

Introduction to Neural Networks

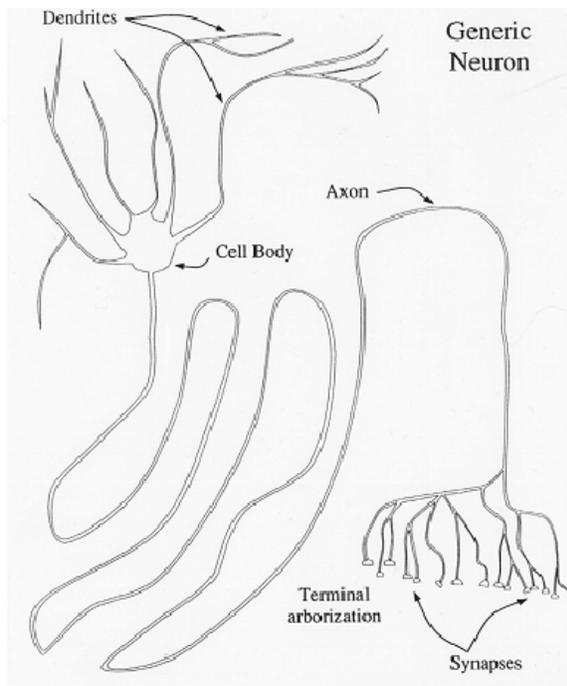
Daniel Kersten

Lecture 2

Getting started with *Mathematica*

Review this section in Lecture 1

The Neuron - overview of structure



From Anderson (1995)

Basic Structure

Information flow: dendrites -> (soma -> axon hillock) -> axon -> terminal buds

Dendrites

The information receiving end of a neuron is a tree-like structure consisting of "dendrites" with special processes or connection sites called synapses. Much computational power is thought to reside in the strength of connections, and in the dendritic tree itself. In this course, we will primarily examine the computational properties of groups of simple neurons, rather than aggregates of dendrites in a single neuron.

Dendrites play the role of wires that convey information through changes in voltage. But they behave rather differently than copper wires. These neural processes are tubes of ionized cytoplasm sitting in a bath of ionized fluid whose ionic composition is not that much different from "seawater". The inside of these tubes during resting state sits at about minus 60-70 millivolts relative to the outside of the cell. The tubes are on the order of microns in diameter, (but other processes, such as the axons discussed below, can reach half a millimeter or so). And for further comparison, the membranes making up the tubes are on the order of 50 Angstroms (50×10^{-10} meters) thick.

Signal transmission is limited by high electrical resistance of the axoplasm, and high electrical capacitance of the neural membrane. Information transmission consequences of these properties are:

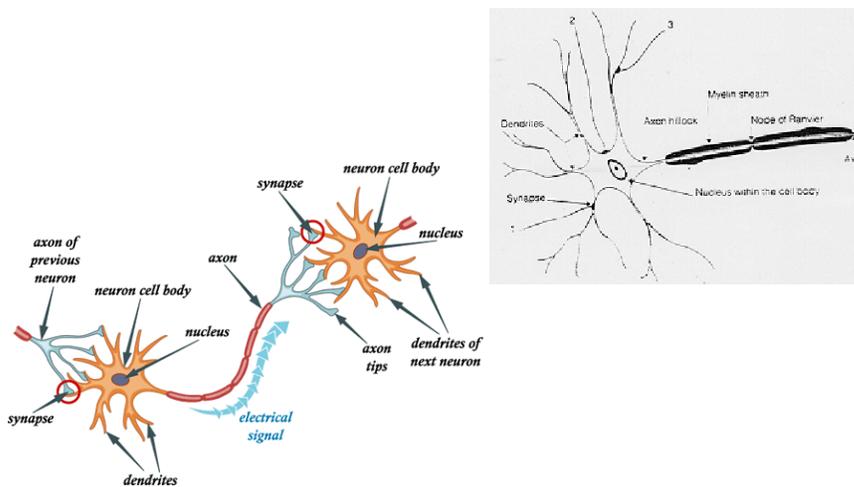
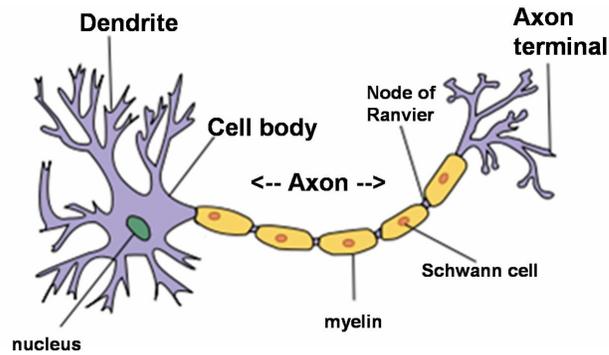
- *the voltage potential changes have a short spatial range of influence, with the amplitude decreasing rapidly as one moves away from the synaptic source.*
- *the signals take time to build up in amplitude, and the signals travel relatively slowly.*

We'll take a quantitative look at these facts shortly.

