Modeling Neurons: A Biological and Mathematical Introduction

Maire Keene and Michael Sheets

May 3, 2016

1 Basic Neuron Biology

As highly evolved organisms, we get to reap the benefits of delicate and networked sensory organs to perceive the world, an intricate and powerful musculoskeletal system that enables us to navigate and manipulate it, and (perhaps most significantly!) an incredibly complex processing center and communications network at the heart of it all to comprehend these inputs and direct our actions. In other words: Neurons are cool! They've been the subject of study and modeling since the early twentieth century both for basic understanding of how this core feature of our personal physiology operates and for the enhancement of a variety of medical interventions such understanding has been able, and will continue, to yield.

Neurons (seen in Fig. 1) transfer information though pulses that travel down the length of the axon to other neurons. These pulses are triggered by chemical or ion signals received through the dendrites from other neurons or external sensations , which travel to the base of the axon, called the hillock. If a large enough signal (which can come from a couple especially large signals, or from a lot of smaller signals) reaches the hillock, it triggers an action potential which jolts down the neuron. When the action potential reaches the end of the neuron, called the neuron terminal, it gets passed on to other neurons and the cycle begins again.

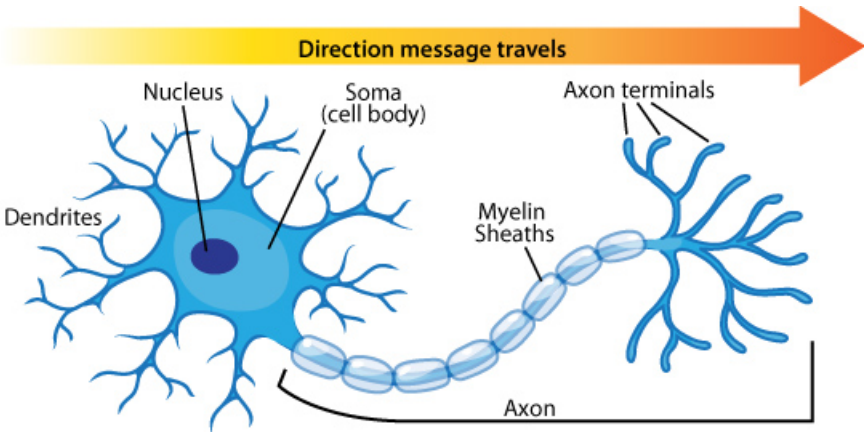

Figure 1: Image of a neuron. Signals are received from dendrites, and if strong enough are carried down the axon in an action potential. (4)

This process is carried out in neurons through diffusing waves of sodium, potassium, and calcium ions, which create a voltage difference that travels down the length of the axon. (We'll mostly care about the first two ions for our purposes). To do this, the neuron has a few key pumps and channels

that prepare sodium and potassium ions for this process. First, we have the sodium-potassium pump, which continually (but slowly) operates to keep a lower concentration of sodium and higher concentration of potassium in the neuron than outside in the surrounding environment. This, along with other ion "leak" channels and organic acids also causes the inside of the neuron to be more electrically negatively charged than the outside, with a resting potential of approximately -60 mV, which is important for the voltage spike that begins the action potential process. Second, we have a voltage-gated sodium channel, which opens at about -50 mV, and closes at about +40 mV. This means that sodium ions are free to flow in (or out, technically) from when the voltage hits -50 until it reaches 40 mV. Third, potassium has a similar voltage gated channel, but in reverse: its channel opens at 40 mV, and closes at about -70 mV. Putting these all together, and adding a large enough trigger signal, we get a cascade of ion flows that create a full action potential spike as seen in Fig. 2.

