Math 360-2: Applied Analysis
Northwestern University, Lecture Notes

Written by Santiago Cañez

These are notes which provide a basic summary of each lecture for Math 360-2, the second quarter of "MENU: Applied Analysis", taught by the author at Northwestern University. The book used as a reference is the 2nd edition of *Differential Equations: A Modeling Perspective* by Borrelli and Coleman. Watch out for typos! Comments and suggestions are welcome.

Contents

Lecture 1: Nonlinear Systems

This quarter we continue our study of differential equations, picking up where we left off. The first chunk of the course will continue to focus on systems of ODEs, where we now move towards studying nonlinear systems (some key notions include stability and chaos), and the rest of the course will focus on *partial differential equations* (PDEs). The study of PDEs will make good use of various other concepts such as Fourier series, eigenfunction expansions, and Green's Functions, which we will develop as needed.

Reactions. As stated above, the main shift is now to studying nonlinear systems, so we begin with a basic scenario where such systems arise. (We saw one example of a nonlinear system last quarter in the Lotka-Volterra predator-prey model, and indeed the scenario we will now consider is born out of similar assumptions.)

Suppose we have three chemicals X, Y, Z undergoing some chemical reaction. Denote the concentrations of each respectively by $x(t), y(t), z(t)$. One simple case is where X undergoes some reaction to produce Y, and then Y undergoes some reaction to produce Z:

$$X \to Y \to Z.$$

We assume that the rate of change in each concentration is proportional to the concentration itself, so that in this case we get a system of linear equations:

$$x' = -k_1 x \qquad y' = k_1 x - k_2 y \qquad z' = k_2 y$$

To be clear, $k_1 x$ is the rate at which X is lost, which then becomes the rate at which Y is introduced, but then Y is lost at a rate of $k_2 y$, which is then the rate at which Z is introduced. This linear system is something we now how to solve using eigenvalues and eigenvectors.

But in an actual chemical reaction, there could be other interactions taking place: maybe X and Y react together to produce Z, or to produce some new amount of X, or some new amounts of Y and Z for instance. The *Chemical Law of Mass Action* states that the rates at which concentrations change are proportional to the *products* of the individual concentrations which are interacting. For instance, if one X molecule reacts with one Y molecule, then there should be a factor of the form $-axy$ in each of the expressions for x' and y' due to "loss". Consider the reaction given by:

$$2X + Y \to 5Y + Z$$

where two X molecules react with one Y molecule to produce five Y molecules and one Z molecule. The concentrations we are interacting are $x(t), x(t), y(t)$, where we have $x(t)$ twice since it is *two* X molecules which are used in the reaction. The various rate equations should then have terms proportional to $x^2 y$. (If we had two X and two Y molecules reacting, we would use $x^2 y^2$.) So, we for instance the following equation for X:

$$x' = -2k_1 x^2 y.$$

Here, X is only being used up (i.e. lost) with no new X introduced, hence the negative, and the factor of two comes from the two X's involved in the reaction: each instance of the X concentration is lost at a rate of $k_1 x^2 y$, so with two instances the overall loss is $2k_1 x^2 y$. The concentration of Y is lost at a rate of $k_1 x^2 y$, but then five new instances of Y are introduced, so we get:

$$y' = \underbrace{-k_1 x^2 y}_{\text{out}} + \underbrace{5k_1 x^2 y}_{\text{in}} = 4k_1 x^2 y.$$

This makes sense, since overall there is a net gain of four instances of Y. Finally, we have

$$z' = k_1 x^2 y$$

since Z is only introduced and now actually used up in the reaction itself.

These nonlinear effects can also be combined with some linear ones, say if X underwent some reaction to become Y, but at the same time X reacted with Z to become something else. Note that the same nonlinearity showed up in the predator-prey model last quarter, and for the same reason: *Volterra's law* then stated at the effect which interactions between the two species have on their populations was proportional to the product of the two populations, which is also what the law for reactions above states if we considering two chemicals reacting as two "species" interacting. Proportionality to products of individual "things" reacting/interacting is a general assumption used in a wide range of phenomena.

Existence & Uniqueness. Now, the resulting nonlinear system we derived above cannot be solved explicitly, but, as we will see, we nonetheless have good ways of deriving much qualitative information about the behavior of solutions. As a first step, we of course have a basic existence and uniqueness theorem:

> If f and $\frac{\partial f}{\partial x_i}$ are continuous on some region in $\mathbb{R}^n$ including $(t_0, \mathbf{x}_0)$, then the IVP $\mathbf{x}' = f(\mathbf{x}, t), \mathbf{x}(t_0) = \mathbf{x}_0$ has a unique solution defined on some interval around t_0.

To be clear, here f is a function of n variables with n outputs (so something like $f : \mathbb{R}^n \to \mathbb{R}^n$ or defined possibly only on some subset of $\mathbb{R}^n$) and $\mathbf{x}(t)$ is a vector function in $\mathbb{R}^n$, so $\mathbf{x}' = f(\mathbf{x}, t)$ is a system of n ODEs. The notation $\frac{\partial f}{\partial x_i}$ denotes the partial derivative of f with respect to x_i, which, since f is vector-valued with multiple components, denotes the vector obtained by differentiating each component of f with respect to x_i:

$$\frac{\partial f}{\partial x_i} = (\frac{\partial f_1}{\partial x_i}, \dots, \frac{\partial f_n}{\partial x_i}).$$

All together then, the requirement that $\frac{\partial f}{\partial x_i}$ be continuous is really a statement about n^2 many functions: n component functions f_j each differentiated with respect to n variables x_i.

