MATH 228 lecture notes for November 29, and December 1, 3, and 5

Russell Milne

November 2021

1 November 29: Hodgkin-Huxley and FitzHugh-Nagumo models; constructing a model from first principles

Now that you've seen some of the theory behind differential equations, I will spend the last few lectures telling you about some examples of how they are used in real life. This should also give you some insight on how mathematical models are created; one of the purposes of this particular lecture is to show how one can build a model from first principles. One of the first, and most important, examples of a mathematical model is the Hodgkin-Huxley model of electrical activity in a neuron. This was created in 1952 by Alan Hodgkin and Andrew Huxley, and was fitted to real experimental data that they obtained by measuring the voltage along the giant axon of a squid. (The axon is giant, not the squid. Action potentials in the neuron that this particular axon is part of cause the squid to contract some of its muscles to expel water from its body, which propels the squid forward very rapidly and can allow it to escape potentially harmful situations.)

Since this is an electrical model, most of the theory behind it is based on electrical circuits. In particular, most of the terms in the Hodgkin-Huxley model consist of voltages and currents. For a brief explanation of how these are related, consider a capacitor in a circuit, which stores charge when a voltage is applied to it. Specifically, if a constant C represents the capacitance (the relative ability of the capacitor to store charge), then the charge Q stored by the capacitor can be expressed as a function of the applied voltage V:

$$Q(t) = C \cdot V(t) \tag{1}$$

The derivative of Q is the current intensity (or just "current") across the capacitor. To get this, we can take the derivative of both sides, which gives an expression in terms of $\frac{dV}{dt}$:

$$I(t) = C \cdot \frac{dV}{dt} \tag{2}$$

Indeed, within the context of the Hodgkin-Huxley model, the lipid bilayer that makes up most of the cell membrane of the squid giant axon is treated as a capacitor. Thus, if V_m represents the axon's "membrane potential" (the difference in electric potential between the inside and outside of the axon), the current flowing across the lipid bilayer I_c can be written as the following:

$$I_c = C_m \frac{dV_m}{dt} \tag{3}$$

However, there are other ways in which ions can move in and out of the cell, which of course have implications for the electrical current. Two of the most important ions in animal cells are sodium and potassium. This is because animal cells maintain a roughly constant electrical potential by exporting sodium ions and importing potassium ions, causing both sodium and potassium ions to usually have a substantial gradient between the inside of the cell and the outside. For a given ion i, the "reversal potential" V_i is the membrane potential at which there is no net flow of the ion into or out of the cell. This means that ions contribute to the electrical current across the cell membrane depending on how much the actual membrane potential differs from each ion's reversal potential. Additionally, the transport of ions across a cell membrane is carried out by ion channels, which can be thought of electrically as conductors, as they facilitate electrical current flowing into or out of the cell. For a given ion i, the conductance of the ion by its associated ion channel can be expressed as some quantity g_i. We know that conductance is equal to current divided by voltage:

$$G = \frac{I}{V} \tag{4}$$

We can therefore model the current generated by a particular ion like this, for i the ion in question and V_m the membrane potential:

$$I_i = g_i(V_m - V_i) \tag{5}$$

Because sodium and potassium are the most important ions in the neuron that the squid giant axon is part of, Hodgkin and Huxley assumed that there would be two "voltage-gated channels" in the cell membrane (one for each of those two ions) that the neuron would use to actively regulate its ion concentrations, as well as a "leak channel" representing the ability of ions to passively diffuse across the cell membrane. (Note that other ions, such as calcium to give one example, are also important, and indeed other neuron models may use different ions than the Hodgkin-Huxley model does.) The leak channel was assumed to have constant conductance, while for the voltage-gated channels, the conductance was assumed to change based on how open the channel in question was. The calculations for the conductances of the sodium and potassium channels were approached differently. Hodgkin and Huxley assumed that the conductance for the potassium channel would be proportional to some variable n associated with the activation of that channel, and that the conductance of the sodium channel would be proportional to variables m and h associated with the

activation and inactivation (respectively) of that channel. The three variables n, m, and h were all assumed to be on a scale from 0 to 1, and all of them are dimensionless quantities rather than concentrations of specific molecules that would activate or inhibit the ion channels. Hence, the assumption was that the neuron would regulate its ion concentrations by undergoing cellular processes that increase or decrease n, m or h. (These would correspond to terms in $\frac{dn}{dt}$, $\frac{dm}{dt}$, and $\frac{dh}{dt}$.) These processes are very complicated, so the actual differential equations in the model for $\frac{dn}{dt}$, $\frac{dm}{dt}$, and $\frac{dh}{dt}$ were assumed to vary directly based on the membrane potential V_m rather than introducing more variables representing molecules involved in cell signaling. Specifically, functions α_n, β_n, α_m, β_m, α_h, and β_h (which all take as input V_m) were introduced to model the increases and decreases of n, m and h. Putting everything together yields the following model (note that $I = I_c + I_{\mathrm{K}} + I_{\mathrm{Na}} + I_l$ is the sum of the currents previously explained):