Soma (or cell body)

- Integrates dendritic signals--*spatial integration* from sites along a dendrite and between dendrites
- The storage of electrical charge across the membrane, and the chemical nature of synaptic transmission leads to *temporal integration* of signals. This observation together with spatial integration of signals from the dendritic tree arriving at the axon hillock will lead to our basic model of the neuron.

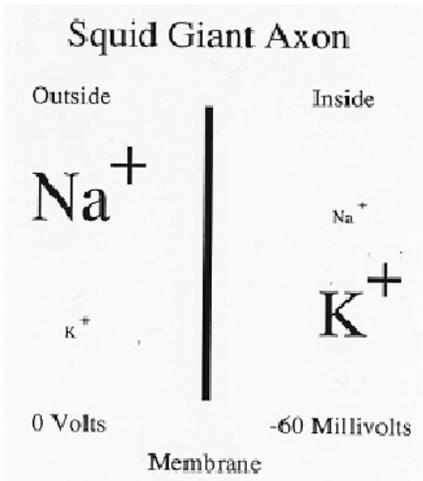
Axon hillock and axon



See <http://www.youtube.com/watch?v=ob5U8zPbAX4>

As seen in the above figures, certain neurons are equipped with a specialized process called an axon that serves to "digitize" the data into all-or-none responses (voltage changes) called *action potentials* or *spikes*. This digitization occurs at the axon hillock near the cell body. There is *passive*, "analog" conduction, called *electrotonic* conduction along the dendrites up to the axon hillock at which point, if there is a sufficient potential change to reach threshold, an *active* process of depolarization kicks in leading to a spike in membrane voltage.

Depolarization means the voltage potential difference across the membrane *decreases*; *hyperpolarization* can also occur, where the voltage difference *increases*.



The potential is maintained by ionic imbalance (excess Na^+ outside, and K^+ inside). The balance between ionic concentration and electric field forces is determined by the Nernst equation (e.g. see Anderson for a derivation).

The action potential signals are carried by rapid (1 msec) voltage depolarizations going from -70 to +40 mV via Na^+ influx, and K^+ outflow through the membrane. From the axon hillock on, a myelin sheath serves to lower the capacitance, increase resistance, and speed up conduction. However it interferes with the regenerative processes that preserve the all-or-none response. At periodic points (Nodes of Ranvier) the myelin sheath is interrupted where high extracellular concentrations of Na^+ ions exist with sodium gates. When a small depolarization arrives, this decreases membrane conductance allowing an increased depolarizing influx of Na^+ , regenerating the spike.

Terminal arborization and terminal buds

Neurons with axons end in a terminal arborization. The terminal buds make synaptic contacts with the dendrites of subsequent neurons, and we have the beginnings of a neural network. Synaptic contacts can either be electrical or chemical, but more about these later.

Basic electrophysiology

A closer look at the passive and active signal transmission properties of the neuron.

Passive properties

We noted two problems: 1) maintaining a signal over a long distance and 2) with sufficient speed. These are BIG problems for an organism that has to transmit signals fast over 10s of centimeters.

Let's take a more quantitative look at these problems to see how they arise from the passive electrical properties of neuronal "electronics".

Membrane electronics can be approximated as a discrete set of sub-circuits, or as a continuous cable.

The advantage of the former is that it is extensible to modeling active, non-linear neuronal properties. The advantage of the second is that there are closed form solutions in simple cases.