Understanding nonlinear solutions. In certain nice cases we can use what we know about single ODEs to understand solutions of nonlinear systems. For instance, consider the system

$$x_1' = x_1^3 \qquad x_2' = -x_2^3.$$

Even though this is written as a system of two equations, this system is *uncoupled*, meaning that actually x_1, x_2 have no effect on each other. We can treat each equation simply as a single ODE, which in this case can be solved using separation:

$$x_1' = x_1^3 \rightsquigarrow \frac{x_1'}{x_1^3} = 1 \rightsquigarrow \left(-\frac{1}{2x_1^2}\right)' = 1 \rightsquigarrow -\frac{1}{2x_1^2} = t + C \text{ and so on,}$$

and similarly for x_2.

But even *coupled* equations, some techniques from last quarter can come in handy. For instance, consider now the system

$$x_1' = x_2^3 \qquad x_2' = -x_1^3.$$

These are coupled, but if we consider the fraction

$$\frac{x_2'}{x_1'} = -\frac{x_1^3}{x_2^3},$$

by recognizing the left side as $\frac{dx_2}{dx_1}$ (implicitly assuming that x_2 can be expressed as a function of x_1), we can turn out system into the single ODE

$$\frac{dx_2}{dx_1} = -\frac{x_1^3}{x_2^3},$$

which is separable. Solving this for x_2 in terms of x_1 gives

$$-x_2^3\frac{dx_2}{dx_1} = x_1^3 \rightsquigarrow \frac{d}{dx_1}\left(-\frac{1}{4}x_2^4\right) = x_1^3 \rightsquigarrow -\frac{1}{4}x_2^4 = \frac{1}{4}x_1^4 + C \text{ and so on.}$$

Now, we do not end up with explicit formulas for $x_1(t), x_2(t)$ here, but rather with an equation which at least tells us how x_1, x_2 relate to one another. If nothing else, we know now that the *orbits* of the original system—which are the curves traced out by $x_1(t), x_2(t)$ in the x_1x_2-plane—lie on curves with equations $x_1^4 + x_2^4 = D$ for constant D.

Not all nonlinear systems can be studied in this way, via some conversion to a single ODE, so in general we will rely on understanding equilibrium points and the qualitative behavior of orbits. We will go into more detail next time.

Lecture 2: Orbits and Nullclines

Warm-Up. We derive a nonlinear system which models the chemical reaction for X and Y given by

$$X \to Y \text{ and } 2X + Y \to X.$$

So, some chemical X undergoes some reaction to become Y, while at the same time two instances of X and one instance of Y react to become a new instance of X. First, the conversion of X to Y in the first portion contributes k_1x terms to x' and y'—subtracted for x' and added for y'. Now, the nonlinear interactions give terms proportional to x^2y since two instances of X are reacting with one instance of Y. This occurs for each instance of X being used, so two k_2x^2y terms are subtracted, but then a new k_2x^2y is added since one instance of X is created, so overall our X equation is:

$$x' = -k_1x - 2k_2x^2y + k_2x^2y = -k_1x - k_2x^2y$$

In other words, there is a net "loss" of one instance of X due to the interaction between X and Y. The Y concentration is only affected by the new Y coming from $X \to Y$ and the Y lost in $2X + Y \to X$, so we get

$$y' = k_1x - k_2x^2y.$$

Thus the system modeling this reaction is

$$\begin{aligned} x' &= -k_1x - k_2x^2y \\ y' &= k_1x - k_2x^2y \end{aligned}$$

for positive constants k_1, k_2.

Autonomous systems. We will focus for now on *autonomous* systems $\mathbf{x}' = f(\mathbf{x})$ which have no explicit dependence on t. As we saw in the single-equation case last quarter, for such systems we can talk about *equilibrium* (or constant) solutions, which come from points at which f is zero.

A key fact about autonomous systems is that orbits never intersect one another, and orbits which intersect themselves must be *cycles*, i.e. closed periodic orbits. Indeed, suppose $\mathbf{x}(t), \mathbf{y}(t)$ describe two orbits of $\mathbf{x}' = f(\mathbf{x})$ such that $\mathbf{x}(t_1) = \mathbf{y}(t_2)$, so that the orbits intersect one another. The fact that our system is autonomous implies that $\mathbf{y}(t - t_1 + t_2)$ is also a solution, but this solution has value at $t = t_1$ of

$$\mathbf{y}(t_1 - t_1 + t_2) = \mathbf{y}(t_2) = \mathbf{x}(t_1).$$

This says that $\mathbf{x}(t)$ and $\mathbf{y}(t - t_1 + t_2)$ satisfy the same IVP, so by uniqueness they must be the same function, meaning that $\mathbf{x}(t), \mathbf{y}(t)$ trace out the same orbit. Taking $\mathbf{y} = \mathbf{x}$, so that $\mathbf{x}(t_1) = \mathbf{y}(x_2)$ becomes $\mathbf{x}(t_1) = \mathbf{x}(t_2)$, covers the case where an orbit intersects itself, in which case we get

$$\mathbf{x}(t) = \mathbf{x}(t - t_1 + t_2),$$

which says that $\mathbf{x}(t)$ is periodic and forms a cycle, or closed orbit.

Nullclines. Apart from identifying equilibrium points, the other key tool in understanding the nature of orbits of nonlinear systems is the use of *nullclines*, which are the curves where at least one of the components of $\mathbf{x}'$ is zero. (A solution of $\mathbf{x}' = f(\mathbf{x})$ occurs when *all* components are zero.) Nullclines describe points at which a component of $\mathbf{x}'$ achieves a maximum or a minimum, and divide our space into regions where a component of $\mathbf{x}'$ maintains the same sign throughout. This is simplest to explain and clarify via examples.