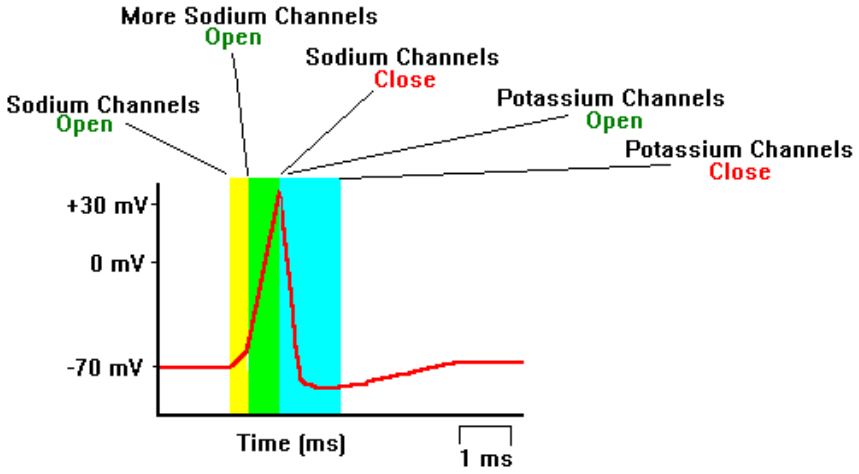

Figure 2: An action potential spike and the ion flows that cause it. This would be the thing we're trying to get our equations to mimic. (5)

Starting at our negative potential once we get a large enough trigger, we see a large spike upwards, where the sodium channel is opened and sodium ions freely flow in to make up for both the voltage imbalance and the sodium ion imbalance, with both electrical forces and diffusion acting in their favor. However, once this flow makes the voltage difference large enough in the positive direction, the sodium channels snap shut and the potassium channels open. Again, charge and diffusion are acting in favor of the potassium ions, causing a sudden drop to a negative potential slightly below the resting potential to ensure that we have an appropriate refractory period during the reset. (Also note that the sodium channels have a *refractory period* while the concentration is lowered by potassium flowing out, which prevents these channels from being pretty much continually open). Finally, after the potassium channels close, our sodium-potassium pump and leak channels help us to reset the ions to their initial state, and the voltage to our resting potential.[1]

Another aspect of the neuron you may notice are the *myelin sheath* that intermittently covers much of the axon, with the nodes of Ranvier being the gaps between it. The sheath is basically

[1]Reading that description may beg the anticipatory question: How do we even *begin* to move from this interconnected system of ions and pumps and channels to some sort of mathematical statement capturing the action-potential phenomenon? Take comfort in simplifying it and recognizing it like this: there is a stable voltage that is determined by ion concentrations and changing those with some incoming, initial disturbance in those concentrations naturally changes the voltage which controls the mechanisms that control the ion concentrations and... this feedback loop starts to sound like a second order ODE, right? It's position and acceleration all over again, and have faith that we can do this, friends.

an extra-thick section of membrane that doesn't have gates or channels in it. This axon structure leads to a electrical and diffusive flow termed *saltatory conduction*, with relatively slow transmission through the nodes of Ranvier and very fast transmission within myelinated regions. This happens for a few reasons, mostly due to how long it takes the neuron to repolarize (reset the ions) between transmission, and the scale of that change. Nodes of Ranvier also serve to boost the signal as it travels along the axon, with sodium gates being activated and sending more sodium and positive charge down the axon, towards the next node. This type of periodic boosting allows for extremely long neurons (in comparison to other cells), reaching up to 4 feet in the human body in the nerves that go from the base of the spine to the tips of the toes. This also means nerve signaling as seen biologically is very complicated mathematically, and is thus often simplified to just the resultant effect: an ionic and electrical signal as seen in Fig. 2 is sent down the entire length of the neuron with minimal decay. This simplification allows us to do most of the useful work we desire with neuron models, without needing variables and terms that are too intensely complicated. (1)

2 Models

2.1 The Hodgkin-Huxley Model

Now that we've developed an understanding of how neurons and action potentials work from a biological perspective, we can apply some (hopefully) pretty familiar basic electrostatics, and start representing what goes on at a particular point in a given axon with some governing equations. Doing this will lead us to one of the first models developed of a neuron's action potential, called the Hodgkin-Huxley model (HH).

If you were to take a lipid bilayer and use it as a component in an electrical circuit (why would we do that? Because it kind of *is*, all the time in our bodies) it would have a certain amount of capacitance, which we'll call C_m. From the definition of capacitance, $Q = CV$ and the definition of a current, $I = \frac{dQ}{dt}$, we can express the current that is traveling through the membrane at a point as a function of the voltage difference (yay, ions!) across the membrane at that point:

$$I_m = C_m \frac{dV_m}{dt} \tag{1}$$

Similarly, for each ion species/ion channel pair involved in the action potential propagation process as well as the leak channels, the current moving through a given ion channel is

$$I_i = g_i(V_m - V_i) \tag{2}$$

where V_i is reversal potential for that ion pump, and g_i is the ion conductance per area.

