The cerebellum learns to predict the physics of our movements

Scott T. Albert and Reza Shadmehr

Think Tank: Forty neuroscientists explore the biological roots of human experience. David J. Linden (editor), Yale University Press, 2018.

To interact with and manipulate the world around us, our bodies must be constantly in motion. Even now, as you read the words on this page, your nervous system is coordinating a series of rapid eye movements (saccades) that allow your eyes to scan the text in front of you. These saccades are an example of a goal-directed movement; your brain chooses an interesting location in space (i.e. the goal) and executes a movement to bring that spatial location into visual focus. Perhaps you might have a cup of coffee nearby. Periodically, your thirst (or your desire for a dose of caffeine) might drive you to interrupt your reading, and reach for your cup. This process of reaching, like moving your eyes, requires you to select a target in the world around you, and move your arm towards that target. Though this process of reaching may require little conscious effort on your part, moving towards an object in the space around you requires many different neural calculations.

Let us consider this reaching movement in greater detail. To grasp the cup, your brain must move your arm from its current position to the position of the cup. Therefore, to plan your movement, your brain first has to figure out the location of the cup in front of you, and also the current position of your arm. This process of localization is a sensory process. In other words, it requires the brain to use sensory information in the form of vision and proprioception (the oft-forgotten sense that allows you to sense the position of your body without vision) to determine the current positions of the cup and your arm. With the starting and ending positions in hand, parietal regions of your cerebral cortex compute the path of the arm that connects these positions in space- the trajectory of the movement. After the trajectory is determined, your primary motor cortex and other associated pre-motor areas then carefully transform this sensory signal into a motor plan, namely the patterns of muscle contraction that will move your arm along the desired path towards the coffee.

So, the process of moving towards a goal begins with sensory measurements about yourself and the world around you, which are then converted into a set of motor actions. In short, the act of moving is guided by a sensation-to-movement map. This type of map is only one component of movement planning. Intriguingly, to execute our movements correctly, this transformation also runs in reverse - our brain uses planned motor actions to predict how the sensory state of our body will change should these actions be performed.

Why does the brain need to predict sensory events that might happen in the future? To answer this question, let us try an experiment. Take a book and place it in your left hand and then ask a friend to pick up the book from your hand. What you will notice is that as the book is lifted off your hand, your hand does not stay perfectly still, but shifts upwards. Now place the book back in your left hand and use your right hand to pick up the book. Something remarkable happens: the left hand that was holding the book remains perfectly still.

This experiment has a practical lesson: in a fancy party when the waiter comes to you with a tray of drinks, don't pick up the glass yourself. Rather, let the waiter pick up the glass and give it to you. It also illustrates two fundamental ideas about our nervous system. First, sensory information (vision and proprioception) is acquired too slowly to allow us to maintain good control of our body. Despite the fact that we can see our friend reaching to pick up the book off our arm, this visual information arrives at the motor planning regions of our brain too late to be used to precisely reduce the muscle activity that was required to support the weight of the book. As a result of this delay, we overcompensate for the weight of the (now absent) book, and our hand shifts upward. Second, when we generate a movement (lift the book with our right hand), our brain predicts the sensory consequences of that movement before the movement occurs. As a result, when it is our own arm that is lifting the book, the brain predicts the sensory consequences that this action will have on our left arm, and changes the muscle activity of our left arm to predictively compensate for the removal of the book. Predicting the sensory consequences of our motor commands allows us to overcome sensory delays that destabilize the control of our movements.

We are all familiar with the consequences of incorrect planning of our motor actions - we screw up. In our coffee cup example, a problem in your brain's knowledge of how to reach to the cup in front of you will likely cause your hand to topple the cup, thus knocking its refreshing contents onto the table below. From our brain's perspective, the act of knocking over the cup, instead of correctly picking it up for a sip, represents a mismatch between the predicted sensory state of our body and its actual sensory state. In neuroscience, we refer to this error as a sensory prediction error. It turns out that our brain continuously monitors these errors, and uses them to fine tune our motor behavior so that we do not make future mistakes. This process of error correction is so fundamental to our biology that distinct circuits of the brain appear to coordinate learning that results from error. One of these error correction pathways involves the cerebellum.[1] The cerebellum receives electrical information from a part of the brainstem called the inferior olive when a sensory prediction error occurs. The cerebellum uses these error signals to predict and correct for errors that might occur in the future.

Why is this process of error-based learning engrained in the physical structure of our nervous system? One answer to this question relates to the very nature of development. As we develop, our brain must learn to understand our bodies and our environment in order to solidify its sensation-to-movement maps that allow us to move accurately, quickly, and smoothly to the stimuli we sense around us. These sensation-to-movement maps require continual updating on many different timescales of life. Over long timescales, the brain must learn to change how it moves in response to the same stimulus because of changes to our physical bodies, e.g. height, weight, and strength. Over short timescales, the brain must learn how to modify its motor actions because of muscle fatigue, or perhaps changes in our environment such as moving through water versus moving on land. The inability to maintain and update sensation-to-movement maps significantly degrades our ability to make accurate movements. For example, individuals who suffer from certain cerebellar disorders (called cerebellar ataxias) exhibit extreme difficulties in the execution of nearly all voluntary movements such as reaching, talking, looking around and walking. These deficits can, in some cases, confine an ataxic patient to a wheelchair.