Example. Consider the system derived in the Warm-Up, with all constants of proportionality set equal to 1:

$$\begin{aligned} x' &= -x - x^2 y \\ y' &= x - x^2 y. \end{aligned}$$

Both x', y' are zero simultaneously whenever $x = 0$, so we get an equilibrium at each point along the y-axis. (This makes sense in the context of the chemical reaction of the Warm-Up: with no X present, nothing will happen regardless of Y, since no new Y is "created" and no Y is "lost" since there is no X available for Y to react with.) These are the only equilibria since for $x \neq 0$, $x' = -x - x^2 y = 0$ becomes $-1 - xy = 0$, while $y' = x - x^2 y = 0$ becomes $1 - xy = 0$, and these equations have no common solutions.

But, the equations $-1 - xy = 0$ and $1 - xy = 0$ (in addition to $x = 0$) do describe the *nullclines* of this system: the x-nullcline is given by $x' = -x - x^2 y = 0$, which excluding the equilibria along $x = 0$ simplifies to $-1 - xy = 0$, and the y-nullcline $y' = x - x^2 y = 0$ simplifies to $1 - xy = 0$. The resulting curves $-1 = xy$ and $1 = xy$ (hyperbolas in this case) describe points at which an orbit will have a vertical ($x' = 0$) or horizontal tangent ($y' = 0$) respectively, so that the orbit possibly "changes direction" at such a point:

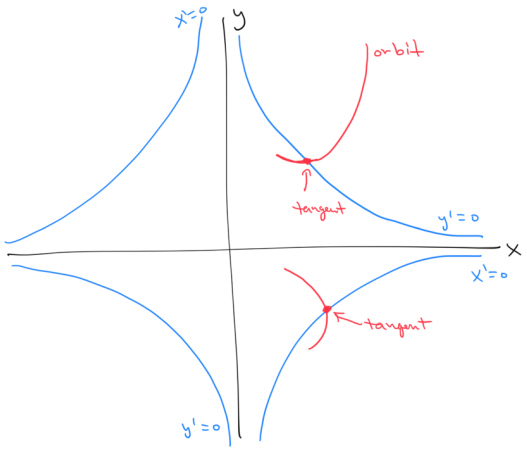

More importantly, the nullclines divide the plane here into regions of key interest. Consider for instance a point in the first-quadrant lying above the hyperbola $xy = 1$. At such a point, $x' = -x - x^2y$ is negative and $y' = x - x^2y$ is also negative. To be clear in the case of y', imagine taking a fixed x and increasing y so that our point moves up: eventually y will be large enough to make $x - x^2y$ negative, and so we must then have $y' < 0$ all throughout this region. (This is an application of the *Intermediate Value Theorem* from calculus/analysis: if there were a point in this region at which y' was negative and another at which y' was positive, then y would have to be zero at some point "in between", which it is not since this region lies above the nullcline $y' = 0$.)

As we pass through the nullcline $y' = 0$, the sign of y' changes so that y' is positive all throughout the portion of the first quadrant lying below the hyperbola $xy = 1$. By checking the signs of x', y' at some specific points, we can thus determine the signs they will each have throughout the regions of the xy-plane cutout by the nullclines:

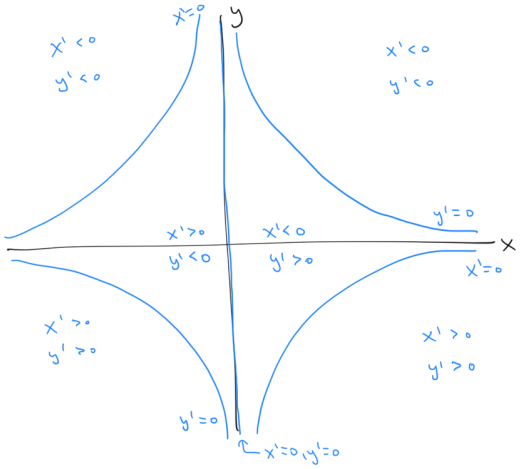

These signs tells us what direction the orbits move in, with the sign of x' telling us whether an orbit moves to the left or to the right, and y' telling us whether an orbit moves up or down:

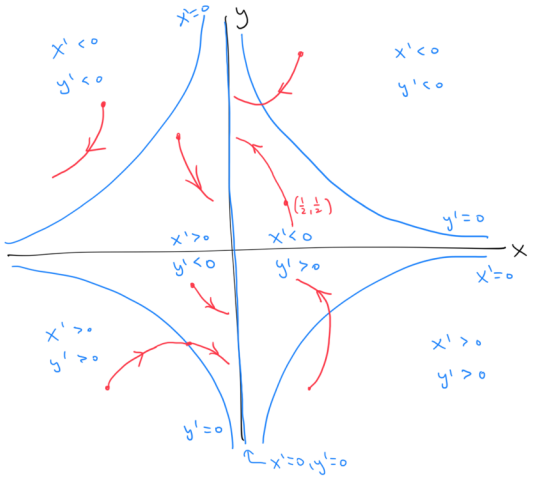

So, for instance, an orbit with initial position at $(1/2, 1/2)$ will move up and to the left, heading towards an equilibrium point on the y-axis as $t \to \infty$. In the context of the chemical reaction in the Warm-Up, this corresponds to a concentration of X which is dying out while the concentration of Y stabilizes; this makes sense since in that reaction X overall was "used up" more than it was "created" (three X "used" overall with only one "created"), while Y is "used up" just as much as it is "created". So, the nullclines do give a fair amount of qualitative information.