So the total current through a membrane (per area) considering the sodium and potassium ions is then the sum of these component currents:

$$I = C_m \frac{\mathrm{d}V_m}{\mathrm{d}t} + g_K(V_m - V_K) + g_{Na}(V_m - V_{Na}) + g_l(V_m - V_l) \tag{3}$$

But here is where we note that the capacitance terms for each of each ions above are dependent on V_m in the real biological system, and we need a more complicated system of ODEs to account for how they all change in relation to each other and how it affects the opening and closing of each of these channels. To include all that, we need the HH in it's full, system-of-four-linked-ODEs-glory:

$$\begin{cases} I = C_m \frac{\mathrm{d}V_m}{\mathrm{d}t} + \bar{g}_{\mathrm{K}} n^4 (V_m - V_K) + \bar{g}_{\mathrm{Na}} m^3 h (V_m - V_{Na}) + \bar{g}_l (V_m - V_l), \\ \frac{dn}{dt} = \alpha_n(V_m)(1-n) - \beta_n(V_m)n \\ \frac{dm}{dt} = \alpha_m(V_m)(1-m) - \beta_m(V_m)m \\ \frac{dh}{dt} = \alpha_h(V_m)(1-h) - \beta_h(V_m)h \end{cases} \tag{4}$$

where α_i and β_i are rate constants for the i-th ion channel (dependent on voltage but not time), $\bar{g}_n$ is the maximum possible conductance, and n, m, & h are constants between 0 and 1 used to regulate sodium and potassium channel activation and sodium channel activation, respectively.

For each of our α and β terms above, we get terms as follows:

$$\alpha_n(V_m) = \frac{n_\infty(V_m)}{\tau_n} \tag{5}$$

$$\beta_n(V_m) = \frac{(1 - n_\infty(V_m))}{\tau_n}, \tag{6}$$

with equivalent expressions for α_m & β_m and α_h & β_h.

This model *is* quite dense to look at, but there are a few especially important things to note. First off, ions affect the voltage difference (and thus current) in the cell, which then in turn affects the voltage. Biologically, this reflects how the ions carry charge with them and are used to change current, and that the voltage differences cause ion channels to open or close, creating a feedback loop that drives the system when active.

Another thing to note would be that the system is indeed very complex and has a lot of terms– many more than can be easily visualized or processed intuitively, giving us many dimensions, or axes, that must be considered at all times and each affect the others. For this reason, simplifications have been made to this model, one of the most famous and most useful being the FitzHugh-Nagumo model. (3)(10)(11)

2.2 The FitzHugh-Nagumo Model

Essentially, we can think the of the FitzHugh-Nagumo model (FHN) as taking the important excitation and recovery behavior of the Hodgkin-Huxley model and simplifying the cellular machinery down to two variables, resulting in a system of nonlinear ODEs that is easier to analyze and visualize:

$$\begin{cases} \dot{V} &= V - \frac{V^3}{3} - W + I \\ \dot{W} &= 0.08(V + 0.7 - 0.8W) \end{cases} \tag{7}$$

Where V is very similar to the voltage difference across the membrane that we saw before with the Hodgkin-Huxley model, and W is the 'recovery variable,' where the different ion species and their movements have been collapsed into one state variable characterizing relevant cellular machinery. I is once again the magnitude a current, in this case the current caused by the stimulating change in local ion concentrations.

Some important things to notice about these equations are that the voltage-like variable has a negative cubic dependence on itself as well as a positive linear term that competes with it, as opposed to the negative, linear relationship (and tiny coefficients in this tuned model!) the recovery variable has with itself. This means that close to zero, a small change in V will have the linear term dominating, and will provide some impetus to keep growing in the direction it was nudged (positive feedback). As it gets far enough away from 0 for the cubic term to start to overpower the linear term, positive values of V will cause some negative feedback and encourage V to move back toward 0, whereas negative values of V will cause V to move further away from zero... dependent on W. This back and forth feedback will be important for oscillation purposes in a minute.

