultaneously recorded P-cells tended to have very similar CS-on directions (Figure 6.7B, between cell difference -4.4±6.3°). However, did the P-cells synchronize their activities during saccades?
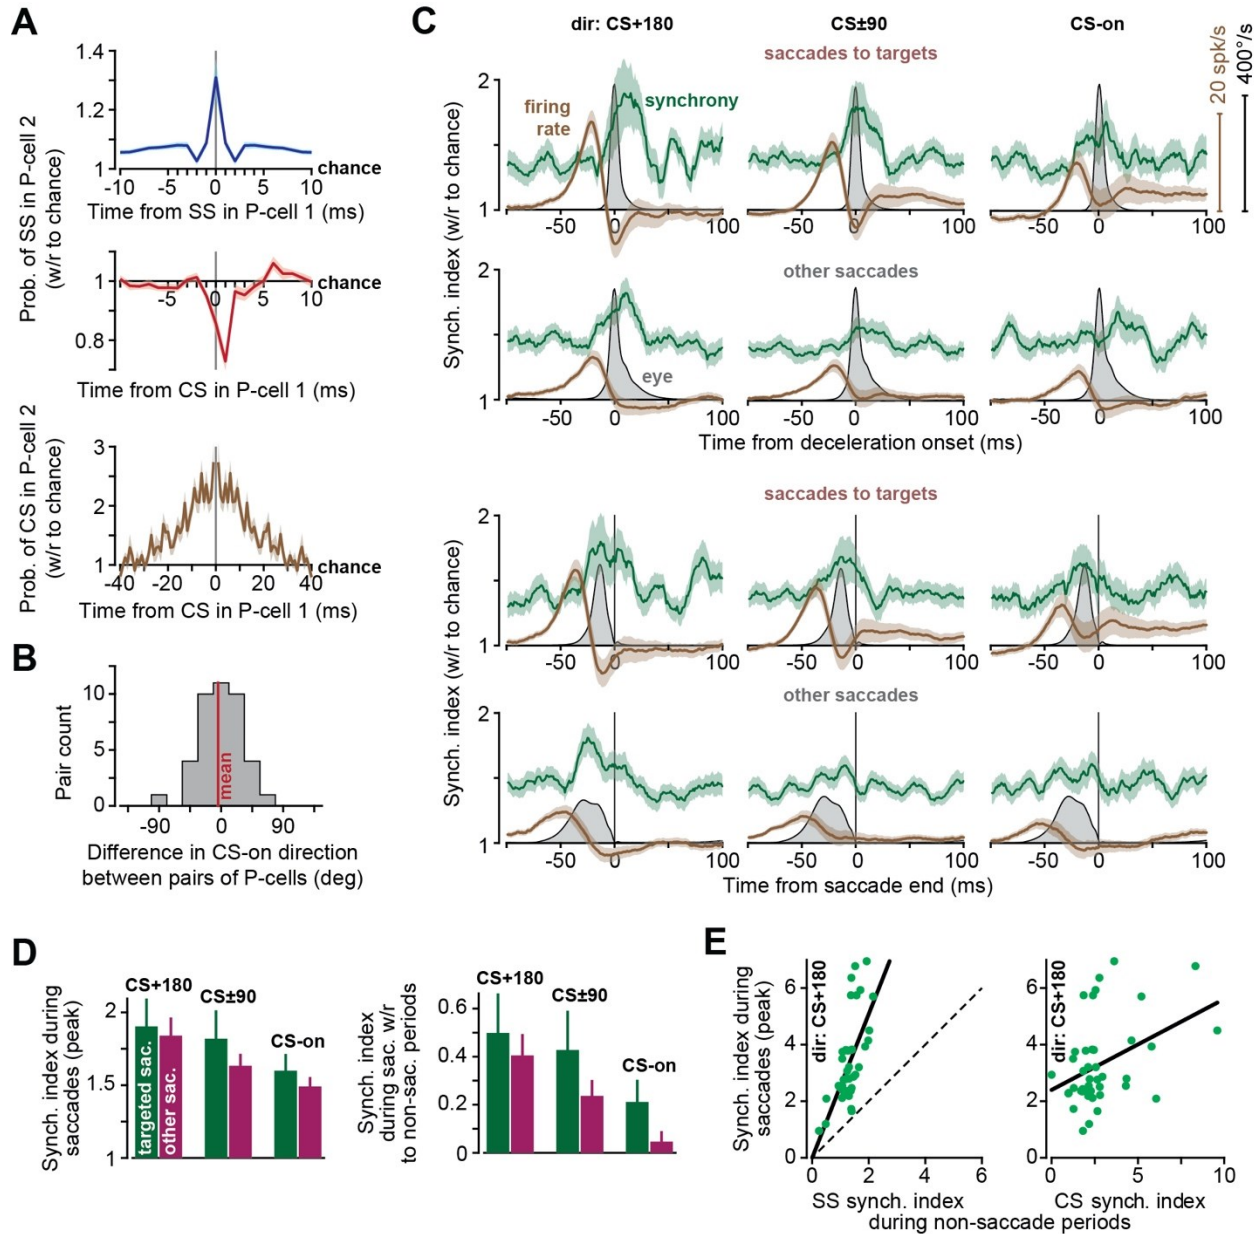
Figure 6.7. P-cells synchronize their spikes during saccade deceleration.
A. Probabilities of spike synchronization in pairs of P-cells during the entire recording session (41±2 minutes, mean±SEM). Top: probability of simple spike in P-cell 2 at time point t (with respect to chance), given that a simple spike occurred in P-cell 1 at time zero. Middle: probability of simple spike in P-cell 2, given that a complex spike was produced in P-cell 1 at time zero. Bottom: probability of complex spike in P-cell 2 given that a complex spike was produced in P-cell 1 at time zero. Bin size is 1 ms. **B.** Difference in CS-on directions among pairs of simultaneously recorded P-cells. **C.** Synchronization index (green) and firing rates (brown) for targeted saccades and other

saccades. In the top two rows, data are aligned to deceleration onset. In the bottom two rows, data are aligned to saccade end. Firing rate is the population response. Bin size is 1 ms. **D.** The magnitude of the synchronization index during saccades (peak value) for saccades in various directions. **E.** Left plot shows the magnitude of the synchronization index during saccades (peak value) in direction CS+180, with respect to SS synchronization index as measured during non-saccade periods. Dashed line is identity. Center plot shows the synchronization index during saccades with respect to complex spike synchronization index as measured during non-saccade periods (1 ms bin for SS and 10 ms bin for CS). Error bars are SEM.

Before saccade onset there was a burst in the P-cell population response. Surprisingly, probability of synchronization remained at baseline (Figure 6.7C, note that before saccade onset, synchronization is greater than chance because at baseline the neighboring P-cells are more synchronous than chance). As the saccade started and then began to decelerate (in direction CS+180) the firing rates fell, but the synchronization index increased, reaching its peak probability after deceleration onset but before saccade end (2.3±3.5 ms after deceleration onset and -5.2±4.1 ms before saccade end). That is, during saccade deceleration the few spikes that remained were significantly more likely than chance to be synchronized. Indeed, synchronization varied with the direction of the saccade: the greatest synchronization occurred in direction CS+180 (Figure 6.7D, Repeated Measure ANOVA, significant effect of direction, $F_{(2,82)}=1269.0$, $p=0.021$). Thus, the synchronization pattern was strongest in the direction associated with the smallest probability of complex spikes.

Like targeted saccades, during other saccades the SS firing rates exhibited a burst before saccade onset, with a magnitude that was largest for direction CS+180. These saccades also had a synchronization that peaked before saccade end (Figure 6.7C), with a probability that was largest for direction CS+180 (Figure 6.7D). Thus, while complex spikes showed strong modulation before targeted saccades but not task-irrelevant saccades, SS firing rates were modulated during all saccades. Indeed, regardless of whether saccades were target-driven or not, simple spikes reached their greatest probability of synchrony as the movement decelerated and came to a stop.

To check whether the increased synchronization during saccades was an artifact of the change in firing rates, we performed a simulation of spiking neurons that burst and paused like cells in our population, but had independent probabilities of spike timing. The changes in rates during simulated saccades produced a synchronization index that remained at chance (Figure 6.8).

Finally, we found that pairs of P-cells that had greater CS synchrony, as measured during non-saccade periods, tended to have greater SS synchrony during saccades (Figure 6.8E, $F_{(1,40)}=7.6$, $p=0.009$). Pairs of P-cells that had greater SS synchrony during non-saccade periods also exhibited greater SS

synchrony during saccades (Figure 6.8E, $F_{(1,40)}$=29.0, p=3x10^{-6}). However, the gain of the relationship between saccade period synchrony and general synchrony was significantly greater than 1 ($F_{(1,40)}$=10.7, p=0.002). Thus, although the timing of simple spikes among nearby P-cells was generally coordinated, during saccades this coordination was greatly enhanced, especially during deceleration.
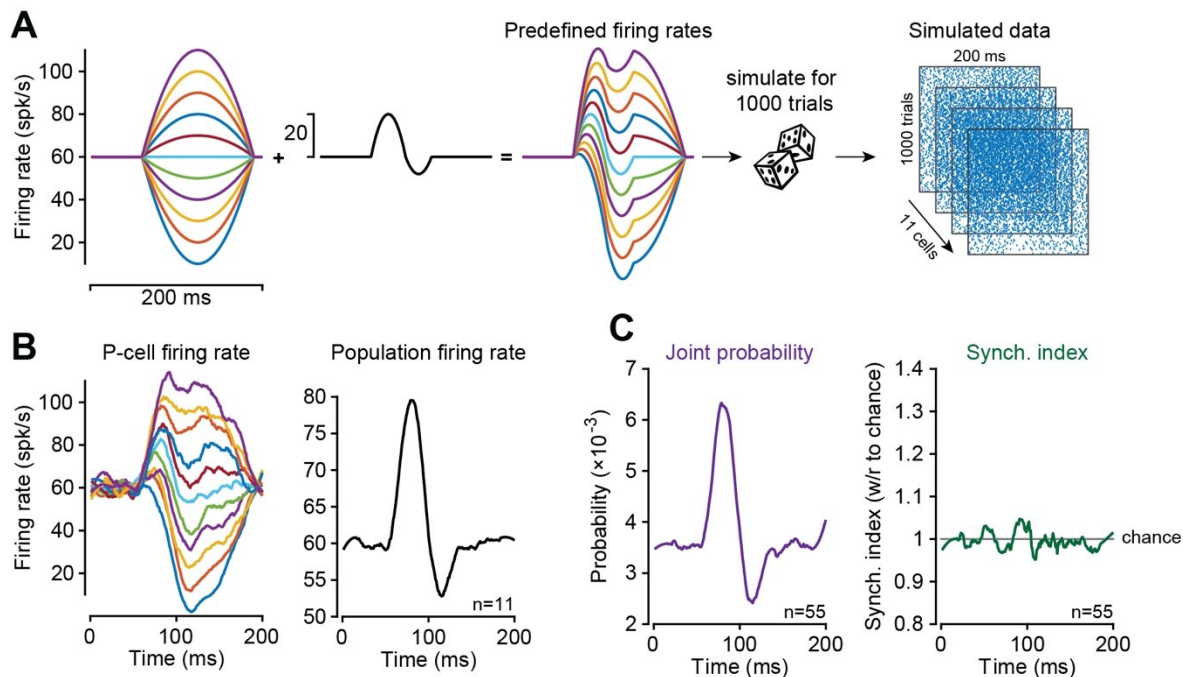
Figure 6.8. Synchronization index among simulation of independent spiking neurons that burst and pause.
    To check whether the increased synchronization index during saccades was an artifact of the change in firing rates of P-cells, we performed a simulation of spiking neurons that burst and paused like the cells in our population. However, the simulated cells had independent probabilities of spiking. **A.** The simulated population consisted of 11 neurons ranging from bursting to pausing in their activity. The plots show the construction of the average firing rates. We started with a unimodal function of various amplitudes, added the response in the middle row to each neuron, producing the data on the right, representing the instantaneous average firing rate of each simulated neuron. **B.** Data from 1000 simulated trials. The left plot shows the average firing rate of each simulated cell. Right plot shows the population firing rate of the simulation. **C.** Joint probability, and synchrony index of 55 pairs of cells (11 cells taken 2 at a time). Joint probability reflected modulation of firing rates, but because spike timing in each neuron was independent, the synchrony index remained at chance level.

