MATH 228 lecture notes for November 15, 17, and 19

Russell Milne

November 2021

1 November 15: Steady states and phase portraits

Previously, I introduced the concept of nonlinear dynamical systems, which will generally be very hard or (usually) impossible to solve by hand. That means that we need to gain information about the state variables in such a system using other means. This week, I'll talk to you about many of the things that you can do with nonlinear dynamical systems without even having to integrate them at all.

One of the most important features of a dynamical system is its steady states. These are points in the system for which all rates of change are zero. If a dynamical system is at a steady state, then it will stay there. (An exception to this is if we are dealing with stochastic differential equations, and some random perturbation knocks the system off of the steady state, but that is beyond the scope of this course.) Other words for a steady state that you may encounter include "fixed point" and "equilibrium point". Mathematically, a steady state of a dynamical system is defined as anywhere at which the derivatives that make up the system are all zero, as at such a point none of the system's state variables can change. In other words, suppose we have the following system:

$$\begin{cases} \frac{dx_1}{dt} = f_1(t, x_1, x_2, \ldots, x_n) \\ \frac{dx_2}{dt} = f_2(t, x_1, x_2, \ldots, x_n) \\ \cdots \\ \frac{dx_n}{dt} = f_n(t, x_1, x_2, \ldots, x_n) \end{cases} \tag{1}$$

A fixed point $\mathbf{x}^* = [x_1^* \; x_2^* \; \ldots \; x_n^*]^T$ is any point in which $f_i = 0 \; \forall i$. This means that no change to any of the state variables $x_1, \; x_2, \; \ldots, \; x_n$ will occur. In practice, this is easiest to get when none of the functions f_i have any dependence on t, i.e. all differential equations in the system are autonomous.

To find where one of the state variables in the system (say x_1) has a rate of change of zero, we just need to set f_1 to zero and solve for values of $x_1, \; x_2, \; \ldots, \; x_n$ that will make this so. If these values also make the rates of change of all of the other state variables zero, then they form a fixed point of the dynamical system (as defined above). Here's an example of this in action. Suppose that we have the following dynamical system:

$$\begin{cases} \frac{dx}{dt} = x^2 - 2y \\ \frac{dy}{dt} = 3x - xy \end{cases} \tag{2}$$

Let's start by setting $\frac{dy}{dt} = 0$. This occurs when either $x = 0$ or $y = 3$, since we get the relation $3x = xy$. Note that these are both lines in the (x, y)-plane rather than individual points. This will typically be the case when finding fixed points, since setting ODEs to zero will result in a relation between several variables (in the two-dimensional case, often a function of one variable with respect to the other). Any line in a 2×2 system on which one of the variables has a rate of change of zero is called a "nullcline"; the term "isocline" is also used, although this more properly refers to any line on which the slope of an ODE takes a specified constant value (not necessarily zero).

Next, we will take $\frac{dx}{dt} = 0$ to get the other condition for a fixed point. This results in the quadratic equation $y = \frac{1}{2}x^2$, which is the nullcline for x. Any point that is on the nullclines of both x and y will be a fixed point of the dynamical system. If you're finding fixed points in a 2×2 system, then you can plot the nullclines on a graph of y versus x to help visualize them and where they intersect. Any such plot where we graph two variables against each other rather than one of them against time is called the "phase plane", and an analogue in higher dimensions is "phase space". Once we have drawn the nullclines, we can also indicate regions of (x, y)-space where x is increasing or decreasing, depending on which side of the x-nullcline it is on, and likewise for y. This will allow us to get a rough approximation of our solutions without doing any integration at all. We'll see more things we can do with a phase plane later on.

In the case that we are working with, there will be three fixed points. This is because $y = \frac{1}{2}x^2$ is a parabola that opens upward, so it will intersect with the line $y = 3$ in two places, and it will intersect $x = 0$ at the point $(x^*, y^*) = (0, 0)$. To find the other two, we just need to find any x such that $\frac{1}{2}x^2 = 3$, which is true when $x = \pm\sqrt{6}$. Therefore, our other two fixed points are $(x^*, y^*) = (\sqrt{6}, 3)$ and $(x^*, y^*) = (-\sqrt{6}, 3)$.

It's possible for a system to have arbitrarily many fixed points. Take, for example, this relatively simple system:

$$\begin{cases} \frac{dx}{dt} = x \\ \frac{dy}{dt} = -2x \end{cases} \tag{3}$$

Here, the nullclines for x and y are the same, specifically $x = 0$, so the point $(x^*, y^*) = (0, y)$ for any real y will be a fixed point. This can be seen by integrating $\frac{dx}{dt}$ by hand. We get $x(t) = Ce^t$, but if we assume that x starts at 0 (for an initial condition), we get $C = 0$ and hence $x = 0$ because e^t cannot be 0 for finite t. Therefore, x is unchanging if it starts at 0, and since $\frac{dy}{dt}$ depends only on x, y will also never change no matter what initial condition is picked for it.

