Progressions

Author reflections on developments since the publication of "Adaptive Representation of Dynamics During Learning of a Motor Task" by R. Shadmehr and F. A. Mussa-Ivaldi. (1994) J Neurosci 14:3208-3224.

Learning to Predict and Control the Physics of Our Movements

Reza Shadmehr

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

When we hold an object in our hand, the mass of the object alters the physics of our arm, changing the relationship between motor commands that our brain sends to our arm muscles and the resulting motion of our hand. If the object is unfamiliar to us, our first movement will exhibit an error, producing a trajectory that is different from the one we had intended. This experience of error initiates learning in our brain, making it so that on the very next attempt our motor commands partially compensate for the unfamiliar physics, resulting in smaller errors. With further practice, the compensation becomes more complete, and our brain forms a model that predicts the physics of the object. This model is a motor memory that frees us from having to relearn the physics the next time that we encounter the object. The mechanism by which the brain transforms sensory prediction errors into corrective motor commands is the basis for how we learn the physics of objects with which we interact. The cerebellum and the motor cortex appear to be critical for our ability to learn physics, allowing us to use tools that extend our capabilities, making us masters of our environment.

Perhaps nothing is so fraught with significance as the human hand, this oldest tool with which man has dug his way from savagery, and with which he is constantly groping forward. Jane Addams

Introduction

Our hands give us the ability to manipulate the environment, but each time that we use them to hold an object, that object's mass alters the physics of our arm. This makes it so that, as we pick up the object, the brain needs to predict and compensate for its physics. Small errors in this prediction produce substantial disturbances in our movement, as evidenced by when you try picking up a can of soda that you thought was full, but in fact is empty: the arm makes a jerky movement. How does the brain learn to predict and control the physics of our body and the objects that we manipulate?

At Emilio Bizzi's laboratory at the Massachusetts Institute of Technology, we had three rooms. In one room, there was a frog that had ones of its legs attached to a small robot while a student stimulated the spinal cord. In a second room there was a monkey that used its hand to hold a medium-sized robot while a student recorded from its motor cortex. And in the third room, a human volunteer held a bigger version of that same robot while a student recorded her reaching movements. Bizzi believed that we should design experiments that not only produced a rich body of behavioral data but could be studied in multiple species. That way, the question would benefit from tools that were specialized for each animal: spinal physiology in frogs (Giszter et al., 1993), cortical neurophysiology in monkeys (Polit and Bizzi, 1979), and behavioral neuroscience in humans (Mussa-Ivaldi et al., 1985; Shadmehr et al., 1993).

In my robotics class, I learned that, although a mechanical arm was much simpler than a human arm, containing fewer degrees of freedom, the robot still required many pages of equations to represent its physics. In Bizzi's laboratory I experienced the reality: a tiny error in one of those equations would cause havoc in the computer's ability to control the robot, making it act violently. This experience of how bad the system could behave when I had the wrong model of physics made it seem like a miracle that our brain routinely and effortlessly controlled our much more complicated biological arm, and did so while interacting with diverse objects. I found the puzzle that I would spend the rest of my life exploring.

The force field paradigm

With Sandro Mussa-Ivaldi, we came up with an experiment: use the computer to control the physics of the robot, making it produce forces that depended on its state (position, velocity, or acceleration). In this way, the robot behaved as a physical object. People and other animals held that object in their hand and learned to control it. If through experience the brain built a model of the novel physics, then that "internal model" would leave its signature in the descending motor commands to the

arm, and one might be able to study the neural basis of this learning process.

To find this signature, Richard Held showed me his prism glasses and explained the concept of "after-effects," movement errors that occurred when there was a mismatch between the physics that the brain predicted and the one that it actually experienced. Learning reduced this mismatch, but the experimenter could behaviorally probe what had been learned by intentionally forcing a mismatch through resetting the physics back to a null condition. Held had explored this by having subjects wear prism glasses. He had found that, when humans practiced movements with the glasses and were then asked to remove them, their very next movement had large errors, opposite in direction to the errors that were induced by the glasses. That is, people knew that the glasses were no longer on their eyes but still produced motor commands that attempted to partially compensate for the errors that they had experienced with the glasses (Held and Freedman, 1963).

Indeed, people produced after-effects as they experienced a force field. Excitedly, we wrote the first draft of the paper and gave a copy to Bizzi. Bizzi asked me to his office and then told me that he thought that the paper "would make my career." To help with that, he asked that I remove his name from the author list. He then sent a copy of the paper to key faculty in the department, requesting their comments. With their help, we improved the paper and then sent it to the *Journal of Neuroscience* (Shadmehr and Mussa-Ivaldi, 1994).

In Bizzi's laboratory, I learned about science, but more importantly, from Bizzi I learned about mentorship.

Internal model of physics

Development of a new experimental tool, called an "error-clamp," allowed exquisite quantification of the process of learning physics by providing a means for probing the after-effects in the absence of error (Scheidt et al., 2000). Measurements suggested that, through experience, the brain learned a model that predicted the motor commands that should be produced to compensate for the novel physics. This "internal model" had interesting properties: it produced an illusion regarding the sensory state of the arm (Ostry et al., 2010); it allowed people to train in one type of movement (straight line reaches), and generalize to another type of movement (drawing ellipses) (Conditt et al., 1997); following training in one part of the workspace, it generalized to reaches in another workspace (Shadmehr and Moussavi, 2000; Green and Labelle, 2015; but see Berniker et al., 2014); following training with the dominant arm, it generalized to reaches with the nondominant arm (Criscimagna-Hemminger et al., 2003). The coordinate system of within arm generalization appeared to depend largely on the neural representation of proprioception, with strong sensitivity to velocity but poor encoding of acceleration (Hwang and Shadmehr, 2005; Hwang et al., 2006). However, the generalization between the two arms appeared to depend on another coordinate system entirely, perhaps vision (Criscimagna-Hemminger et al., 2003; Malfait and Ostry, 2004; Joiner et al., 2013).

These diverse patterns of generalization suggested that the internal model was unlikely to be a single neural entity, but rather a combination of entities possibly involving distinct regions of the brain. Together, various neural streams appeared to receive a copy of the motor commands as the movement was executed, then predicted in real-time the visual and proprioceptive sensory consequences. That is, a critical element of learning physics was

to predict the sensory consequences of the motor commands in the coordinate system of the sensors.