Neuroscientists who study motor learning think that one key to understanding movement disorders like cerebellar ataxias, lies in elucidating the process by which the brain updates its sensation-to-movement maps over time, or in other words, how it uses sensory prediction errors to correct for movement errors in the future. Over many decades, the motor learning community has devised clever experimental protocols that allow researchers to carefully control errors that an individual experiences and precisely measure how individuals correct their movements in the future. One of these protocols, known as reach adaptation, focuses specifically on how people learn to correct reaching movements of their arm, or in more prosaic terms, how people avoid knocking over cups of coffee.

To study how people adapt their reaching movements over time, we ask participants to reach from one point to another while holding the handle of a robotic arm. As the individual makes their reach, the experimenter can perturb the movement of the subject's arm by imposing forces on their hand via the robot. In other words, by pushing on the subject's hand we can convert their normally straight and ideal movements, into curved, distorted movements.[2] In order to reach the target in a specified amount of time, the subject must learn to produce additional forces that predict and counteract the robot's intervention. Remarkably, healthy adults can almost completely eliminate the errors caused by the force field in just a few dozen trials.

How does the nervous system learn the appropriate corrections so rapidly? To answer this question, it is useful to consider how people learn things in everyday life – we use teachers. If we want to get better at basketball, we take basketball lessons. If we want to learn a new language we seek a language instructor. We refer to a teacher-driven learning process as supervised learning. Is it possible that a supervised learning process is used for the correction of our movements? If so, who is the teacher? A good teacher is someone who already knows how to do the task we seek to learn. So for movement correction, a good teacher would be someone who knows how to eliminate a sensory error by changing our patterns of muscle activation. Fortunately, healthy individuals have such a teacher; it is called a reflex.

Our reflexes are part of our automatic error correction circuit, and they naturally coordinate the sensation-to-movement map we seek: observing a sensory error and rapidly correcting it with a motor response. For example, when we slip, our sensory imbalance causes us to throw our hands out in front of us without thinking, to catch our imminent fall. Everyone is likely familiar with the stretch-reflex. When a physician abruptly hits the patellar ligament of the knee with a dense object, our quadriceps muscle stretches rapidly, and our spinal reflexes quickly counteract this unintended stretch through a reciprocal contraction that causes us to kick our leg out in front of us. These stretch reflexes are also present in our arm, and can be measured in the laboratory using a technique called electromyography (EMG). With EMG, we can record muscle activity to precisely measure the muscle signals that move our arm in space during both voluntary movement as well as subconscious reflexive movement.

Let us now return to our question regarding force field adaptation. Is it possible that the reflexes that counteract the unexpected force field, also instruct the brain on how to produce better motor commands in the future? Preliminary evidence suggests that the answer to this question is 'yes'. The EMG patterns that our reflexes add to our movement during a force field perturbation, closely resemble the changes we make to our future muscle activation patterns. In other words, corrective muscle activity during a single movement, appears to instruct future changes in our motor actions on the next reach. For example, if we contracted our biceps and relaxed our triceps in response to the imposing force field, our motor system appears to incorporate this feedback correction into a new sensation-to-movement map, i.e. the brain thinks that to achieve the desired straight-line movement in the future, the biceps needs to be more active and the triceps needs to be less active than they were during the original movement.

To summarize, the movements we make are a bit more complicated than they might seem at first glance- they involve complex conversions between sensory predictions and motor actions. And the conversions themselves are constantly being tweaked as we learn more about ourselves and our environment. This learning is triggered by sensory prediction errors, or mismatches between what we expected to feel after moving, and what we actually feel. Here we explored this process of learning, and culminated with the idea that to improve our future movements, our brain likely seeks advice from teaching systems embedded within our brain and spinal cord, namely our reflexes.

Insights into the control of movement, as we have made here, can have large implications in the understanding of our behavior both in healthy and diseased populations. For example, our discussion may inform the question of why some individuals learn motor skills faster than others; perhaps they have a better system of reflexes. More importantly, perhaps the motor deficits that accompany various movement disorders, such as stroke and cerebellar ataxia, are in part caused by a patient's inability to execute appropriate corrections to their movements, thus depriving their brain of an extremely knowledgeable teacher. This research suggests that encouraging patients to make mistakes while moving and reinforcing patient-driven feedback corrections to their movement errors may be one path towards neurorehabilitation.

In other words, as neuroscientists, we hope that the fundamental insights we make concerning the brain will be translated into methods for improving human life. At the very least, hopefully in the not-so-distant future, we will never spill our coffee again.

References

Reza Shadmehr, Maurice A. Smith, and John W. Krakauer (2010) Error correction, sensory prediction, and adaptation in motor control. *Annual Reviews of Neuroscience* 33:89-108.

Reza Shadmehr and Steven P. Wise (2005) *Computational neurobiology of reaching and pointing: A foundation for motor learning*. MIT Press.

Reza Shadmehr and John W. Krakauer (2008) A computational neuroanatomy for motor control. *Experimental Brain Research* 185:359-381.

Reza Shadmehr (2016) Learning to predict the physics of our movements. *Journal of Neuroscience*, 37:1663-1671.

Scott Albert and Reza Shadmehr (2016) The neural feedback response to error as a teaching signal for the motor learning system. *Journal of Neuroscience* 36:4832-4845.

---

[1] The cerebellum (Latin for 'little brain') is a brain region located at the back of the head that contains roughly half of all neurons in our brain.

[2] A common way to do this is by applying a force field to the arm – one of the most common force fields used in reach adaptation is called a curl field. In a curl field, the robot will push the subject's arm in a direction that is perpendicular to their intended direction of motion