6.4. Discussion

In describing symptoms of cerebellar damage, Holmes (1939) noted that "the most obvious errors are seen toward the end of the movement during which the speed of the affected limb is often unchecked until the object is reached or even passed." For example, during an outward reach, many interposed

nucleus neurons of the cerebellum produce their greatest discharge during deceleration, with spiking activity that plays a causal role in stopping the movement (Becker & Person, 2019). Similarly, following inactivation of the fastigial nucleus, extraocular motoneurons that act as saccade agonists produce an abnormally large amount of activity during the deceleration period of ipsilateral movements (Kojima et al., 2014), resulting in saccades that overshoot the target (Buzunov et al., 2013; Robinson et al., 1993). Thus, the computations that are performed by the cerebellar cortex are critical for monitoring the ongoing motor commands and predicting when the movement should be stopped. Yet, P-cell simple spikes are often modulated long after the movement ends (Hewitt et al., 2015; Ishikawa et al., 2014; Kase et al., 1980; Thier et al., 2000; Tomatsu et al., 2016).

Here, we found that if P-cells were organized based on their complex spike response to visual stimuli, their population simple spikes produced a burst-pause pattern that started before saccade onset and ended with saccade termination. Changes in saccade velocity affected the timing and magnitude of the burst, but the pause remained time-locked to deceleration onset. Critically, in simultaneously recorded P-cells, during the pause period the probability of spike synchronization reached a maximum value. The resulting inhibition-disinhibition pattern of firing rates, coupled with spike synchronization, hints that the P-cells attempted to entrain the nucleus neurons specifically at the onset of deceleration (Gauck & Jaeger, 2000; Özcan et al., 2020; Person & Raman, 2012).

What might be the behavioral consequence of this synchronization? The synchronization probability was greatest for saccades that were in the direction that coincided with the least probability of complex spikes (CS+180). For both saccades and limb movements, the CS tuning of a P-cell is likely aligned with the direction of action of the downstream nucleus neuron (Ekerot et al., 1995; Herzfeld et al., 2018). For example, trial-to-trial analysis of the effects of complex spikes on simple spikes and behavior suggests that P-cells that have CS-on tuning to the left project to nucleus neurons that have a downstream direction of action that indirectly promotes production of leftward forces (Herzfeld et al., 2018). This implies that during a saccade in direction CS+180, the increased synchrony combines with disinhibition (peak pause) to entrain the nucleus neurons during deceleration (Gauck & Jaeger, 2000), producing downstream forces that are aligned with the CS-on of the parent P-cells. As a result, the effect of synchronization of P-cells, coupled with disinhibition of the nucleus, is likely the production of forces that oppose the direction of movement, bringing it to a stop.

To our knowledge, one earlier work had reported increased spike synchrony among P-cells during movements. Using multiunit signals (i.e., not single unit isolation of spikes), Heck et al. (2007)

found increased covariance between P-cells during reaching movements (in rats). That work found that synchrony was most prominent as the hand approached the target, i.e., during deceleration. Here, we found that P-cell firing rates and spike synchrony were coordinated, especially among populations that had a common CS tuning.

Our results were obtained in the marmoset, a New World primate that like macaques and humans relies on saccades to explore its environment. Like macaques, individual P-cells in lobule VI and VII of marmosets produced simple spikes that were bursting, pausing, or a combination of the two, with no obvious relationship to the direction of the saccade or its velocity. However, following onset of a visual stimulus, the P-cells received information from the olive regarding the location of the stimulus with respect to the fovea, producing complex spikes that were tuned to the direction of the target. In both species, this tuning was anatomically organized, with P-cells in the right vermis showing highest CS probability for targets to the left. A similar anatomical representation of contralateral stimuli/movements of the arm has recently been noted in the vermis of mice (Wagner et al., 2021).

In both marmosets and macaques, when we organized P-cells based on their complex spike tuning properties, the simple spikes produced firing rates that varied strongly with direction and velocity of the movement (Herzfeld et al., 2015, 2018). Indeed, the gain of the response with respect to velocity was highest when saccades were in direction CS+180, and in both species the response was a burst followed by a pause that ended as the movement came to a halt. The consistency of these results across species suggests that viewing P-cell activity through the lens of population coding (Shadmehr, 2020), i.e., a lens in which the climbing fibers organize the P-cells into olivo-cerebellar modules (Chris I. De Zeeuw et al., 2011), may provide a key for unlocking the language with which the cerebellar cortex encodes information.

How does synchronization arise during a specific phase of a movement? P-cells that show elevated synchrony in their complex spikes also tend to fire simple spikes more synchronously (Wise et al., 2010). Indeed, here we found that P-cells with greater complex spike synchrony tended to have greater simple spike synchrony during saccades. In addition, we found that following a CS in one P-cell, after a 1 ms delay there was a 1-2 ms period of simple spike suppression in the neighboring P-cell, confirming recent findings in mice (Han et al., 2020). One possibility is that P-cells that receive a common input from the olive generate a synchronized CS, leading to SS suppression, which may then be followed by a synchronized resumption of SS firing (Bengtsson et al., 2011). However, SS synchrony was greatest in direction CS+180, i.e., the direction for which there was the smallest probability of CS.

Furthermore, complex spikes showed modulation before the onset of targeted saccades but not task-irrelevant saccades, whereas simple spikes showed synchronization for both types of saccades. These observations make it seem unlikely that during saccades, presence of complex spikes played a role in synchronization of simple spikes.

Synchrony is also present in P-cells that are likely to have common parallel fiber inputs (on-beam) but different climbing fibers (Ebner & Bloedel, 1981; Heck et al., 2007). Indeed, here we found that nearby P-cells not only had greater than chance levels of synchrony, but that cells with greater SS synchrony in general had much greater than expected SS synchrony during saccades. Thus, it is possible that SS synchrony arises from a shared input from ascending granule cell axons that are positioned directly beneath the P-cells and receive inputs from the same mossy fibers (Heck et al., 2007). However, why this synchrony would be focused during the deceleration phase, particularly for saccades in direction CS+180, is unclear. Because granule cells also recruit molecular layer interneurons (Dizon & Khodakhah, 2011), synchronization of P-cells may engage a network wide organization to overcome inhibition by the basket and stellate cells (de Solages et al., 2008).

In a typical artificial neural network, information transfer from one layer to the next is via firing rates of neurons, and learning modifies synaptic weights to change the activity of each neuron and minimize error in the output layer. The cerebellum resembles a 3-layer network where learning is at least partially guided by the climbing fibers (Raymond & Medina, 2018; Shadmehr, 2020). Our results demonstrate that the information that is transmitted from the P-cells to the nucleus is encoded in an exquisite coordination of firing rates and synchronization. This implies that when there is error in performance, cerebellar learning cannot simply focus on changing the P-cell firing rates. Rather, learning must also alter network wide synchronization. This conjecture predicts that complex spikes that arise following movement errors not only promote learning via changes in the activity of individual P-cells (Herzfeld et al., 2018; Medina & Lisberger, 2008; Yang & Lisberger, 2014a), but may also alter the synchronization patterns of populations of P-cells. Learning to transfer information via synchronization is an exciting new direction with which to explore the function of the cerebellum.

Chapter 7. Conclusions and future directions

To move accurately, the brain relies on internal models that predict the consequences of motor commands. How does the brain adaptively control our movements? To answer this question, we chose saccadic eye movements as our model movement. Due to the ballistic nature of saccades, the brain relies on its internal models to ensure the accuracy of this type of movement. In this dissertation, through behavioral and neural measurements, we studied the underlying mechanisms of the adaptive control of saccades in the brain.

In Chapter 2, we asked how does dopamine release in milliseconds before the onset of a movement serve to invigorate the ensuing movement? We designed an innovative behavioral paradigm in which the reward prediction error (RPE) events occurred just before the onset of a movement. Participants were presented with face or noise images and then during the execution of the primary saccade, we probabilistically changed the position and content of images, encouraging a secondary saccade. The change in the content produced an RPE event before the onset of the secondary saccade. We produced a combination of large- or small-, negative- or positive-RPE events. We found that the reaction time of the secondary saccade was affected in an orderly fashion by the magnitude and direction of the preceding RPE event: the most vigorous saccades followed the largest positive-RPE, whereas the least vigorous saccades followed the largest negative-RPE. This suggests that RPE events, which are thought to transiently alter the release of dopamine, modulate the vigor of the ensuing movement.

Our interpretations regarding RPE events relied on the assumption that the opportunity to gaze at an image served as a proxy for reward acquisition. One potential future direction to our work is to look at the effect of history on RPE and movement vigor. For example, if participants experience a change in the image content from noise to face for tens of trials, the event will become predictable, and it will no longer serve as a prediction error. One can measure saccade reaction time throughout this process and assess how the history of image content affects movement vigor. Another future direction is to limit the number of available images. In this way, after experiencing each image for a dozen trials, their content will not be salient anymore and this will potentially break our initial assumption that gazing at images are rewarding. Measuring saccade vigor as participants experience the same image multiple times can provide insight into how humans adjust the value of what they gaze at.

In Chapter 3, we used saccade adaptation as a model of sensorimotor learning and explored whether imposing a cost on the time spent correcting for error could modulate learning from error. We designed a novel paradigm and combined random dot motion discrimination with saccade adaptation to impose a cost on movement error. In the resulting paradigm, movement errors resulted in corrective saccades, but those corrections took time away from acquiring information in the discrimination task. By changing the coherence of the random dot motion, we modulated the error cost in our task. Our results demonstrated that when the error cost was large, pupil diameter increased, and the brain learned more from error. However, when the error cost was small, the pupil constricted, and the brain learned less from the same error. Thus, during sensorimotor adaptation, the act of correcting for error carries an implicit cost for the brain. Modulating the landscape of this implicit loss regulates the rates of sensorimotor learning.

In our task, we kept the size of the target jump (perturbation) consistent and controlled the statistics of the error that the subjects experienced. In addition to the cost of error, earlier works showed that the brain's sensitivity to error can be influenced by factors such as the size of error (Hanajima et al., 2015; Marko et al., 2012), stability of error in an environment (Herzfeld, Vaswani, et al., 2014; Leow et al., 2016), and variability in the trial-by-trial sequence of errors (Albert et al., 2021). One potential future direction to our work is to expand the current experiment design and incorporate a systematic variation in the statistics of the error. This way, one can evaluate the interaction between different factors on the brain's sensitivity to error. Do these factors engage similar brain circuitries and their effects are additive or do they engage different mechanisms and their effects are orthogonal to each other?

In Chapter 2 and Chapter 3 we focused on behavioral correlates of adaptive control of saccades. We proposed brainstem dopaminergic system modulates movement vigor. Next, we showed that modulation of implicit error cost can modulate the rate of saccade adaptation. Animal lesion studies and results from cerebellar patients have established the cerebellum as the main locus for control of saccade accuracy. How do the principal cells of the cerebellar cortex, Purkinje cells (P-cells), control the accuracy of saccades? We introduced marmoset monkeys as a new animal model to study motor control and motor learning and recorded from their cerebellum while subjects performed a saccadic eye movement task.