The way biology had approached the problem of physics was quite different from how a roboticist would go about doing the same. If I wanted to build a robot that held objects in its gripper and moved them, I would code an internal model based on Isaac Newton's insights: linearly separate velocity from acceleration in the equations of motion. Indeed, we build robots with sensors that separately measure acceleration and velocity. Biology, however, had sensors like muscle spindles that combined position, velocity, and acceleration in a form not structurally optimized to represent equations of motion of inertial objects. Although our arms were inertial objects, our sensors had the evolutionary baggage of muscle spindles, which could not measure acceleration independent of velocity. Internal models that people learned appeared consistent with entities that represented motion in the coordinate system of the proprioceptive sensors (Hwang and Shadmehr, 2005; Sing et al., 2009).

However, learning physics of an object required not just internal models that predicted the sensory consequences of the motor commands that moved that object, but also the motor commands that were required so that the object moved as one intended. That is, to learn control of an object, one ultimately needed to learn the motor commands that were necessary to produce an intended sensory consequence. When a force field altered the physics of the task, the motor commands produced sensory prediction errors (the arm did not go where it was supposed to go). To improve performance, the brain needed to transform the sensory representation of error into better motor commands. How did the transformation from sensory coordinates of error to muscle coordinates of commands take place?

As a reaching movement unfolded, errors in performance engaged sensorimotor feedback pathways, producing reflexive and voluntary corrections after a delay. These corrections were too late to compensate for the novel physics but represented a transformation that took prediction errors in sensory coordinates and produced a motor response in muscle coordinates. As theory had predicted (Kawato et al., 1987), experiments demonstrated that the feedback response acted as a template, teaching the brain the motor commands that it should produce to reduce the error (Thoroughman and Shadmehr, 1999; Franklin et al., 2003; Milner and Franklin, 2005; Albert and Shadmehr, 2016). That is, the neural feedback response to error was co-opted by the internal model and used as a template from which to learn. Individuals who learned physics faster appeared to have a better teacher in their existing feedback control system (Albert and Shadmehr, 2016).

Neural correlates of internal models in humans

Suzanne Corkin invited us to try the force field task with Henry Molaison, better known as H.M., the severely amnestic individual who years earlier had regions of his medial temporal lobe bilaterally removed. We put our robot in my wife's station wagon and drove up to Boston to see H.M. He sat quietly in the experiment chair and, like every other volunteer, did not touch the robotic arm that was in front of him. I asked him to hold the robot's handle and move it around, and he did so while looking at his hand. I asked him to look at the monitor that was suspended in front of him, where there was a cursor that moved with the robot handle. A target appeared, and I instructed him to move the cursor to the target. He did so slowly, at which time the computer produced a low tone, meaning "too slow." I encouraged him to reach faster, and after a few trials he reached the target in the

appropriate amount of time, causing the target to "explode" — our primitive animation. After a couple of more explosions, he recalled a story: When he was a kid, he had two pistols and a rifle. His back yard had rabbits and birds that he enjoyed hunting. With pleasure he recalled the details of his hunting days and repeated the story many times throughout the 2 day experiment. In my notebook, I wrote this quote from him: "this is like target shooting."

The task became more difficult as the robot produced a force field, but H.M. learned to compensate for the forces, frequently experiencing those target explosions and producing after-effects (Shadmehr et al., 1998). After some training, he left the room and returned 4 h later. I asked him whether he had seen me or the robot before, and he said "no." Now something interesting happened: he sat in the chair and, without any instruction, grabbed the robot handle, brought it toward him, and began moving it as he looked up at the monitor. I engaged the force field and gave him the first target. He reached while producing motor commands that partially compensated for the field. He retrained in the field for ~30 min, getting those explosions, during which he joyfully repeated the story of his childhood hunting days. He left for the night to return 18 h later. Upon his return, without instruction he once again grabbed the robot and sought out the target. Now the robot was programmed to not produce a field. However, H.M. produced motor commands that expected the field, resulting in after-effects.

The observations in H.M. suggested that the ability to learn physics could proceed despite profound damage to the medial temporal lobe. The fact that the after-effects were present 1 d after training suggested that the motor memory was stored and could be recalled upon seeing and/or holding the object (the robot) without conscious knowledge of having seen the object before. That is, experiencing novel physics engaged an automatic process of forming an internal model that could be recalled despite having little or no accompanying declarative memory.

H.M.'s behavior also demonstrated that, despite lack of declarative memory, he remembered the purpose of the robot and its relationship with the motion of the cursor. In addition to forming a model of the robot's physics, the experience produced a memory that associated the robot with the possibility of a rewarding outcome (target explosions). These observations demonstrated that, despite severe damage to the declarative memory system, the remaining brain could do the following: (1) learn the purpose of a novel tool (the robot's purpose was to move a cursor on the screen, so that targets could explode); and (2) learn to control the physics of that tool (to acquire the target, produce forces that compensated for the robot's physics).

Tom Thach and colleagues had demonstrated that damage to the cerebellum impaired the ability of people to adapt their movements in response to donning of prism glasses (Martin et al., 1996). It was soon discovered that people with cerebellar damage were also impaired in the force field task (Maschke et al., 2004; Smith and Shadmehr, 2005): whereas damage to the basal ganglia through Huntington's disease produced no significant deficits in learning the field, damage to the cerebellum profoundly impaired this ability.

We built an MRI-compatible robot that used pneumatic valves that actuated the arm and produced force fields, and then used that robot to look for regions of the cerebellum that were activated when people experienced novel physics during reaching. We found two distinct cerebellar regions: one in the anterior cerebellum, lobule V, and another in the posterior cerebellum, lobule VIII (Diedrichsen et al., 2005), both ipsilateral to the mov-

ing arm. Studies of people with cerebellar damage that measured learning deficits as a function of the location of damage in their cerebellum largely confirmed the imaging results: the critical regions inferred from lesion studies were in the anterior arm region of the cerebellar cortex, lobules IV and V, and to a lesser extent the lateral posterior lobe (Crus I), ipsilateral to the moving arm (Rabe et al., 2009; Donchin et al., 2012; Burciu et al., 2014). Imaging of the human brain using PET demonstrated that, during weeks of training in a force field, anterior cerebellar cortex ipsilateral to the moving arm showed consistent reductions in regional blood flow (Nezafat et al., 2001).

With development of noninvasive stimulation, it became possible to temporarily alter the function of the cerebellum in healthy people. A study using transcranial direct current stimulation observed that cathodal cerebellar stimulation impaired learning, whereas anodal cerebellar stimulation made the subjects superlearners, improving their ability to learn physics beyond those who received sham stimulation (Herzfeld et al., 2014a).