$$\begin{cases} I = C_m \frac{dV_m}{dt} + \bar{g}_{\mathrm{K}} n^4 (V_m - V_{\mathrm{K}}) + \bar{g}_{\mathrm{Na}} m^3 h (V_m - V_{\mathrm{Na}}) + \bar{g}_l (V_m - V_l) \\ \frac{dn}{dt} = (1-n)\alpha_n(V_m) - n\beta_n(V_m) \\ \frac{dm}{dt} = (1-m)\alpha_m(V_m) - m\beta_m(V_m) \\ \frac{dh}{dt} = (1-h)\alpha_h(V_m) - h\beta_h(V_m) \end{cases} \tag{6}$$

The functions causing the increases and decreases in the three dimensionless variables n, m and h were fitted by Hodgkin and Huxley to real data that they obtained in lab experiments. They were empirically determined to have the following forms, for V_r the resting potential of the neuron:

$$\begin{aligned} \alpha_n(V_m) &= \tfrac{0.01(10+V_r-V_m)}{\exp(0.1(10+V_r-V_m))-1} & \beta_n(V_m) &= 0.125\exp(\tfrac{V_r-V_m}{80}) \\ \alpha_m(V_m) &= \tfrac{0.1(25+V_r-V_m)}{\exp(0.1(25+V_r-V_m))-1} & \beta_m(V_m) &= 4\exp(\tfrac{V_r-V_m}{18}) \\ \alpha_h(V_m) &= 0.07\exp(\tfrac{V_r-V_m}{20}) & \beta_h(V_m) &= \tfrac{1}{\exp(0.1(30+V_r-V_m))+1} \end{aligned} \tag{7}$$

Note how the forms of the α and β functions are different for h compared to the other two, due to the slightly different biological role played by h. This system is 4-dimensional and highly nonlinear, so analytical solutions of it are impossible. However, a fixed point can be found, and the Jacobian of the system can be evaluated at it. Doing so reveals that there are two negative eigenvalues and two complex eigenvalues, which means that the preconditions for a Hopf bifurcation in the Hodgkin-Huxley model are satisfied. Most of the parameters in the model are biologically determined and therefore not easily manipulated, but taking the current I as a bifurcation parameter can cause the real parts of the two imaginary eigenvalues to switch from negative to positive. This Hopf bifurcation represents a change in model behaviour from the membrane potential reaching an equilibrium to the membrane potential exhibiting a periodic solution with sharp spikes in voltage. Biologically, these represent action potentials,

also known as when the neuron is firing. The spikes in voltage bear almost no resemblance to sine waves. Periodic oscillations of this type, with sharp rather than gradual increases or decreases, are known as "pulse-relaxation oscillations" or simply "relaxation oscillations", following electrical engineering terminology. (If parameters are chosen such that the periodic solution does not exist, but the system is sufficiently close to the Hopf bifurcation, then the transient dynamics of V_m before reaching equilibrium will resemble one period of the periodic solution; biologically speaking, this is a single action potential.) Note that in general, pulse-relaxation oscillations occur when there is some separation in timescales between different variables in a model.

Owing to the success of the Hodgkin-Huxley model, many other mathematical neuron models have been proposed over the years. This includes models in which many neurons are coupled together in one system; since ions can be transmitted between neurons via their synapses, the electrical potential in one neuron in such a system will depend on the electrical potential in any neurons that are upstream from it. For instance, a dynamical system modelling three neurons might look like this, for $i = 1, 2, 3$:

$$\begin{cases} I_i = C_m \frac{dV_{m,i}}{dt} + \bar{g}_{\mathrm{K}} n_i^4 (V_{m,i} - V_{\mathrm{K}}) + \bar{g}_{\mathrm{Na}} m_i^3 h_i (V_{m,i} - V_{\mathrm{Na}}) + \bar{g}_l (V_{m,i} - V_l) + I_{\mathrm{syn},i} \\ \frac{dn_i}{dt} = (1 - n_i)\alpha_n(V_{m,i}) - n_i \beta_n(V_{m,i}) \\ \frac{dm_i}{dt} = (1 - m_i)\alpha_m(V_{m,i}) - m_i \beta_m(V_{m,i}) \\ \frac{dh_i}{dt} = (1 - h_i)\alpha_h(V_{m,i}) - h_i \beta_h(V_{m,i}) \end{cases} \tag{8}$$

If we take $i = 1, 2, 3$, then this is a twelve-dimensional system, since each of the three neurons in the system has its own internal dynamics (and hence its own values for V_m, n, m, and h). However, the voltage in a given neuron may increase or decrease based on the voltages of other neurons that are connected to it via a synapse. Here, the term $I_{\mathrm{syn},i}$ for $i = 1, 2, 3$ (in other words, $I_{\mathrm{syn},1}$, $I_{\mathrm{syn},2}$, and $I_{\mathrm{syn},3}$) represents the amount of current that neuron i receives by virtue of having a synaptic connection with the other neurons. This is the way that electrical activity can propagate along many different neurons, causing signals to be transmitted through the entire nervous system.

