Introduction

As you are reading this book, you are moving your eyes rapidly from one point to another. Each movement is a saccade. There is regularity in how you move your eyes in that the speed of your eye movements, the duration of each movement, and the places where you choose to move your eyes are similar to how another person reading this book would move their eyes. In fact, if you look closely at how we move our arms during reaching, and how we move our legs during walking, there are also strong similarities across people. It is useful to notice regularity in nature because occasionally, regularity leads to discovery of fundamental mechanisms. Think of regularity in the motion of celestial objects: Kepler gave an approximate description of the orbits of the planets around the sun (the elliptical path, with the square of the orbital period being equal to the cube of the semi-major axis of the orbit), and Newton was able to account for this regularity by proposing that gravity produces a centripetal force inversely proportional to the square of the radius. In this book, we will focus on the regularity in how our brain perceives the world around us, the regularity in how our brain reacts to sensory stimuli, and most crucially, the regularity in how our brain controls our movements. Our goal is to first describe this regularity, and then attempt to make sense of it using theory, in particular, theory of the kind that relies on mathematics.

Let us start with an example of regularity in perception. Children sometimes play a game where a mischievous kid licks his *index* finger while a victim is watching and then runs over and wipes the *middle* finger on the arm of the second child. One of our students (Molly Marko) told us of this story and we had to try it. Indeed, the victim was certain that his arm felt wet. Why should our brain feel that the arm is wet when in fact it is perfectly dry? To account for this, we will build a theoretical framework in which perception is a combination of two streams of information: one due to prediction, and the other due to observation. Prediction is due to internal models that describe what should happen. Measurements are due to reports from a sensory system that describes what actually did happen. Perception is a combination of the two, describing our belief about what happened.

But perhaps a more interesting question is ‘what is good about sensing the world this way’? What do we gain by making predictions and then combining the predictions with our observations? By presenting the problem in a theoretical framework, we will see that our ability to sense the world is significantly improved by our ability to predict the world. In other words, by relying on predictions, our brain can dramatically improve upon the limitations of our sensors. Because of our brain's ability to predict, we can see what's around us better than what our eyes can tell us, and we can feel where our body is in space better than what our proprioception can tell us.

If our brain is going to rely on its own predictions to form beliefs about sensory observations, then a fundamental problem is how to maintain the accuracy of the predictions. To cope with this, the brain continuously learns from prediction errors. For example, in America soft drinks come in cans that are sold in colorful cardboard boxes. In a grocery store, if you were to pick up one of these boxes and it were empty, your arm would jerk upward and you may lose your balance and fall backward. The reason is that your brain expects the box to be full and weigh a significant amount. The motor commands that you send to your arms (lift the box), and your legs (stabilize the body) are based on this expectation. If the box is empty, the result of the action is a prediction error, i.e., a difference between the expected sensory consequences of the motor commands and the actual sensory consequences. This error produces a change in our prediction. If we were asked to lift another box, the brain predicts a smaller weight than before, and the motor commands reflect this change in prediction.

Now consider the regularity in our movements. The next time that you open a web page on your computer, pay attention to where you look first. Say that the page has some text on it, and also some graphics, and one of the graphics contains a human face. It is quite likely that the first place that your eyes will look at is the face. Advertisers know this, and that's why they place graphics that contains faces in whatever they sell. Why should you look at the face first? Perhaps the reason is that every movement that you make is a reflection of an economic decision: among the potential things that I can look at, which is the most valuable? Attention is indeed a resource that comes in limited supply. If you could measure your eye movements, you would note that the motor commands that your brain sends to your eyes to direct the gaze at a particular location are such that the resulting movement has a particular velocity and duration. For example, a 15 degree displacement of your eyes will take about 0.06 seconds and reach a peak velocity of 400 degrees per second. Among healthy people of a given age these numbers are fairly consistent. So there seems to be some regularity in where we tend to move our eyes, as well as in how our brain controls the motion of the eyes during that movement.