Discrete model - sequence of RC circuits

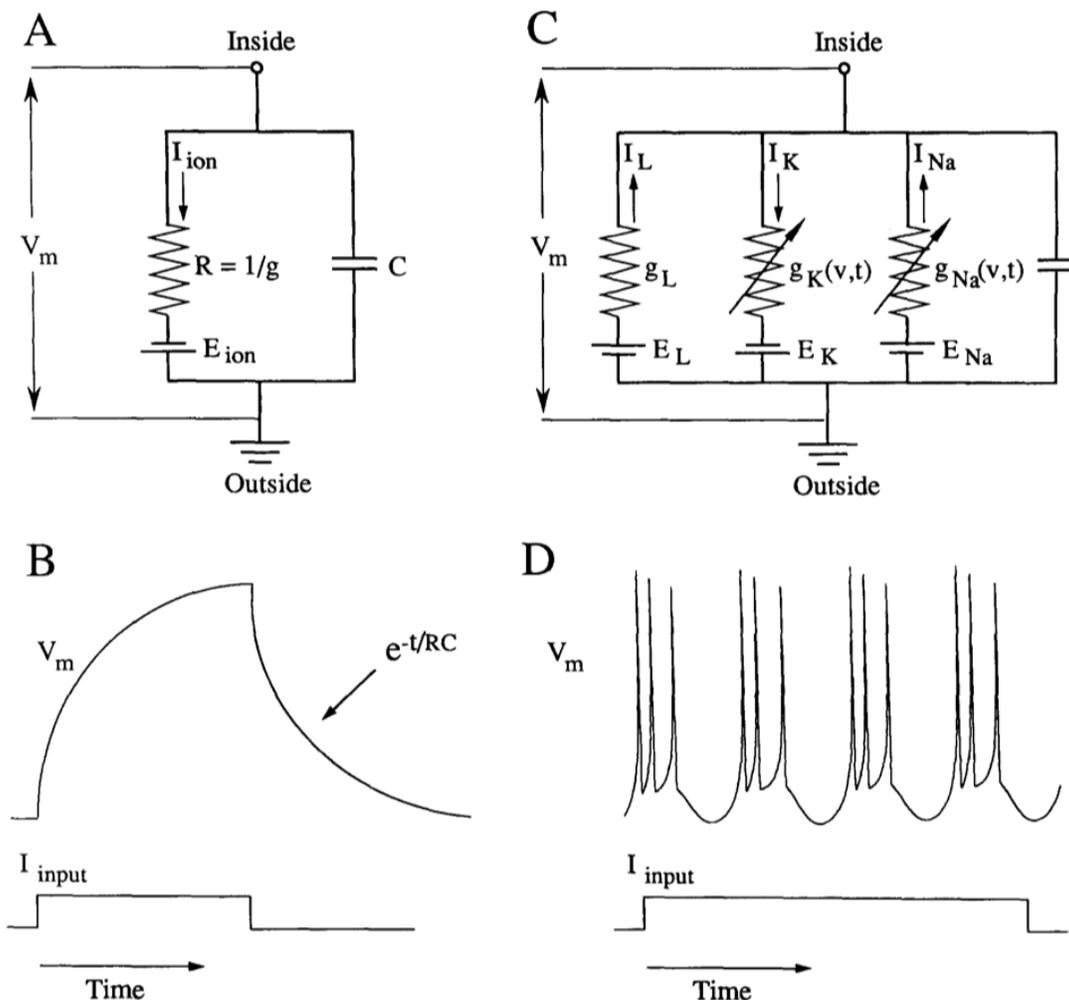
Modeling the passive electrical properties as a function of **time**.

- A small portion of the membrane can be modeled by an “RC circuit” where R stands for resistance, and C for capacitance.

RC-circuits are “low pass temporal filters”, i.e. favor signals with low temporal frequencies. We’ll see shortly that there is a temporal delay in voltage response characterized by time constant τ or rise time.

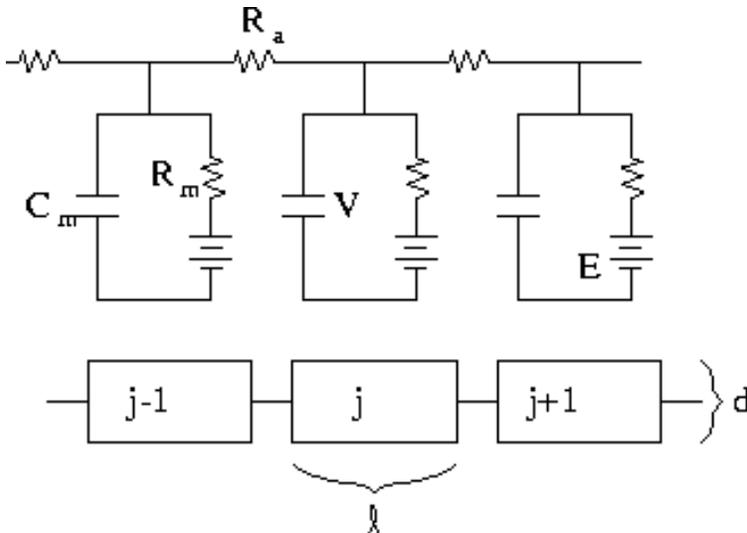
We can also model the passive electrical properties as a function of **space**.