In Chapter 4, we presented the details of our animal protocol that motivated the subjects to perform more than 1000 trials of reward-driven saccade tasks on a daily basis. We explained how we

designed and surgically implanted CT-guided and subject-specific head-posts. Having control over the head-post design was an integral part of our system because the design of our electrode guidance system was based on defining a coordinated system that was attached to the head-post. We explained the details of designing the alignment tool of the electrode guidance system based on pre-operative and post-operative CT and MRI images of each animal. This system enabled us to target the cerebellar cortex and robustly record simultaneous P-cells during saccade tasks.

The design of our system provided an accurate approach to target deep brain areas and robustly record from those areas. In our case, we were mostly interested in recording from the cerebellar cortex. In our design, we placed the recording chamber over the visual cortex, and to reach the cerebellum we penetrated the visual cortex and then tentorium to access the cerebellar P-cells. The motivation for this placement was to prevent any intrusion of the neck muscle attached to the back of the animal's head. In this design, we avoided the use of thick electrodes (thicker than 150 μm) that could damage the visual cortex or thin electrodes (thinner than 70 μm) that mechanically could not penetrate the tentorium. In practice, we found that electrodes that have 70-150 μm thickness work the best with our setup. One major modification for our design would be to place the recording chamber directly above the cerebellar cortex and bypass the visual cortex. In this scenario, one can perform a micro-incision on the tentorium and directly access the cerebellar cortex without the risk of damaging the visual cortex or getting stopped by the tentorium.

After establishing the techniques to record from P-cells, we faced the next challenge. Solving the detection and attribution problem for sorting complex and simple spikes. In Chapter 5, we explained how we designed and implemented an open-source software package, named P-sort, to address the challenges of cerebellar neurophysiology. We presented examples of different complex spike waveforms and how to sort them using the P-sort. We analyzed data from P-cells recorded in three species: marmosets, macaques, and mice. We validated and compared our results with expert manual curation and state-of-the-art automatic spike sorters.

As discussed in Chapter 5, P-sort is semi-automated and relies on the input from the user to further curate the detected simple and complex spikes. P-sort extracts and visualizes different linear and non-linear features to help the user gain better insights about the data and the best approach for sorting it. With the rapid evolution of silicon probe technology, more labs are collecting data over many channels simultaneously. Sorting this type of data needs automatic algorithms that need minimal input from the user. To improve the development of automatic algorithms for cerebellar neurophysiology, we

have provided labeled data for over 300 P-cells in three species, along with statistical characteristics of their waveforms and firing patterns. Moreover, investigators can use their current spike sorting algorithm on our dataset and uncover the unforeseen limitations of the P-sort, contributing to further develop this software to help the cerebellar community.

Finally, in Chapter 6, we started to systematically analyze the recorded data from n=149 P-cells (42 pairs of simultaneous neurons) during saccades. We organized P-cells into populations that shared the same preference for error as measured via their complex spike response. While simple spike activity of individual P-cells was quite diverse: pausing, bursting, or the mixture of the two in different directions, the population activity produced a burst-pause pattern that started before saccade onset and ended with saccade termination. This reveals that P-cells transmitted information via a rate code that predicted the direction and velocity of the impending saccade. However, if P-cells synchronize the timing of their simple spikes, they could entrain downstream nucleus neurons and presumably control when to stop the saccade. We computed the probability of synchronized firing in our population of simultaneously recorded P-cells. Our results demonstrated that before saccade onset synchronization remained at baseline, but as the saccade started and then began to decelerate the synchronization index increased, reaching its peak probability after deceleration onset but before saccade end. That is, during saccade deceleration the few spikes that remained became synchronized. Thus, the combination of synchronization of spikes, with the disinhibition produced by the firing rates, is likely to produce spiking in the nucleus, signaling when to stop the saccade.

Our results showed that the cerebellar cortex transmits the information to the cerebellar nucleus in an exquisite coordination of firing rates and synchronization. Previous works established that complex spikes that arise following movement errors promote learning via changes in the activity of individual P-cells (Herzfeld et al., 2018; Medina & Lisberger, 2008; Yang & Lisberger, 2014a). Results of our study predict that the production of error-related complex spikes would also alter the synchronization pattern among population of P-cells. This prediction remains to be tested by future studies.

With cerebellar damage, movement speed is left unchecked until the goal is reached or even passed. A major difficulty has been to understand how the cerebellum contributes to adaptive control of movements. The new techniques and findings presented in this dissertation have paved the way to systematically study the function of the cerebellum using marmoset monkeys as the animal model and saccadic eye movements as the model movement. A new approach to explore the function of the

cerebellum can indeed be the effect of motor learning on transfer of the information via synchrony and rate code.

Bibliography

Abe, M., Schambra, H., Wassermann, E. M., Luckenbaugh, D., Schweighofer, N., & Cohen, L. G. (2011). Reward improves long-term retention of a motor memory through induction of offline memory gains. *Current Biology : CB*, *21*(7), 557–562. https://doi.org/10.1016/j.cub.2011.02.030

Achilly, N. P., He, L.-J., Kim, O. A., Ohmae, S., Wojaczynski, G. J., Lin, T., Sillitoe, R. V, Medina, J. F., & Zoghbi, H. Y. (2021). Deleting Mecp2 from the cerebellum rather than its neuronal subtypes causes a delay in motor learning in mice. *ELife*, *10*. https://doi.org/10.7554/eLife.64833

Aharon, I., Etcoff, N., Ariely, D., Chabris, C. F., O'Connor, E., & Breiter, H. C. (2001). Beautiful faces have variable reward value: fMRI and behavioral evidence. *Neuron*, *32*(3), 537–551. https://doi.org/10.1016/s0896-6273(01)00491-3

Albert, S. T., Jang, J., Sheahan, H. R., Teunissen, L., Vandevoorde, K., Herzfeld, D. J., & Shadmehr, R. (2021). An implicit memory of errors limits human sensorimotor adaptation. *Nature Human Behaviour*, *5*(7), 920–934. https://doi.org/10.1038/s41562-020-01036-x

Albert, S. T., & Shadmehr, R. (2018). Estimating properties of the fast and slow adaptive processes during sensorimotor adaptation. *Journal of Neurophysiology*, *119*(4), 1367–1393. https://doi.org/10.1152/jn.00197.2017

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. *The European Journal of Neuroscience*, *12*(1), 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. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *19*(24), 10931–10939. http://www.ncbi.nlm.nih.gov/pubmed/10594074

Batton, R. R., Jayaraman, A., Ruggiero, D., & Carpenter, M. B. (1977). Fastigial efferent projections in the monkey: an autoradiographic study. *The Journal of Comparative Neurology*, *174*(2), 281–305. https://doi.org/10.1002/cne.901740206

Bayer, H. M., & Glimcher, P. W. (2005). Midbrain dopamine neurons encode a quantitative reward prediction error signal. *Neuron*, *47*(1), 129–141. https://doi.org/10.1016/j.neuron.2005.05.020

Becker, M. I., & Person, A. L. (2019). Cerebellar Control of Reach Kinematics for Endpoint Precision. *Neuron*, *103*(2), 335-348.e5. https://doi.org/10.1016/j.neuron.2019.05.007

Behnel, S., Bradshaw, R., Citro, C., Dalcin, L., Seljebotn, D. S., & Smith, K. (2011). Cython: The Best of Both Worlds. *Computing in Science & Engineering*, *13*(2), 31–39. https://doi.org/10.1109/MCSE.2010.118

Bell, C. C., & Grimm, R. J. (1969). Discharge properties of Purkinje cells recorded on single and double microelectrodes. *Journal of Neurophysiology*, *32*(6), 1044–1055. https://doi.org/10.1152/jn.1969.32.6.1044

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*(4), e18822.

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. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *31*(41), 14708–14720. https://doi.org/10.1523/JNEUROSCI.3323-11.2011

Bruno, R. M., & Sakmann, B. (2006). Cortex is driven by weak but synchronously active thalamocortical synapses. *Science (New York, N.Y.)*, *312*(5780), 1622–1627. https://doi.org/10.1126/science.1124593

Buccino, A. P., Hurwitz, C. L., Garcia, S., Magland, J., Siegle, J. H., Hurwitz, R., & Hennig, M. H. (2020). SpikeInterface, a unified framework for spike sorting. *ELife*, *9*. https://doi.org/10.7554/eLife.61834

Burroughs, A., Wise, A. K., Xiao, J., Houghton, C., Tang, T., Suh, C. Y., Lang, E. J., Apps, R., & Cerminara, N. L. (2017). The dynamic relationship between cerebellar Purkinje cell simple spikes and the spikelet number of complex spikes. *The Journal of Physiology*, *595*(1), 283–299. https://doi.org/10.1113/JP272259

Buzunov, E., Mueller, A., Straube, A., & Robinson, F. R. (2013). When during horizontal saccades in monkey does cerebellar output affect movement? *Brain Research*, *1503*(1872-6240 (Electronic)), 33–42. https://doi.org/10.1016/j.brainres.2013.02.001

Campello, R. J. G. B., Moulavi, D., & Sander, J. (2013). *Density-Based Clustering Based on Hierarchical Density Estimates* (pp. 160–172). https://doi.org/10.1007/978-3-642-37456-2_14

Catz, N., Dicke, P. W., & Thier, P. (2005). Cerebellar complex spike firing is suitable to induce as well as to stabilize motor learning. *Current Biology : CB*, *15*(24), 2179–2189. https://doi.org/10.1016/j.cub.2005.11.037

Chen-Harris, H., Joiner, W. M., Ethier, V., Zee, D. S., & Shadmehr, R. (2008). Adaptive control of saccades via internal feedback. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *28*(11), 2804–2813. https://doi.org/10.1523/JNEUROSCI.5300-07.2008

Chung, J. E., Magland, J. F., Barnett, A. H., Tolosa, V. M., Tooker, A. C., Lee, K. Y., Shah, K. G., Felix, S. H., Frank, L. M., & Greengard, L. F. (2017). A Fully Automated Approach to Spike Sorting. *Neuron*, *95*(6), 1381-1394.e6. https://doi.org/10.1016/j.neuron.2017.08.030

Codol, O., Holland, P. J., & Galea, J. M. (2018). The relationship between reinforcement and explicit control during visuomotor adaptation. *Scientific Reports*, *8*(1), 9121. https://doi.org/10.1038/s41598-018-27378-1

Coltman, S. K., Cashaback, J. G. A., & Gribble, P. L. (2019). Both fast and slow learning processes contribute to savings following sensorimotor adaptation. *Journal of Neurophysiology*, *121*(4), 1575–1583. https://doi.org/10.1152/jn.00794.2018

Courchesne, E., Karns, C. M., Davis, H. R., Ziccardi, R., Carper, R. A., Tigue, Z. D., Chisum, H. J., Moses, P., Pierce, K., Lord, C., Lincoln, A. J., Pizzo, S., Schreibman, L., Haas, R. H., Akshoomoff, N. A., & Courchesne, R. Y. (2001). Unusual brain growth patterns in early life in patients with autistic disorder: an MRI study. *Neurology*, *57*(2), 245–254. https://doi.org/10.1212/wnl.57.2.245

da Silva, J. A., Tecuapetla, F., Paixão, V., & Costa, R. M. (2018). Dopamine neuron activity before action initiation gates and invigorates future movements. *Nature*, *554*(7691), 244–248. https://doi.org/10.1038/nature25457