If internal models of physics formed in the cerebellum, they likely influenced reach motor commands through the cerebello-thalamo-cortical pathway. One way to test this hypothesis was to examine essential tremor patients that had undergone surgery and implanted a deep brain stimulator in the thalamic region that received inputs from the cerebellum. With the stimulator turned on, the output of the cerebellum was disrupted. This produced an immediate reduction in the symptoms of the disease (tremor). However, the patients became impaired in learning of the force field (Chen et al., 2006). That is, despite the fact that thalamic stimulation improved tremor, it impaired the patient's ability to learn physics, providing further evidence that the human cerebellum, and its projections to the motor cortex via the thalamus, were critical for learning of physics.

Another region that appeared important was the primary motor cortex (M1). People who learned to reach in Field A formed an internal model that helped them if they encountered Field A again, but that same internal model impaired performance when they encountered the opposite Field B (Brashers-Krug et al., 1996). Remarkably, this impaired performance in Field B could be rescued if M1 was disrupted via repetitive transcranial magnetic stimulation following learning of Field A (Cothros et al., 2006). The results suggested that a component of the internal model remained present in M1 after completion of training.

Because M1 is a recipient of inputs from the cerebellum (via the thalamus), and the cerebellum is critical for learning physics, it was possible that disruption of M1 produced learning deficits because it blocked the ability of the cerebellum to express its contributions. To consider this, a line of work attempted to excite M1 during learning and measure its downstream effects. Depending on the location of stimulation, a single pulse of TMS to M1 evoked an EMG response in specific arm muscles. This evoked potential was a measure of excitability of the cortical network engaged by the stimulation. During practice of reaching movements that required forces perpendicular to the direction of reach, the brain learned to activate muscles that produced those forces (Thoroughman and Shadmehr, 1999). As training proceeded, excitability of M1 in the preparatory period before a reach increased for the field-specific muscles (Orban de Xivry et al., 2013). When the field was removed during a period labeled "washout," the brain no longer activated the field-specific muscles. However, despite the fact that EMG had returned to baseline, M1 measurements of excitability persisted and were now joined by increased excitability of antagonist muscles. These re-

sults appeared similar to neurophysiological measurements taken from single cells in M1 of monkeys during force field learning: training in Field A coincided with changes in the activity of many cells, but these changes tended to persist during washout and were joined by changes (often in the opposite direction) in another group of cells (Li et al., 2001; Arce et al., 2010b; Mandelblat-Cerf et al., 2011).

The idea that the human motor cortex played an important role in learning of physics gained further support by a series of creative experiments pioneered by Paul Gribble. When people watched other people train in a force field, they learned from the errors that they observed in the other person's movements (Mattar and Gribble, 2005; Brown et al., 2009), suggesting that a sensory prediction error (in visual space) could guide learning in both the person that made the movement and the person who observed that movement. While people who practiced in Field A were impaired in subsequent exposure to Field B, people who observed an actor learn Field A were also impaired in their actual learning of Field B. However, repetitive transcranial magnetic stimulation of M1 in people who had observed an actor learn Field A rescued their learning in Field B (Brown et al., 2009), providing evidence that learning of physics produced a memory that partially depended on the motor cortex.

In summary, imaging, lesion, and stimulation studies in humans suggested that the cerebellum and the motor cortex were two regions critical for learning of physics. In the cerebellum, force field adaptation engaged two distinct regions, one in the anterior lobe and the other in the posterior lobe, both ipsilateral to the moving arm. Following training, a component of the memory of the internal model remained in the motor cortex, acting as a prior that facilitated performance during reexposure to the same field while impairing performance during exposure to the opposite field.

Neural correlates of learning physics in the cerebellum

Purkinje cells (P-cells) are the principal cells in the cerebellar cortex and are a key site of plasticity that may underlie cerebellar contributions to motor learning. Tim Ebner and colleagues performed a series of studies that focused on the activity of P-cells in the lateral zones of lobules V and VI while monkeys held the handle of a robotic arm and moved it along a random path (moving target) or a straight line (stationary target). In some studies, they varied the speed of the moving target and its direction, fitting P-cell simple spikes to kinematics of hand motion (Coltz et al., 1999). They observed that speed scaled the depth of modulation of firing with respect to position and direction of the movement (Roitman et al., 2005, 2009). The simple spikes predicted position and velocity in some cells but lagged these variables in other cells (Hewitt et al., 2011). During reaching movements in the pre-
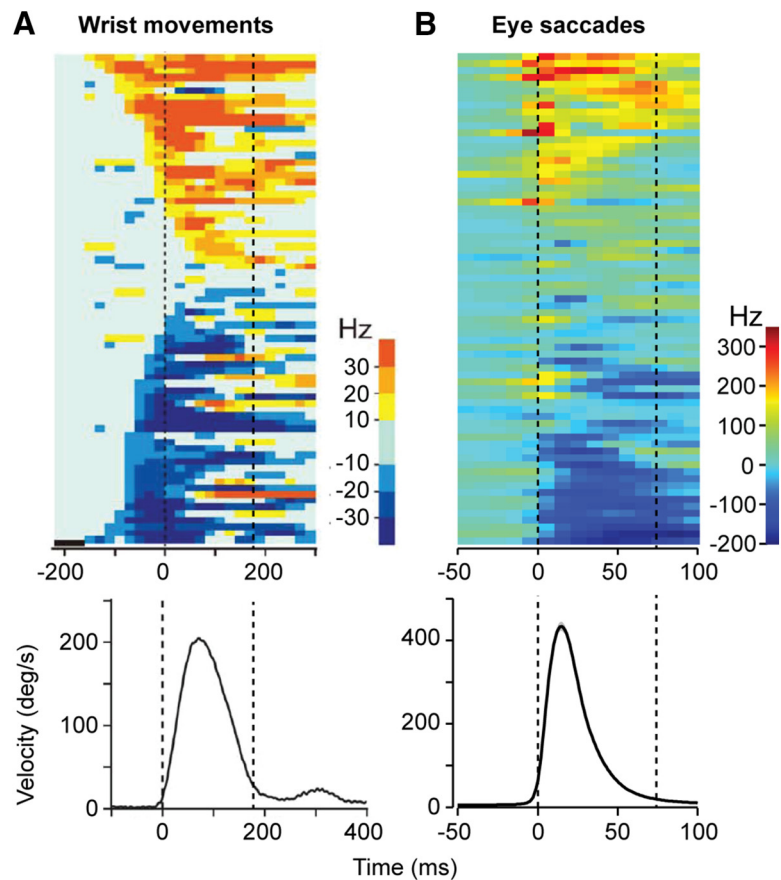