Another example. Consider the system

$$x' = x - y$$
$$y' = x - y^2.$$

This has two equilibria: $(0, 0)$ and $(1, 1)$. The x-nullcline is $x = y$ and the y-nullcline is $x = y^2$. By testing the signs of $x - y$ and $x - y^2$ at various points (in particular ones with large/small values of x, y) we get the following picture of the behavior of orbits:

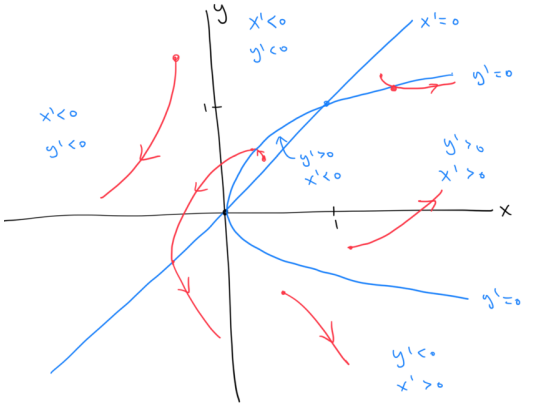

Note the interesting behavior of an orbit with initial point in the region enclosed by the line $y = x$ and parabola $x = y^2$ where $y' < 0$ and $x' > 0$: this initially moves up and left, hits the parabola at a horizontal tangent, moves down and left, hits the line at a vertical tangent, and finally continues to move down and right as $t \to \infty$.

Polar conversions. Apart from using nullclines or other means to determine some qualitative behavior, one might still hope to be able to find explicit solutions of nonlinear systems. As stated

before, there is no general method which will work here, although some ideas work in some cases. We already mentioned the idea of describing orbits (so not really explicit solutions) by turning a system into a single ODE for y in terms of x for instance, or the fact that for a decoupled system we can solve for each desired function separately using techniques from last quarter.

But here is a new technique for other nonlinear systems. Motivated by the fact that we have alternate coordinate-systems we can use for describing points in the xy-plane, we look at converting a given system into one expressed using *polar coordinates*. For $x = r\cos\theta, y = r\sin\theta$, taking derivatives with respect to t gives:

$$\begin{aligned} x' &= r'\cos\theta - r(\sin\theta)\theta' \\ y' &= r'\sin\theta + r(\cos\theta)\theta'. \end{aligned}$$

To be clear, r and θ here are still dependent on t since they vary just as x, y do, so differentiating something like $\cos\theta$ requires the chain rule to get $(-\sin\theta)\theta'$. We can write this more compactly using matrix notation:

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} \cos\theta & -r\sin\theta \\ \sin\theta & r\cos\theta \end{bmatrix} \begin{bmatrix} r' \\ \theta' \end{bmatrix},$$

and then use an inverse to express r', θ' in terms of x', y':

$$\begin{bmatrix} r' \\ \theta' \end{bmatrix} = \begin{bmatrix} \cos\theta & -r\sin\theta \\ \sin\theta & r\cos\theta \end{bmatrix}^{-1} \begin{bmatrix} x' \\ y' \end{bmatrix} = \frac{1}{r} \begin{bmatrix} r\cos\theta & r\sin\theta \\ -\sin\theta & \cos\theta \end{bmatrix} \begin{bmatrix} x' \\ y' \end{bmatrix}.$$

Thus, the autonomous system $x' = f(x, y), y' = g(x, y)$ becomes the following in terms of r, θ:

$$\begin{aligned} r' &= x'\cos\theta + y'\sin\theta = f(r\cos\theta, r\sin\theta)\cos\theta + g(r\cos\theta, r\sin\theta)\sin\theta \\ \theta' &= -x'\tfrac{\sin\theta}{r} + y'\tfrac{\cos\theta}{r} = -f(r\cos\theta, r\sin\theta)\tfrac{\sin\theta}{r} + g(r\cos\theta, r\sin\theta)\tfrac{\cos\theta}{r}. \end{aligned}$$

Here, the $f(x, y)$ and $g(x, y)$ functions in the end are simply expressed in terms of r, θ using $x = r\cos\theta$ and $y = r\sin\theta$. Converting to polar coordinates in this way will not help to understand all nonlinear systems, but it will definitely help for some.

Polar example. Consider the system

$$\begin{aligned} x' &= -y(x^2 + y^2) \\ y' &= x(x^2 + y^2). \end{aligned}$$

In polar coordinates, $f(x, y) = -y(x^2 + y^2)$ becomes $f(r\cos\theta, r\sin\theta) = -r^3\sin\theta$ and $g(x, y) = x(x^2 + y^2)$ becomes $g(r\cos\theta, r\sin\theta) = r^3\cos\theta$. Thus, the polar equivalent of our system is:

$$\begin{aligned} r' &= -r^3\sin\theta\cos\theta + r^3\cos\theta\sin\theta = 0 \\ \theta' &= -(-r^3\sin\theta)\tfrac{\sin\theta}{r} + r^3\cos\theta\tfrac{\cos\theta}{r} = r^2. \end{aligned}$$

This is now much simpler to solve: we get that $r = r_0$ is constant from the first equation, and then from the second $\theta' = r_0^2$ that $\theta = r_0^2 t + \theta_0$ for some constant θ_0. The equations

$$r(t) = r_0 \qquad \theta(t) = r_0^2 t + \theta_0$$

parametrize circles centered at the origin in polar coordinates, which thus form the orbits of our system. In this case, solving the given system explicitly is not-so-straightforward in terms of x and y, but becomes more straightforward in polar coordinates.

Lecture 3: Interacting Species

Warm-Up. We determine the long-term behavior of the orbits of the nonlinear system

$$x' = x(1 - \sqrt{x^2 + y^2}) - y$$
$$y' = y(1 - \sqrt{x^2 + y^2}) + x.$$

This system is not simple to analyze as-is, but converting to polar coordinates produces the following decoupled system:

$$r' = [r\cos\theta(1-r) - r\sin\theta]\cos\theta + [r\sin\theta(1-r) + r\cos\theta]\sin\theta = r(1-r)$$
$$\theta' = -[r\cos\theta(1-r) - r\sin\theta]\tfrac{\sin\theta}{r} + [r\sin\theta(1-r) + r\cos\theta]\tfrac{\cos\theta}{r} = 1.$$

Thus for θ we get $\theta(t) = t + \theta_0$, where θ_0 is constant, which indicates counterclockwise motion.