It's also possible for a system to have no fixed points. Consider this 3×3 system:

$$\begin{cases} \frac{dx}{dt} = x - z \\ \frac{dy}{dt} = y - x \\ \frac{dz}{dt} = z - x - 2 \end{cases} \tag{4}$$

If we take $\frac{dx}{dt} = 0$ and $\frac{dz}{dt} = 0$, we get two parallel planes in (x, y, z)-space. These can never intersect, so we can never get all three rates of change to be zero. There are places where two of the three state variables will be fixed, since the plane $y = x$ that we find by setting $\frac{dy}{dt} = 0$ intersects both of the other planes; the line $x = y = z$ is one of these.

What if our dynamical system is large and highly nonlinear, and the equations that we get when setting each rate of change equal to zero are hard to solve? In that case, we can use root-finding methods to obtain accurate approximations. You may have seen Newton's method in the past, most likely as a way to find roots of functions of a single variable. If you haven't, then it is defined as follows. Suppose that we want to find a root of a function $f(t)$ whose derivative exists. Then, starting at some initial guess t_0, we can apply the following formula to get a (usually) better guess t_1:

$$t_1 = t_0 - \frac{f(t_0)}{f'(t_0)} \tag{5}$$

This is a recurrence relation, so can be repeated additional times to come closer to the root and eventually get a very good approximation of it. We would like to have a way to generalize this to functions of multiple variables, as if we have this, we can find the roots of the functions $\frac{dx_1}{dt} = f_1(t, x_1, x_2, \ldots, x_n)$, $\frac{dx_2}{dt} = f_2(t, x_1, x_2, \ldots, x_n)$, and so forth. Luckily, such a way exists. The multidimensional analogue of the derivative is the Jacobian. Suppose that we have an autonomous dynamical system of the following form:

$$\begin{cases} \frac{dx_1}{dt} = f_1(x_1, x_2, \ldots, x_n) \\ \frac{dx_2}{dt} = f_2(x_1, x_2, \ldots, x_n) \\ \cdots \\ \frac{dx_n}{dt} = f_n(x_1, x_2, \ldots, x_n) \end{cases} \tag{6}$$

We can treat $f_1, f_2, \ldots, f_n$ as the entries of a vector-valued function, which we will call $\mathbf{F}$. The Jacobian of this function looks like this:

$$J_{\mathbf{F}}(x_1, x_2, \ldots, x_n) = \begin{bmatrix} \frac{\partial f_1}{\partial x_1} & \cdots & \frac{\partial f_1}{\partial x_n} \\ \vdots & \ddots & \vdots \\ \frac{\partial f_n}{\partial x_1} & \cdots & \frac{\partial f_n}{\partial x_n} \end{bmatrix} \tag{7}$$

If this is evaluated at a single point in $(x_1, \ldots, x_n)$-space, then it becomes a matrix full of constants. If this matrix is invertible, then we can use Newton's method, albeit a modified version in which we left-multiply F by the inverse of the Jacobian, $J_{\mathbf{F}}^{-1}$, instead of multiplying a univariate function f by $\frac{1}{f'}$. Thus, if $\mathbf{x}$ represents our vector of guesses for the fixed point of the system and $\mathbf{x}_0$ our initial guess, we can use the following formula to iterate towards a fixed point:

$$\mathbf{x}_1 = \mathbf{x}_0 - J_{\mathbf{F}}^{-1}(\mathbf{x}_0)\mathbf{F}(\mathbf{x}_0) \tag{8}$$

Note that taking the inverse of a matrix is unbelievably computationally expensive, i.e. it takes an extremely long time for a computer to do it (especially for big matrices). Therefore, faster methods have been developed to solve linear equations such as the one above. However, that is a topic for a different day.

Now that we know our fixed points, what can we say about them? Since fixed points represent the zeros of the vector-valued function with entries being the rates of change of our state variables, it would make sense that the rates of change of the state variables in the area around a fixed point would be very small. This is correct, but more important is which direction those small rates of change are in, particularly if they point towards or away from the fixed point. In the former case, any solution that starts sufficiently close to the fixed point will be drawn in towards it as time increases. In the latter case, solutions will be pushed away from it as time increases, although they would be drawn towards it if time is run backwards to $-\infty$. One example of this is the two differential equations $\frac{dx}{dt} = x$ and $\frac{dx}{dt} = -x$. Both of these have a single fixed point, namely $x = 0$. For $\frac{dx}{dt} = -x$, this is an "attractive" fixed point, which can be seen as the analytical solution to that ODE ($x(t) = e^{-t}$) tends towards the fixed point $x^* = 0$ as t increases. However, for $\frac{dx}{dt} = x$, $x^* = 0$ is instead a "repelling" fixed point as its solution ($x(t) = e^t$) moves away from the fixed point as time increases.

