

1 Granular cell math

We posit that nearby mossy fibers encode related information, like similar lengths of the same muscle, or particular lengths of muscles that are near to each other and therefore likely to be utilized in the same action. Other than this, we posit that each granular cell makes its connections at random - that is, it selects 4-5 random mossy fibers within a particular distance d of the granular cell soma.

2 Intro inputs

Inputs to the cerebellum from the proprioceptive system

One important class of inputs is that of the muscle spindles. These are sensory neurons that fire most rapidly when a particular muscle has a particular length. Given a set of length assignments, for each muscle, it is possible to calculate the angles of each joint in the body, and thus to know the exact position of each limb, and all the parts of each limb, like fingers.

3 Intro to negative pathway

The negative pathway functions as a classifier - given the input from the mossy fibers, it attempts to classify each potential elementary action at each time step as either *positive* (should be allowed through, although it will only take place if it is produced by the positive pathway at that time step) or *negative* (should be suppressed, so that it will not take place even if it is produced by the positive pathway).

The negative pathway reencodes its input as a very large collection of product features, in a way similar to a support vector machine with a polynomial kernel. This encoding is implemented by the granular cells. Each granular cell is active only when all of the inputs in its *active set* are active. The active set is a subset of the 4-5 inputs the granular cell receives. The active set is learned by the granular cell - it is the largest subset of the granular cell's inputs that is likely to be active at the same time.

The purpose of the active set is to make each granular cell be maximally informative to the classification system. If we choose the active set to be too large, then the granular cell will almost never fire, and thus will convey almost no information almost all of the time.

If we choose the active set to be too small, the granular cell will not give us very specific information. This can be seen in the prototypical example where all of the mossy fiber inputs are from muscle spindles, which fire most rapidly when a particular muscle has a particular length. Knowing only a single muscle length does not give us as much information as knowing the lengths of multiple muscles simultaneously. For example, during the walking cycle, a particular muscle in the leg may have the same length at two different points in the cycle, and thus making a decision based on the muscle being this length is likely to

cause us to make an incorrect decision at one of the two points. (It is extremely likely that the two different points in the walking cycle will require different muscle activations for the walking cycle to function correctly.)

4 Consolidation

The theory we've described is sufficient to learn skills through practice. However, it is well-known that skill at many tasks improves in the periods between practice. This phenomenon is called "consolidation". Experimental evidence suggests that during consolidation, the cerebellum transfers learning from the negative pathway to the positive pathway.

It seems likely that consolidation takes place through a sort of rehearsal, where the brain simulates a particular activity without actually performing the correct muscle movements. This would provide a functional explanation for the sort of procedural imagery that many people report after performing some skilled action (a classic example is playing Tetris). This theory is appealing because we don't need to postulate any additional learning mechanisms - all we need to postulate is some mechanism for preventing the actual expression of muscle movements. This could be as simple as inhibitory neurons suppressing the activity of neurons downstream from the cerebellum. Alternatively, it could be the case that, during rehearsal, commands are issued to the cerebellum at a lower firing rate, which would result in the cerebellar outputs being issued at lower firing rates, which would result in muscle twitches too small to be detected.

5 Postulate long-term and short-term learning in the negative pathway

It seems likely that there are actually two learning mechanisms in the negative pathway, a long-term mechanism and a short-term mechanism. Both would be implemented in terms of synapse strength. The simplest method would be for there to be two weights at each synapse, the long-term weight w_{ij}^{long} and the short-term weight w_{ij}^{short} . At each time step, we would apply the standard learning rules to w_{ij}^{short} in the form of an update Δ_{ij} , and then move w_{ij}^{short} closer to w_{ij}^{long} . At the same time, we would move w_{ij}^{long} towards w_{ij}^{short} more slowly:

$$\begin{aligned}w_{ij}^{\text{short}}(t+1) &= (1 - \alpha) (w_{ij}^{\text{short}}(t) + \Delta_{ij}(t)) + \alpha w_{ij}^{\text{long}} \\w_{ij}^{\text{long}}(t+1) &= (1 - \beta) w_{ij}^{\text{long}}(t) + \beta w_{ij}^{\text{short}}(t),\end{aligned}$$

with $0 < \beta < \alpha \ll 1$.

This would ensure that valuable information isn't lost unnecessarily, but also that any skill could be permanently modified through sufficient practice.

6 Talk about learned reflexes and other weird things that come from the positive pathway

Since there is no distinction between command mossy fibers and context mossy fibers in the naive Bayes model, that means that we are able to acquire “learned reflexes” - if a particular sensory signal is reliably followed by a particular action, the positive pathway will learn to trigger the action automatically upon receiving that sensory signal.

7 Talk about the positive pathway speeding up movements and the negative pathway preventing it from speeding things up too much

8 IO fires on sensory input that predicts a false negative

One of the central problems of the cerebellum is that it is trying to learn to perform actions quickly by observing the behavior of a slower system. It is thus the case that error signals will always come too late. One method of addressing this is to make learning strongest when the error signal comes slightly late, which is observed in the synapses between parallel fibers and Purkinje cells.