Figure 1. Activity of Purkinje cells during wrist movements and saccadic eye movements display a diversity of patterns, with some cells exhibiting a burst, some cells exhibiting a pause, and change in activity often outlasting the movement. Each row represents activity in a single cell. Activity was measured as rate of discharge over a 10 ms period of time with respect to average rate produced by the same cell in the baseline period. For example, 300 Hz implies that, over a 10 ms period, the cell produced 3 more spikes than baseline. *A*, Activity of Purkinje cells (*n* = 76) in the lateral zone of lobules V and VI of the right cerebellum during 20° movements of the ipsilateral wrist. Bottom trace, Kinematics of a single movement. Vertical dashed lines indicate average movement onset and offset for all trials. Data from Ishikawa et al. (2014). *B*, Activity of Purkinje cells (*n* = 72) in the oculomotor vermis region of the cerebellum (midline regions of lobule VI and VII) during 10° saccades. Bottom trace, Average kinematics. Error bars indicate SEM. Data reanalyzed from Herzfeld et al. (2015).

ferred direction of the cell (direction that produced biggest change with respect to baseline), some cells showed a reduction in their activity (with respect to a baseline period before movement onset), whereas other cells showed an increase (Hewitt et al., 2015). As the monkey trained in various force fields, more than half of the task-related P-cells changed their discharge. However, the patterns of change were diverse, with some cells increasing their activity and others showing a decrease (Hewitt et al., 2015). Regression analysis of simple spikes in individual P-cells demonstrated a diverse encoding of various kinematic parameters of the arm movements, at various time delays. Learning a force field changed both the strength of how each movement parameter was represented in the simple spikes of individual P-cells and the time delay that related that activity to the movement parameter.

Modulation of P-cell simple spikes during force field learning was also noted by Shigeru Kitazawa and colleagues. Monkeys performed elbow flexion and extension while holding a robotic arm and learned to compensate for either a resistive or an assistive force field. Activity in P-cells began changing ~100 ms before movement onset but often persisted long after movement had ended (Yamamoto et al., 2007). The two fields produced diverse patterns of change in the P-cells, further illustrating that, as the

physics of the task changed, P-cells responded and changed their activity during the period of adaptation.

Despite these important studies, we still did not know how P-cells encoded the internal model that might represent the physics of the movement. The problem, in my view, was exacerbated by the limited knowledge that we had on how the simple spikes of P-cells related to control of movements during motion of the arm. To illustrate this point, consider a point-to-point movement performed by the wrist (Fig. 1A, bottom plot). Such a movement coincided with a burst of activity in some P-cells (Fig. 1A) and a pause of activity in other P-cells (Mano and Yamamoto, 1980; Ishikawa et al., 2014). Importantly, the period of modulation for many cells persisted long after the movement had completed. These bidirectional changes in P-cell activity, and seemingly inappropriate durations of discharge, complicated our attempts to develop a framework of how movements were encoded in the cerebellum.

To consider this problem, we turned our attention to cerebellar control of a much simpler goal-directed movement, saccadic eye movements. Like wrist movements, P-cell activity during saccadic eye movements exhibited a bewildering assortment of responses, including cells that exhibited a burst, cells that produced a pause, and cells that did a combination of the two, as illustrated in Figure 1B (Soetedjo et al., 2008; Kojima et al., 2010). When we divided the cells into two groups, we found that the activities of both the burst and pause groups tended to outlast the saccade (Fig. 2A). Therefore, similar to the activities present during arm and wrist movements, during saccadic eye movements the discharge of individual P-cells could not be easily decoded in terms of the kinematics of the ongoing motion.

The key to unlocking this puzzle was the fact that the cerebellum is anatomically organized in such a way that ∼50 P-cells project onto a single deep cerebellar nucleus neuron (Person and Raman, 2011). Therefore, what mattered was the convergence of the simple spikes that were produced by the population of P-cells onto a nucleus neuron. Indeed, Peter Thier and colleagues had earlier demonstrated that a population coding was often a better predictor of eye motion during saccades than activity of individual P-cells (Thier et al., 2000). However, the critical problem was to identify the membership of the P-cells in the micro-cluster that projected onto a single nucleus neuron. That is, was there something in common among the 50 P-cells of the micro-cluster that together projected onto a single nucleus neuron?

We considered the possibility that these P-cells within a micro-cluster were not selected randomly but were organized by their inputs from the inferior olive (Herzfeld et al., 2015). That is, we imagined that the olive projections divided the P-cells where, within a micro-cluster, all the P-cells shared a common input from the olive (Fig. 2B). The inputs from the olive produced complex spikes in the P-cells. Therefore, to compute the population response of the hypothesized micro-cluster, we organized the simple spikes of the P-cells based on a coordinate system that depended on their complex spikes (CS).

Working with our colleagues Yoshiko Kojima and Robi Soetedjo who had collected P-cell data from many monkeys over the course of a decade in the laboratory of Albert Fuchs, my student David Herzfeld reanalyzed these data using the micro-cluster hypothesis (Herzfeld et al., 2015). The approach required us to initially ignore the simple spikes and instead focus on the complex spikes of the P-cells. Kitazawa et al. (1998) had shown that, for arm movements, reach errors produced complex spikes. Importantly, each P-cell had a preference for a specific error direction. During saccades, Kojima et al. (2010) had quantified the
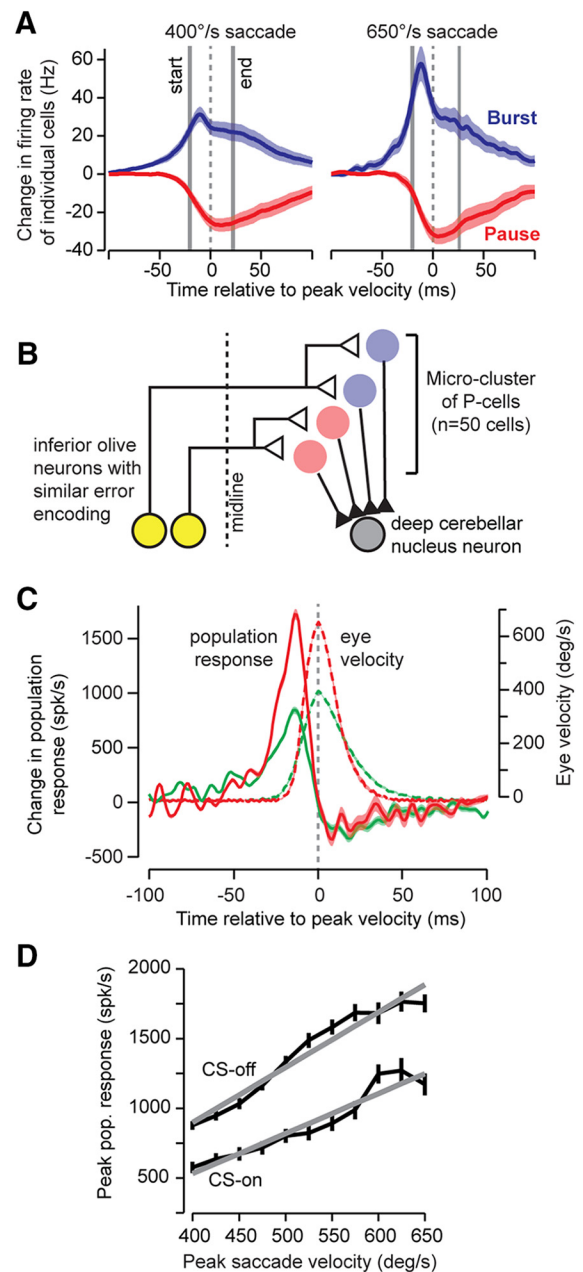