How do we tell whether a fixed point is attractive or repelling? It depends on the slopes of the functions f_1, f_2, and so on, specifically their slopes evaluated at the fixed point. This is because a fixed point $\mathbf{x}^*$ of the system is a place where the vector-valued function $\mathbf{F}$ (using the notation that we introduced above) is zero for all entries in the vector. Suppose we go a slight distance away from the fixed point, say to some point $\mathbf{x}^* + \mathbf{h}$ for some vector $\mathbf{h}$ whose entries are very small. The value of $\mathbf{F}$ is just the rates of change of all of our state variables. Therefore, if $\mathbf{x}^*$ is an attractive fixed point, then we want the slope of $\mathbf{F}$ to be in the opposite direction as our perturbation $\mathbf{h}$, so the action of $\mathbf{F}$ pushes us back into the fixed point. Similarly, for a repelling fixed point, we want $\mathbf{F}$ to act in the the same direction as the perturbation, so perturbing a solution away from the fixed point causes it to move even further away. I'll illustrate this with a one-dimensional example, for simplicity. Suppose we have this very straightforward ODE:

$$\frac{dx}{dt} = f(x) = x \tag{9}$$

Here, the only fixed point is $x^* = 0$. Therefore, we will be interested in the behaviour of $f(x)$ at some point $x = 0 + h = h$. If we take h to be positive, then $f(h)$ will also be positive, since it will just be h. Likewise, if we take h to be negative, then $f(h)$ will also be negative. In this way, perturbing some hypothetical solution $x(t)$ a small distance away from the equilibrium at 0 causes $\frac{dx}{dt} = f(x)$ to push it further away from the equilibrium, so $x^* = 0$ is a repelling fixed point. This can be confirmed by integrating this DE by hand and noticing that solutions tend to move away from 0 as time increases. Similarly, if we took $\frac{dx}{dt} = -x$, we would get exactly the opposite result. There, 0 is an attractive fixed point, consistent with the observed behaviour of exponential decay. For a general one-dimensional ODE $\frac{dx}{dt} = f(x)$, the sign of the derivative $f'(x)$ evaluated at the fixed point $x = x^*$ determines whether the fixed point is attractive or repelling. If $f'(x^*)$ is negative, the fixed point is attractive. If it's positive, then the fixed point is repelling. This can be verified for the cases mentioned above, as $\frac{d}{dx}(x) = 1 > 0$ and $\frac{d}{dx}(-x) = -1 < 0$ regardless of which point they are evaluated at.

2 November 17: Stability of fixed points in higher dimensions, including a modelling example

Previously, we looked at how to determine if a fixed point for a one-dimensional system (i.e. a single differential equation) was attractive or repelling. Today, we'll do the same for larger systems. Before we do, I have a final note on the one-dimensional case. We showed that for a system $x' = f(x)$ and a fixed point x^*, x^* is attractive, or "stable", if $f'(x^*) < 0$, and it is repelling, or "unstable", if $f'(x^*) > 0$. What if $f'(x^*) = 0$? In that case, we don't have enough information to tell whether or not the fixed point is stable or unstable, and we will need to look at the slope of f around the fixed point rather than just at it. It could even be "semi-stable", in which solution trajectories that start on one side of the fixed point flow towards it and those that start on the other side flow away from it. To visualize this, consider the following ODE:

$$\frac{dx}{dt} = x^2 \tag{10}$$

This has a fixed point at $x^* = 0$, but we can see that $f'(x) = x$, which also takes the value of 0 at $x = x^*$. However, $f'(x)$ takes negative values for negative values of x, and positive values for positive values of x. Based on what we said about the sign of f', we would thus expect a solution $x(t)$ that started with a negative value of x to flow towards the fixed point, and a solution $x(t)$ that started with a positive value of x to flow away from the fixed point. This is, in fact, true. If you integrate the above ODE by hand, you will get the solution $x(t) = \frac{1}{C-t}$ for C a constant of integration. This solution is a hyperbola with its singularity at $t = C$ and the horizontal axis as an asymptote; it increases towards ∞ when it is positive and decreases towards $-\infty$ when it is negative.

As a sidenote, the process of evaluating the stability of a fixed point of a nonlinear dynamical system by evaluating the slopes of the functions making up the system is called "linearization". The reason for the name is linked to the reason why it works. Suppose we have a one-dimensional ODE, $x' = f(x)$, which has a fixed point x^*. Without loss of generality, we can assume that this fixed point is $x^* = 0$, since if it is some other value of x we can just do a coordinate transform to make it 0 (e.g. $u = x - 3$ for a fixed point $x^* = 3$). If this is a linear ODE, i.e. $f(x) = kx$ for some k, then determining the stability is easy, because we can integrate by hand and get either exponential growth or exponential decay. (In this case, the derivative f' is also just a constant, making it easy by this method as well.) However, for nonlinear f, integration by hand might be difficult. On the other hand, all we want to know is the behaviour of a solution in some neighbourhood of the fixed point, not over all possible values of t. Therefore, we can make a linear approximation to the function by using the Taylor series:

$$f(x) = x^* + f'(x^*)(x - x^*) + \mathcal{O}(x^2) \tag{11}$$

Note that if $x^* = 0$, we get a linear function in x with only one term, which can easily be integrated by hand. In general, this will be a pretty good approximation close to the fixed point, since the linear term of the Taylor expansion will dominate all of the others close to 0. However, the more nonlinear our function is, the more error will accumulate as we move away from the fixed point. If we are far away from the fixed point under consideration, the solution might do something that its linear approximation wouldn't, such as converge to a different fixed point.