- If we imagine cascading a series of RC-circuits, each connected by additional resistances, we have a discrete model of a section of neural membrane. This kind of model is good for computer simulation. But a continuous model can be solved exactly under simple assumptions.



From Segev (1992). A. illustrates an RC-circuit at a single point of *passive* membrane. B is the temporal response to a step current input. C illustrates additional variable conductance components that model

the electrical processes of spike generation (panel D)—the *active* properties. To model the active properties, one needs a more complicated set of differential equations: the Hodgkin-Huxley equations. (See Claude Meunier and Idan Segev, 2002 for an overview and critique). In the next lecture, we'll use *Mathematica* to derive a simplification of the Hodgkin-Huxley equations.



From Ermentrout (<http://www.cnbc.cmu.edu/~bard/passive2/passive2.html>).

Continuous model - the cable equation

A continuous model over time and short lengths is the “Cable equation” (Anderson, pages 25-32), made famous by Lord Kelvin (Sir William Thomson, 1824-1907) in the context of submarine telegraph cables (Ireland to Newfoundland in 1858.). http://en.wikipedia.org/wiki/Cable_theory

We'll see below how the cable equation predicts an exponential drop-off of voltage with distance for constant current. The length constant λ (distance to the $1/e$ drop-off point or 63% drop) is on the order of millimeters. By solving the cable equation governing the voltage change over distance and time, we can get a quantitative idea of how voltage drops with distance, and how voltage changes with time—the basic message being that *change is not instantaneous*.

Solutions of the cable equation

The cable equation is given by:

$$V = \lambda^2 \frac{\partial^2 V}{\partial^2 x} - \tau \frac{\partial V}{\partial t}$$

where $V = V(x,t)$ is the voltage across the membrane as a function of distance x along the membrane and time t .

($\text{L}^{\text{A}}\text{T}_{\text{E}}\text{X}$ format: $V = \lambda^2 \frac{\partial^2 V}{\partial^2 x} - \tau \frac{\partial V}{\partial t}$)

Let's look separately at the space and time properties.

Space.

In a later Notebook, we'll see how to use *Mathematica* to find solutions to equations, including differential equations for passive and active membranes. For now, let's take the steady-state solution for a fixed voltage, V_0 , at a specific place, say $x = 0$, and see how the voltage drops away from zero. Steady-state

means that the voltage is no longer changing with time, or in another words that $\frac{\partial V}{\partial t}=0$.

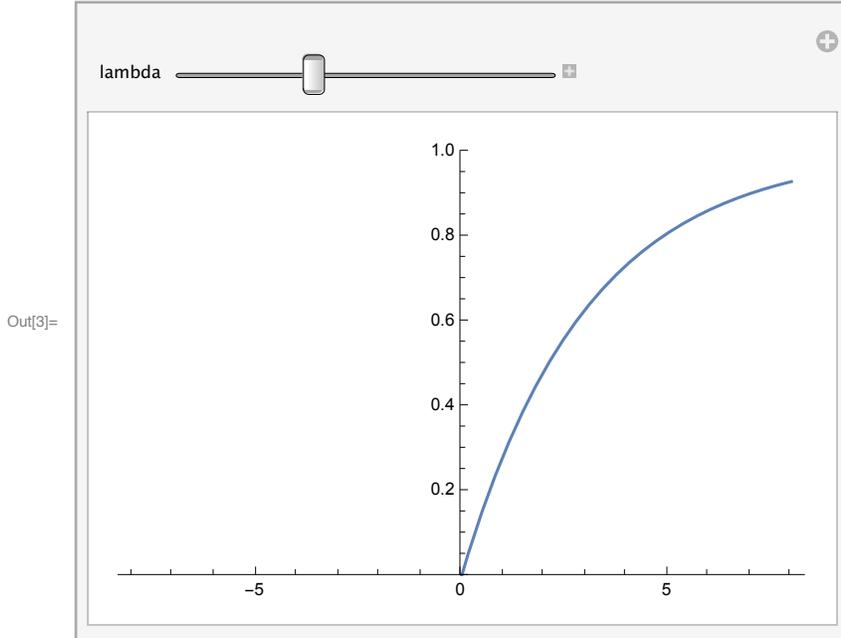
$$V = \lambda^2 \frac{\partial^2 V}{\partial^2 x}$$