But where did this two-equation simplification come from?
Did the original authors just pull this form out of thin air and it somehow worked out that it had the rapid excitation and slower recovery that they were looking for? With Hodgkin and Huxley's model, it kind of makes sense and ties directly back to the biology, and we can use

electrostatics and the actual physical voltage limits of the ion gates to construct our equations, but what is this?
So the answer is kind of that FitzHugh got really annoyed trying to implement the HH model in his lab's giant, analog, computer one day and at someone else's suggestion, just modified the equations for another, existing simple nonlinear relaxation oscillation model developed in the 20s– the van der Pol Oscillator. Those equations were pretty similar to those governing a circuit Nagumo was working on with a cubic current-voltage curve. This is why you may see equivalent circuit diagrams for the model, or see the FHN referred to as the "Bonhoeffer-van der Pol model."
The less-history-infused version is basically when the frustrated scientists decided to cast about for a simpler model with similar behaviors, they decided the important bits to keep were that all the cellular machinery needed to be represented by this one variable, W, which needed to change relatively slowly, and the voltage, which changes relatively suddenly depending on what the cell machinery is doing at a given moment, got its own variable V. They needed a model with a fast component, a slow component, and sudden spikes with a somewhat slower recovery. Looking for something that already did that, they found a type of electrical circuit with a particular arrangement of capacitors, inductors, and diodes that was being actively studied somewhere else and figured out the equivalent damping for their system. Mathematical abstraction of complicated systems and transferring those abstractions for the win!

We can also note that because of the cubic term, when you get far enough away from 0, a small change in V is going to produce a much bigger (read, 'faster') change in itself than a similar small change in W (read, it's the 'slower' variable, especially with the specific coefficients we have on it right now). The recovery variable also always provides strictly negative feedback for itself, meaning it's a kind of slow relaxation back to equilibrium...dependent on V.

So, these equations indeed respond to themselves in interesting ways, and are of course linked, so combined with the non-linearity and difference in variable 'speed' or the 'fast and slow timescales' they operate on serve to balance each other and create an oscillating effect.

2.2.1 'Fast' and 'Slow' Time

Although the FHN is vastly simplified compared to the Hodgkin-Huxley, some further mathematical analysis is needed before we can do much of the useful work the model is intended for. For instance, it would be very difficult to conceptualize how the fast variable and slow variable moves cause changes in relation to each other without knowing about each separately. A key breakdown that can be done is looking at the model in "fast time" for the fast variable and "slow time" for the slow variable by introducing a new time variable:

$$\tau = \epsilon t, \tag{8}$$

where t is now our fast time variable, and τ is now our slow time variable for some small ϵ.

Applying ϵ to each of our equations, calling a few of our coefficients 1 for now because they're a little irrelevant to understanding how the model behaves, and incorporating in τ, we get the following two sets of equations. One set for fast time:

$$\begin{cases} \frac{dV}{dt} = \dot{V} & = V - V^3 - W + I \\ \frac{dW}{dt} = \dot{W} & = \epsilon(V - W), \end{cases} \tag{9}$$

and one set for slow time:

$$\begin{cases} \frac{dV}{d\tau} = \frac{dV}{dt}\frac{dt}{d\tau} = \frac{1}{\epsilon}\dot{V} & = \frac{1}{\epsilon}(V - V^3 - W + I) \\ \frac{dW}{d\tau} & = \frac{1}{\epsilon}\dot{W} = V - W. \end{cases} \tag{10}$$

As we make ϵ smaller and smaller, we eventually set it to zero, giving us two singular systems (thus making them one-dimensional ODEs). One where fast time is instantaneous, freezing the slow part and allowing us to only track the fast part:

$$\begin{cases} \dot{V} & = V - V^3 - W + I \\ \dot{W} & = 0, \end{cases} \tag{11}$$

and one where the fast component just a blur, and we are able to follow the slow component over time:

$$\begin{cases} 0 & = V - V^3 - W + I \\ \dot{W} & = (V - W). \end{cases} \tag{12}$$

This allows us to create two pictures as in Fig. 3 to describe the system, one that examines how the fast component of the system moves with a stationary slow component, and one that tracks the slow component while letting the fast component move quickly on its own. You might hear these referred to as the fast and slow manifolds of a system.