Figure 2. Simple spikes of P-cells organized into a population based on their complex spike properties predicted real-time motion of the eye during saccades via a gain field. *A*, Change in firing rates (with respect to baseline) in the bursting and pausing P-cells for two saccade speeds. Gray bars represent onset and termination of the saccade (width is SEM). *B*, Hypothesized organization of the P-cells into micro-clusters. Each micro-cluster houses P-cells that receive the same error information, and is composed of approximately equal number of pause and burst cells, all projecting to a single nucleus neuron. Open triangles represent excitatory synapses. Filled triangles represent inhibitory synapses. *C*, Population response, as computed via the sum of simple spikes generated by a micro-cluster of P-cells. To compute a population response, we measured the simple spikes of each P-cell as a function of saccade direction with respect to the CS-on direction of that cell. Saccade is in direction CS-off. Population response appears to predict in real-time the velocity of the eye. *D*, Encoding of direction. Peak population response grows linearly with saccade peak velocity but has a higher gain for saccades in CS-off direction. Data from Herzfeld et al. (2015).

direction of error that produced the maximum probability of complex spikes in each P-cell, and then labeled that direction of error as CS-on. We used this complex spike-based coordinate system of sensory prediction error to organize the P-cells into micro-clusters.

When a saccade took place in some direction, we estimated the population response (sum of simple spikes) by computing the direction of movement with respect to CS-on of each P-cell. We found that when the simple spikes of the P-cells were organized based on this hypothesized anatomy, a pattern was unmasked: if the saccade was in the CS-off direction of a group of P-cells, the collection of cells together produced a total number of simple spikes that predicted in real-time the velocity of the eye (Fig. 2C). When the saccade direction changed so it aligned with the CS-on direction of the same group of P-cells, the population still predicted eye velocity in real-time, but now with a lower gain (Fig. 2D). Therefore, within each micro-cluster, the P-cells produced a population response that predicted real-time velocity of the eye, with a gain that multiplicatively depended on direction of motion (i.e., a type of encoding called "gain field"). Encoding via a gain field in the cerebellum was intriguing because a similar type of encoding had been found among the cells in the posterior parietal cortex (Andersen et al., 1985) as well as the motor cortex (Paninski et al., 2004).

The new idea was that the fundamental computational element of the cerebellum was a micro-cluster of P-cells, all sharing the same preferred direction of error (Fig. 2B). This anatomical prediction provided a clue regarding an important question: how were motor memories protected from erasure?

Consider a typical experiment in which a perturbation is imposed on movements, resulting in behavioral adaptation. For example, Kojima et al. (2004) trained a monkey to make saccadic eye movements, and using perturbations produced an increase in the gain of the saccade (Fig. 3A). Following this gain-up period of training, they reversed the perturbation direction, resulting in a return of behavior back to baseline. However, despite the behavioral washout, the reexposure to the previous perturbation resulted in faster learning than before, a phenomenon called "savings." Savings is a fundamental property of motor memory, illustrating that washout training does not erase the previously acquired memory.

The idea that P-cells may be organized based on their preference for error suggested one way with which cerebellar-dependent memories were protected from erasure. Suppose that the perturbations are imposed on only rightward saccades. As training begins, the perturbations are always positive (Fig. 3A), which means that, when the primary saccade ends, the visual stimulus is to the right of the fovea. This visual input activates the left superior colliculus, which excites the right inferior olive, producing complex spikes in the left cerebellum in a few micro-clusters of P-cells, resulting in plasticity. When the perturbation reverses direction, so do the errors. The new errors engage the opposite side of the inferior olive, resulting in complex spikes on the right cerebellum. Therefore, washout does not erase the memory. Rather, washout encourages new learning because the errors engage a new group of P-cells, those that have a preference for leftward errors. Learning followed by washout lays down two anatomically distinct memories, not one. Behavioral experiments in humans provided evidence that supported some of the predictions of this "memory of errors" model (Herzfeld et al., 2014b).

It remains to be seen whether the anatomy described in Figure 2B is indeed present in the cerebellum. If so, organizing the P-cell simple spikes based on their complex spikes might shed light on how the cerebellum encodes movements of the wrist and the arm. Given the similarities in the patterns of discharge during wrist movements and saccadic eye movements (Fig. 1), it is plausible that the encoding present for eye movements may be shared by other types of movements. It also remains to be seen whether this
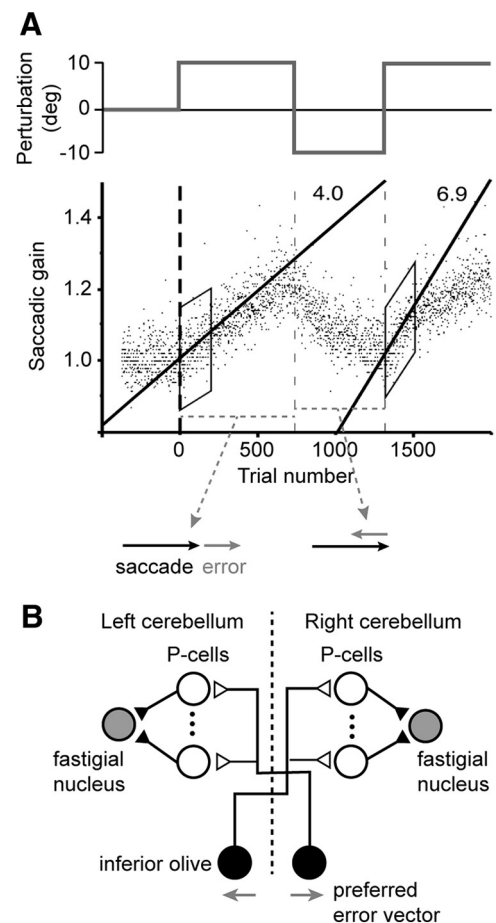