Another important consequence of mathematical research on ODEs in neuroscience is the development of a way to reduce the differential equations for oscillatory quantities (i.e. any state variable in a dynamical system that has a periodic solution) to differential equations describing how far along in their periods these quantities are. (In other words, we would be describing the phase of each state variable.) This is well beyond the scope of this course, although the synchronization of voltages in different neurons as well as other oscillatory variables (i.e. determining when the differences in their phases go to zero) is an active area of current research.

2 December 1: SIR model; fitting a model to data

The event with the highest global impact in recent years is arguably the COVID-19 pandemic. Throughout the world, much of the response to Covid has been driven by mathematical models. Specifically, variations on one particular dynamical system model have been used to predict caseloads and deaths; this model is the SIR model. In addition to Covid, the SIR model has been used to make predictions regarding many other infectious diseases. Constructing this model from first principles is not very challenging. We start from the assumption that everyone in the population belongs to one of three categories, which are susceptible to infection, infected, and recovered from or resistant to infection (hence the name "SIR"). We therefore have three state variables S, I, and R, which represent the proportions of the population that belong to each group. The rates of change of each of these model components represent ways that a person can transition between categories, such as a susceptible person becoming infected or an infected person recovering. This style of dynamical system model, in which the state variables are categories of some kind, is called a "compartmental model".

So, how do we construct such a model? To start, we will look at one of the most basic cases: the standard SIR model without any added bells and whistles. We will assume that the total population of wherever we're studying is constant, because population growth takes place over a longer timescale than disease spread. (In other words, we are assuming that total population is at equilibrium relative to S, I and R.) We will also assume only two events of interest in the system: a susceptible person catching the disease, and an infected person recovering. Assuming that the disease spreads by person-to-person contact, the rate at which a susceptible person can become infected is proportional to how often a person in the category S encounters someone in the category I. As we have previously seen with predator-prey models, this is an interaction term and hence scales with both S and I. Contact is also more frequent if the total population is smaller (and vice versa) as a smaller population means more opportunities for the same people to bump into each other. Taking the population size to be $S + I + R = N$, we get the interaction term describing the process of a susceptible person becoming infected to be $\frac{\beta}{N}SI$. This term is added to $\frac{dI}{dt}$ and subtracted from $\frac{dS}{dt}$, as it represents someone leaving the susceptible category and entering the infected category. Finally, we will assume that infected people recover at some rate γ. This means that the total number of people leaving the infected category and entering the recovered category per unit time is γI. Putting these together yields the following system:

$$\begin{cases} \frac{dS}{dt} = \frac{-\beta SI}{N} \\ \frac{dI}{dt} = \frac{\beta SI}{N} - \gamma I \\ \frac{dR}{dt} = \gamma I \end{cases} \tag{9}$$

One thing that we can immediately notice (and that we assumed during the derivation of the model) is that the total population is some constant N, and hence a conserved quantity. Therefore, we can reduce the dimensionality of this system by 1, which we can easily do by considering $R = N - S - I$ since R does not occur in any of the three ODEs making up the model. Taking $N = 1$, as is often done, means that the state variables represent percentages of a total population. This also means that the model does not output fractional numbers of people, although the presence of this behaviour is not typically viewed as a problem since all models are approximations in the first place. We therefore can get the following, the simplest form of the SIR model:

$$\begin{cases} \frac{dS}{dt} = -\beta SI \\ \frac{dI}{dt} = \beta SI - \gamma I \end{cases} \tag{10}$$

This is simpler than the Hodgkin-Huxley neuron model, or even the Lotka-Volterra predator-prey model, and we can determine a lot about its behaviour analytically. For instance, we can find its fixed points. Based on the equation for $\frac{dI}{dt}$, we need $\beta SI = \gamma I$, which is true if $I = 0$ or $S = \frac{\gamma}{\beta}$. However, if $I \neq 0$, then we get a nonzero value for $\frac{dS}{dt}$ (and also $\frac{dR}{dt}$). This means that $I = 0$ is a requirement for a fixed point of this model, but there are no other requirements, so $(S^*, I^*, R^*) = (k, 0, 1 - k)$ is a fixed point for any $k \in [0, 1]$. The interpretation of this is that the epidemic will stop when there are no more infectious people to spread the disease, which is very intuitive.