The first-order nonlinear equation $r' = r(1-r)$ for r is a logistic equation which can be solved using separation, but this is not needed if all we seek is to determine qualitative behavior. This has equilibrium solutions $r = 0$ and $r = 1$, which correspond to the origin and unit circle respectively. (To be clear, $r = 1$ is an equilibrium point of the ODE $r' = r(1-r)$, but it does NOT correspond to an "equilibrium point" of the original xy-system since it is not a constant solution. In this case, the unit circle $r = 1$ would be called a "limit cycle" of the system, which is a concept we'll talk about later on.) Now, based on the type of analysis we did last quarter, we can see that $r \to 1$ from above for $r(0) > 1$ and from below for $0 < r(0) < 1$:

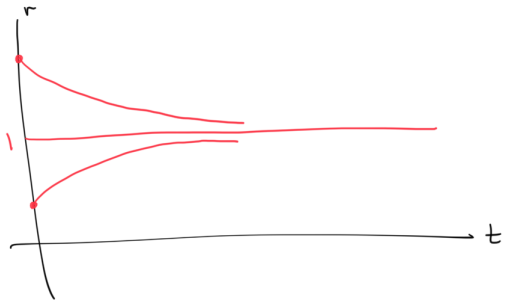

Thus, for $r(0) > 1$ we get orbits which are counter-clockwise moving spirals which are converging towards the unit circle from the outside, while for $0 < r(0) < 1$ we get spirals approaching the unit circle from the inside:

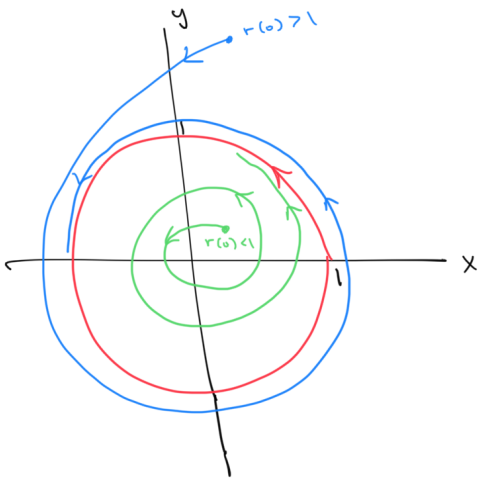

For $r(0) = 1$, the orbit is just the unit circle itself, while for $r(0) = 0$ the orbit is the origin.

Interactions. Recall the Lotka-Volterra predator-prey model:

$$x' = -ax + bxy = (-a + by)x$$
$$y' = cy - kxy = (c - kx)y$$

where $a, b, c, k > 0$. Here, x corresponded to the predator species and y the prey, and the "interaction" terms came from the assumption that the effect these interactions had on each species was proportional to the product of the two populations which interact, which as we've said is similar-in-spirit to the assumption modeling chemical reactions. In this case, interactions have a positive effect on the predator species (hence the $+bxy$ term) but a negative effect ($-kxy$) on the prey.

But we can also consider more general types of interactions, say ones modeling cooperation or other forms of competition, by allowing other expressions. For instance, possibly the "proportionality constants" should not be constants after all, but might themselves depend on the present population. One example of this is due to overcrowding, where an increasing population has a more and more detrimental effect on the popluation going forward; in the case of logistic equations, overcrowding was modeled by a quadratic term, not simply a linear one.

So, more generally we will consider autonomous systems of the form

$$x' = R_1(x, y)x$$
$$y' = R_2(x, y)y,$$

where R_1, R_2 are used to describe different types of effects affecting the "proportionality" to x and$y$ respectively. The Lotka-Volterra case of $R_1 = -a + by, R_2 = c - kx$ models linear death and birth rates for the predator and prey species respectively, and interactions with affect predator and prey positive and negatively respectively. Overcrowding can then be incorporated with quadratics x^2, y^2—or linear terms in R_1 and R_2 which become quadratics after multiplying through by x or y as in $R_1(x, y)x$ and $R_2(x, y)y$—so something like

$$x' = (-a - dx + by)x$$
$$y' = (c - kx - fy)y$$

models predator-prey interactions with overcrowding, where the positive constants d, f describe the effect overcrowding has on each population. Harvesting/restocking rates which are proportional to the populations can be incorporated with additional constants:

$$x' = (-a - dx + by - H_1)x$$
$$y' = (c - kx - fy + H_2)y,$$

so in this case x is being harvested and y restocked, or with H_1, H_2 outside the coefficients of X, y we get constant-rate harvesting/restocking

$$x' = (-a - dx + by)x - H_1$$
$$y' = (c - kx - fy)y + H_2.$$

(So not in the form $x' = R_1(x, y)x, y' = R_2(x, y)y$ in this case.)

Cooperation, where interactions between the two species have a positive effect on *both*, is obtained by mixed xy terms which are being added in both cases, and *competition*, where interactions

affect both species negatively, by xy terms which are subtracted. One final type of effect, which we will not really consider in any depth going forward but is worth pointing out, is that of *satiation*, where the benefit which, say, a prey species provides to a predator species can only get so high, and *decreases* as the populations increase. For instance, consider something like

$$x' = -x + \left(\frac{1}{1+y}\right)xy$$
$$y' = y - \left(\frac{1}{1+y}\right)xy.$$

Here, the measure of the positive effect $\frac{1}{1+y}$ which interactions have on the predator species x gets smaller as y increases, and correspondingly the negative effect which interactions have on the prey species gets smaller; this makes sense if for instance the predator species can only consume so much prey, so that having more prey available does not provide that much more of a benefit. The term *satiation* comes from the fact that the predator's hunger is satiated.