So, let's dive in to higher-dimensional systems. Using the notation that we introduced on Monday, suppose that we have the following dynamical system:

$$\frac{d\mathbf{x}}{dt} = \mathbf{F}(\mathbf{x}) \tag{12}$$

As with the one-dimensional case, we want to find out what the slope of $\mathbf{F}$ is at values very close to the fixed point, or in other words $\mathbf{F}(\mathbf{x}^* + \mathbf{h})$ for small $\mathbf{h}$. We will still do this by linearization, but in multiple dimensions, this is more complicated than just looking at the sign of the derivative of f, like we did in the one-dimensional case. Instead, we'll look at the Jacobian of our system (which we talked about earlier), since the linear term in the Taylor expansion for a vector-valued function involves the Jacobian:

$$\mathbf{F}(\mathbf{x}) \approx \mathbf{F}(\mathbf{x}^*) + J_{\mathbf{F}}(\mathbf{x}^*) \cdot (\mathbf{x} - \mathbf{x}^*) \tag{13}$$

Another way to look at this is that the Jacobian encompasses all of the partial derivatives of the functions in $\mathbf{F}$, and therefore the slopes of all of the functions in $\mathbf{F}$ with respect to all of the variables in the system. Note also that if our system isn't autonomous, then we might have some terms in the Jacobian with t, based on how the partial derivatives turn out. In that case, linearization requires some extra justifications, which I won't go into here.

So, we have that the linearization of our function $\mathbf{F}$ involves taking the Jacobian of $\mathbf{F}$ at the fixed point $\mathbf{x}^*$. How do we use this to evaluate the stability of $\mathbf{x}^*$? Well, instead of just taking the sign of f' in the one-dimensional case, we will take the signs of all of the eigenvalues of $J_{\mathbf{F}}$. More specifically, we will take the signs of all of the real parts of the eigenvalues of $J_{\mathbf{F}}$, This is because real-valued matrices can have complex eigenvalues (as we have seen), but the complex parts of these don't affect convergence to the fixed point $\mathbf{x}^*$, because we have been operating under the assumption that $\mathbf{x}^*$ is real and $\mathbf{F}$ is real-valued.

If all of the eigenvalues have negative real parts, then perturbing our solution by some slight amount in any direction in phase space away from the fixed point will cause the trajectory of the solution to fall back into the fixed point. This means that the fixed point will be stable, or a "sink". If at least one eigenvalue has a positive real part, then solutions will eventually escape along the eigendirection in phase space associated with that eigenvalue, making it unstable. However, there are two different ways that this can happen. If every eigenvalue has a positive real part, then solution trajectories will escape from the fixed point in any direction in phase space. Such a fixed point is called a "source"; note that if time is run backwards, sources become sinks and vice versa. If some eigenvalues have positive real parts and some have negative real parts, then solution trajectories may approach the fixed point along an eigendirection associated with a negative eigenvalue, but then move away from it along an eigendirection associated with a positive eigenvalue. These kinds of fixed points are called "saddle points". To see why, picture some small object sitting in the middle of a saddle, perfectly balanced. If you perturb it forward or backward with no lateral motion, then it would theoretically roll back into the fixed point in the middle. However, if you perturbed it to the left or right, it would roll off. If instead the small object started out somewhere other than the very middle of the saddle, it would initially roll towards the fixed point at the middle, but its ultimate fate would be to fall off. Plotting a solution trajectory in phase space would reveal a similar pattern, with "forward or backward" corresponding to a stable eigendirection and "left or right" corresponding to an unstable one.

I will illustrate this with a few simple examples. Suppose we have the following system, of which the solution should be obvious:

$$\begin{cases} \frac{dx}{dt} = x \\ \frac{dy}{dt} = 2y \end{cases} \tag{14}$$

Here, we have one fixed point, which is the origin. This system is already linear, so finding the Jacobian of it is trivial. We get the following:

$$J_{\mathbf{F}} = \begin{bmatrix} 1 & 0 \\ 0 & 2 \end{bmatrix} \tag{15}$$

This is a diagonal matrix, so its eigenvalues are just the diagonal entries, namely $r = 1$ and $r = 2$. (If you really want to calculate the characteristic polynomial, it's $r^2 - 3r + 2$.) Both of these are positive, so we get that the origin is a source, which lines up with what we know about exponential functions. If, instead, we had the following dynamical system:

$$\begin{cases} \frac{dx}{dt} = -x \\ \frac{dy}{dt} = -2y \end{cases} \tag{16}$$

then we would get the origin to be a sink, which once again confirms what we can find analytically. What about a saddle point? Well, suppose we have the following dynamical system:

$$\begin{cases} \frac{dx}{dt} = x + y \\ \frac{dy}{dt} = x - y \end{cases} \tag{17}$$