Davie, J. T., Clark, B. A., & Häusser, M. (2008). The origin of the complex spike in cerebellar Purkinje cells. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *28*(30), 7599–7609. https://doi.org/10.1523/JNEUROSCI.0559-08.2008

de Solages, C., Szapiro, G., Brunel, N., Hakim, V., Isope, P., Buisseret, P., Rousseau, C., Barbour, B., & Léna, C. (2008). High-frequency organization and synchrony of activity in the purkinje cell layer of the cerebellum. *Neuron*, *58*(5), 775–788. https://doi.org/10.1016/j.neuron.2008.05.008

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. *Journal of Neurophysiology*, *77*(4), 1747–1758. https://doi.org/10.1152/jn.1997.77.4.1747

De Zeeuw, C I, 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*(4), 981–986. https://doi.org/10.1016/s0306-4522(97)00249-2

De Zeeuw, Chris I., Hoebeek, F. E., Bosman, L. W. J., Schonewille, M., Witter, L., & Koekkoek, S. K. (2011). Spatiotemporal firing patterns in the cerebellum. *Nature Reviews. Neuroscience*, *12*(6), 327–344. https://doi.org/10.1038/nrn3011

Deubel, H. (1987). Adaptivity of gain and direction in oblique saccades. In J. K. O'Regan & A. Levy-Schoen (Eds.), *Eye Movements from Physiology to Cognition* (pp. 181–190). Elsevier. https://doi.org/10.1016/B978-0-444-70113-8.50030-8

Dizon, M. J., & Khodakhah, K. (2011). The role of interneurons in shaping Purkinje cell responses in the cerebellar cortex. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *31*(29), 10463–10473. https://doi.org/10.1523/JNEUROSCI.1350-11.2011

Donchin, O., Rabe, K., Diedrichsen, J., Lally, N., Schoch, B., Gizewski, E. R., & Timmann, D. (2012). Cerebellar regions involved in adaptation to force field and visuomotor perturbation. *Journal of Neurophysiology*, *107*(1), 134–147. https://doi.org/10.1152/jn.00007.2011

Ebner, T. J., & Bloedel, J. R. (1981). Correlation between activity of Purkinje cells and its modification by natural peripheral stimuli. *Journal of Neurophysiology*, *45*(5), 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*(2), 268–296. https://doi.org/10.1113/jphysiol.1966.sp007824

Ekerot, C. F., Jörntell, H., & Garwicz, M. (1995). Functional relation between corticonuclear input and movements evoked on microstimulation in cerebellar nucleus interpositus anterior in the cat. *Experimental Brain Research*, *106*(3), 365–376. https://doi.org/10.1007/BF00231060

Eliades, S. J., & Miller, C. T. (2017). Marmoset vocal communication: Behavior and neurobiology. *Developmental Neurobiology*, *77*(3), 286–299. https://doi.org/10.1002/dneu.22464

Eliades, S. J., & Wang, X. (2008). Chronic multi-electrode neural recording in free-roaming monkeys. *Journal of Neuroscience Methods*, *172*(2), 201–214. https://doi.org/10.1016/j.jneumeth.2008.04.029

Eliades, S. J., & Wang, X. (2013). Comparison of auditory-vocal interactions across multiple types of vocalizations in marmoset auditory cortex. *Journal of Neurophysiology*, *109*(6), 1638–1657.

https://doi.org/10.1152/jn.00698.2012

Ethier, V., Zee, D. S., & Shadmehr, R. (2008a). Spontaneous recovery of motor memory during saccade adaptation. *Journal of Neurophysiology*, *99*(5), 2577–2583. https://doi.org/10.1152/jn.00015.2008

Ethier, V., Zee, D. S., & Shadmehr, R. (2008b). Changes in control of saccades during gain adaptation. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *28*(51), 13929–13937. https://doi.org/10.1523/JNEUROSCI.3470-08.2008

Fedorov, A., Beichel, R., Kalpathy-Cramer, J., Finet, J., Fillion-Robin, J.-C., Pujol, S., Bauer, C., Jennings, D., Fennessy, F., Sonka, M., Buatti, J., Aylward, S., Miller, J. V, Pieper, S., & Kikinis, R. (2012). 3D Slicer as an image computing platform for the Quantitative Imaging Network. *Magnetic Resonance Imaging*, *30*(9), 1323–1341. https://doi.org/10.1016/j.mri.2012.05.001

Fuchs, A. F., Kaneko, C. R., & Scudder, C. A. (1985). Brainstem control of saccadic eye movements. *Annual Review of Neuroscience*, *8*, 307–337. https://doi.org/10.1146/annurev.ne.08.030185.001515

Fujita, H., Oh-Nishi, A., Obayashi, S., & Sugihara, I. (2010). Organization of the marmoset cerebellum in three-dimensional space: lobulation, aldolase C compartmentalization and axonal projection. *The Journal of Comparative Neurology*, *518*(10), 1764–1791. https://doi.org/10.1002/cne.22301

Galea, J. M., Mallia, E., Rothwell, J., & Diedrichsen, J. (2015). The dissociable effects of punishment and reward on motor learning. *Nature Neuroscience*, *18*(4), 597–602. https://doi.org/10.1038/nn.3956

Gao, H., Solages, C. de, & Lena, C. (2012). Tetrode recordings in the cerebellar cortex. *Journal of Physiology, Paris*, *106*(3–4), 128–136. https://doi.org/10.1016/j.jphysparis.2011.10.005

Gauck, V., & Jaeger, D. (2000). The control of rate and timing of spikes in the deep cerebellar nuclei by inhibition. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *20*(8), 3006–3016. http://www.ncbi.nlm.nih.gov/pubmed/10751453

Glaser, J. I., Wood, D. K., Lawlor, P. N., Ramkumar, P., Kording, K. P., & Segraves, M. A. (2016). Role of expected reward in frontal eye field during natural scene search. *Journal of Neurophysiology*, *116*(2), 645–657. https://doi.org/10.1152/jn.00119.2016

Golla, H., Tziridis, K., Haarmeier, T., Catz, N., Barash, S., & Thier, P. (2008). Reduced saccadic resilience and impaired saccadic adaptation due to cerebellar disease. *The European Journal of Neuroscience*, *27*(1), 132–144. https://doi.org/10.1111/j.1460-9568.2007.05996.x

Goodale, M. A., & Milner, A. D. (1992). Separate visual pathways for perception and action. *Trends in Neurosciences*, *15*(1), 20–25. https://doi.org/10.1016/0166-2236(92)90344-8

Haith, A. M., Reppert, T. R., & Shadmehr, R. (2012). Evidence for hyperbolic temporal discounting of reward in control of movements. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *32*(34), 11727–11736. https://doi.org/10.1523/JNEUROSCI.0424-12.2012

Hall, N. J., Herzfeld, D. J., & Lisberger, S. G. (2021). Evaluation and resolution of many challenges of neural spike-sorting: a new sorter. *BioRxiv*. https://doi.org/10.1101/2021.01.19.427297

Han, K.-S., Chen, C. H., Khan, M. M., Guo, C., & Regehr, W. G. (2020). Climbing fiber synapses rapidly and transiently inhibit neighboring Purkinje cells via ephaptic coupling. *Nature Neuroscience*, *23*(11), 1399–1409. https://doi.org/10.1038/s41593-020-0701-z

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*(3), 564-578.e3. https://doi.org/10.1016/j.neuron.2018.09.018

Hanajima, R., Shadmehr, R., Ohminami, S., Tsutsumi, R., Shirota, Y., Shimizu, T., Tanaka, N., Terao, Y., Tsuji, S., Ugawa, Y., Uchimura, M., Inoue, M., & Kitazawa, S. (2015). Modulation of error-sensitivity during a prism adaptation task in people with cerebellar degeneration. *Journal of Neurophysiology*, *114*(4), 2460–2471. https://doi.org/10.1152/jn.00145.2015

Hanes, D. P., Smith, M. K., Optican, L. M., & Wurtz, R. H. (2005). Recovery of saccadic dysmetria following localized lesions in monkey superior colliculus. *Experimental Brain Research*, *160*(3), 312–325. https://doi.org/10.1007/s00221-004-2013-z

Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N. J., Kern, R., Picus, M., Hoyer, S., van Kerkwijk, M. H., Brett, M., Haldane, A., del Río, J. F., Wiebe, M., Peterson, P., … Oliphant, T. E. (2020). Array programming with NumPy. *Nature*, *585*(7825), 357–362. https://doi.org/10.1038/s41586-020-2649-2

Harting, J. K. (1977). Descending pathways from the superior collicullus: an autoradiographic analysis in the rhesus monkey (Macaca mulatta). *The Journal of Comparative Neurology*, *173*(3), 583–612. https://doi.org/10.1002/cne.901730311

Hashimoto, T., Tayama, M., Murakawa, K., Yoshimoto, T., Miyazaki, M., Harada, M., & Kuroda, Y. (1995). Development of the brainstem and cerebellum in autistic patients. *Journal of Autism and Developmental Disorders*, *25*(1), 1–18. https://doi.org/10.1007/BF02178163

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. *Proceedings of the National Academy of Sciences of the United States of America*, *104*(18), 7658–7663. https://doi.org/10.1073/pnas.0609966104

Heiney, S. A., Ohmae, S., Kim, O. A., & Medina, J. F. (2018). Single-unit extracellular recording from the cerebellum during eyeblink conditioning in head-fixed mice. *Neuromethods*, *134*, 39–71. https://doi.org/10.1007/978-1-4939-7549-5_3

Heitz, R. P., & Schall, J. D. (2012). Neural mechanisms of speed-accuracy tradeoff. *Neuron*, *76*(3), 616–628. https://doi.org/10.1016/j.neuron.2012.08.030

Helmchen, C., & Büttner, U. (1995). Saccade-related Purkinje cell activity in the oculomotor vermis during spontaneous eye movements in light and darkness. *Experimental Brain Research*, *103*(2), 198–208. https://doi.org/10.1007/BF00231706

Herzfeld, D. J., Kojima, Y., Soetedjo, R., & Shadmehr, R. (2015). Encoding of action by the Purkinje cells of the cerebellum. *Nature*, *526*(7573), 439–442. https://doi.org/10.1038/nature15693

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. *Nature Neuroscience*, *21*(5), 736–743. https://doi.org/10.1038/s41593-018-0136-y

Herzfeld, D. J., Pastor, D., Haith, A. M., Rossetti, Y., Shadmehr, R., & O'Shea, J. (2014). Contributions of the cerebellum and the motor cortex to acquisition and retention of motor memories. *NeuroImage*, *98*(1095-9572 (Electronic)), 147–158. https://doi.org/10.1016/j.neuroimage.2014.04.076

Herzfeld, D. J., Vaswani, P. A., Marko, M. K., & Shadmehr, R. (2014). A memory of errors in sensorimotor learning. *Science (New York, N.Y.)*, *345*(6202), 1349–1353. https://doi.org/10.1126/science.1253138

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. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *35*(3), 1106–1124. https://doi.org/10.1523/JNEUROSCI.2579-14.2015

Hewitt, A. L., Popa, L. S., Pasalar, S., Hendrix, C. M., & Ebner, T. J. (2011). Representation of limb kinematics in Purkinje cell simple spike discharge is conserved across multiple tasks. *Journal of Neurophysiology*, *106*(5), 2232–2247. https://doi.org/10.1152/jn.00886.2010

Holmes, G. (1939). The cerebellum of man. *Brain*, *62*(1), 1–30. https://doi.org/10.1093/brain/62.1.1

Hopp, J. J., & Fuchs, A. F. (2004). The characteristics and neuronal substrate of saccadic eye movement plasticity. *Progress in Neurobiology*, *72*(1), 27–53. https://doi.org/10.1016/j.pneurobio.2003.12.002