Positive populations. If we modeling actual populations, we really only care about what happens for x, y in the first quadrant where x, y are both positive. A basic fact is that if $x(0), y(0)$ are initially positive, then $x(t), y(t)$ remain positive for all time. This comes from the fact that distinct population orbits cannot intersect, so that no orbits in the first quadrant can ever cross an axis.

Example 1. Consider the system

$$x' = (5 - x + y)x$$
$$y' = (10 + x - 5y)y,$$

which models cooperation (positive xy terms) between two species with overcrowding. The positive equilibrium occurs at $(8.75, 3.75)$, which comes from solving

$$5 - x + y = 0$$
$$10 + x - 5 = 0.$$

The nullclines are $y = x - 5$ and $y = \frac{1}{5}x + 2$, which give the following orbit behavior:

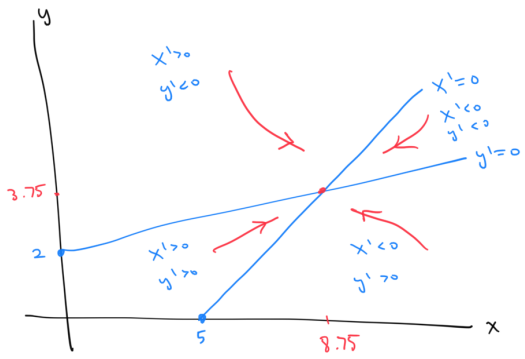

In this case, all orbits seem to tends towards the positive equilibrium, meaning that regardless of the initial populations, the populations in the long-term will stabilize towards $(8.75, 3.75)$. This behavior can be confirmed more carefully using a computer to plot the direction field—i.e. the field of tangent vectors defined by (x', y')—and some varying orbits:

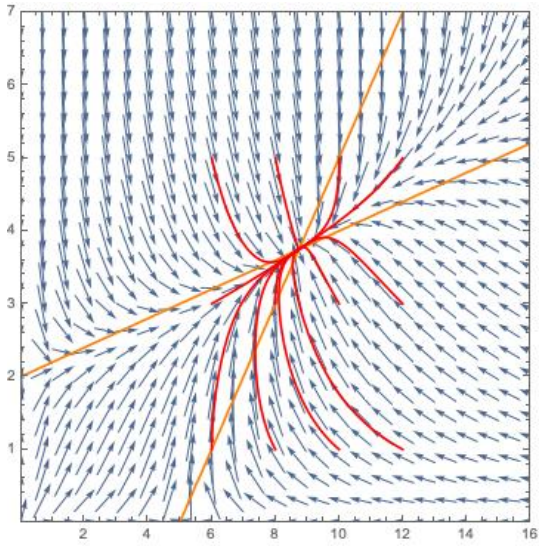

In this case, even with overcrowding, the mutual cooperation leads to controlled behavior.

Example 2. Consider now the system

$$x' = (10 - x + 5y)x$$
$$y' = (5 + x - y)y$$

There is an equilibrium point at $(-8.75, -3.75)$, which is outside the first quadrant where populations are realistic. The nullclines are $y = \frac{1}{5}x - 2$ and $y = x + 5$, and give the following behaviors:

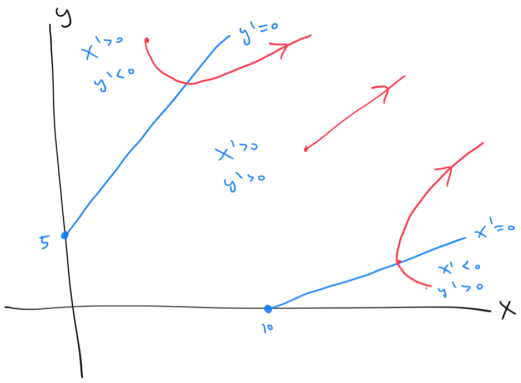

Here, populations do not seem to stabilize and appear to grow without restriction. In other words, the cooperation seems to greatly outweigh the overcrowding, so that both species grow in an uncontrolled manner. This can be confirmed with a more careful plot:

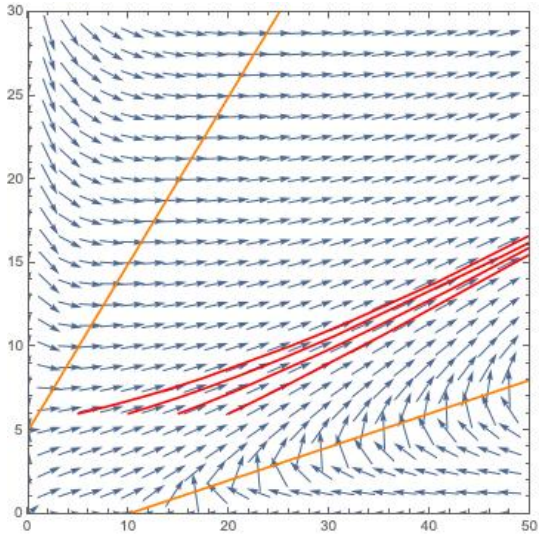

Moreover, the y-population appears to grow much more slowly than the x-population.