You might say that this regularity has a lot to do with the fact that across people, eyes are biomechanically very similar. Maybe the commands are similar because the biomechanics are similar. Biomechanics are indeed important, but consider the fact that our movements change from the teenage years to young adulthood, and continue to change as we grow older. Our movements change due to diseases like Parkinson's disease, schizophrenia, and depression. Our movements have changed due to evolution: movements made by monkeys are different than those made by humans. A clear example is how the brain moves our eyes during a saccade. When teenagers move their eyes, their saccades have very high velocities, and this speed declines at every decade of life. In Parkinson's disease, saccade speeds are abnormally slow while in schizophrenia saccade speeds are abnormally fast. Monkeys have saccades that are nearly twice the speed of humans, despite nearly identical eye biomechanics. Why?

By presenting this puzzle in a theoretical framework, we will try to make sense of the facts. We will suggest that motor commands that move our body reflect an economic decision that our brain makes regarding reward and effort. In particular, if we view the goal of any movement as to acquire a more rewarding state for our body, and further hypothesize that the brain discounts reward with the passage of time (we would rather receive reward sooner rather than later), then the theory predicts that as temporal discounting of reward changes in the brain (for example due to development, disease, or evolution), the shape of our movements will also change.

Temporal discounting of reward is typically measured in decision making tasks by asking people and other animals to make decisions associated with reward and time. For example, given a choice between $20 now and $80 in one year, you may wait for a year to get $80. Impulsivity is often assessed by quantifying the rate at which an individual discounts reward as a function of time. A person that is more impulsive will discount that $80 more steeply as a function of time, and may therefore pick $20 now. Diseases that affect the reward system can produce changes in impulsivity. Development produces changes in impulsivity, as has evolution. It is a peculiar fact that changes in impulsivity as measured in these decision making tasks are sometimes correlated with changes in movement kinematics (Shadmehr et al., 2010). By attempting to build a theory that links movements with reward, we will build a story that has the potential to explain why development and aging alter the way we move, why diseases that affect the reward system alter the way we move, and why evolution that has brought about different primate species have also made them different in the way they move.

It may seem a bit odd that we are going to be using fairly sophisticated mathematics to study movements like saccades and reaching. These movements may not seem like much in a world where movies show robots fighting robots and enhancing man's power as exoskeletons. Surely moving an eye ball from one place to another or an arm from here to there couldn't be that interesting. What's the big deal?

Well, consider that the motors in robots can produce the same force for a given input over and over again, whereas our muscles fatigue quickly and alter their responses from one movement to the next. The sensors that record motion of a robot do so with far more precision than one finds in the response of our proprioceptive neurons. The transmission lines that connect a robot's motors and sensors to its controller move information at the speed of light, and the controller can process sensory information to issue commands in microseconds. In contrast, our transmission lines (axons) move information slower than the speed of sound, and neural computations often require tens of milliseconds. Indeed, our ability to produce a lifetime of accurate movements is not because we are born with an invariant set of actuators, precise set of sensors, or fast transmission lines. Once we see the problem in this way, we cannot help but be amazed by the consistency and fidelity of a lifetime of movements. There is really nothing ordinary about our most ordinary movements.

How do we maintain our ability to move accurately for our entire life? The answer seems to be that we are born with a nervous system that adapts to these limitations and continuously compensates for them. Internal models are representations that our nervous system builds of our own body, and the world around us. Internal models allow the brain to predict the sensory consequences of motor commands. Our brain uses these predictions to plan actions so that we can convert goals and desired rewarding states into motor commands that achieve them in some efficient way. Internal models change constantly because our body is changing from moment to moment and because the world around us provides us with innumerably rewarding states.