Huxter, J., Burgess, N., & O'Keefe, J. (2003). Independent rate and temporal coding in hippocampal pyramidal cells. *Nature*, *425*(6960), 828–832. https://doi.org/10.1038/nature02058

Ishikawa, T., Tomatsu, S., Tsunoda, Y., Lee, J., Hoffman, D. S., & Kakei, S. (2014). Releasing dentate nucleus cells from Purkinje cell inhibition generates output from the cerebrocerebellum. *PloS One*, *9*(10), e108774. https://doi.org/10.1371/journal.pone.0108774

Ito, M., & Simpson, J. I. (1971). Discharges in Purkinje cell axons during climbing fiber activation. *Brain Research*, *31*(1), 215–219. https://doi.org/10.1016/0006-8993(71)90648-2

Iwamoto, Y., & Yoshida, K. (2002). Saccadic dysmetria following inactivation of the primate fastigial oculomotor region. *Neuroscience Letters*, *325*(3), 211–215. https://doi.org/10.1016/s0304-3940(02)00268-9

Izawa, J., Criscimagna-Hemminger, S. E., & Shadmehr, R. (2012). Cerebellar contributions to reach adaptation and learning sensory consequences of action. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *32*(12), 4230–4239. https://doi.org/10.1523/JNEUROSCI.6353-11.2012

Jackson, N., & Muthuswamy, J. (2008). Artificial dural sealant that allows multiple penetrations of implantable brain probes. *Journal of Neuroscience Methods*, *171*(1), 147–152. https://doi.org/10.1016/j.jneumeth.2008.02.018

Jiang, Z., Huxter, J. R., Bowyer, S. A., Blockeel, A. J., Butler, J., Imtiaz, S. A., Wafford, K. A., Phillips, K. G., Tricklebank, M. D., Marston, H. M., & Rodriguez-Villegas, E. (2017). TaiNi: Maximizing research output whilst improving animals' welfare in neurophysiology experiments. *Scientific Reports*, *7*(1), 8086. https://doi.org/10.1038/s41598-017-08078-8

Johnston, K. D., Barker, K., Schaeffer, L., Schaeffer, D., & Everling, S. (2018). Methods for chair restraint and training of the common marmoset on oculomotor tasks. *Journal of Neurophysiology*, *119*(5), 1636–1646. https://doi.org/10.1152/jn.00866.2017

Joshi, S., Li, Y., Kalwani, R. M., & Gold, J. I. (2016). Relationships between pupil diameter and neuronal activity in the locus coeruleus, colliculi, and cingulate cortex. *Neuron*, *89*(1), 221–234.

https://doi.org/10.1016/j.neuron.2015.11.028

Joshua, M., & Lisberger, S. G. (2012). Reward action in the initiation of smooth pursuit eye movements. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *32*(8), 2856–2867. https://doi.org/10.1523/JNEUROSCI.4676-11.2012

Kahneman, D., & Beatty, J. (1966). Pupil diameter and load on memory. *Science (New York, N.Y.)*, *154*(3756), 1583–1585. https://doi.org/10.1126/science.154.3756.1583

Kase, M., Miller, D. C., & Noda, H. (1980). Discharges of Purkinje cells and mossy fibres in the cerebellar vermis of the monkey during saccadic eye movements and fixation. *The Journal of Physiology*, *300*(0022-3751 (Print)), 539–555. https://doi.org/10.1113/jphysiol.1980.sp013178

Kawagoe, R., Takikawa, Y., & Hikosaka, O. (1998). Expectation of reward modulates cognitive signals in the basal ganglia. *Nature Neuroscience*, *1*(5), 411–416. https://doi.org/10.1038/1625

Kawagoe, R., Takikawa, Y., & Hikosaka, O. (2004). Reward-predicting activity of dopamine and caudate neurons--a possible mechanism of motivational control of saccadic eye movement. *Journal of Neurophysiology*, *91*(2), 1013–1024. https://doi.org/10.1152/jn.00721.2003

Keller, E. L., Gandhi, N. J., & Shieh, J. M. (1996). Endpoint accuracy in saccades interrupted by stimulation in the omnipause region in monkey. *Visual Neuroscience*, *13*(6), 1059–1067. https://doi.org/10.1017/s0952523800007719

Khaliq, Z. M., & Raman, I. M. (2005). Axonal propagation of simple and complex spikes in cerebellar Purkinje neurons. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *25*(2), 454–463. https://doi.org/10.1523/JNEUROSCI.3045-04.2005

Kim, H. E., Parvin, D. E., & Ivry, R. B. (2019). The influence of task outcome on implicit motor learning. *ELife*, *8*, 363606. https://doi.org/10.7554/eLife.39882

Kishi, N., Sato, K., Sasaki, E., & Okano, H. (2014). Common marmoset as a new model animal for neuroscience research and genome editing technology. *Development, Growth & Differentiation*, *56*(1), 53–62. https://doi.org/10.1111/dgd.12109

Kojima, Y., Robinson, F. R., & Soetedjo, R. (2014). Cerebellar fastigial nucleus influence on ipsilateral abducens activity during saccades. *Journal of Neurophysiology*, *111*(8), 1553–1563. https://doi.org/10.1152/jn.00567.2013

Kojima, Y., & Soetedjo, R. (2017a). Selective reward affects the rate of saccade adaptation. *Neuroscience*, *355*(0306-4522 (Linking)), 113–125. https://doi.org/10.1016/j.neuroscience.2017.04.048

Kojima, Y., & Soetedjo, R. (2017b). Change in sensitivity to visual error in superior colliculus during saccade adaptation. *Scientific Reports*, *7*(1), 9566. https://doi.org/10.1038/s41598-017-10242-z

Kojima, Y., & Soetedjo, R. (2018). Elimination of the error signal in the superior colliculus impairs saccade motor learning. *Proceedings of the National Academy of Sciences of the United States of America*, *115*(38), E8987–E8995. https://doi.org/10.1073/pnas.1806215115

Kojima, Y., Soetedjo, R., & Fuchs, A. F. (2010a). Changes in simple spike activity of some Purkinje cells in the oculomotor vermis during saccade adaptation are appropriate to participate in motor learning. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *30*(10), 3715–

3727. https://doi.org/10.1523/JNEUROSCI.4953-09.2010

Kojima, Y., Soetedjo, R., & Fuchs, A. F. (2010b). Behavior of the oculomotor vermis for five different types of saccade. *Journal of Neurophysiology*, *104*(6), 3667–3676. https://doi.org/10.1152/jn.00558.2010

Kording, K. P., Tenenbaum, J. B., & Shadmehr, R. (2007). The dynamics of memory as a consequence of optimal adaptation to a changing body. *Nature Neuroscience*, *10*(6), 779–786. https://doi.org/10.1038/nn1901

Körding, K. P., & Wolpert, D. M. (2004). The loss function of sensorimotor learning. *Proceedings of the National Academy of Sciences of the United States of America*, *101*(26), 9839–9842. https://doi.org/10.1073/pnas.0308394101

Kostadinov, D., Beau, M., Blanco-Pozo, M., & Häusser, M. (2019). Predictive and reactive reward signals conveyed by climbing fiber inputs to cerebellar Purkinje cells. *Nature Neuroscience*, *22*(6), 950–962. https://doi.org/10.1038/s41593-019-0381-8

Lankarany, M., Al-Basha, D., Ratté, S., & Prescott, S. A. (2019). Differentially synchronized spiking enables multiplexed neural coding. *Proceedings of the National Academy of Sciences of the United States of America*, *116*(20), 10097–10102. https://doi.org/10.1073/pnas.1812171116

Lee, E. K., Balasubramanian, H., Tsolias, A., Anakwe, S. U., Medalla, M., Shenoy, K. V, & Chandrasekaran, C. (2021). Non-linear dimensionality reduction on extracellular waveforms reveals cell type diversity in premotor cortex. *ELife*, *10*. https://doi.org/10.7554/eLife.67490

Leigh, R. J., & Zee, D. S. (2015). *The neurology of eye movements* (5th ed.). Oxford University Press.

Leow, L.-A., de Rugy, A., Marinovic, W., Riek, S., & Carroll, T. J. (2016). Savings for visuomotor adaptation require prior history of error, not prior repetition of successful actions. *Journal of Neurophysiology*, *116*(4), 1603–1614. https://doi.org/10.1152/jn.01055.2015

Leow, L.-A., Marinovic, W., de Rugy, A., & Carroll, T. J. (2020). Task errors drive memories that improve sensorimotor adaptation. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *40*(15), 3075–3088. https://doi.org/10.1523/JNEUROSCI.1506-19.2020

Lu, T., Liang, L., & Wang, X. (2001a). Neural representations of temporally asymmetric stimuli in the auditory cortex of awake primates. *Journal of Neurophysiology*, *85*(6), 2364–2380. https://doi.org/10.1152/jn.2001.85.6.2364

Lu, T., Liang, L., & Wang, X. (2001b). Temporal and rate representations of time-varying signals in the auditory cortex of awake primates. *Nature Neuroscience*, *4*(11), 1131–1138. https://doi.org/10.1038/nn737

MacDougall, M., Nummela, S. U., Coop, S., Disney, A., Mitchell, J. F., & Miller, C. T. (2016). Optogenetic manipulation of neural circuits in awake marmosets. *Journal of Neurophysiology*, *116*(3), 1286–1294. https://doi.org/10.1152/jn.00197.2016

Magland, J. F., & Barnett, A. H. (2015). *Unimodal clustering using isotonic regression: ISO-SPLIT*. http://arxiv.org/abs/1508.04841

Manohar, S. G., Chong, T. T., Apps, M. A. J., Batla, A., Stamelou, M., Jarman, P. R., Bhatia, K. P., & Husain, M. (2015). Reward pays the cost of noise reduction in motor and cognitive control. *Current*

Biology : CB, 25(13), 1707–1716. https://doi.org/10.1016/j.cub.2015.05.038

Marino, R. A., Levy, R., & Munoz, D. P. (2015). Linking express saccade occurance to stimulus properties and sensorimotor integration in the superior colliculus. Journal of Neurophysiology, 114(2), 879–892. https://doi.org/10.1152/jn.00047.2015

Markanday, A., Bellet, J., Bellet, M. E., Inoue, J., Hafed, Z. M., & Thier, P. (2020). Using deep neural networks to detect complex spikes of cerebellar Purkinje cells. Journal of Neurophysiology, 123(6), 2217–2234. https://doi.org/10.1152/jn.00754.2019

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 : A Journal of Neurology, 138(Pt 3), 784–797. https://doi.org/10.1093/brain/awu394

Marko, M. K., Haith, A. M., Harran, M. D., & Shadmehr, R. (2012). Sensitivity to prediction error in reach adaptation. Journal of Neurophysiology, 108(6), 1752–1763. https://doi.org/10.1152/jn.00177.2012

Maschke, M., Gomez, C. M., Ebner, T. J., & Konczak, J. (2004). Hereditary cerebellar ataxia progressively impairs force adaptation during goal-directed arm movements. Journal of Neurophysiology, 91(1), 230–238. https://doi.org/10.1152/jn.00557.2003

Mathôt, S. (2018). Pupillometry: psychology, physiology, and function. Journal of Cognition, 1(1), 16. https://doi.org/10.5334/joc.18

Matsumoto, M., & Hikosaka, O. (2007). Lateral habenula as a source of negative reward signals in dopamine neurons. Nature, 447(7148), 1111–1115. https://doi.org/10.1038/nature05860

Matsumoto, M., & Hikosaka, O. (2009). Two types of dopamine neuron distinctly convey positive and negative motivational signals. Nature, 459(7248), 837–841. https://doi.org/10.1038/nature08028