Another way to address the problem is to use any other means available to predict actions that will be taken in the future. The most obvious such signal is pain - we can predict that the eye will try to close immediately upon the perception of pain, so we can use the pain itself as the training signal in the inferior olive. Similarly, a limb should generally be withdrawn towards the body when pain is perceived in it, so we can use the pain itself as a training signal that triggers the withdrawal motion.

9 Merging of the first half of one movement with the second half of another

There are other ways to learn a transition between state cells. For example, suppose two commands are consistently issued at the same time. Each command triggers a set of state cells at subsequent time steps; for simplicity, assume that each command triggers a linear sequence of state cells, say s_1, s_2, s_3, s_4 for the first command and s_5, s_6, s_7, s_8 for the second command. If the two commands are consistently issued at the same time, we may learn a transition between s_2 and s_7 , since s_2 reliably fires on the time step before s_7 fires.

10 prediction: Should be pairs of DCN neurons with similar firing patterns, one of which projects to IO, one of which projects as output

Prediction 10.1. There should be pairs of deep nuclear cells with similar firing patterns, one of which projects to the inferior olive and is inhibitory, and one of which projects to premotor areas or wherever else the output of the cerebellum projects to.

11 many zones are active during learning, reduces to a single zone

Imagine the cerebral cortex planning a sequence of movements. It has at its disposal all of the commands which the cerebellum is capable of executing, as well as the ability to issue motor commands directly. The cerebral cortex must have some model of what each command does, in order for it to use each one appropriately. When the cerebral cortex encounters a new task, it can make guesses as to which existing commands are appropriate for the new task, or it can try to create a new command. In either case, it will initially be issuing a random set of commands, which will over time be narrowed down to a smaller set of commands which have proven themselves useful. This explains the finding that, when presented with a novel task, the cerebellum initially shows activity in many regions, but that eventually the activity is confined to a small region.

12 State mossy fibers interspersed with others

This helps make the negative pathway maximally effective.

13 Bayes' Theorem

Let A and B be two random variables. Bayes' Theorem states that:

$$P(A = a|B = b) = \frac{P(B = b|A = a)P(A = a)}{P(B = b)}.$$

The probability $P(A = a)$ is called the *prior probability* - it is the probability that A takes on the value a before we know the value of B . The probability $P(B = b|A = a)$ is called the *likelihood*. The probability $P(A = a|B = b)$ is called the *posterior probability* - it is the probability that A takes on the value a after we know that the value of B is b .

Bayes' Theorem is often used to calculate optimal predictions given evidence. Specifically, suppose that we want to predict the most likely value of A , and

that we have observed the evidence that $B = b$. The most likely value for A is the a that maximizes the posterior probability $P(A = a|B = b)$.

For a random variable A that only takes on two values (say, 0 and 1), we can write Bayes' Theorem in a different way:

$$\frac{P(A = 1|B = b)}{P(A = 0|B = b)} = \frac{P(B = b|A = 1)}{P(B = b|A = 0)} \cdot \frac{P(A = 1)}{P(A = 0)}.$$

Thus, the ratio of the posterior probabilities $P(A = 1|B = b)$ and $P(A = 0|B = b)$ is the product of the ratio of the likelihoods $P(B = b|A = 1)$ and $P(B = b|A = 0)$; and the ratio of the prior probabilities $P(A = 1)$ and $P(A = 0)$.

14 Redoing Naive Bayes

We wish to estimate $P(s_i(t)|s_1(t-1), \dots, s_n(t-1))$.

Using Bayes' Theorem, this can be rewritten as

$$P(s_i(t)|s_1(t-1), \dots, s_n(t-1)) = \frac{P(s_i(t))P(s_1(t-1), \dots, s_n(t-1)|s_i(t))}{P(s_1(t-1), \dots, s_n(t-1))}.$$

We only care about the numerator of this fraction, since the denominator is the same for $s_i(t) = 1$ and $s_i(t) = 0$.

Using the chain rule for conditional probability, we can write

$$P(s_1(t-1), \dots, s_n(t-1)|s_i(t)) = P(s_1(t-1)|s_i(t)) \cdot P(s_2(t-1)|s_i(t), s_1(t-1)) \cdots P(s_n(t-1)|s_i(t), s_1(t-1), \dots, s_{n-1}(t-1))$$

We now make the assumption (known as the "naive Bayes" assumption) that each $s_j(t-1)$ is conditionally independent of every other $s_k(t-1)$ given $s_i(t)$. That is, once we know the value of $s_i(t)$, the variables $s_j(t-1)$ and $s_k(t-1)$ are independent. This simplifies the above to

$$P(s_1(t-1), \dots, s_n(t-1)|s_i(t)) = P(s_1(t-1)|s_i(t)) \cdots P(s_n(t-1)|s_i(t)).$$

We therefore have

$$P(s_i(t)|s_1(t-1), \dots, s_n(t-1)) = \frac{1}{Z} P(s_i(t)) \prod_j P(s_j(t-1)|s_i(t)).$$