Lecture 4: Linear Stability

Warm-Up. We analyze the populations modeled by the system

$$x' = (1 - 2x - 4y)x$$
$$y' = (1 - 4x - 2y)y.$$

This describes two competing populations with overcrowding. There is an equilibrium point at $(\frac{1}{6}, \frac{1}{6})$ with nullclines $1 - 2x - 4y = 0, 1 - 4x - 2y = 0$ in the first-quadrant. By checking the signs of x' and y', we get the following orbital picture:

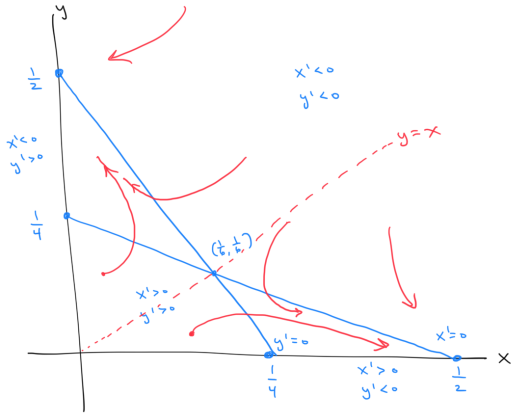

The key observation here is that some orbits tends towards one equilibrium at $(\frac{1}{2}, 0)$, but others tend towards the equilibrium at $(0, \frac{1}{2})$. In particular, the behavior is different below $y = x$ versus above. This is termed *competitive exclusion*, where competition benefits one species and harms the other, but which is which depends on the initial sizes.

Note that the given system remains the same if we swap x and y, which accounts for the symmetry across $y = x$ in the picture above: if $(x(t), y(t))$ is an orbit, then so is $(y(t), x(t))$. Moreover, for $x = y$ both equations for x' and y' become the same, which says that if $x(t)$ alone

satisfies the resulting equation for x', then $(x(t), x(t))$ is an orbit. This accounts for the orbits along $y = x$ which actually *stabilize* towards the equilibrium at $(\frac{1}{6}, \frac{1}{6})$. Deviating a bit away from this orbit produces the behaviors described previously.

Linear stability. A *linear* autonomous system is of the form

$$\mathbf{x}' = A\mathbf{x} + \mathbf{F}$$

where $\mathbf{F}$ is a constant vector and A a constant matrix. We studied such systems in detail last quarter, where we derived explicit solutions phrased in terms of the eigenvalues and eigenvectors of A, or more compactly in terms of the matrix exponential e^{At}.

For now, let us consider only the case $\mathbf{F} = \mathbf{0}$. (The behavior for nonzero $\mathbf{F}$ can be obtained by translation, as we will see.) Recall that in the 2-dimensional case we had the following types of behavior for the equilibrium at the origin, depending on the eigenvalues of A:

- stable improper or star node (two negative eigenvalues or one negative and non-deficient):

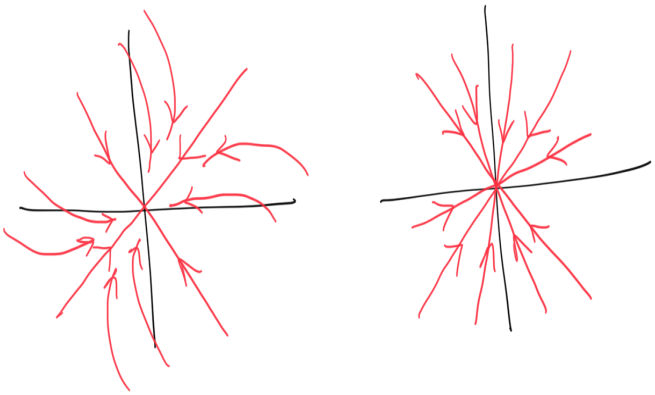

- unstable saddle (one positive, one negative eigenvalue):

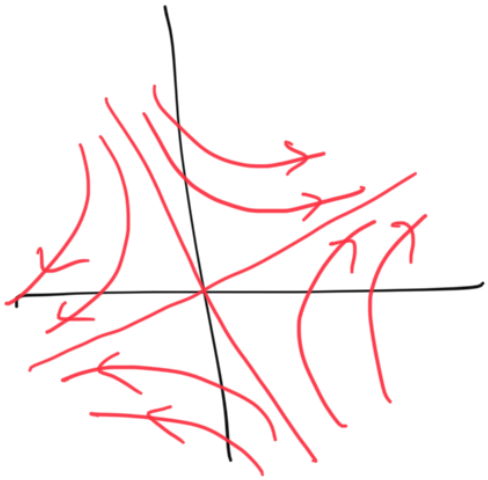

- unstable improper or star node (two positive eigenvalues or one positive and non-deficient):

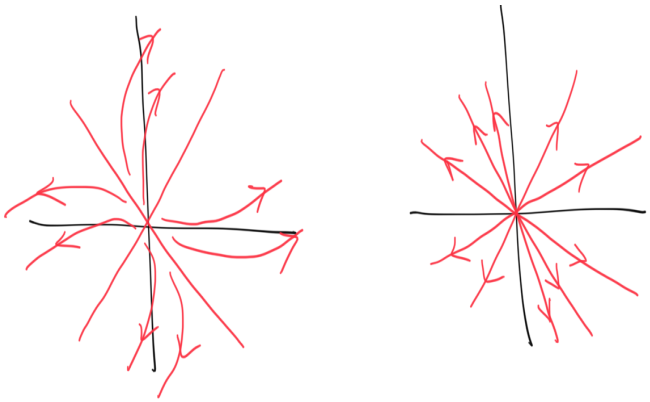

- stable deficient node (one negative deficient eigenvalue):

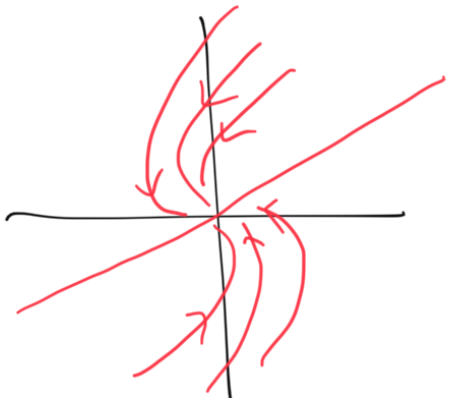

- unstable deficient node (one positive deficient eigenvalue):