May, P. J., Warren, S., Bohlen, M. O., Barnerssoi, M., & Horn, A. K. E. (2016). A central mesencephalic reticular formation projection to the Edinger-Westphal nuclei. Brain Structure & Function, 221(8), 4073–4089. https://doi.org/10.1007/s00429-015-1147-z

McInnes, L., Healy, J., & Melville, J. (2018). UMAP: uniform manifold approximation and projection for dimension reduction. ArXiv. http://arxiv.org/abs/1802.03426

McKee, S. P., & Nakayama, K. (1984). The detection of motion in the peripheral visual field. Vision Research, 24(1), 25–32. https://doi.org/10.1016/0042-6989(84)90140-8

Medina, J. F., & Lisberger, S. G. (2007). Variation, signal, and noise in cerebellar sensory-motor processing for smooth-pursuit eye movements. The Journal of Neuroscience : The Official Journal of the Society for Neuroscience, 27(25), 6832–6842. https://doi.org/10.1523/JNEUROSCI.1323-07.2007

Medina, J. F., & Lisberger, S. G. (2008). Links from complex spikes to local plasticity and motor learning in the cerebellum of awake-behaving monkeys. Nature Neuroscience, 11(10), 1185–1192. https://doi.org/10.1038/nn.2197

Meermeier, A., Gremmler, S., & Lappe, M. (2016). The influence of image content on oculomotor plasticity. Journal of Vision, 16(8), 17. https://doi.org/10.1167/16.8.17

Miller, C. T., Freiwald, W. A., Leopold, D. A., Mitchell, J. F., Silva, A. C., & Wang, X. (2016). Marmosets: a

neuroscientific model of human social behavior. *Neuron*, *90*(2), 219–233. https://doi.org/10.1016/j.neuron.2016.03.018

Milstein, D. M., & Dorris, M. C. (2007). The influence of expected value on saccadic preparation. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *27*(18), 4810–4818. https://doi.org/10.1523/JNEUROSCI.0577-07.2007

Mitchell, J. F., Reynolds, J. H., & Miller, C. T. (2014). Active vision in marmosets: a model system for visual neuroscience. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *34*(4), 1183–1194. https://doi.org/10.1523/JNEUROSCI.3899-13.2014

Monsivais, P., Clark, B. A., Roth, A., & Häusser, M. (2005). Determinants of action potential propagation in cerebellar Purkinje cell axons. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *25*(2), 464–472. https://doi.org/10.1523/JNEUROSCI.3871-04.2005

Murakami, J. W., Courchesne, E., Press, G. A., Yeung-Courchesne, R., & Hesselink, J. R. (1989). Reduced cerebellar hemisphere size and its relationship to vermal hypoplasia in autism. *Archives of Neurology*, *46*(6), 689–694. https://doi.org/10.1001/archneur.1989.00520420111032

Mustoe, A. C., Cavanaugh, J., Harnisch, A. M., Thompson, B. E., & French, J. A. (2015). Do marmosets care to share? Oxytocin treatment reduces prosocial behavior toward strangers. *Hormones and Behavior*, *71*(0018-506X (Linking)), 83–90. https://doi.org/10.1016/j.yhbeh.2015.04.015

Najafi, F., Giovannucci, A., Wang, S. S.-H., & Medina, J. F. (2014). Sensory-driven enhancement of calcium signals in individual Purkinje cell dendrites of awake mice. *Cell Reports*, *6*(5), 792–798. https://doi.org/10.1016/j.celrep.2014.02.001

Nikooyan, A. A., & Ahmed, A. A. (2015). Reward feedback accelerates motor learning. *Journal of Neurophysiology*, *113*(2), 633–646. https://doi.org/10.1152/jn.00032.2014

O'Doherty, J., Winston, J., Critchley, H., Perrett, D., Burt, D. M., & Dolan, R. J. (2003). Beauty in a smile: the role of medial orbitofrontal cortex in facial attractiveness. *Neuropsychologia*, *41*(2), 147–155. https://doi.org/10.1016/s0028-3932(02)00145-8

Ohmae, S., & Medina, J. F. (2015). Climbing fibers encode a temporal-difference prediction error during cerebellar learning in mice. *Nature Neuroscience*, *18*(12), 1798–1803. https://doi.org/10.1038/nn.4167

Optican, L. M., & Miles, F. A. (1985). Visually induced adaptive changes in primate saccadic oculomotor control signals. *Journal of Neurophysiology*, *54*(4), 940–958. https://doi.org/10.1152/jn.1985.54.4.940

Optican, L. M., & Robinson, D. A. (1980). Cerebellar-dependent adaptive control of primate saccadic system. *Journal of Neurophysiology*, *44*(6), 1058–1076. https://doi.org/10.1152/jn.1980.44.6.1058

Osmanski, M. S., Song, X., & Wang, X. (2013). The role of harmonic resolvability in pitch perception in a vocal nonhuman primate, the common marmoset (Callithrix jacchus). *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *33*(21), 9161–9168. https://doi.org/10.1523/JNEUROSCI.0066-13.2013

Osmanski, M. S., & Wang, X. (2011). Measurement of absolute auditory thresholds in the common marmoset (Callithrix jacchus). *Hearing Research*, *277*(1–2), 127–133. https://doi.org/10.1016/j.heares.2011.02.001

Özcan, O. O., Wang, X., Binda, F., Dorgans, K., De Zeeuw, C. I., Gao, Z., Aertsen, A., Kumar, A., & Isope, P. (2020). Differential coding strategies in glutamatergic and GABAergic neurons in the medial cerebellar nucleus. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *40*(1), 159–170. https://doi.org/10.1523/JNEUROSCI.0806-19.2019

Pachitariu, M., Steinmetz, N. A., Kadir, S. N., Carandini, M., & Harris, K. D. (2016). Fast and accurate spike sorting of high-channel count probes with KiloSort. *Advances in Neural Information Processing Systems*, 4448–4456.

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*(2050-084X (Linking)). https://doi.org/10.7554/eLife.37102

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., & others. (2011). Scikit-learn: Machine learning in Python. *The Journal of Machine Learning Research*, *12*, 2825–2830.

Pekny, S. E., Criscimagna-Hemminger, S. E., & Shadmehr, R. (2011). Protection and expression of human motor memories. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *31*(39), 13829–13839. https://doi.org/10.1523/JNEUROSCI.1704-11.2011

Person, A. L., & Raman, I. M. (2011). Purkinje neuron synchrony elicits time-locked spiking in the cerebellar nuclei. *Nature*, *481*(7382), 502–505. https://doi.org/10.1038/nature10732

Person, A. L., & Raman, I. M. (2012). Synchrony and neural coding in cerebellar circuits. *Frontiers in Neural Circuits*, *6*(1662-5110 (Electronic)), 97. https://doi.org/10.3389/fncir.2012.00097

Pilly, P. K., & Seitz, A. R. (2009). What a difference a parameter makes: a psychophysical comparison of random dot motion algorithms. *Vision Research*, *49*(13), 1599–1612. https://doi.org/10.1016/j.visres.2009.03.019

Quaia, C., Lefèvre, P., & Optican, L. M. (1999). Model of the control of saccades by superior colliculus and cerebellum. *Journal of Neurophysiology*, *82*(2), 999–1018. https://doi.org/10.1152/jn.1999.82.2.999

Quaia, C., Paré, M., Wurtz, R. H., & Optican, L. M. (2000). Extent of compensation for variations in monkey saccadic eye movements. *Experimental Brain Research*, *132*(1), 39–51. https://doi.org/10.1007/s002219900324

Quattrocchi, G., Monaco, J., Ho, A., Irmen, F., Strube, W., Ruge, D., Bestmann, S., & Galea, J. M. (2018). Pharmacological dopamine manipulation does not alter reward-based improvements in memory retention during a visuomotor adaptation task. *ENeuro*, *5*(3), ENEURO.0453-17.2018. https://doi.org/10.1523/ENEURO.0453-17.2018

Rabe, K., Livne, O., Gizewski, E. R., Aurich, V., Beck, A., Timmann, D., & Donchin, O. (2009). Adaptation to visuomotor rotation and force field perturbation is correlated to different brain areas in patients with cerebellar degeneration. *Journal of Neurophysiology*, *101*(4), 1961–1971. https://doi.org/10.1152/jn.91069.2008

Raymond, J. L., & Medina, J. F. (2018). Computational principles of supervised learning in the cerebellum. *Annual Review of Neuroscience*, *41*(1), 233–253. https://doi.org/10.1146/annurev-neuro-080317-061948

Remington, E. D., Osmanski, M. S., & Wang, X. (2012). An operant conditioning method for studying auditory behaviors in marmoset monkeys. *PloS One*, *7*(10), e47895. https://doi.org/10.1371/journal.pone.0047895

Reppert, T. R., Lempert, K. M., Glimcher, P. W., & Shadmehr, R. (2015). Modulation of saccade vigor during value-based decision making. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *35*(46), 15369–15378. https://doi.org/10.1523/JNEUROSCI.2621-15.2015

Ritchie, L. (1976). Effects of cerebellar lesions on saccadic eye movements. *Journal of Neurophysiology*, *39*(6), 1246–1256. https://doi.org/10.1152/jn.1976.39.6.1246

Robinson, F. R., Straube, A., & Fuchs, A. F. (1993). Role of the caudal fastigial nucleus in saccade generation. II. Effects of muscimol inactivation. *Journal of Neurophysiology*, *70*(5), 1741–1758. https://doi.org/10.1152/jn.1993.70.5.1741

Roitman, A. V., Pasalar, S., & Ebner, T. J. (2009). Single trial coupling of Purkinje cell activity to speed and error signals during circular manual tracking. *Experimental Brain Research*, *192*(2), 241–251. https://doi.org/10.1007/s00221-008-1580-9

Roitman, A. V., Pasalar, S., Johnson, M. T. V, & Ebner, T. J. (2005). Position, direction of movement, and speed tuning of cerebellar Purkinje cells during circular manual tracking in monkey. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *25*(40), 9244–9257. https://doi.org/10.1523/JNEUROSCI.1886-05.2005

Roth, M. J., Synofzik, M., & Lindner, A. (2013). The cerebellum optimizes perceptual predictions about external sensory events. *Current Biology*, *23*(10), 930–935. https://doi.org/10.1016/j.cub.2013.04.027

Roy, S., Miller, C. T., Gottsch, D., & Wang, X. (2011). Vocal control by the common marmoset in the presence of interfering noise. *The Journal of Experimental Biology*, *214*(Pt 21), 3619–3629. https://doi.org/10.1242/jeb.056101

Roy, S., & Wang, X. (2012). Wireless multi-channel single unit recording in freely moving and vocalizing primates. *Journal of Neuroscience Methods*, *203*(1), 28–40. https://doi.org/10.1016/j.jneumeth.2011.09.004

Ruigrok, T. J., & Voogd, J. (2000). Organization of projections from the inferior olive to the cerebellar nuclei in the rat. *The Journal of Comparative Neurology*, *426*(2), 209–228. https://doi.org/10.1002/1096-9861(20001016)426:2<209::aid-cne4>3.0.co;2-0

Sasaki, E., Suemizu, H., Shimada, A., Hanazawa, K., Oiwa, R., Kamioka, M., Tomioka, I., Sotomaru, Y., Hirakawa, R., Eto, T., Shiozawa, S., Maeda, T., Ito, M., Ito, R., Kito, C., Yagihashi, C., Kawai, K., Miyoshi, H., Tanioka, Y., … Nomura, T. (2009). Generation of transgenic non-human primates with germline transmission. *Nature*, *459*(7246), 523–527. https://doi.org/10.1038/nature08090