A solution of this equation (which you can verify by differentiating twice, see exercise below) is a standard result. The voltage drops exponentially with distance from the point of stimulation. Lambda (λ) is the "space constant", which for an axon would be about $\lambda = 3$ to 5 mm.

$$V(x) = V_0 e^{-x/\lambda}$$

```
In[1]:= V0 = 1; lambda0 = 3; (*Space constant*)
V[x_, lambda_] := V0 Exp[-Abs[x]/lambda];

In[3]:= Manipulate[Plot[V[x, lambda], {x, -8, 8}, PlotRange -> {0, 1}],
  {{lambda, lambda0}, 2, 5}]
```



The space constant $\lambda = \sqrt{(d/4) R_m / R_a}$ where d is the diameter of the cable and R_m and R_a are the membrane and axial resistances, respectively.

Use *Mathematica's* derivative function `D[]` on `V[x]` re-defined below to verify the solution. Differentiate $V(x)$ twice with respect to x , where $V(x)$ is re-defined over positive x values (it is simpler to treat positive and negative x separately). You can test it with: `V[x] == lambda^2 * V''[x]`

```
Clear[V];
V[x_] := V0 Exp[-x/lambda];
```

Time.

Now let's see how membrane voltage changes with time at a single location by plotting a time solution to the cable equation. Let's assume some initial conditions. Suppose a 1 volt step is applied (through some resistance to generate a step current change) across the neuron membrane at time $t=1$. How

does voltage change with time at a fixed point x ? Later, we'll see how to calculate temporal dependence more generally, but for this simple case:

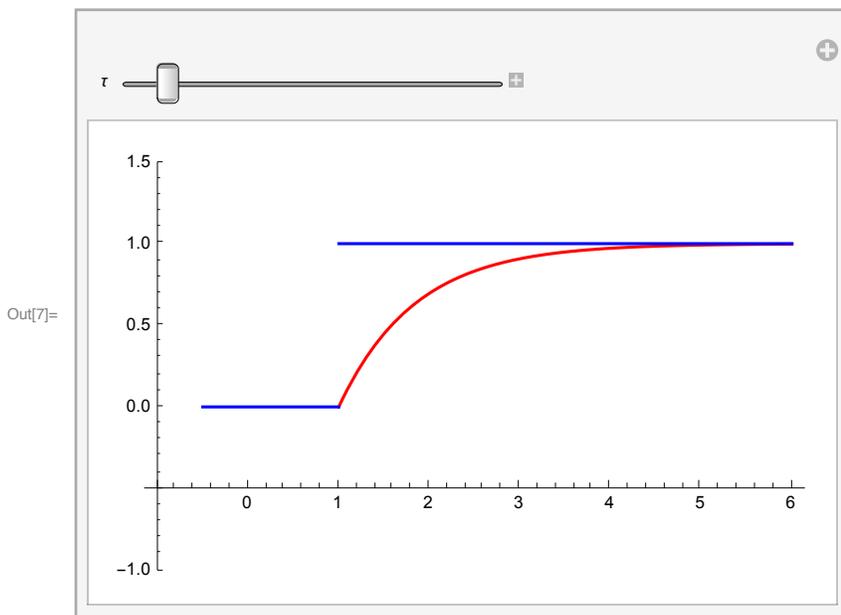
$$V(t) = 1 - e^{-t}$$

The answer is that the output voltage doesn't immediately change to 1 volt, but grows gradually. The time constant is proportional to the product of membrane resistance and capacitance. A typical time constant τ is on the order of 1 or 2 msec.

```
In[4]:= Clear[V, τ, g1];
V[t_, τ_] = 1 - e-t/τ

g1[τ_] := Plot[V[t - 1, τ], {t, 1, 6}, PlotRange → {-1, 1.5},
  AxesOrigin → {-1, -.5}, PlotStyle → RGBColor[1, 0, 0]];
g2 = Plot[UnitStep[t - 1], {t, -.5, 6}, PlotRange → {-1, 1.5},
  AxesOrigin → {-1, -.5}, PlotStyle → {RGBColor[1, 0, 0], RGBColor[0, 0, 1]};
Manipulate[Show[{g1[τ], g2}], {{τ, 1.5}, .1, 10}]
```