This is a coupled linear system, and we can solve it analytically using the methods that we have already learned. It has the origin as its only fixed point, as that is the only place where the lines $x + y = 0$ and $x - y = 0$ cross. Taking the Jacobian of this system (which is the same as the coefficient matrix, as the system is linear), we get the following:

$$J_{\mathbf{F}} = \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \tag{18}$$

The characteristic polynomial of this matrix is $r^2 - 2 = 0$, meaning that our eigenvalues will be $r = \pm\sqrt{2}$. From what we know about the theory behind linear systems, we can say that our general solution will include a term with $e^{\sqrt{2}t}$ and a term with $e^{-\sqrt{2}t}$. As t increases, depending on the coefficients on the terms (which we can determine using the initial conditions), we could see trajectories in phase space (i.e. the (x, y)-plane) in which the solution (x, y) first appears to approach the origin, but then moves away from it. (A time series of $y(t)$ would show something similar.) This can be seen, of course, by the fact that one of the eigenvalues of the Jacobian is positive and the other is negative.

What if the eigenvalues are complex? If that is the case, then the behaviour of our solutions in phase space will involve rotation. (This is intuitive to see. Think of the behaviour of e^{kt} when k is real versus when k is complex or purely imaginary, and the fact that the sine and cosine functions in (x, y)-space are linked to rotation around a unit circle.) In this case, if the fixed point that we are evaluating the Jacobian at is a sink or a source, then solutions will travel towards or away from the fixed point (respectively) in a spiral manner in phase space. (Saddle points won't see much change in practice.) If the eigenvalues are purely imaginary, then solutions will neither converge to nor escape from the fixed point, instead rotating around it in a circle, at least theoretically. In practice, this is more accurate for linear systems, since for these the Jacobian is the same as the coordinate matrix that is used when solving by hand, as seen above. For nonlinear systems, the extra terms present in the Taylor expansion that get ignored during linearization may cause the eigenvalues to have nonzero real parts. We might still get cyclic patterns, but if we do, it's highly unlikely that they'll be circular.

Since we have just looked at linear systems so far, let's see an example of the linearization process for evaluating the stability of a fixed point of a nonlinear system. Suppose we have the following system, which we used last week as a predator-prey model:

$$\begin{cases} \frac{dN}{dt} = rN - \alpha NP \\ \frac{dP}{dt} = \beta NP - mP \end{cases} \tag{19}$$

We will assume that all constants in the model are non-negative, since the system becomes non-biological otherwise (e.g. "I eat you and then there's more of you"). The first step here is to find the fixed points. $\frac{dN}{dt} = 0$ when either $N = 0$ or $P = \frac{r}{\alpha}$, and $\frac{dP}{dt} = 0$ when either $P = 0$ or $N = \frac{m}{\beta}$. Therefore, there are two locations in phase space where an N-nullcline intersects with a P-nullcline, which are $(N^*, P^*) = (0, 0)$ and $(N^*, P^*) = (\frac{m}{\beta}, \frac{r}{\alpha})$. The first of these represents extinction of both species, and the second represents coexistence.

Now, let's evaluate the Jacobian at these points. If we take the partial derivatives of both of the functions making up the dynamical system, we get the following:

$$J_{\mathbf{F}} = \begin{bmatrix} r - \alpha P & -\alpha N \\ \beta P & \beta N - m \end{bmatrix} \tag{20}$$

At $(0, 0)$, most of the terms in the Jacobian reduce to zero, and we are left with the following:

$$J_{\mathbf{F}} = \begin{bmatrix} r & 0 \\ 0 & -m \end{bmatrix} \tag{21}$$

If we assume that both r and m are positive, then the origin is therefore a saddle point and hence unstable. In particular, the positive eigenvalue will be r; you can calculate the corresponding eigenvector (and hence the eigendirection that solutions will escape from $(0, 0)$ on) yourself. A biological interpretation of this is that so long as the two species aren't both completely extinct, they will continue to survive in the long term. What about the other fixed point, $(\frac{m}{\beta}, \frac{r}{\alpha})$? In this case, we get the following for our Jacobian:

$$J_{\mathbf{F}} = \begin{bmatrix} 0 & \frac{-\alpha m}{\beta} \\ \frac{r\beta}{\alpha} & 0 \end{bmatrix} \tag{22}$$

Note that we have zeros on the diagonal, and that the other two terms have opposite signs. That's a clue that we'll have purely imaginary eigenvalues, and indeed our characteristic polynomial is $r^2 + rm$, which has the roots $r = i\sqrt{rm}$. We had previously said that this might indicate periodic orbits of the fixed point, although as the system is nonlinear, more analysis is needed. In order to see whether or not these periodic solutions exist, since we can't readily integrate the functions in this system with respect to t, we will instead look at the movement of solutions within the phase plane, i.e. in (N, P)-space. Consider the following differential equation:

$$\frac{dP}{dN} = \frac{dP}{dt}\left(\frac{dN}{dt}\right)^{-1} = \frac{\beta NP - mP}{rN - \alpha NP} = \frac{P(\beta N - m)}{N(r - \alpha P)} \tag{23}$$