Sato, Y., Miura, A., Fushiki, H., & Kawasaki, T. (1992). Short-term modulation of cerebellar Purkinje cell activity after spontaneous climbing fiber input. *Journal of Neurophysiology*, *68*(6), 2051–2062. https://doi.org/10.1152/jn.1992.68.6.2051

Schmolesky, M. T., Weber, J. T., De Zeeuw, C. I., & Hansel, C. (2002). The making of a complex spike: ionic composition and plasticity. *Annals of the New York Academy of Sciences*, *978*(1), 359–390. https://doi.org/10.1111/j.1749-6632.2002.tb07581.x

Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science (New York, N.Y.)*, *275*(5306), 1593–1599. https://doi.org/10.1126/science.275.5306.1593

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 Research : Official Journal of the International Society for Autism Research*, *2*(5), 246–257. https://doi.org/10.1002/aur.97

Scudder, C A. (1988). A new local feedback model of the saccadic burst generator. *Journal of Neurophysiology*, *59*(5), 1455–1475. https://doi.org/10.1152/jn.1988.59.5.1455

Scudder, Charles A, & McGee, D. M. (2003). Adaptive modification of saccade size produces correlated changes in the discharges of fastigial nucleus neurons. *Journal of Neurophysiology*, *90*(2), 1011–1026. https://doi.org/10.1152/jn.00193.2002

Sedaghat-Nejad, E., Fakharian, M. A., Pi, J., Hage, P., Kojima, Y., Soetedjo, R., Ohmae, S., Medina, J. F., & Shadmehr, R. (2021). P-sort: an open-source software for cerebellar neurophysiology. *Journal of Neurophysiology*, *126*(4), 1055–1075. https://doi.org/10.1152/jn.00172.2021

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. *Journal of Neurophysiology*, *122*(4), 1502–1517. https://doi.org/10.1152/jn.00389.2019

Sedaghat-Nejad, E., Herzfeld, D. J., & Shadmehr, R. (2019). Reward prediction error modulates saccade vigor. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *39*(25), 5010–5017. https://doi.org/10.1523/JNEUROSCI.0432-19.2019

Sedaghat-Nejad, E., Pi, J. S., Hage, P., Fakharian, M. A., & Shadmehr, R. (2021). Synchronous spiking of cerebellar Purkinje cells during control of movements. *BioRxiv*. https://doi.org/10.1101/2021.09.16.460700

Sedaghat-Nejad, E., & Shadmehr, R. (2021). The cost of correcting for error during sensorimotor adaptation. *Proceedings of the National Academy of Sciences of the United States of America*, *118*(40), e2101717118. https://doi.org/10.1073/pnas.2101717118

Servais, L., Bearzatto, B., Hourez, R., Dan, B., Schiffmann, S. N., & Cheron, G. (2004). Effect of simple spike firing mode on complex spike firing rate and waveform in cerebellar Purkinje cells in non-anesthetized mice. *Neuroscience Letters*, *367*(2), 171–176. https://doi.org/10.1016/j.neulet.2004.05.109

Shadmehr, R. (2020). Population coding in the cerebellum: a machine learning perspective. *Journal of Neurophysiology*, *124*(6), 2022–2051. https://doi.org/10.1152/jn.00449.2020

Shadmehr, R., & Ahmed, A. A. (2020). *Vigor: Neuroeconomics of movement control*. MIT Press.

Shadmehr, R., Orban de Xivry, J. J., Xu-Wilson, M., & Shih, T.-Y. (2010). Temporal discounting of reward and the cost of time in motor control. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *30*(31), 10507–10516. https://doi.org/10.1523/JNEUROSCI.1343-10.2010

Shadmehr, R., Reppert, T. R., Summerside, E. M., Yoon, T., & Ahmed, A. A. (2019). Movement vigor as a reflection of subjective economic utility. *Trends in Neurosciences*, *42*(5), 323–336. https://doi.org/10.1016/j.tins.2019.02.003

Shin, S.-L., & De Schutter, E. (2006). Dynamic synchronization of Purkinje cell simple spikes. *Journal of Neurophysiology*, *96*(6), 3485–3491. https://doi.org/10.1152/jn.00570.2006

Siegle, J. H., López, A. C., Patel, Y. A., Abramov, K., Ohayon, S., & Voigts, J. (2017). Open Ephys: an open-source, plugin-based platform for multichannel electrophysiology. *Journal of Neural Engineering*, *14*(4), 045003. https://doi.org/10.1088/1741-2552/aa5eea

Smalianchuk, I., Jagadisan, U. K., & Gandhi, N. J. (2018). Instantaneous midbrain control of saccade velocity. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *38*(47), 10156–10167. https://doi.org/10.1523/JNEUROSCI.0962-18.2018

Smith, M. A., Ghazizadeh, A., & Shadmehr, R. (2006). Interacting adaptive processes with different timescales underlie short-term motor learning. *PLoS Biology*, *4*(6), e179. https://doi.org/10.1371/journal.pbio.0040179

Smith, M. A., & Shadmehr, R. (2005). Intact ability to learn internal models of arm dynamics in Huntington's disease but not cerebellar degeneration. *Journal of Neurophysiology*, *93*(5), 2809–2821. https://doi.org/10.1152/jn.00943.2004

Soetedjo, R., & Fuchs, A. F. (2006). Complex spike activity of purkinje cells in the oculomotor vermis during behavioral adaptation of monkey saccades. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *26*(29), 7741–7755. https://doi.org/10.1523/JNEUROSCI.4658-05.2006

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? *Journal of Neurophysiology*, *100*(4), 1949–1966. https://doi.org/10.1152/jn.90526.2008

Soetedjo, R., Kojima, Y., & Fuchs, A. F. (2019). How cerebellar motor learning keeps saccades accurate. *Journal of Neurophysiology*, *121*(6), 2153–2162. https://doi.org/10.1152/jn.00781.2018

Solomon, S. G., & Rosa, M. G. P. (2014). A simpler primate brain: the visual system of the marmoset monkey. *Frontiers in Neural Circuits*, *8*, 96. https://doi.org/10.3389/fncir.2014.00096

Song, Y., & Smiley-Oyen, A. L. (2017). Probability differently modulating the effects of reward and punishment on visuomotor adaptation. *Experimental Brain Research*, *235*(12), 3605–3618. https://doi.org/10.1007/s00221-017-5082-5

Spampinato, D. A., Satar, Z., & Rothwell, J. C. (2019). Combining reward and M1 transcranial direct current stimulation enhances the retention of newly learnt sensorimotor mappings. *Brain Stimulation*, *12*(5), 1205–1212. https://doi.org/10.1016/j.brs.2019.05.015

Sparks, D. L. (2002). The brainstem control of saccadic eye movements. *Nature Reviews. Neuroscience*, *3*(12), 952–964. https://doi.org/10.1038/nrn986

Steel, A., Silson, E. H., Stagg, C. J., & Baker, C. I. (2016). The impact of reward and punishment on skill learning depends on task demands. *Scientific Reports*, *6*(July), 36056. https://doi.org/10.1038/srep36056

Sugihara, I. (2011). Compartmentalization of the deep cerebellar nuclei based on afferent projections and aldolase C expression. *Cerebellum (London, England)*, *10*(3), 449–463. https://doi.org/10.1007/s12311-010-0226-1

Sugihara, I., Marshall, S. P., & Lang, E. J. (2007). Relationship of complex spike synchrony bands and climbing fiber projection determined by reference to aldolase C compartments in crus IIa of the rat cerebellar cortex. *The Journal of Comparative Neurology*, *501*(1), 13–29. https://doi.org/10.1002/cne.21223

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. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *27*(36), 9696–9710. https://doi.org/10.1523/JNEUROSCI.1579-07.2007

Summerside, E. M., Shadmehr, R., & Ahmed, A. A. (2018). Vigor of reaching movements: reward discounts the cost of effort. *Journal of Neurophysiology*, *119*(6), 2347–2357. https://doi.org/10.1152/jn.00872.2017

Tachibana, Y., & Hikosaka, O. (2012). The primate ventral pallidum encodes expected reward value and regulates motor action. *Neuron*, *76*(4), 826–837. https://doi.org/10.1016/j.neuron.2012.09.030

Takagi, M., Zee, D. S., & Tamargo, R. J. (1998). Effects of lesions of the oculomotor vermis on eye movements in primate: saccades. *Journal of Neurophysiology*, *80*(4), 1911–1931. https://doi.org/10.1152/jn.1998.80.4.1911

Takahashi, D. Y., Liao, D. A., & Ghazanfar, A. A. (2017). Vocal learning via social reinforcement by infant marmoset monkeys. *Current Biology : CB*, *27*(12), 1844-1852.e6. https://doi.org/10.1016/j.cub.2017.05.004

Takikawa, Y., Kawagoe, R., & Hikosaka, O. (2002). Reward-dependent spatial selectivity of anticipatory activity in monkey caudate neurons. *Journal of Neurophysiology*, *87*(1), 508–515. https://doi.org/10.1152/jn.00288.2001

Tang, T., Blenkinsop, T. A., & Lang, E. J. (2019). Complex spike synchrony dependent modulation of rat deep cerebellar nuclear activity. *ELife*, *8*(2050-084X (Linking)). 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 (London, England)*, *15*(1), 10–13. https://doi.org/10.1007/s12311-015-0743-z

Thach, W. T. (1967). Somatosensory receptive fields of single units in cat cerebellar cortex. *Journal of Neurophysiology*, *30*(4), 675–696. https://doi.org/10.1152/jn.1967.30.4.675

Thiele, A., Henning, P., Kubischik, M., & Hoffmann, K.-P. (2002). Neural mechanisms of saccadic suppression. *Science (New York, N.Y.)*, *295*(5564), 2460–2462. https://doi.org/10.1126/science.1068788

Thier, P., Dicke, P. W., Haas, R., & Barash, S. (2000). Encoding of movement time by populations of cerebellar Purkinje cells. *Nature*, *405*(6782), 72–76. https://doi.org/10.1038/35011062

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. *Journal of Neurophysiology*, *115*(1), 255–270. https://doi.org/10.1152/jn.00530.2015

van der Maaten, L., & Hinton, G. (2008). Visualizing data using t-SNE. *Journal of Machine Learning Research*, *9*, 2579–2605. http://www.jmlr.org/papers/v9/vandermaaten08a.html

Van Gisbergen, J. A., Robinson, D. A., & Gielen, S. (1981). A quantitative analysis of generation of saccadic eye movements by burst neurons. *Journal of Neurophysiology*, *45*(3), 417–442. https://doi.org/10.1152/jn.1981.45.3.417

van Gompel, R. P. G., Fischer, M. H., Murray, W. S., & Hill, R. L. (2007). *Eye movements: A window on mind and brain.* Elsevier.