Figure 3. A possible neural basis for the phenomenon of savings in the cerebellum. *A*, Saccadic adaptation. The monkey's rightward saccades are perturbed via intrasaccadic movement of the target. Positive perturbations produce rightward errors, resulting in an increase in saccadic gain. Negative perturbations produce leftward errors, resulting in a decrease in saccadic gain. After gain-up, gain-down training, the animal exhibits savings in the gain-up period, demonstrating that reversal of errors did not erase the memory. The lines and the numbers indicate the slope of the behavioral data. Data from Kojima et al. (2004). *B*, Hypothesized model of the oculomotor vermis region of the cerebellum. Rightward errors during gain-up training engage the right inferior olive, producing complex spikes in the left cerebellum. Gain-down training reverses the direction of error, producing complex spikes in the right cerebellum. Perturbation reversal changes the direction of error. If P-cells are organized based on their preference for error, reversal of error engages a new micro-cluster of P-cells, laying down a competing memory. As a result, training followed by "washout" produces two anatomically distinct memories: one for the errors experienced during training and one for the errors experienced during washout. Unfilled triangles represent excitatory synapses. Filled triangles represent inhibitory synapses.

hypothesis can shed light on the neural process of cerebellar adaptation. Because behavior during adaptation exhibits numerous characteristics that are shared in both arm and eye movements, including multiple time-scales of learning, spontaneous recovery (Kojima et al., 2004; Smith et al., 2006; Ethier et al., 2008), and modulation of error sensitivity (Hanajima et al., 2015), these steps, if successful, may put our field in a position to better understand how the P-cells of the cerebellum participate in the problem of learning physics.

Neural correlates of learning physics in the motor cortex

Teams led by Emilio Bizzi and Eilon Vaadi systematically studied the events that took place in the frontal motor regions of the cerebral cortex during force field training. Both teams characterized activity of cortical cells by computing discharge as a function

of direction of movement. Recordings from muscles during force field learning had demonstrated that, when activity was represented as a function of movement direction, the preferred direction of muscles rotated with training (Thoroughman and Shadmehr, 1999). Bizzi and colleagues trained monkeys to reach and place a cursor at a target in an 8-direction task that involved a baseline period, a force field period, and then a washout period (Li et al., 2001). In the force epoch, many cells in M1 showed a shift of preferred direction in the direction of the applied force. The magnitude of change in the cell preferred direction was similar to change in muscle preferred direction. Interestingly, in the washout period, some cells maintained their change in preferred direction, whereas others rotated in the direction opposite to the field. This provided the first neural evidence that behavioral washout in force field learning was not erasure of the changes that had taken place during training, but addition of new learning that masked the older changes.

Arce et al. (2010a) extended these results by showing that during training in a force field, M1 and premotor cells increased their activity if their preferred direction was opposite to the direction of field. In contrast, M1 and premotor cells decreased their activity if the preferred direction was parallel to the direction of the field. These results demonstrated that, with training in Field A, two groups of cortical cells with preferred directions opposite or parallel to the field altered their activities to counter the field. They tracked these cells as the animal first trained in Field A, and then washout, and observed that cells that had increased or decreased their activities during Field A training did not return to baseline during washout trials (Arce et al., 2010b). That is, training produced changes that depended on the direction of error with respect to the preferred direction of the cortical cell: errors opposite to the preferred direction produced an increase in discharge, and errors parallel to the preferred direction produced a decrease in discharge.

As a result, when training transitioned from Field A to washout, errors changed in direction, and the group of motor cortical cells that had changed their activity during Field A largely maintained their activity, and were now joined by a new group of cells. Because cortical changes were dependent on the coincidence of error direction and preferred direction of the cell, a change in error direction did not erase the previous training, but rather recruited additional cells into the internal model. These results suggested that understanding the neural coding of sensory prediction errors by individual cells might be a critical factor in decoding activity of cells in the cerebellum (Herzfeld et al., 2015) as well as the motor cortex (Inoue et al., 2016).

Overview

When we move an object that has novel physics, we experience a sensory prediction error. The coordinate system of that error is that of the primary sensors, vision and proprioception. These sensors are anatomically linked to reflex pathways that can produce corrective motor commands. Learning physics appears to be in part a process of transforming reflex-generated corrective motor commands into models with which the brain predicts physics: the delayed reflex-dependent responses become predictions.

Sensory prediction errors engage the cortical motor regions as well as the cerebellum. For example, motor cortical neurons respond to errors that are visually sensed at the end of a reaching movement, and stimulation of these cortical neurons produces trial-to-trial reach adaptation (Inoue et al., 2016). In a similar fashion, neurons in the superior colliculus respond to errors that

are visually sensed at the end of a saccadic eye movement, and stimulation of these collicular neurons produces trial-to-trial saccade adaptation (Kaku et al., 2009; Soetedjo et al., 2009). For saccadic eye movements, the visual errors are transmitted from the colliculus to the inferior olive, which then produce complex spikes in the P-cells of the cerebellum, resulting in plasticity. It seems likely that a similar anatomy connects the motor cortex to the inferior olive, and then to the cerebellum, transforming reach errors into cerebellar-dependent motor memories.

Computational models suggest that control of movements may be represented as a process in which the sensory goal forms a utility that depends on reward and effort (Rigoux and Guigon, 2012; Shadmehr et al., 2016). The problem of movement control may be viewed as one in which the utility associated with the goal is transformed into a policy that maximizes that utility via a feedback controller (Todorov and Jordan, 2002). For saccades, the superior colliculus appears to be critical for implementing this feedback controller (Goossens and Van Opstal, 2006), whereas for reaching, the motor cortex may play the role of the feedback controller (Pruszynski et al., 2011). In this framework, the role of the cerebellar cortex may be to transform a copy of the motor commands into a prediction regarding sensory consequences. The output of the cerebellum, however, is not what the P-cells are predicting, but an interaction between the P-cell predictions and mossy fiber inputs to the deep nucleus neurons. Unfortunately, we currently know very little about the mossy fiber inputs to the deep nucleus, which makes it very difficult to understand this final step of computation in the cerebellum. Once this has been worked out, our field should be poised to decipher the seemingly alien language used by the brain to predict and control the physics of our body.