One important part about the model is the rate of change of the infected component, since that will determine whether the epidemic under consideration will spread or die out. To examine this, we will consider a population that starts out entirely susceptible (i.e. $S(t = 0) = 1$). Now, suppose that a disease is stochastically introduced into this population. In other words, suppose that we perturb S and I by $-\varepsilon$ and $+\varepsilon$, respectively, where ε is a number very close to zero. This means that S is still approximately 1, but I is now nonzero and the rates of change in the model involving I are also nonzero. In particular, we now have the following for the rate of change of I:

$$\frac{dI}{dt} \approx (\beta - \gamma)I \tag{11}$$

From this, we can tell that the infection will spread if $\beta > \gamma$, since $\frac{dI}{dt}$ will be positive in that case. Likewise, if $\gamma > \beta$, the infection will die out, as the rate at which people get infected will be less than the rate at which infected people recover. Since 2020, you will have undoubtedly heard the term "R-naught", or "R_0", many times in describing how capable a disease is of spreading. It is actually a term taken directly from this model:

$$R_0 = \frac{\beta}{\gamma} \tag{12}$$

As you may already know, a value of R_0 above 1 means that a disease will spread, while a value of R_0 below 1 means that it will die out. This makes

R_0 the most important quantity represented in the model; in epidemiology, determining R_0 for a given disease is often the main goal of research. Another important feature of the model is that since γ represents the rate at which infected individuals recover per unit time, the quantity γ^{-1} represents the average length of time that a person stays infected for. Likewise, βS can be thought of as the number of people that an infected person can themselves infect per unit time, and hence an interpretation for β is the number of contacts an infected person has over a given length of time multiplied by the probability of each contact becoming infected. This means that both β and γ can be input into the model from real data, which I'll explain more about later.

One important thing about the SIR model is the fact that its simplicity makes it endlessly customizable. Diseases are very heterogeneous in terms of their effects and characteristics, and this can be reflected by adding new terms to the basic SIR model. For example, suppose that recovery from a particular disease does not grant immunity to that disease going forward. This can be reflected by turning the SIR model into what might be referred to as an SIS model:

$$\begin{cases} \frac{dS}{dt} = -\beta SI + \gamma I \\ \frac{dI}{dt} = \beta SI - \gamma I \end{cases} \tag{13}$$

This is even simpler, and can actually be solved analytically after reducing the model to one dimension based on the conserved quantity $S + I = 1$. (You can do this on your own time if you want.)

So far, we have assumed that nobody actually dies from the disease, which is obviously an inaccurate assumption to make. What if they do? This can be represented by adding a new state variable D (for "dead"), and an additional term to $\frac{dI}{dt}$:

$$\begin{cases} \frac{dS}{dt} = -\beta \frac{SI}{N} \\ \frac{dI}{dt} = \beta \frac{SI}{N} - \gamma I - \mu I \\ \frac{dR}{dt} = \gamma I \\ \frac{dD}{dt} = \mu I \end{cases} \tag{14}$$

Here, μ is (as you might be able to guess) the mortality rate of the disease. Note that this version of the model also explicitly brings back $N = S + I + R$, since the assumption that the population is constant over time is now broken.

Alternatively, suppose that when someone is exposed to a particular disease, there is some latency period during which they do not show symptoms of the disease and cannot infect other people. This additional state can be accounted for by introducing another state variable (E for "exposed") to the model. We could then end up with something like this:

$$\begin{cases} \frac{dS}{dt} = -\beta \frac{SI}{N} \\ \frac{dE}{dt} = \beta \frac{SI}{N} - \alpha E - \gamma E \\ \frac{dI}{dt} = \alpha E - \mu I - \gamma I \\ \frac{dR}{dt} = \gamma E + \gamma I \\ \frac{dD}{dt} = \mu I \end{cases} \tag{15}$$

Note that here, the calculations that go into finding R_0 will be quite different compared to what they are in the simple SIR model. Additionally, further possibilities for building a model are essentially endless. (That's one of the big issues in mathematical modelling: knowing which of the millions of possible model formulations to use.)

So, how do we use such a model in practice? We have all of these parameters that we need in order to determine the exact dynamics of the model, so how do we find their values? The answer is based on the available data. If we are trying to determine how an actual disease will spread throughout an actual population, we might have data points representing case counts, deaths and recoveries over time in that population. This gives us a set of points that our model should replicate at least fairly accurately, if it is a good model for predicting future dynamics of the same disease in the same population. We can therefore pick several sets of values for the model parameters, then simulate the chosen parameter sets in the model, to see which set corresponds to the best fit of the model output for the real-life data. (This can be done using least squares, for example.)