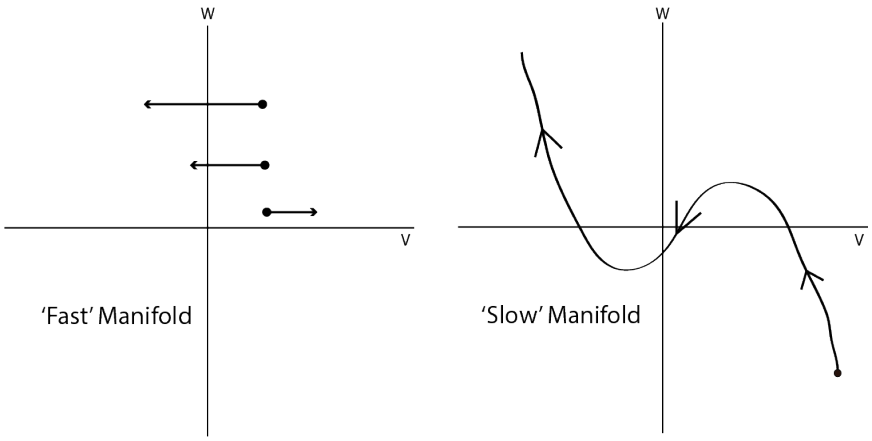

Figure 3: Pictures exemplifying how the fast variable moves with various set values of W, and how the slow variable moves. Also sometimes referred to as the fast and slow manifolds of a system.

After we have observed the behaviors on either end of the spectrum, we can use both of these to construct a more complete image of how the system acts as a whole, when neither variable is actually frozen but just operating in its own timescale, resulting in a system that follows the fast time behavior at certain points, and the slow time behavior at others. For the FHN and the specific tuned parameters we wrote down above, this results in a 'phase portrait' (don't worry about that terminology, we'll explain in a second) as in Fig. 4

The many black lines are time-traces of this system's behavior with different sized initial stimulating currents; the red notes are physiologically relevant notes; and the blue lines are the particular 'nullclines', the lines along which $\dot{W}$ and $\dot{V}$ are 0. Notice how there *might* just be a smidgen of relationship between the manifolds and nullclines. We'll soon see more images like this one.

This concept of mapping out how the variables change in relation to one another is an incredibly useful tool for mathematical analysis of these complicated systems of non-linear equations, and if you notice the images above are plotted as V vs W, or in this system's "Phase Space."

2.3 Phase Space

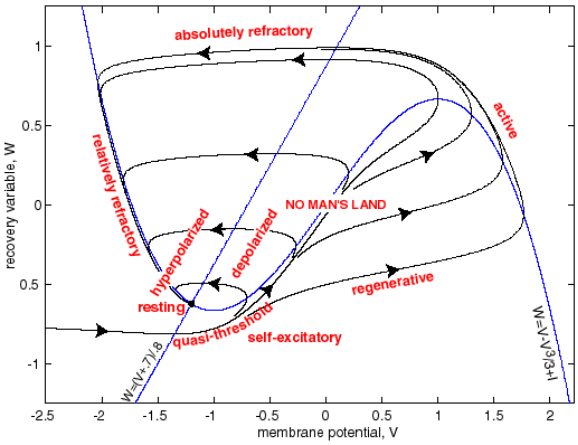

Figure 4: Phase portrait of the FHN with several different time traces of different magnitudes of initial conditions traced out, along with physiological notes of what those magnitude responses correspond with in a firing neuron.(9)

The Pendulum: A simpler, familiar example of phase space.
Phase what-now? Usually, we're used to seeing variables such as position, velocity, or voltage (hint hint) plotted against time. However, sometimes it can reveal more about a given system to plot these variables against each other. For instance, in the case of the pendulum (Fig. 5), we can see that the system behaves in a cyclic manner when not interfered with, and that position and velocity vary regularly in accordance not only with time but with each other. This type of phase portrait can give us key insights into the system, such as if you started at a higher or lower position (bigger and smaller circles, respectively), that might not be immediately apparent from the time series graph. This concept will be used to further investigate the FHN model of neuronal action potentials.

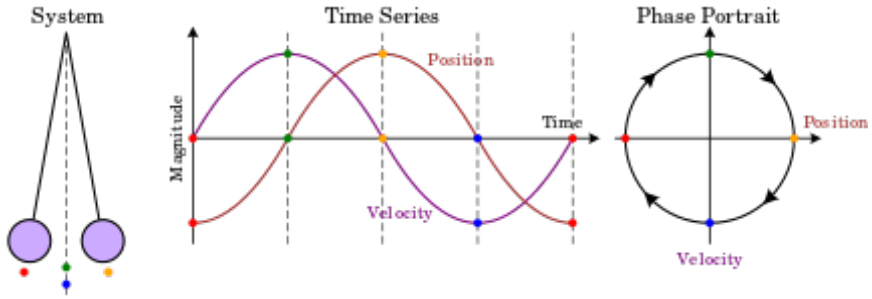

Figure 5: System, time series, and phase portrait images of a swinging pendulum. As the pendulum swings, velocity and position vary cyclically with respect to each other, which can be concisely visualized in a phase portrait of position versus velocity.(2)