Vazey, E. M., Moorman, D. E., & Aston-Jones, G. (2018). Phasic locus coeruleus activity regulates cortical encoding of salience information. *Proceedings of the National Academy of Sciences of the United States of America*, *115*(40), E9439–E9448. https://doi.org/10.1073/pnas.1803716115

Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., … Vázquez-Baeza, Y. (2020). SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nature Methods*, *17*(3), 261–272. https://doi.org/10.1038/s41592-019-0686-2

Wächter, T., Lungu, O. V, Liu, T., Willingham, D. T., & Ashe, J. (2009). Differential effect of reward and punishment on procedural learning. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *29*(2), 436–443. https://doi.org/10.1523/JNEUROSCI.4132-08.2009

Wagner, M. J., Savall, J., Hernandez, O., Mel, G., Inan, H., Rumyantsev, O., Lecoq, J., Kim, T. H., Li, J. Z., Ramakrishnan, C., Deisseroth, K., Luo, L., Ganguli, S., & Schnitzer, M. J. (2021). A neural circuit state change underlying skilled movements. *Cell*, *184*(14), 3731-3747.e21. https://doi.org/10.1016/j.cell.2021.06.001

Wakabayashi, M., Koketsu, D., Kondo, H., Sato, S., Ohara, K., Polyakova, Z., Chiken, S., Hatanaka, N., & Nambu, A. (2018). Development of stereotaxic recording system for awake marmosets (Callithrix jacchus). *Neuroscience Research*, *135*(0168-0102 (Linking)), 37–45. https://doi.org/10.1016/j.neures.2018.01.001

Wang, C.-A., Boehnke, S. E., White, B. J., & Munoz, D. P. (2012). Microstimulation of the monkey superior colliculus induces pupil dilation without evoking saccades. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *32*(11), 3629–3636. https://doi.org/10.1523/JNEUROSCI.5512-11.2012

Wang, X. (2018). Cortical coding of auditory features. *Annual Review of Neuroscience*, *41*, 527–552. https://doi.org/10.1146/annurev-neuro-072116-031302

Wang, X., Lu, T., Snider, R. K., & Liang, L. (2005). Sustained firing in auditory cortex evoked by preferred stimuli. *Nature*, *435*(7040), 341–346. https://doi.org/10.1038/nature03565

Warnaar, P., Couto, J., Negrello, M., Junker, M., Smilgin, A., Ignashchenkova, A., Giugliano, M., Thier, P., & De Schutter, E. (2015). Duration of Purkinje cell complex spikes increases with their firing frequency. *Frontiers in Cellular Neuroscience*, *9*(1662-5102 (Linking)), 122. https://doi.org/10.3389/fncel.2015.00122

Wei, K., & Körding, K. (2009). Relevance of error: what drives motor adaptation? *Journal of Neurophysiology*, *101*(2), 655–664. https://doi.org/10.1152/jn.90545.2008

Welsh, J. P., Lang, E. J., Suglhara, I., & Llinás, R. (1995). Dynamic organization of motor control within the olivocerebellar system. *Nature*, *374*(6521), 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 (London, England)*, *7*(3), 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. *The Journal of Physiology*, *588*(Pt 13), 2373–2390. https://doi.org/10.1113/jphysiol.2010.189704

Wong, R. O., Meister, M., & Shatz, C. J. (1993). Transient period of correlated bursting activity during development of the mammalian retina. *Neuron*, *11*(5), 923–938. https://doi.org/10.1016/0896-6273(93)90122-8

Xu-Wilson, M., Chen-Harris, H., Zee, D. S., & Shadmehr, R. (2009). Cerebellar contributions to adaptive control of saccades in humans. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *29*(41), 12930–12939. https://doi.org/10.1523/JNEUROSCI.3115-09.2009

Xu-Wilson, M., Tian, J., Shadmehr, R., & Zee, D. S. (2011). TMS perturbs saccade trajectories and unmasks an internal feedback controller for saccades. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *31*(32), 11537–11546. https://doi.org/10.1523/JNEUROSCI.1584-11.2011

Xu-Wilson, M., Zee, D. S., & Shadmehr, R. (2009). The intrinsic value of visual information affects saccade velocities. *Experimental Brain Research*, *196*(4), 475–481. https://doi.org/10.1007/s00221-009-1879-1

Y, D., JM, A., WM, U., & RC, R. (1998). Coding of visual information by precisely correlated spikes in the lateral geniculate nucleus. *Nature Neuroscience*, *1*(6), 501–507. https://doi.org/10.1038/2217

Yang, Y., & Lisberger, S. G. (2014a). Role of plasticity at different sites across the time course of cerebellar motor learning. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *34*(21), 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*(7506), 529–532. https://doi.org/10.1038/nature13282

Yasuda, M., Yamamoto, S., & Hikosaka, O. (2012). Robust representation of stable object values in the oculomotor Basal Ganglia. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *32*(47), 16917–16932. https://doi.org/10.1523/JNEUROSCI.3438-12.2012

Yger, P., Spampinato, G. L., Esposito, E., Lefebvre, B., Deny, S., Gardella, C., Stimberg, M., Jetter, F., Zeck, G., Picaud, S., Duebel, J., & Marre, O. (2018). A spike sorting toolbox for up to thousands of electrodes validated with ground truth recordings in vitro and in vivo. *ELife*, *7*. https://doi.org/10.7554/eLife.34518

Yoon, T., Geary, R. B., Ahmed, A. A., & Shadmehr, R. (2018). Control of movement vigor and decision making during foraging. *Proceedings of the National Academy of Sciences of the United States of America*, *115*(44), E10476–E10485. https://doi.org/10.1073/pnas.1812979115

Yoon, T., Jaleel, A., Ahmed, A. A., & Shadmehr, R. (2020). Saccade vigor and the subjective economic value of visual stimuli. *Journal of Neurophysiology*, *123*(6), 2161–2172. https://doi.org/10.1152/jn.00700.2019

Zee, D. S. (2012). What the Future Holds for the Study of Saccades. *Biocybernetics and Biomedical*

Engineering, *32*(2), 65–76. https://doi.org/10.1016/S0208-5216(12)70037-2

Zee, D. S., Yee, R. D., Cogan, D. G., Robinson, D. A., & Engel, W. K. (1976). Ocular motor abnormalities in hereditary cerebellar ataxia. *Brain : A Journal of Neurology*, *99*(2), 207–234. https://doi.org/10.1093/brain/99.2.207

Zhou, H., Lin, Z., Voges, K., Ju, C., Gao, Z., Bosman, L. W. J., Ruigrok, T. J. H., Hoebeek, F. E., De Zeeuw, C. I., & Schonewille, M. (2014). Cerebellar modules operate at different frequencies. *ELife*, *3*(2050-084X (Electronic)), e02536. https://doi.org/10.7554/eLife.02536

Zur, G., & Joshua, M. (2019). Using extracellular low frequency signals to improve the spike sorting of cerebellar complex spikes. *Journal of Neuroscience Methods*, *328*, 108423. https://doi.org/10.1016/j.jneumeth.2019.108423

Curriculum Vitae

Ehsan Sedaghat-Nejad

CONTACT INFORMATION	Laboratory for Computational Motor Control, Johns Hopkins University 720 Rutland Ave, 416 Traylor Building, Baltimore, MD 21205. E-mail: esedaghatnejad@gmail.com

EDUCATION

Ph.D., Biomedical Engineering   *Aug 2015 – Nov 2021*
Johns Hopkins School of Medicine (JHU), Baltimore, USA

- Area of study: Systems Neuroscience and Neuroengineering
- Thesis: Adaptive control of saccades
- Mentor: Professor Reza Shadmehr

M.Sc., Mechanical Engineering   *Sep 2009 – May 2012*
Sharif University of Technology (SUT), Tehran, Iran

- Area of study: Mechanical Engineering – Applied Design
- Thesis: Muscle synergies as a control strategy of the neuro-muscular system
- Supervisor: Professor Mohamad Parnianpour

B.Sc., Mechanical Engineering   *Sep 2005 – Aug 2009*
Iran University of Science and Technology (IUST), Tehran, Iran

- Area of study: Mechanical Engineering
- Thesis: Evaluation of vehicle ride comfort in varying speed
- Supervisor: Professor Morteza Montazeri

LEADERSHIP ROLES

- **Primary Co-Investigator** of protocol PR19M403, *Neural correlates of motor learning in the marmoset monkey*, Principal Investigator: Dr. Reza Shadmehr, JHU. *Feb 2018 – Nov 2021*.
- **Professional Development Chair** of JHU Translational Neuroengineering Technologies Network (TNT), JHU. *Dec 2020 – Nov 2021*.

SELECTED AWARDS

- MBL Methods in Computational Neuroscience Course Trainee Award, MBL. *2021.*
- BRAIN Initiative Investigators Trainee Award, NIH. *2021.*
- Society for the Neural Control of the Movement Scholarship Award, NCM. *2021.*
- Computational Sensory-Motor Neuroscience Course (CoSMo) Trainee Award. *2016.*

TEACHING EXPERIENCES

- Lead teaching assistant for "Neuromatch Academy", managing a team of 7 teaching assistants and 70 students. *July 2020.*
- Teaching assistant for "Mathematical Foundations of BME", Instructor: Dr. Reza Shadmehr, JHU. *Spring 2018.*
- Teaching assistant for "Principles of Design of BME Instrumentation", Instructor Dr. Nitish Thakor, JHU. *Fall 2017.*

PEER REVIEWED JOURNAL PUBLICATIONS	• **Sedaghat-Nejad, E.**, Pi, J.S., Hage, P., Fakharian, M.A., & Shadmehr, R. (2021). Synchronous spiking of cerebellar Purkinje cells during control of movements. (*bioRxiv*, 2021.09.16.460700). • **Sedaghat-Nejad, E.** & Shadmehr, R. (2021). The cost of correcting for error during sensorimotor adaptation. ***PNAS***, 118(40), e2101717118. (*bioRxiv*, 2021.01.13.426535). • **Sedaghat-Nejad, E.**, Fakharian, M.A., Pi, J.S., Hage, P., Kojima, Y., Soetedjo, R., Ohmae, Sh., Medina, J., & Shadmehr, R. (2021). P-sort: an open-source software for cerebellar neurophysiology. ***Journal of Neurophysiology***, 126(4), 1055–1075. (*bioRxiv*, 2021.03.16.435644). • **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. ***Journal of Neurophysiology***, *122(4)*, 1502–1517. (*bioRxiv*, 683706). • **Sedaghat-Nejad, E.**, Herzfeld, D.J., & Shadmehr, R. (2019). Reward Prediction Error Modulates Saccade Vigor. ***The Journal of Neuroscience***, *39(25)*, 5010–5017. (*bioRxiv*, 555573). • Reppert, T. R., Rigas, I., Herzfeld, D.J., **Sedaghat-Nejad, E.**, Komogortsev, O., & Shadmehr, R. (2018). Movement vigor as a traitlike attribute of individuality. ***Journal of Neurophysiology***, *120(2)*, 741–757.
PEER REVIEWED CONFERENCE PROCEEDINGS	• **Sedaghat-Nejad, E.**, Pi, J.S., Hage, P., Fakharian, M.A., & Shadmehr, R. (2021). Synchronous spiking of cerebellar Purkinje cells during control of movements. ***Proceedings of Advances in Motor Learning and Motor Control***. • **Sedaghat-Nejad, E.**, Hage, P., Pi, J., Shadmehr, R. (2020). Population coding of saccadic eye movements by Purkinje cells of the marmoset cerebellum. ***Proceedings of Advances in Motor Learning and Motor Control***. • **Sedaghat-Nejad, E.**, Shadmehr, R. (2019). Cost of error: mechanisms that modulate sensorimotor learning. ***Proceedings of Advances in Motor Learning and Motor Control***.
INVITED TALKS	• Synchronous spiking of cerebellar Purkinje cells during control of movements. (2021). Invited talk at: *Advances in Motor Learning and Motor Control*. • Control of movement deceleration by the Purkinje cells of the cerebellum. (2021). Invited talk at: *Society for the Neural Control of Movement*. • Population coding of saccadic eye movements by Purkinje cells of the marmoset cerebellum. (2020). Invited talk at: *Advances in Motor Learning and Motor Control*. • Encoding of saccadic eye movements by the Purkinje cells of the cerebellum. (2020). Invited talk at: *Neuromatch Conference*. • Cost of error: mechanisms that modulate sensorimotor learning. (2019). Invited talk at: *Advances in Motor Learning and Motor Control*, Chicago, IL. • Value of error: mechanisms that modulate sensorimotor learning. (2019). Invited talk at: *Gordon Eye Movement Conference*, Lewiston, ME.