But how do we choose the parameter values to test? One way to do this is by random sampling. If our model doesn't have very many parameters that we need to determine, then we can simulate all combinations of values within whatever ranges we think those parameters are likely to be. For instance, the original SIR model has only two parameters, β and γ. If (for example) we thought that β was between 2 and 3, and γ was between 0.5 and 1.5, then we could test value pairs (β, γ) featuring regularly spaced values of β in the interval $[2, 3]$ (e.g. 2, 2.1, 2.2, et cetera), and likewise for $\gamma \in [0.5, 1.5]$. However, if we have a lot of parameters to fit, then this will take a lot of time. One alternative is to use what is called "Latin hypercube sampling", which significantly cuts down on computation time, and is easy to implement (it's actually a bit like sudoku). Suppose that for each parameter in our model, we have n possible choices of values for that parameter, which we can number from 1 to n. (So, in the SEIRD model above with parameters β, α, γ and μ, we would have $\beta_1, \ldots, \beta_n$ being the potential choices for β, $\alpha_1, \ldots, \alpha_n$ being the choices for α, and likewise for γ and μ.) A Latin hypercube sample is one in which we take n different parameter sets overall, so that each value from 1 to n only occurs once for each parameter. So, for instance, if we had a standard SIR model with two parameters β and γ, and four different choices for each parameter, then one possible Latin hypercube sample would be $\{(\beta_1, \gamma_2), (\beta_2, \gamma_4), (\beta_3, \gamma_3), (\beta_4, \gamma_1)\}$. Another advantage of a Latin hypercube sample is that it portrays a relatively accurate picture of the variability in parameters; sampling four points (β, γ) at

random might lead to something like $\{(\beta_3, \gamma_3), (\beta_3, \gamma_4), (\beta_4, \gamma_3), (\beta_4, \gamma_4)\}$ where the samples are all clustered together.

But how do we know these ranges that we can choose from? The best way to do this is do draw values directly from the literature related to the problem that we're trying to model. This is where knowledge of some subject besides math comes in handy, since building a model will invariably involve reading papers published in different fields. However, this is well beyond the scope of this course.

3 December 3: Blood glucose model

Another field that mathematical models are often used in (that we haven't seen very many examples from yet) is physiology. The human body can be thought of as a vast network of interacting processes, in which different types of cells and chemicals are produced, perform their functions, and are consumed or destroyed. You've seen with the SIR model that dynamical system models can be constructed based on diagrams of boxes and arrows, in which the boxes represent different categories of people (susceptible to a disease, currently infected, recovered from the disease, et cetera) and the arrows represent the movement of people between categories. In physiology, similar principles are used in order to construct mathematical models from process diagrams. The simplest form of this, which you've already seen, is reaction kinetics, in which an enzyme converts its substrate into a reaction product. However, many biological processes are not as simple as one chemical being turned into another. For instance, in the human pancreas, there exists a certain kind of cell called the β cell whose job it is to produce insulin, but this is obviously not the same as actual β cells being converted into insulin. In general, this means that we can still start off a mathematical model in physiology by drawing boxes and arrows, but the boxes and arrows in question might have different meanings than we previously used in the SIR model.

Since the human body is so complex, most mathematical models used in physiology cover specific processes. Models on a larger scale do exist: for instance, Prof. Robert Hester at the University of Mississippi works on a very large mathematical model called HumMod, which simulates a variety of different physiological processes on the scale of the entire human body. The full range of processes in that model obviously cannot be fully explained within the span of one lecture, so I will instead focus on something on a smaller scale, namely a model of insulin production that has been used to predict the onset of diabetes. Since diabetes is associated with elevated blood glucose levels and decreased insulin production, the model tracks three different variables: blood concentrations of glucose and insulin, and mass of β cells (which produce insulin). This means that the model will look like this, for G glucose concentration, I insulin concentration, and β the mass of β cells in the pancreas:

$$\begin{cases} \frac{dG}{dt} = u_1(G, I, \beta) \\ \frac{dI}{dt} = u_2(G, I, \beta) \\ \frac{d\beta}{dt} = u_3(G, I, \beta) \end{cases} \tag{16}$$

In order to get the specific terms in the model, we can think all the way back to the tank problems that we saw at the beginning of the course, where the rate of change of the contents of a tank was defined as the rate in minus the rate out. Likewise, the rate of change for blood glucose will be the rate at which it is produced minus the rate at which various parts of the body uptake it for their own use:

$$\frac{dG}{dt} = \text{Production} - \text{Uptake} \tag{17}$$

Since we are dealing with an actual physical quantity (i.e. glucose concentration), it is highly important to get the units right for G, $\frac{dG}{dt}$, and the terms defining $\frac{dG}{dt}$. (If we built all of the terms in the model correctly, but we evaluated the model dynamics where G is on some highly implausible scale such as millions of kilograms per liter, then we wouldn't get any results that are biologically meaningful.) A common scale used to track glucose concentration in medical settings is milligrams per deciliter, so we will take those to be the units for G. Additionally, this model specifically looks at changes in fasting glucose levels over long time scales (days to years), we will measure time t in days. Therefore, the units for $\frac{dG}{dt}$ are milligrams per deciliter per day, and therefore all terms that make up the differential equation $\frac{dG}{dt}$ must be in terms of milligrams per deciliter per day as well.