As we've said, one of the key simplifications of the FitzHugh-Nagumo model is that it reduces the system to two dimensions, which makes it much easier to visualize, and we can use many of these visualizations to draw conclusions about the behavior of the system. Graphing our voltage and recovery against time is one way that we can gain information about our system (Fig. 6 right). Another particularly useful way to look at our system is in phase space, with voltage plotted against

the recovery variable (Fig. 6 left). This also can allow us to build up an intuition on the visual 2D system to better understand how higher-level models operate.

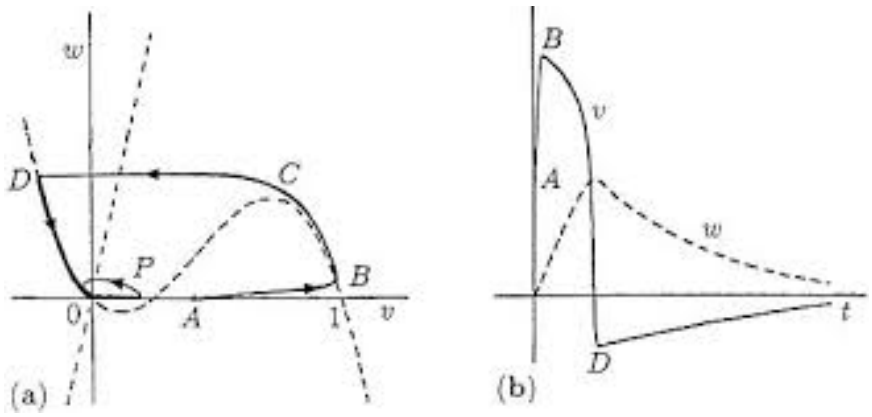

Figure 6: Phase space, left, and state variable over time, right, representations of the FHN model. We can gain important and different insights into the behavior of our system of equations from both types of representation. (3)

Phase space allows a representation of all possible states of a system's voltage in relation to recovery and vice versa, allowing us to translate between time trace and voltage/recovery possibilities (see the pendulum example in the insert for a simpler case). Examining these possibilities and the associated directionalities allows us to follow both variables from one set of coordinates to the next throughout neuronal spikes.

In phase space diagrams, it is often helpful to plot the nullclines, shown as dashed lines in the phase space plot in Fig. 6. Often, our phase space path will follow and jump between nullclines. As our nullclines line up with the fast and slow manifolds seen in the system (see the "Fast and Slow Time" section), our variables will follow these when they can, as they are the locally "flat" places, or indents in our system that are easier to follow than to break away from. Our equilibrium points also happen to be at the intersection(s) of the nullclines, the point where neither variable is changing. Disruptions to the system (such as an action potential spike) push our variables out of this equilbrium, causing them to follow the spike path set up by the nullclines.

2.4 Limits of FHN

Although the FitzHugh-Nagumo model vastly simplifies the mathematics necessary to visualize and analyze a neural action potential as compared to Hodgkin-Huxley, it does move the mathematics slightly further from the biology. However, these losses are in most cases well worth the usability the FHN model provides.

One example of this is the "slow time" arc seen at the top of the FHN action potential spike, which is not present in the HH model or the biological system, as seen in Figure 7.

As the FHN model collapses multiple ion potentials down into one recovery variable, it is no longer mathematically capable of the fast drop seen immediately after the fast rise. This has a few implications for the usefulness of this model, however, the minimization of the slow part at the top means that for most cases, FHN would be an appropriate choice. When studying the effect of ion channel flow rates or activation voltages, this may become relevant, and the HH model may be the more appropriate choice, or in situations where the slow portion may affect impulse stability.