References

Albert ST, Shadmehr R (2016) The neural feedback response to error as a teaching signal for the motor learning system. J Neurosci 36:4832–4845. CrossRef Medline

Andersen RA, Essick GK, Siegel RM (1985) Encoding of spatial location by posterior parietal neurons. Science 230:456–458. CrossRef Medline

Arce F, Novick I, Mandelblat-Cerf Y, Israel Z, Ghez C, Vaadia E (2010a) Combined adaptiveness of specific motor cortical ensembles underlies learning. J Neurosci 30:5415–5425. CrossRef Medline

Arce F, Novick I, Mandelblat-Cerf Y, Vaadia E (2010b) Neuronal correlates of memory formation in motor cortex after adaptation to force field. J Neurosci 30:9189–9198. CrossRef Medline

Berniker M, Franklin DW, Flanagan JR, Wolpert DM, Kording K (2014) Motor learning of novel dynamics is not represented in a single global coordinate system: evaluation of mixed coordinate representations and local learning. J Neurophysiol 111:1165–1182. CrossRef Medline

Brashers-Krug T, Shadmehr R, Bizzi E (1996) Consolidation in human motor memory. Nature 382:252–255. CrossRef Medline

Brown LE, Wilson ET, Gribble PL (2009) Repetitive transcranial magnetic stimulation to the primary motor cortex interferes with motor learning by observing. J Cogn Neurosci 21:1013–1022. CrossRef Medline

Burciu RG, Reinold J, Rabe K, Wondzinski E, Siebler M, Müller O, Theysohn N, Gerwig M, Donchin O, Timmann D (2014) Structural correlates of motor adaptation deficits in patients with acute focal lesions of the cerebellum. Exp Brain Res 232:2847–2857. CrossRef Medline

Chen H, Hua SE, Smith MA, Lenz FA, Shadmehr R (2006) Effects of human cerebellar thalamus disruption on adaptive control of reaching. Cereb Cortex 16:1462–1473. CrossRef Medline

Coltz JD, Johnson MT, Ebner TJ (1999) Cerebellar Purkinje cell simple spike discharge encodes movement velocity in primates during visuomotor arm tracking. J Neurosci 19:1782–1803. Medline

Conditt MA, Gandolfo F, Mussa-Ivaldi FA (1997) The motor system does not learn the dynamics of the arm by rote memorization of past experience. J Neurophysiol 78:554–560. Medline

Cothros N, Köhler S, Dickie EW, Mirsattari SM, Gribble PL (2006) Proactive interference as a result of persisting neural representations of previ-

ously learned motor skills in primary motor cortex. J Cogn Neurosci 18:2167–2176. CrossRef Medline

Criscimagna-Hemminger SE, Donchin O, Gazzaniga MS, Shadmehr R (2003) Learned dynamics of reaching movements generalize from dominant to nondominant arm. J Neurophysiol 89:168–176. CrossRef Medline

Diedrichsen J, Hashambhoy Y, Rane T, Shadmehr R (2005) Neural correlates of reach errors. J Neurosci 25:9919–9931. CrossRef Medline

Donchin O, Rabe K, Diedrichsen J, Lally N, Schoch B, Gizewski ER, Timmann D (2012) Cerebellar regions involved in adaptation to force field and visuomotor perturbation. J Neurophysiol 107:134–147. CrossRef Medline

Ethier V, Zee DS, Shadmehr R (2008) Spontaneous recovery of motor memory during saccade adaptation. J Neurophysiol 99:2577–2583. CrossRef Medline

Franklin DW, Osu R, Burdet E, Kawato M, Milner TE (2003) Adaptation to stable and unstable dynamics achieved by combined impedance control and inverse dynamics model. J Neurophysiol 90:3270–3282. CrossRef Medline

Giszter SF, Mussa-Ivaldi FA, Bizzi E (1993) Convergent force fields organized in the frog spinal cord. J Neurosci 13:467–491. Medline

Goossens HH, Van Opstal AJ (2006) Dynamic ensemble coding of saccades in the monkey superior colliculus. J Neurophysiol 95:2326–2341. CrossRef Medline

Green AM, Labelle JP (2015) The influence of proprioceptive state on learning control of reach dynamics. Exp Brain Res 233:2961–2975. CrossRef Medline

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) Spared modulation of error-sensitivity during a prism adaptation task in people with cerebellar degeneration. J Neurophysiol 114:2460–2471. CrossRef Medline

Held R, Freedman SJ (1963) Plasticity in human sensorimotor control. Science 142:455–462. CrossRef Medline

Herzfeld DJ, Pastor D, Haith AM, Rossetti Y, Shadmehr R, O'Shea J (2014a) Contributions of the cerebellum and the motor cortex to acquisition and retention of motor memories. Neuroimage 98:147–158. CrossRef Medline

Herzfeld DJ, Vaswani PA, Marko MK, Shadmehr R (2014b) A memory of errors in sensorimotor learning. Science 345:1349–1353. CrossRef Medline

Herzfeld DJ, Kojima Y, Soetedjo R, Shadmehr R (2015) Encoding of action by the Purkinje cells of the cerebellum. Nature 526:439–442. CrossRef Medline

Hewitt AL, Popa LS, Pasalar S, Hendrix CM, Ebner TJ (2011) Representation of limb kinematics in Purkinje cell simple spike discharge is conserved across multiple tasks. J Neurophysiol 106:2232–2247. CrossRef Medline

Hewitt AL, Popa LS, Ebner TJ (2015) Changes in Purkinje cell simple spike encoding of reach kinematics during adaption to a mechanical perturbation. J Neurosci 35:1106–1124. CrossRef Medline

Hwang EJ, Shadmehr R (2005) Internal models of limb dynamics and the encoding of limb state. J Neural Eng 2:S266–S278. CrossRef Medline

Hwang EJ, Smith MA, Shadmehr R (2006) Adaptation and generalization in acceleration-dependent force fields. Exp Brain Res 169:496–506. CrossRef Medline

Inoue M, Uchimura M, Kitazawa S (2016) Error signals in motor cortices drive adaptation in reaching. Neuron 90:1114–1126. CrossRef Medline

Ishikawa T, Tomatsu S, Tsunoda Y, Lee J, Hoffman DS, Kakei S (2014) Releasing dentate nucleus cells from Purkinje cell inhibition generates output from the cerebrocerebellum. PLoS One 9:e108774. CrossRef Medline

Joiner WM, Brayanov JB, Smith MA (2013) The training schedule affects the stability, not the magnitude, of the interlimb transfer of learned dynamics. J Neurophysiol 110:984–998. CrossRef Medline

Kaku Y, Yoshida K, Iwamoto Y (2009) Learning signals from the superior colliculus for adaptation of saccadic eye movements in the monkey. J Neurosci 29:5266–5275. CrossRef Medline

Kawato M, Furukawa K, Suzuki R (1987) A hierarchical neural-network model for control and learning of voluntary movement. Biol Cybern 57:169–185. CrossRef Medline