So, what factors influence the production and uptake of glucose? Based on experimental data, we know that these rates are based on the concentrations of insulin and glucose itself in the blood. When insulin concentration is held constant, increasing blood glucose causes less additional glucose to be produced by the body, and more of the glucose in the blood to be uptaken. This means that $\frac{dG}{dt}$ might take the form $a - \sigma G$, for a the net rate of glucose production when blood glucose levels are zero, since higher levels of G put downward pressure on G. (This is known as a "negative feedback loop"; a "positive feedback loop" is when higher levels of one variable cause the rate of change of that variable to increase even further.) We also know that higher blood insulin concentration causes greater uptake of glucose; diabetes is associated with the breakdown of this relationship, where the body's sensitivity to insulin is low. We can thus formulate $\frac{dG}{dt}$ as follows:

$$\frac{dG}{dt} = a - (b + cI)G \tag{18}$$

Here, b is the rate at which glucose is utilized by the body independent of insulin concentration, and c is insulin sensitivity. By our assumptions, the units for a need to be the same as those for $\frac{dG}{dt}$, namely milligrams per deciliter per day. Since the term $(b + cI)$ multiplies G, it needs to be in terms of day^{-1} so

that the units line up. Therefore, b has units of day^{-1}, and the units of c will be whatever is needed to cancel out the units of I (see below) to get cI also on the scale of day^{-1}.

For the differential equation governing insulin, $\frac{dI}{dt}$, we can start from similar assumptions. First of all, insulin concentration is often measured in thousandths of international units of insulin per milliliter, or μU ml^{-1}. Since we already know that we will be measuring time in days, that gives us units of μU ml^{-1} day^{-1}. We know that insulin is secreted by β cells, and that it is cleared by being uptaken by the liver, kidneys and various insulin receptors. We can therefore assume the following "rate in minus rate out"-style dynamics for $\frac{dI}{dt}$:

$$\frac{dI}{dt} = \text{Secretion} - \text{Clearance} \tag{19}$$

We know that β cells secrete insulin in response to high levels of blood glucose concentration. More specifically, experimental work has shown that the relationship between glucose concentration and the rate at which insulin is secreted by β cells is sigmoidal. We will therefore assume that the secretion term in $\frac{dI}{dt}$ is a sigmoidal saturation function of G (which will also depend on how many β cells there are). One relatively standard form for a sigmoidal function is $u(x) = \frac{x^2}{1+x^2}$. We will additionally assume that the maximum rate of insulin production is some constant d, and that the half-saturation constant that affects the shape of the sigmoidal curve is some value e. Furthermore, we will also assume that insulin is cleared at a constant rate f. This yields the following form for $\frac{dI}{dt}$:

$$\frac{dI}{dt} = \frac{d\beta G^2}{e + G^2} - fI \tag{20}$$

Once again, we will need our parameters to take units so that all terms in $\frac{dI}{dt}$ are measured in μU ml^{-1} day^{-1}. It should be relatively obvious that f has the units of day^{-1}. You can figure out the units for the rest of the parameters yourself (note that β measures the mass of existing β cells and is measured in milligrams).

What about $\frac{d\beta}{dt}$? Since β represents the mass of β cells, we can assume that the rates of change of β are related to β cells replicating and dying. We know that the replication of β cells increases with blood glucose concentration, for the same reasons that blood glucose concentration increases the rate of insulin production by existing β cells. However, extremely high levels of blood glucose concentration have been shown experimentally to decrease β cell replication. We will therefore assume that the rate at which new β cells are produced is a quadratic function of G, something like $\alpha G - \gamma G^2$, making the mass of β cells produced per unit time something like $(\alpha G - \gamma G^2)\beta$. As for β cell death, this can happen in two different ways, namely apoptosis (planned, or "natural", cell death) and necrosis (unregulated cell death due to harmful conditions). Experimental results suggest that β cell death also varies nonlinearly with glucose concentration, albeit in the opposite directions compared to replication.

Therefore, we will assume that the death rate for β cells is another quadratic polynomial in G, namely $k - \delta G + \eta G^2$. This makes the mass of β cells that die over a given time interval equal to $(k - \delta G + \eta G^2)\beta$, and hence the β cell death term will be $-(k - \delta G + \eta G^2)\beta = (-k + \delta G - \eta G^2)\beta$. Adding these together gives us our equation for $\frac{d\beta}{dt}$:

$$\frac{d\beta}{dt} = (-k + hG - mG^2)\beta \tag{21}$$

Therefore, we get the following for our model:

$$\begin{cases} \frac{dG}{dt} = a - (b + cI)G \\ \frac{dI}{dt} = \frac{d\beta G^2}{e+G^2} - fI \\ \frac{d\beta}{dt} = (-k + hG - mG^2)\beta \end{cases} \tag{22}$$