Another limitation of the FHN model is lack of periodic boosting between nodes of Ranvier, which may be relevant when attempting to look at the effects of myelination, axon length and diameter,

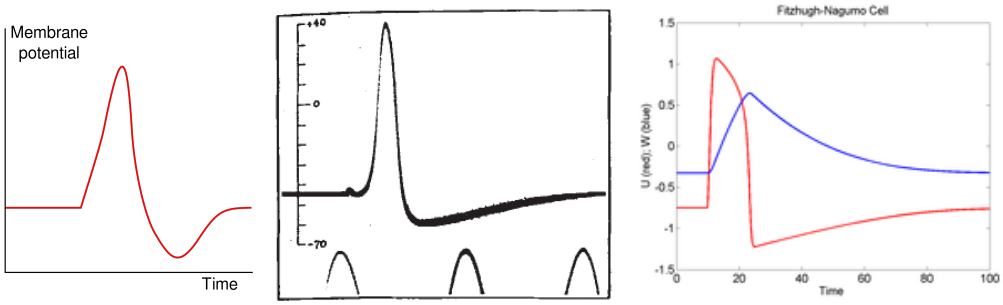

Figure 7: Biological (6), HH (7), FHN (8) neural spike images. Note that HH closely mimics the biological system, while the FHN introduces a slight change with the "slow time" necessitated at the top of the spike.

and transmission throughout an entire neuron. This can be addressed by incorporating a spatially variant capacitance–however, to even begin looking at this, we must extend the model beyond a point and into multiple dimensions, which requires moving from ordinary differential equations to (dun dun dun) *partial* differential equations.

2.5 Extending FHN Into Multiple Dimensions: Time *and* Space!

So far we've discussed how to model the action potential spike and recovery at a single point in the neuron, which is pretty interesting in it's own right.[2] But with neurons we are also somewhat inherently interested how these signals *travel* down the axon, since that's kind of the whole point of these signals – they travel down the axon rather rapidly and deliver the signal to the next neuron in the chain to actually achieve things in the body. Mathematically, this rather sounds like something we've seen before: a traveling wave. Perhaps, then, we can incorporate traveling wave behavior into our existing oscillator by transforming our voltage-variable equation from an ODE into a PDE with the addition of the double spatial derivative:

$$\dot{V} = V_{xx} + \underbrace{V - \frac{V^3}{3} - W + I}_{\text{familiar nonlinear bits, f(V)}} \tag{13}$$

This move might look eerily familiar to those who recall how equations for modeling the diffusion of heat or particles... which perhaps makes sense from a biological perspective since it is in make diffusing/electrically shoved ions that are responsible for the signal moving in the real world. But back to math land.

Ok, so we've modified our equation and are hoping that there is actually a traveling wave solution to our new PDE, and that it's physically reasonable. In other words, we want

$$V(t, x) = \Phi(x - ct). \tag{14}$$

Mathematically, the good news for saying something intelligible and useful about this solution we want is that if we can solve the equation

$$-c\Phi' = \Phi'' + f(\Phi)$$

then there is indeed a traveling wave solution to the equation and if it's bounded then that solution is reasonably non-pathological and can make sense in a physical system.

[2]You could also view this as what was happening along a length of axon when time was frozen, but we couldn't play a movie showing us the progression of the action potential at every point as it progressed. We like movies.

Unfortunately, given what we mentioned above with the whole periodic-boosting, non-uniform capacitance of the axon in the spatial dimension thing, our coefficients become spatially dependent and actually getting that exact solution becomes incredibly hairy. Like it does with most PDEs. Which is why it's fortunate we have tools to estimate the solutions numerically, such as MATLAB and COMSOL. Our new and annoying term can be approximated in a couple different ways that are worth mentioning quickly, like with discretization or convolution, but you probably don't have to worry too much about it:

$$\begin{aligned}(K(x)V_x)_x &= \int V(y)K(x-y)dy \\ &= V_{n+1} + V_{n-1} - 2V_n\end{aligned}$$

This PDE version of the FHN can be modified and used for a number of different "reaction-diffusion" systems in the world such as blood clotting and crystal formation, and you may very well run into it again and again in the future.

2.6 Are the Travelling Waves Stable?

Another use of neuronal modeling that can have real-life implications is the stability of neuronal travelling waves. This can allow us to look at when things start to break down in neuron function, as well as the bridge between our model and an actual neuron. For instance, if a type of ion channel is malfunctioning at the hillock or a node of Ranvier, will the signal still be sent? Or on the model side, how finely tuned do our parameters have to be to allow normal action potential signalling?

Finding the stability of a travelling wave necessitates we first have found a travelling wave solution, which we'll continue to call ϕ. Having found ϕ, we then ask: is it stable? To do this ,we define our function for V_t as an arbitrary function G with linear and nonlinear components, and create a small function of space, ψ, to act as our nudging variable, or as a way to "poke" V and see if it causes any dramatic changes in the behavior of V.