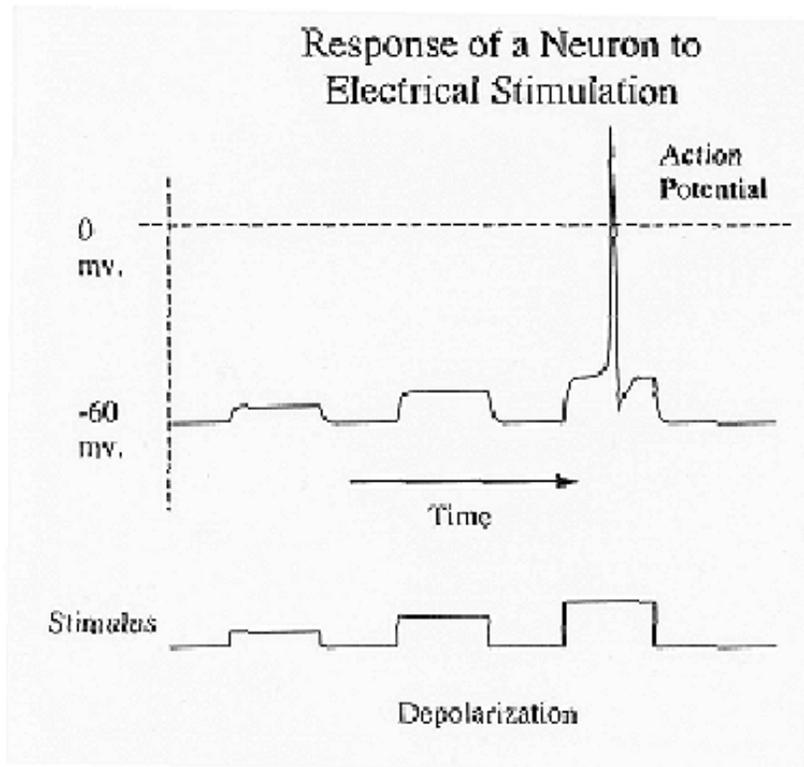
Out[5]= $1 - e^{-\frac{t}{\tau}}$



(UnitStep is a built-in *Mathematica* function. You could define your own as: **myUnitStep[x_] := If[x < 0, 0, 1];**)

In summary, 1) with passive properties there is a rapid (exponential) drop-off in signal strength with distance; 2) the response to an abrupt input takes time to develop.

Active (non-linear) properties



Action potentials (spike trains) are nature's solution to the problem of long distance signalling. Action potential generation is analogous to a "toilet flush" - small push of the handle and a little leak, but reach a threshold, and the whole tank empties. To quantitatively model these voltage changes, one needs to add extra terms to the RC circuit (panels C and D above). We'll see a *Mathematica* implementation later, and also see Meunier and Segev (2002).

Effect of active propagation on *space* properties

What is the neuron's solution to the rapid decline in voltage signal over distance due to passive properties? Action potential at one location provides the depolarization stimulus at a nearby spatial location, travels like a lit fuse down the membrane. Unlike the passive voltage change, the size of the action potential voltage stays the same.

Effect of active propagation on *time* properties

And how about the speed problem?

myelin sheath (insulation increases resistance and reduces current leakage)

Nodes of Ranvier, action potentials jump from node to node, up to **30x faster**.

Refractory period: absolute, and relative

The ion pumps need time to restore the ionic imbalance. *Absolute refractory* period is a brief time (~1 msec) right after the depolarization where no strength of input current would be sufficient for another spike.

Relative refractory period- Immediately after an action potential, one can elicit a spike, but it requires a stronger input signal. This is one of the factors that leads to the idea of *frequency coding*. (Later in the

course, we'll spend time on the question of the "neural code"--i.e. what information is represented by the pattern of action potentials in a neural network).

A constant step input leads to a series of spikes of a particular frequency (e.g. if the absolute refractory period is 1 msec, we'd expect a maximum spike frequency of 1000 spikes/second), but usually much lower (e.g. 20 spikes/second). To get a faster rate, the input voltage would have to be increased. This leads to the idea of the neuron as a "voltage-to-frequency" converter. But it gets a little more complicated because neurons often show "adaptation" and the firing rate declines for a fixed step input. More on that later when we develop the leaky integrate-and-fire model of the neuron.