So, what can we do with this model? For one, we can determine which parameter values are likely to lead to normal versus diabetic blood glucose levels. This can be demonstrated by finding fixed points in the model, which we can do analytically. There exists one which corresponds to healthy levels of G, I and β, one that represents a hyperglycemic state in which G is pathologically high and I and β are both zero, and one additional fixed point between them. The pathological equilibrium is just $(G^*, I^*, \beta^*) = (\frac{a}{b}, 0, 0)$, which can be obtained trivially. The other two take the form $(G^*, I^*, \beta^*) = (G_i, I_i, \beta_i)$, for $i = 1, 2$ and the following specific parameter values:

$$G_{1,2} = \frac{h \pm \sqrt{h^2 - 4mk}}{2m} \tag{23}$$

$$I_i = \frac{a}{cG_i} - \frac{b}{c} \tag{24}$$

$$\beta_i = \frac{fI_i(e + G_i^2)}{dG_i^2} \tag{25}$$

The Jacobian can be calculated for each of these points, and the stability of them can therefore be determined. When taking parameter values corresponding to experimental results, the healthy and pathological fixed points are both stable, while the fixed point between them is a saddle point. However, if we perturb different model parameters, we can cause bifurcations to happen. For instance, if we decrease h, corresponding to a drop in β cell production, the healthy fixed point and the saddle point move closer together. For a critical value of h, these two fixed points collide and cause a fold bifurcation, annihilating one another. The effect of this is that if h falls below that critical threshold, the only equilibrium will be the pathological equilibrium. Another important analysis of this model has been to incorporate periodic oscillations in blood glucose and β cell count. The original version of the model always considered fasting levels of glucose, insulin and β cells, whereas in real life, these levels may vary over the course of the day. Allowing different model parameters to change

based on the human circadian rhythm caused much more complicated dynamics to emerge, including additional ways for hyperglycemic conditions to develop. The research on this was actually just recently performed by a friend of mine, and the paper containing the important findings is still under review, so I can't really reveal too much more about it. However, there are still many more results pertaining to diabetes waiting to be discovered, which almost certainly include some that can be obtained by using mathematical models.

4 December 5: Spatially explicit systems of ordinary differential equations

Since Calculus 3 is a corequisite for this course, and this is the last lecture that I'm giving this semester, I'm sure that you know all about partial derivatives by now. Therefore, you might think that we can construct differential equations using partial derivatives as well. This is correct; these are known as partial differential equations. For instance, you might run into this one a lot:

$$\frac{\partial u}{\partial t} = k\frac{\partial^2 u}{\partial x^2} \tag{26}$$

Or this one:

$$\frac{\partial^2 u}{\partial t^2} = k\frac{\partial^2 u}{\partial x^2} \tag{27}$$

As this is a class on ordinary differential equations rather than partial ones, the analysis of these is outside the scope of it. However, it is certainly possible to represent spatial patterns within a framework of ordinary differential equations. Instead of representing space as another variable, this is typically done by assuming different copies of the same dynamical system, representing the dynamics in several different locations. The state variables in these systems then influence each other based on how the objects that they represent interact across space. I will show you some examples of this, based on models that you have already seen.

Let's start with a simple one. Suppose we have a predator-prey system, specifically the Lotka-Volterra model that we saw in class earlier:

$$\begin{cases} \frac{dN}{dt} = rN - \alpha NP \\ \frac{dP}{dt} = \beta NP - mP \end{cases} \tag{28}$$

What if we have two different locations that the predators and prey both live in? We can first extend this model to represent the predators and prey in both locations:

$$\begin{cases} \frac{dN_1}{dt} = r_1 N_1 - \alpha_1 N_1 P_1 \\ \frac{dP_1}{dt} = \beta_1 N_1 P_1 - m_1 P_1 \\ \frac{dN_2}{dt} = r_2 N_2 - \alpha_2 N_2 P_2 \\ \frac{dP_2}{dt} = \beta_2 N_2 P_2 - m_2 P_2 \end{cases} \tag{29}$$

Note that I have introduced the indices 1 and 2 for the predators and prey in locations 1 and 2. I also use these indices for the parameters which represent conditions in each location that might differ. For example, if we think of the prey as a herbivore, which eats grass, the local prey growth rates r_1 and r_2 might be different depending on which area is more suitable for grass to grow.

So far, with this setup, the predators and prey in each location (also called a "patch") are completely independent of each other. However, in real life, one or both species might migrate between patches. Suppose that over some unit of time, the proportion of prey that migrate from one patch to the other is μ_N, which can be taken to be some value between 0 (no migration happens at all) and 1 (the entire population migrates during that length of time). It could theoretically also be greater than 1 if the time it takes for the entire population to move between patches is less than what t is defined as. Likewise, suppose that some proportion of predators also moves between patches over the unit of time. We will call this μ_P. This allows us to rewrite our system to include the effects of migration:

$$\begin{cases} \frac{dN_1}{dt} = r_1 N_1 - \alpha_1 N_1 P_1 + \mu_N N_2 - \mu_N N_1 \\ \frac{dP_1}{dt} = \beta_1 N_1 P_1 - m_1 P_1 + \mu_P P_2 - \mu_P P_1 \\ \frac{dN_2}{dt} = r_2 N_2 - \alpha_2 N_2 P_2 + \mu_N N_1 - \mu_N N_2 \\ \frac{dP_2}{dt} = \beta_2 N_2 P_2 - m_2 P_2 + \mu_P P_1 - \mu_P P_2 \end{cases} \tag{30}$$

The functional form for migration that I have used here is often called "passive dispersal", because it is essentially identical to diffusion or Brownian motion. We have also made the assumption here that the rate of migration from Patch 1 into Patch 2 is the same as the rate of migration from Patch 2 into Patch 1 for both species. This is not necessarily true: if a species prefers one patch or the other, the migration rates will not be symmetric. We could then include terms like $\mu_{N_{1,2}}$, $\mu_{N_{2,1}}$, $\mu_{P_{1,2}}$, and $\mu_{P_{2,1}}$ to describe the specific rates of migration from one given patch to another.