Model is a loaded word. Some models are abstract representations of the world. There are economical models of how prices fluctuate in free markets based on laws such as supply and demand. There are celestial models that allow us to know when and where a comet will appear in the sky. There are models of neuronal networks to explain how a collection of simple interconnected element can do something as difficult as recognizing a familiar face. All these models have in common the use of mathematics for understanding reality but- perhaps more importantly- for predicting the future. Predictions may take place over geologic eras or over fractions of a second. In all cases, the possibility to predict rests on the observation of some ongoing external influence, such as a gravitational force, combined with the knowledge of some relevant state variables that evolve through time. Another important meaning of the word "model" is the physical reproduction of an object, such as the model of ship or of a car. This is also another family of human-made artifacts. When we reproduce an object into a model, we capture some of the object's essential properties. These can be shape features, as in a doll, but also functional features, as in a paper airplane. The essence of all models may be less in what they represent than in what they omit. Modern physics was built upon ignoring large amounts of empirical facts to focus on those few that could bear a clear relation to ideas. This theoretical neglect is consistent with the way in which our brains perceive the external world. Our retinas have a tiny spot, called the fovea, where visual information is captured with high resolution. The rest is a blur. Yet, our visual perception is quite uniform. As we look around, we are constructing the very scene that we observe by combining a mosaic of fragments in our memories. You may be familiar with a video clip by Daniel Simmons, where a group of students play by passing a ball to each other. At some point an actor dressed as a gorilla walks in the middle of the players and engages in gorilla-like gestures before leaving the scene. In Simmons' experiment, subjects look at the scene and count the number of times the ball changes hands. Almost invariably, when at the end they were asked if they noticed the gorilla they were astonished. "What gorilla?" But when one sees the same clip after having been informed of the trick, then one is astonished for not having seen the gorilla the first time. Far from being a flaw, this kind of neglect is a foundation of our ability to operate in a world that presents us with a constant torrent of facts, among which only few are really important.

There is another type of model, the internal models that are formed by the brain as we act upon the world around us. Internal models are not human artifacts but products of biology. They are not made of clay or mathematical equations. However, they carry out many of the same functions of all other models: they constitute a form of fundamental knowledge of the physical world, of the space around us and of the consequences of actions. Most importantly they provide the brain with one of the most essential skills for survival, the ability to predict based on past observations. In the following chapters we will present mathematical concepts that relate to the formation of internal models in movement control and perception. In a way, this book is about mathematical models of brain models. The mathematical notation is our own instrument for describing how representations are formed, what structure they may have and how our own theories can be tested.

We do not expect that the brain carries out computations in the way we do. But, hopefully, to the extent that we are capturing some truth about information processing in the brain, the behavior and the structure of our models have a degree of equivalence with the way in which biology operates.

Regularity in how people and other animals adapt to their environment provides us with clues as to how internal models are represented in the brain. We are going to be using primarily behavioral data from experiments in humans and other animals in order to describe the regularity, and then use mathematics to ask whether the behaviors make sense. Because we are not going to be talking about the 'how' question with biology, e.g. we are not going to be asking about how neurons in the brain might actually perform the computations that are implied by our mathematics, our approach is going to be fraught with danger. After all, it is quite possible that we cannot know why an organism does something until we know how it was built. That is, we may not be able to successfully theorize about the regularity in the way that our brain moves our body until we are much farther along in understanding the basic facts regarding the biology of our brain and its evolutionary history. Can we have any hope for success in using a purely theoretical approach?

Stephen J. Gould (1995), the eminent paleontologist considered this question in a broader sense and wrote the following regarding the role of theory:

> Geology, in the late 18th century, had been deluged with a rash of comprehensive, but mostly fatuous, 'theories of the earth' --- extended speculations about everything, generated largely from armchairs. When the Geological Society of London was inaugurated in the early 19th century, the founding members overreacted to this admitted blight by banning theoretical discussion from their proceedings. Geologists, they ruled, should first establish the facts of our planet's history by direct observation --- and then, at some future time when the bulk of accumulated information becomes sufficiently dense, move to theories and explanations. … In mid career, in 1861, in a letter to Henry Fawcett, Darwin reflected on the false view of earlier geologists. In so doing, he outlined his own conception of proper scientific procedure in the best one-liner ever penned. "About thirty years ago there was much talk that geologists ought only to observe and not theorize; and I well remember someone saying that at this rate a man might as well go into a gravel pit and count the pebbles and describe the colors.

> How odd it is that anyone should not see that all observations must be for or against some view if it is to be of any service!"

Admittedly, theories can also act as straight-jackets, focusing attention on observations that agree with the theory and blurring views that can challenge it. In telling our story, our goal will not to be to sell one view or another, but to point out that there is true wonder in the data, and that on occasion theory allows one to make sense of a good deal of it.

Reference List

Gould SJ (1995) Dinosaur in a haystack: reflections in natural history. Harmony Books.

Shadmehr R, Orban de Xivry JJ, Xu-Wilson M, Shih TY (2010) Temporal discounting of reward and the cost of time in motor control. J Neurosci 30:10507-10516.