You can verify the refractory properties yourself with the *Mathematica* model of the solutions to the Hodgkin-Huxley equations in the course notes. Also see: <http://demonstrations.wolfram.com/NeuralImpulsesTheActionPotentialInAction/>

Synaptic Integration--Combining information from incoming signals to the neuron

Terminal buds do not make "direct" contact with dendrites. There is a gap between the two neurons, with information resulting in a voltage change being passed either electrically, or via chemicals called neurotransmitters.

Some definitions:

pre-synaptic potential --voltage across the membrane of the terminal bud of the "transmitting" neuron

post-synaptic potential -- voltage across the membrane of the dendrite of the "receiving" neuron

Two types of connections: **excitatory** (make the cell more likely to fire) and **inhibitory** (less likely to fire) post-synaptic potentials.

Post-synaptic potentials have a relatively long duration - even a fast excitatory post-synaptic potential or **EPSP** may have a 1 to 2 msec rise time and 3 to 5 msec decay (recall: action potential 1-2 msec). (IPSP is the corresponding short-hand for inhibitory post synaptic potential). The synaptic delay, information transfer across the synapse, itself is rather short, about 500 micro-seconds.

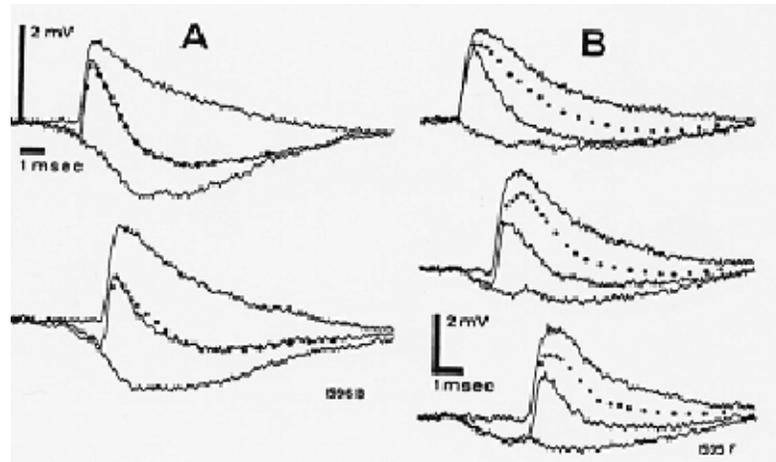
Interaction of synaptic potentials

Imagine 1000's of synaptic inputs to a single neuron. How do they interact? This question is fundamental to understanding neural network modeling.

When the graded potentials arrive at the axon hillock, do they combine algebraically? I.e. is the cumulative effect at the hillock an algebraic sum? If so, this would greatly simplify the modeling, because of the rich set of tools and understanding we have of "linear systems".

Our "generic neuron model" introduced later will assume linear algebraic summation, and only introduce non-linearities as needed.

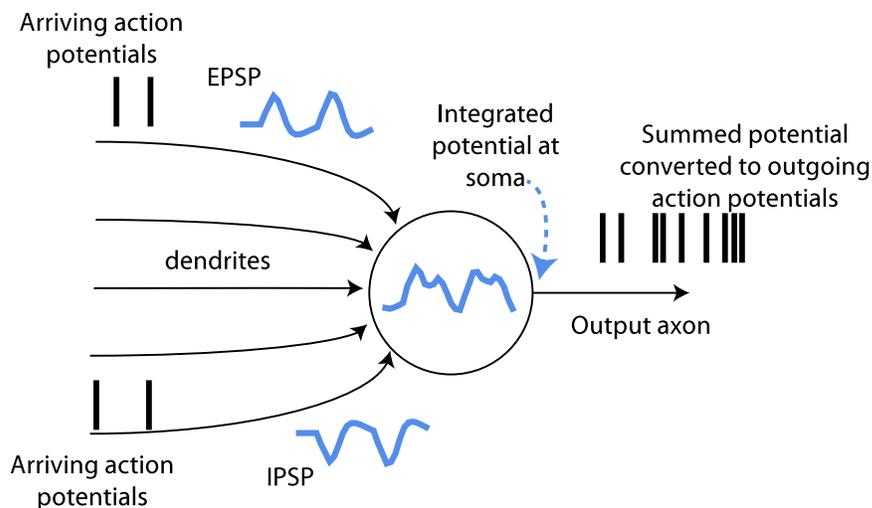
However, it is important to keep in mind that linear summation properties are sometimes but not always seen (see Koch and Segev, 2000). Linearity is usually for the convenience of the modeler. The figure below shows examples of some empirical results. The top and bottom traces are inputs, the middle trace is the result of combining these two inputs. The dotted trace is the algebraic sum prediction. Panel A shows simple algebraic summation, but Panel B shows measurements in which the integrated signal is not a simple sum.