This describes the movement of P relative to N, and luckily for us it is separable. Separating and integrating yields the following:

$$\int \frac{r - \alpha P}{P}\, dP = \int \frac{\beta N - m}{N}\, dN \tag{24}$$

This evaluates to the following expression, after moving around some symbols and combining constants of integration:

$$r \ln P - \alpha P - \beta N + m \ln N = C \tag{25}$$

We have from before that $N = 0$ is an N-isocline and that $P = 0$ is a P-isocline. From this, we can conclude that if N starts off positive, then it can't become negative, and likewise for P. In other words, any solution that starts in the first quadrant of the (N, P)-plane will stay there. This is reassuring, since if it weren't true then this model wouldn't be very realistic. Additionally, we just derived that the expression $r \ln P - \alpha P - \beta N + m \ln N$ must always be equal to a constant, which is finite. This prevents N or P from blowing up to infinity, so long as both species exist. (Since $\ln N$ and $\ln P$ could be $-\infty$ if N or P is zero, we could get the prey going to ∞ if there are no predators present. If there is no prey, $\frac{dP}{dt}$ will be strictly negative for initial conditions in the first quadrant, so the predators cannot escape to ∞.) Based on these conclusions, we can confidently say that both the predator and prey populations will follow bounded, periodic solutions.

3 November 19: Introduction to bifurcations

Previously, we worked through finding the fixed points of Lotka-Volterra predator-prey model and their stability. This model was a bit different from the other dynamical systems that you may have seen before, since it included some parameters. These are any constant in the model that is left unspecified (but is not a state variable or t). In the Lotka-Volterra model, these were r, α, β, and m. We made the assumption that all of these were positive, which caused us to make some conclusions about our fixed points, namely that $(N^*, P^*) = (0, 0)$ is a saddle point and that $(N^*, P^*) = (\frac{m}{\beta}, \frac{r}{\alpha})$ admits periodic orbits. However, in other dynamical systems, the parameters might not be defined so strictly. Another way of saying this is that there might be a fixed point for which the stability changes depending on the values of certain parameters in the system (even if we make an assumption that all parameters are positive or something similar). It's also possible for a dynamical system to have fixed points that only exist for certain ranges of a given parameter; if the parameter is outside these ranges, then $\mathbf{F}(\mathbf{x})$ evaluated at that "fixed" point might be nonzero. These occurrences, as well as other significant qualitative changes in the behaviour of a dynamical system when one of its parameters is changed, are called "bifurcations".

I will demonstrate with a simple example. Suppose we have the following one-dimensional ODE, for k a constant:

$$\frac{dx}{dt} = f(x) = k - x^2 \tag{26}$$

Let's try finding the fixed points of this DE. They occur when $x^2 = k$, which means that we will see different behaviour when $k > 0$, $k = 0$ and $k < 0$. For $k > 0$, the equation $x^2 = k$ has two real solutions, which means that there will be two fixed points, $x^* = \sqrt{k}$ and $x^* = -\sqrt{k}$. In order to evaluate the stability of these, we will take the derivative of $f(x)$, which is $f'(x) = -2x$. Evaluated at $\sqrt{k}$, this is a negative number (since $\sqrt{k}$ is positive), which means that the fixed point $x^* = \sqrt{k}$ is stable. Likewise, we get that the fixed point $x^* = -\sqrt{k}$ is unstable by using the same method.

On the other hand, if we let $k < 0$, then there are no fixed points in this system, because the values of x that make $\frac{dx}{dt} = 0$ will be purely imaginary. Taking $k = 0$ makes the system reduce to $\frac{dx}{dt} = -x^2$, which has just one fixed point at $x^* = 0$ rather than the two that appear when $k > 0$. In this case, $f'(x)$ evaluated at the single fixed point $x^* = 0$ is also zero, which means that we can't draw any conclusions about its stability from evaluating $f'(x)$. (We saw previously that $\frac{dx}{dt} = x^2$ has a semi-stable fixed point at $x^* = 0$, and $\frac{dx}{dt} = -x^2$ is the same, albeit with the areas in which the fixed point attracts or repels solutions being swapped.) Because our ODE $\frac{dx}{dt} = k - x^2$ has a major change to its fixed points at the parameter value $k = 0$, we say that a bifurcation happens at $k = 0$.

A good way to visualize the fixed points of a system and how they can change with different values of a certain parameter is by drawing a bifurcation plot. This is a graph with a parameter on the horizontal axis and the fixed point of some state variable on the vertical axis. In other words, given a dynamical system, we are plotting the value taken by one of the variables in the system at one of the system's fixed points as a function of a parameter. For the system we saw previously, we can plot x^* as a function of k. On the right-hand side of this bifurcation plot, where $k > 0$, we will get two lines showing the locations of our two fixed points. (By convention, on a bifurcation plot, stable fixed points are shown as solid lines, while unstable ones are shown as dashed lines.) Since the locations of these two fixed points are $x^* = \pm\sqrt{k}$, the graph looks like what we would get if we plotted the two equations $y(t) = \sqrt{t}$ and $y(t) = -\sqrt{t}$ as part of a standard time series. These two lines representing the fixed points will collide with each other at $k = 0$, and on the left-hand side of the graph no lines representing fixed points will exist. This kind of bifurcation, in which a stable fixed point and an unstable fixed point come together and annihilate one another, is called a "fold bifurcation", as on a bifurcation plot it looks like the line representing the fixed point is being folded over. Fold bifurcations are something to watch out for in mathematical models, as they can indicate the possibility of drastic changes in the trajectory of a solution if the underlying parameter is altered in some way.