So, what are the effects of migration in this system? Well, consider the case without it, but where the parameters in each patch are different. (So, assume that $r_1 \neq r_2$, et cetera.) In such a system, the dynamics of the predators and prey in the two different locations will be completely different, with the two pairs of solutions (N_1, P_1) and (N_2, P_2) each following their own trajectories. However, increasing the migration parameters (each μ) will cause the solutions in each patch to resemble each other more and more. The reason for this is that we have added a term $\mu_N(N_2 - N_1)$ to N_1 that subtracts a proportion of the prey in patch 1 and adds a proportion of the prey in patch 2 (the other patch), and so on for P_1, N_2 and P_2. As these proportions increase towards 0.5, the

effect becomes more and more like subtracting half of N_1 from N_1 and adding back half of N_2, which has a similar effect as averaging the populations in the two patches.

We also know that populations in the Lotka-Volterra model tend to oscillate, due to the limit cycle that exists in the (N, P)-plane. What happens when we introduce migration? If the values of the μ parameters are high enough, the oscillations of N_1 and N_2 will start to synchronize, and likewise with P_1 and P_2. This happens even if their frequencies had been very different in the absence of migration due to different parameter values in patch 1 and patch 2. In fact, the more the solution trajectories in each patch differ from each other in terms of shape, the higher the migration threshold above which synchrony occurs.

What if we have even more patches? (Real-life ecological networks are often fairly large.) Then, we could assume that each species in our system can migrate between some or all of the pairs of patches, and derive equations like this:

$$\begin{cases} \frac{dN_i}{dt} = r_i N_i - \alpha_i N_i P_i + \sum_{j \neq i} \mu_{N_{i,j}} N_j - \mu_{N_{j,i}} N_i \\ \frac{dP_i}{dt} = \beta_i N_i P_i - m_i P_i + \sum_{j \neq i} \mu_{P_{i,j}} P_j - \mu_{P_{j,i}} P_i \end{cases} \tag{31}$$

Note that some of the μ constants might be zero, if there is no migration between one or more pairs of patches. For instance, two patches might be separated by impassable terrain such as a mountain range or a large body of water. Additionally, we might consider a case in which migration only happens one way, such as populations in a river. Since the flow of the river only goes one way, we might have something like $\mu_{N_{1,2}} = 0$ but $\mu_{N_{2,1}} \neq 0$. In other words, patch 1 can send organisms to patch 2, but cannot receive organisms from patch 2. In general, this framework allows us to construct arbitrarily large networks with arbitrarily many patches, of which each might have its own different parameter values corresponding to different ecological conditions. As a result, we can model spatially complex scenarios without leaving the boundaries of ODEs.

For another example, in a previous lecture, I alluded to the fact that the Hodgkin-Huxley neuron model can be used to simulate the dynamics in multiple different neurons that are connected by synapses. In that version of the Hodgkin-Huxley model, we assumed that the variables in each neuron that represented activation or inactivation of the neuron's ion channels would not depend on the conditions in other neurons. (In other words, their dynamics would be local.) On the other hand, we assumed that the voltage in a given neuron would be affected by the voltages in the neurons that it is connected to. This is because there is a biological mechanism that allows voltage to be carried from one neuron to the next, namely the release of neurotransmitters, while there is no biological mechanism that would cause the ion channels in one neuron to open and close based on whether the ion channels in a different neuron are open or closed.

As it turns out, we can actually model the effects of specific synapses and

neurotransmitters. Synapses can either be excitatory or inhibitory: an excitatory synapse increases the voltage in the postsynaptic neuron, increasing the possibility of a spike, whereas an inhibitory synapse decreases voltage in the postsynaptic neuron and reduces the possibility of it spiking. Regardless of whether the synapse is excitatory or inhibitory, the current produced by it follows the same general pattern, similar to the current produced by an ion channel:

$$I_{\text{syn}} = \bar{g}_{\text{syn}} f(V_m, \ldots)(V_m - V_{\text{syn}}) \tag{32}$$

In other words, the current produced by a given synapse has a maximum value $\bar{g}_{\text{syn}}$, and depends on the reversal potential V_{syn} of the neurotransmitter that the synapse uses. Apart from this, synapses can get very biologically complex (hence the function f that I have left unknown). Explaining these would take a long time to do, and the biology involved is probably beyond the scope of this course, so I'll leave the problem of constructing a model of multiple neurons as an exercise for those interested.