Interaction of **inhibitory postsynaptic potentials (IPSP)** and **excitatory postsynaptic potentials (EPSP)** can be linear (left column, A) or non-linear (right column, B). From Rall (1967). Dotted line shows linear prediction.

Qualitative summary of slow potential neuron model

Let's summarize the essential qualitative features of signal integration and transmission of a neuron with what is sometimes called the "slow potential model" (Anderson, 1995).



Slow potential at axon hillock waxes and wanes (because of low-pass temporal characteristics and the spatial distribution of the inputs) depending on the number of active inputs, whether they are excitatory or inhibitory, and their arrival times. (Temporal "low-pass" filtering simply means that fast changes don't get through because their amplitudes are so strongly damped.)

The slow integrated voltage potential now and then exceeds threshold producing an axon potential. Further, if the slow potential goes above threshold, frequency of firing is related to size of slow potential.

We'll develop a quantitative description of this, called the "integrate and fire" model, in the next Lecture.

Caveat: Not all signal transmission in neural computation is done through action potentials. For example, of the more than 6 classic types (there are more!) of cells in the retina of your eye, it is primarily the ganglion cells which use action potentials, the others communicate via slow potentials. Ganglion cells send their messages over a range of centimeters. The other cells on the order of millimeters.

Spike generation isn't a strictly deterministic process. There is "**noise**" or **random fluctuation**. This means one isn't guaranteed of getting exactly the same spike discharge for identical inputs. This is because:

- ion channels open and close probabilistically
- neurotransmitter is released in discrete packages, adding uncertainty in the signal to the next neural in the signalling pathway
- sensory receptors can produce spontaneous signals (not to mention that the physical causes of sensory stimulation have noise too).

One way to model the noisiness of a neuron's discharge is to treat spike train frequency as roughly like a Poisson process (better--an interval Gamma distribution) whose mean is modulated by the already noisy slow potential. (See PoissonDistribution[.])

We haven't covered what happens at the synapse connections yet. But for a lively overview of the basic properties of a neuron, including synaptic transmission, see: <http://www.youtube.com/watch?v=eZundDVPIYw>

Next time

Brief history of neural models

Various classes of models

Develop the "leaky integrate and fire" model

Appendix

Solve the space-dependent portion of the cable equation subject to initial conditions: $V[x=0]=V_0, V'[x=0]=-V_0/\lambda$:

`Remove[V, V0, λ];`

`DSolveValue[{V[x] == (λ^2) * V''[x], V[0] == V0, V'[0] == -V0/λ}, V[x], x]`

$e^{-\frac{x}{\lambda}} V_0$

Solve the time-dependent portion of the cable equation subject to initial conditions: $V[t]=0, V''[x]=1/\lambda^2$:

`Remove[V, V0, λ];`

`DSolveValue[{V[t] == -τ V'[t] + 1, V[0] == 0}, V[t], t]`

$e^{-t} (-1 + e^t)$

References

Ermentrout, G. Bard <http://www.cnbc.cmu.edu/~bard/passive2/passive2.html>

Hodgkin, A. L., & Huxley, A. F. (1952). A quantitative description of membrane current and its application to conduction and excitation in nerve. *J. Physiol. (Lond.)*, 117, 500-544.

Koch, C. (1998). Biophysics of Computation : Information Processing in Single Neurons . Oxford: Oxford Univ. Press.

Koch, C., & Segev, I. (1998). Methods in Neuronal Modeling : From Ions to Networks. Cambridge, MA: MIT Press, 671 pages.

Koch, C., & Segev, I. (2000). The role of single neurons in information processing. *Nat Neurosci*, 3 *Suppl*, 1171-1177.

Koch, C., & Poggio, T. (1985). A simple algorithm for solving the cable equation in dendritic trees of arbitrary geometry. *Journal of Neuroscience Methods*, 12(4), 303–315.

Meunier, C., & Segev, I. (2002). Playing the devil's advocate: is the Hodgkin-Huxley model useful? *Trends Neurosci*, 25(11), 558-563.

Segev, I. (1992). Single neurone models: oversimple, complex and reduced. Trends in Neuroscience, 15(11), 414-421.