What are some other kinds of bifurcations? Let's look at another one-dimensional ODE, once again for k a constant:

$$\frac{dx}{dt} = f(x) = kx - x^2 \tag{27}$$

This one has fixed points when $kx = x^2$, or in other words when $x = k$ and $x = 0$. Taking the derivative of f results in $f'(x) = k - 2x$. When we evaluate the stability of $x^* = k$, we can notice that f' reduces to $-k$. This means that if $k > 0$, $x^* = k$ is stable, and if $k < 0$ it is unstable. Meanwhile, if we evaluate the stability of $x^* = 0$, we arrive at the opposite conclusion, because there f' reduces instead to k. Therefore, at $k = 0$, the stability of both of the fixed points changes, with $x^* = k$ going from unstable to stable and $x^* = 0$ going from stable to unstable. This also corresponds to the intersection of $x^* = 0$ and $x^* = k$ on a bifurcation plot. (Hence, at $k = 0$ only one fixed point will exist, and we won't be able to ascertain its stability by taking f' because f' will be zero.) However, unlike in the example we saw previously, the two lines on the bifurcation plot cross at an angle rather than colliding head-on, so we get a change in stability rather than them obliterating each other. This kind of bifurcation, in which two fixed points cross each other and change each other's stability, is called a "transcritical bifurcation".

Are there more kinds of bifurcations? Of course there are. Consider the following system, for k a constant:

$$\frac{dx}{dt} = kx - x^3 \tag{28}$$

This has fixed points whenever $x^3 = kx$, which happens when $x = 0$ and when $x = \pm\sqrt{k}$. Therefore, there can be a maximum of three fixed points for this dynamical system, but this only occurs when $k > 0$. For $k < 0$, $\pm\sqrt{k}$ will both be imaginary, so $x^* = 0$ will be the only fixed point. So far, we already know that something interesting happens at $k = 0$, namely the creation of two additional fixed points. What happens when we evaluate their stability? Well, $f'(x) = k - 3x^2$. When we test $x^* = 0$, this reduces to just k, so $x^* = 0$ is stable if $k < 0$ and unstable if $k > 0$. For $x^* = \pm\sqrt{k}$, f' evaluates to $k - 3k$, which is negative for the only values of k for which those two fixed points exist (i.e. $k > 0$). Therefore, at $k = 0$, a lot of different things happen. The existing fixed point in the system switches from being stable to unstable, but two new stable fixed points are created on either side of it. If you plot all of the fixed points on a bifurcation plot, the shape of all of the lines resembles a pitchfork, and hence this kind of bifurcation is called a "pitchfork bifurcation". Note that you can also get pitchfork bifurcations where a fixed point goes from unstable to stable and two new unstable fixed points are created on either side of it. For an example of this, look at $\frac{dx}{dt} = kx + x^3$. A pitchfork bifurcation which has one stable fixed point that branches into two stable and one unstable ones is called "supercritical", whereas a pitchfork bifurcation which has one unstable fixed point that branches into two unstable and one stable fixed points is called "subcritical".

There's another kind of bifurcation that only occurs in systems of two dimensions or higher. This is called the Hopf bifurcation, and it is related to a dynamical system having periodic solutions. Before I begin explaining it, I would first like to point out that having a dynamical system with at least two state variables is a prerequisite for obtaining a periodic solution. You can see this by noting that for a one-dimensional ODE $\frac{dx}{dt} = f(x)$, the value of f' evaluated at any real fixed point will itself be a real number, but for a dynamical system of dimension 2 or higher, the eigenvalues of the Jacobian evaluated at a fixed point might be imaginary. A periodic solution to a dynamical system is also referred to as an "orbit", as that's what a drawing of it in a phase plane will resemble. If it serves as the limit for some solution trajectories as either time goes to ∞ or time goes to $-\infty$, then it is referred to as a "limit cycle" (stable and unstable, respectively). Periodic solutions that are solely composed of sines and cosines are not limit cycles, since a dynamical system with sines and cosines as its general solution will have each individual solution be a circle in phase space (and thus no solution curve will ever converge to another one). However, if you have a nonlinear dynamical system that admits periodic solutions, these will be limit cycles.