To determine when a function is stable, we solve:

$$V_t = G(V), V(0,x) = \phi(x) + \psi(x), \tag{15}$$

and ask whether $V(t,x)$ is close to $\phi(x-ct)$.

Another helpful way to define this is to introduce the function r (think r for residual), where:

$$r(t,x) = V(t,x-ct) - \phi(x) \tag{16}$$

and ask if is $r(t,x)$ is small. By the way, we're using $V(t,x-ct)$ and $\Phi(x)$ because we're interested in the residual in the moving reference frame of the traveling wave solution, or in other words how well the shape of the wave going down the axon is matching the shape of the wave in the absence of the "poke."

Introducing the r term allows us to require that r satisfies:

$$r_t = \underbrace{\frac{\partial G}{\partial V}(\phi)}_{\text{linear operator} L} r + \underbrace{G_{nonlinear}}_{\text{remaining terms of G}}. \tag{17}$$

Defining r in this way allows us to create a linear operator, L, for the linear part of r_t. Because we are interested in whether the residual goes to zero in response to little "pokes," i.e. in the area of 0, the nonlinear terms will be small and dominated by the term with the linear operator. This is very neat, because it means if the linear portion of r_t has a stable equilibrium at 0, *the whole thing is stable.* Determining the stability of a linear operator about zero is actually something that might sound familiar as well– we just need to find the eigenvalues of L, and they will tell us. Getting the

r PDE then in turn tells us about the stability about ϕ for the V PDE, which is the answer to our original question of if-we-poke-this-wave-a-bit-does-it-collapse-or-get-better. If (the real components of) the eigenvalues of $L < 0$, then we know that the the change due to our poke is decaying over time, and thus the function is stable. If the eigenvalues of $L \geq 0$, then our function is not stable and will have expanding (or at 0, other kinds of unstable) deviations from our original function over time.

3 Some Applications

As we continue to develop our understanding and models of how neurons work, there are a few ends that many researchers are working towards:

- How networks of neurons operate. Some symmetries between how individual neurons operate and how their networks function, and this often becomes interesting mathematically and when developing models for neuronal networks. This research can often be especially relevant to epilepsy.
- Neural prostheses. Much current and former neural modelling can be applied to prostheses such as retinal devices, cochlear implants, and neurally-controlled artificial limbs. How neurons send signals to the brain and each other is mimicked in these devices, and developments in our understanding of how neurons work has been vital in their development.

As our understanding of the basic biology increases, so does our ability to model and artificially manipulate these systems, helping us understand the systems at an even deeper level and use these understandings to improve health.

Acknowledgements

Many thanks to the illustrious Aaron Hoffman for much assistance and many office hours spent helping us understand the information presented above.

References

[1] Khan Academy. "Advanced nervous system physiology." Khan Academy. https://www.khanacademy.org/science/health-and-medicine/nervous-system-and-sensory-infor Citation for an overview of video series contained, most notably Nervous System Introduction and Neuron Membrane Potentials series.

[2] "Phase Space." Wikipedia. https://en.wikipedia.org/wiki/Phase_space

[3] Murray, James. Mathematical Biology I. an Introduction. N.p.: Springer Verlag, 2014. Print.

[4] Arizona State University. "A Nervous Journey." Ask A Biologist. https://askabiologist.asu.edu/neuron-anatomy

[5] Chudler, E. H. "Lights, Camera, Action Potential." Neuroscience of Kids. https://faculty.washington.edu/chudler/ap.html

[6] "Action Potential." Wikipedia. https://en.wikipedia.org/wiki/Action_potential

[7] Simpson, W. A. "Hodgkin & Huxley in Plymouth." University of Plymouth. http://www.psy.plymouth.ac.uk/research/Wsimpson/misc/hh.html

[8] "Temporal interaction of excitation and inhibition (Fitzhugh-Nagumo Model)." Sounder Sleep System. http://soundersleep.com/fitzhughNagumo.php

[9] "FitzHugh-Nagumo Model." Scholarpedia. http://www.scholarpedia.org/article/FitzHugh-Nagumo_model

[10] "Hodgkin-Huxley Model." Wikipedia. https://en.wikipedia.org/wiki/Hodgkin%E2%80%93Huxley_model

[11] "Hodgkin-Huxley Model and FitzHugh-Nagumo Model." http://phaser.com/modules/maa2005/h-h_model.pdf