Kitazawa S, Kimura T, Yin PB (1998) Cerebellar complex spikes encode both destinations and errors in arm movements. Nature 392:494–497. CrossRef Medline

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

Kojima Y, Soetedjo R, Fuchs AF (2010) Changes in simple spike activity of some Purkinje cells in the oculomotor vermis during saccade adaptation are appropriate to participate in motor learning. J Neurosci 30:3715–3727. CrossRef Medline

Li CS, Padoa-Schioppa C, Bizzi E (2001) Neuronal correlates of motor performance and motor learning in the primary motor cortex of monkeys adapting to an external force field. Neuron 30:593–607. CrossRef Medline

Malfait N, Ostry DJ (2004) Is interlimb transfer of force-field adaptation a cognitive response to the sudden introduction of load? J Neurosci 24:8084–8089. CrossRef Medline

Mandelblat-Cerf Y, Novick I, Paz R, Link Y, Freeman S, Vaadia E (2011) The neuronal basis of long-term sensorimotor learning. J Neurosci 31:300–313. CrossRef Medline

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

Martin TA, Keating JG, Goodkin HP, Bastian AJ, Thach WT (1996) Throwing while looking through prisms: I. Focal olivocerebellar lesions impair adaptation. Brain 119:1183–1198. CrossRef Medline

Maschke M, Gomez CM, Ebner TJ, Konczak J (2004) Hereditary cerebellar ataxia progressively impairs force adaptation during goal-directed arm movements. J Neurophysiol 91:230–238. CrossRef Medline

Mattar AA, Gribble PL (2005) Motor learning by observing. Neuron 46:153–160. CrossRef Medline

Milner TE, Franklin DW (2005) Impedance control and internal model use during the initial stage of adaptation to novel dynamics in humans. J Physiol 567:651–664. CrossRef Medline

Mussa-Ivaldi FA, Hogan N, Bizzi E (1985) Neural, mechanical and geometric factors subserving arm posture in humans. J Neurosci 5:2732–2743. Medline

Nezafat R, Shadmehr R, Holcomb HH (2001) Long-term adaptation to dynamics of reaching movements: a PET study. Exp Brain Res 140:66–76. CrossRef Medline

Orban de Xivry JJ, Ahmadi-Pajouh MA, Harran MD, Salimpour Y, Shadmehr R (2013) Changes in corticospinal excitability during reach adaptation in force fields. J Neurophysiol 109:124–136. CrossRef Medline

Ostry DJ, Darainy M, Mattar AA, Wong J, Gribble PL (2010) Somatosensory plasticity and motor learning. J Neurosci 30:5384–5393. CrossRef Medline

Paninski L, Shoham S, Fellows MR, Hatsopoulos NG, Donoghue JP (2004) Superlinear population encoding of dynamic hand trajectory in primary motor cortex. J Neurosci 24:8551–8561. CrossRef Medline

Person AL, Raman IM (2011) Purkinje neuron synchrony elicits time-locked spiking in the cerebellar nuclei. Nature 481:502–505. CrossRef Medline

Polit A, Bizzi E (1979) Characteristics of motor programs underlying arm movements in monkeys. J Neurophysiol 42:183–194. Medline

Pruszynski JA, Kurtzer I, Nashed JY, Omrani M, Brouwer B, Scott SH (2011) Primary motor cortex underlies multi-joint integration for fast feedback control. Nature 478:387–390. CrossRef Medline

Rabe K, Livne O, Gizewski ER, 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. J Neurophysiol 101:1961–1971. CrossRef Medline

Rigoux L, Guigon E (2012) A model of reward- and effort-based optimal decision making and motor control. PLoS Comput Biol 8:e1002716. CrossRef Medline

Roitman AV, Pasalar S, Johnson MT, Ebner TJ (2005) Position, direction of movement, and speed tuning of cerebellar Purkinje cells during circular manual tracking in monkey. J Neurosci 25:9244–9257. CrossRef Medline

Roitman AV, Pasalar S, Ebner TJ (2009) Single trial coupling of Purkinje cell activity to speed and error signals during circular manual tracking. Exp Brain Res 192:241–251. CrossRef Medline

Scheidt RA, Reinkensmeyer DJ, Conditt MA, Rymer WZ, Mussa-Ivaldi FA (2000) Persistence of motor adaptation during constrained, multi-joint, arm movements. J Neurophysiol 84:853–862. Medline

Shadmehr R, Moussavi ZM (2000) Spatial generalization from learning dynamics of reaching movements. J Neurosci 20:7807–7815. Medline

Shadmehr R, Mussa-Ivaldi FA (1994) Adaptive representation of dynamics during learning of a motor task. J Neurosci 14:3208–3224. Medline

Shadmehr R, Mussa-Ivaldi FA, Bizzi E (1993) Postural force fields of the human arm and their role in generating multi-joint movements. J Neurosci 13:45–62. Medline

Shadmehr R, Brandt J, Corkin S (1998) Time-dependent motor memory processes in amnesic subjects. J Neurophysiol 80:1590–1597. Medline

Shadmehr R, Huang HJ, Ahmed AA (2016) A representation of effort in decision-making and motor control. Curr Biol 26:1929–1934. CrossRef Medline

Sing GC, Joiner WM, Nanayakkara T, Brayanov JB, Smith MA (2009) Primitives for motor adaptation reflect correlated neural tuning to position and velocity. Neuron 64:575–589. CrossRef Medline

Smith MA, Shadmehr R (2005) Intact ability to learn internal models of arm dynamics in Huntington's disease but not cerebellar degeneration. J Neurophysiol 93:2809–2821. CrossRef Medline

Smith MA, Ghazizadeh A, Shadmehr R (2006) Interacting adaptive processes with different timescales underlie short-term motor learning. PLoS Biol 4:e179. CrossRef Medline

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

Soetedjo R, Fuchs AF, Kojima Y (2009) Subthreshold activation of the superior colliculus drives saccade motor learning. J Neurosci 29:15213–15222. CrossRef Medline

Thier P, Dicke PW, Haas R, Barash S (2000) Encoding of movement time by populations of cerebellar Purkinje cells. Nature 405:72–76. CrossRef Medline

Thoroughman KA, Shadmehr R (1999) Electromyographic correlates of learning internal models of reaching movements. J Neurosci 19:8573–8588. Medline

Todorov E, Jordan MI (2002) Optimal feedback control as a theory of motor coordination. Nat Neurosci 5:1226–1235. CrossRef Medline

Yamamoto K, Kawato M, Kotosaka S, Kitazawa S (2007) Encoding of movement dynamics by Purkinje cell simple spike activity during fast arm movements under resistive and assistive force fields. J Neurophysiol 97:1588–1599. CrossRef Medline