So, what is a Hopf bifurcation? Well, consider the following dynamical system, for k a parameter:

$$\begin{cases} \frac{dx}{dt} = y \\ \frac{dy}{dt} = -x + (k - x^2)y \end{cases} \tag{29}$$

This is what we get when we transform a second-order ODE called the Liénard equation (which is $x'' - (k - x^2)x' + x = 0$) into a system of first-order ODEs. The only fixed point for this system is the origin, as you can see by setting up the nullclines. The Jacobian of this system evaluated at the origin is as follows:

$$J_{\mathbf{F}} = \begin{bmatrix} 0 & 1 \\ -1 & (k - x^2) \end{bmatrix} \implies J_{\mathbf{F}}(0,0) = \begin{bmatrix} 0 & 1 \\ -1 & k \end{bmatrix} \tag{30}$$

The characteristic equation for this matrix is $r^2 - kr + 1 = 0$, so we get solutions of the following form:

$$r = \frac{k}{2} \pm \frac{\sqrt{k^2 - 4}}{2} = \frac{k}{2} \pm \frac{i\sqrt{4 - k^2}}{2} \tag{31}$$

When the absolute value of k is small, the eigenvalues will be complex conjugates. If k is negative, then the real part of both eigenvalues will be negative, so the origin will be a stable node (specifically a spiral because of the imaginary parts). However, when $k = 0$, the real part disappears and we are left with periodic orbits for our solutions. What about when $k > 0$? We would expect the origin to become an unstable node. However, this is only part of what actually happens. In addition to the origin switching from stable to unstable, a limit cycle appears around the origin. This new limit cycle is stable, which means that it soaks up any solution trajectory that starts somewhere near it. (In this case, as the origin is the only fixed point and it's unstable, any solution to our dynamical system that starts anywhere other than the origin will converge to this new limit cycle as time goes to ∞.) This is the result of the Hopf bifurcation. The general criteria for a Hopf bifurcation are as follows. Suppose that we have a dynamical system including some parameter k, which has a fixed point $\mathbf{x}^*$. Suppose also that there is a value of k (which we will call k_0) for which all of the eigenvalues of the Jacobian evaluated at $\mathbf{x}^*$ have negative real parts, except for one pair of eigenvalues that are purely imaginary. In other words, suppose that this pair of eigenvalues is of the form $g(k) \pm ih(k)$ in general, but for $k = k_0$ we get $g(k_0) = 0$ and $h(k_0) \neq 0$. Now, suppose further that $\frac{dg}{dk}(k = k_0) > 0$. Then, there exists some k_1 such that the dynamical system in question has a periodic orbit surrounding $\mathbf{x}^*$ for $k_0 < k < k_1$. (If $\frac{dg}{dk}(k = k_0) < 0$, then we instead get an orbit for $k_1 < k < k_0$, as the values of k for which $g(k) > 0$ and $g(k) < 0$ will be reversed.)

In practice, the existence of Hopf bifurcations means that an unstable equilibrium in 2×2 systems or higher will often be surrounded by a limit cycle. Here's an example of a mathematical model in which they occur. Consider the following dynamical system:

$$\begin{cases} \frac{dx}{dt} = -x + ay + x^2 y \\ \frac{dy}{dt} = b - ay - x^2 y \end{cases} \tag{32}$$

This is the Sel'kov model of glycolysis, which is a simplified version of a more complicated enzyme kinetic model (similar to the one that you saw earlier) after a few biological assumptions were made about some reactions being very fast compared to others. This model has a fixed point at $(x^*, y^*) = (b, \frac{b}{a+b^2})$. The Jacobian is as follows:

$$J_{\mathbf{F}} = \begin{bmatrix} -1 + 2xy & a + x^2 \\ -2xy & -a - x^2 \end{bmatrix} \implies J_{\mathbf{F}}(x^*, y^*) = \begin{bmatrix} -1 + \frac{2b^2}{a+b^2} & a + b^2 \\ \frac{-2b^2}{a+b^2} & -a - b^2 \end{bmatrix} \tag{33}$$

Calculating the eigenvalues of this matrix is quite heavy on the algebra, so I'll omit it here. (Of course, you could also calculate them numerically.) They will take the form of a conjugate pair, with the real part as follows:

$$r = g(a, b) \pm ih(a, b) \implies g(a, b) = \frac{a + a^2 - b^2 + 2ab^2 + b^4}{-2(a + b^2)} \tag{34}$$

If we want this to be zero, we will need a and b such that everything in the numerator drops. Once again, this is algebraically rather involved, but you will eventually end up with the following values for b as a function of a:

$$b(a) = \sqrt{\frac{1}{2}\left(1 - 2a \pm \sqrt{1 - 8a}\right)} \tag{35}$$

Note that there are two values here, in keeping with the fact that $g(a, b)$ has higher powers of both a and b. This means that there will be two Hopf bifurcations. Indeed, if you plot the solutions for the Sel'kov model for a fixed value of a and varying values of b, you will see your solutions tend towards a stable node to start, followed by a stable limit cycle, then back to a stable node. While the limit cycle exists, there will be an unstable node inside it (at the same location that the stable node would be otherwise). However, this unstable node will be essentially impossible to hit if you're doing numerical simulations, because the nature of floating-point arithmetic and finite step sizes means that you will always be making slight perturbations away from any unstable fixed point in a system that you're simulating.