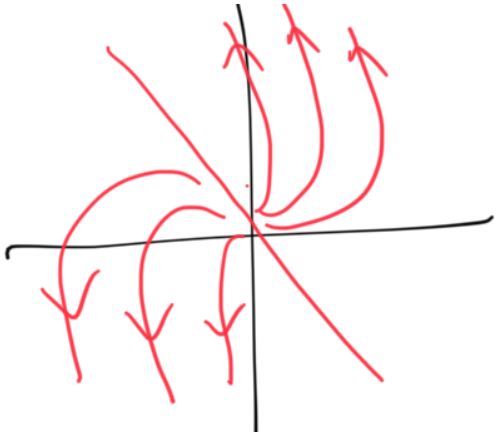

- stable spiral (negative real part non-real eigenvalues):

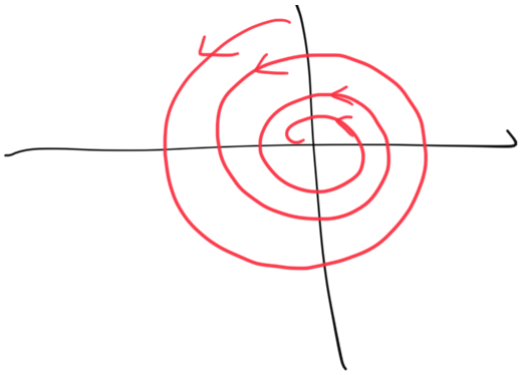

- unstable spiral (positive real part non-real eigenvalues):

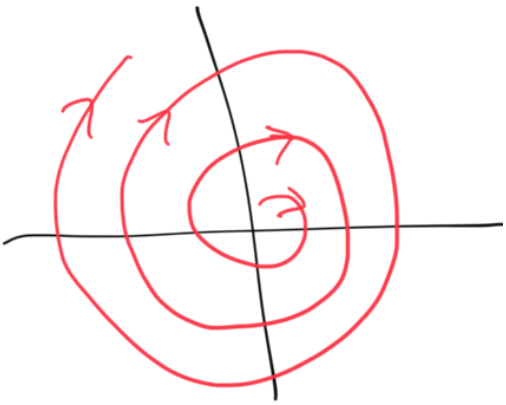

- neutrally stable center (imaginary eigenvalues):

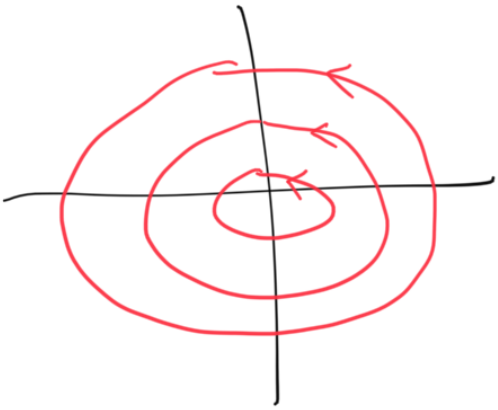

(The lines in the first few pictures come from the eigenvector directions.) These characterizations came from the types of exponentials $e^{\lambda t}$ used in the explicit solutions, or more precisely the possible signs of λ and whether such exponentials were present at all. Here, "stable" (or "asymptotically stable") refers to the fact that orbits *approach* the equilibrium as $t \to \infty$, and "unstable" the fact that orbits move away instead. (Later we will give more precise definitions of these terms, so that

they apply to nonlinear systems as well.) The "imaginary eigenvalues" case is where we get complex solutions of the form

$$e^{ibt}(\mathbf{v} + i\mathbf{w}) = [(\cos bt)\mathbf{v} - (\sin bt)\mathbf{w}] + i[(\cos bt)\mathbf{w} + (\sin bt)\mathbf{v}]$$

without any real exponentials, and it is this lack of real exponentials which causes the "neutrally stable" behavior; the presence of a real exponential here causes the "spiraling" effect.

The same types of behaviors occur for higher-dimensional linear autonomous systems, with one wrinkle we will soon clarify which is not noticeable in the 2-dimensional case. Here is our theorem on linear stability:

Consider the system $\mathbf{x}' = A\mathbf{x}$. If A has:

- all eigenvalues with negative real part, the origin is an asymptotically stable equilibrium point;
- at least one eigenvalue with a positive real part OR all imaginary eigenvalues at least one of which is deficient, the origin is an unstable equilibrium point; or
- all imaginary eigenvalues, none of which are deficient, the origin is a neutrally stable equilibrium point.

We will clarify the "imaginary eigenvalues at least one of which is deficient" criteria in a bit. Note that the case where A has purely real eigenvalues is accounted for here is as well, since in that case the eigenvalue is its own real part, so that either the first or second conditions above apply. As was the case last quarter, we are not treating the case of a zero eigenvalue.

Deficient eigenvalues. An eigenvalue of A is *deficient* if a basis for its eigenspace consisting of eigenvectors alone cannot be found, so that the use of generalized eigenvectors is required. There are numerous equivalent ways of saying this, for instance to borrow some terminology from linear algebra: the dimension of its eigenspace is strictly less than the multiplicity of the eigenvalue, where "multiplicity" is the number of times the eigenvalue occurs as a root of the characteristic polynomial. (For those of who have taken Math 334, we can also phrase this as saying that there is a Jordan block of size larger than 1 corresponding to that eigenvalue in the Jordan form.)

Recall that for a real deficient eigenvalue, in addition to solutions of the form $e^{\lambda t}\mathbf{v}$, we also looked for solutions of the form

$$e^{\lambda t}\mathbf{w} + te^{\lambda t}\mathbf{v}$$

when trying to form the general solution. In this case, the t term was outweighed by the $e^{\lambda t}$ term, so that the solution would approach the origin or blow up depending on the sign of λ. Deficient real eigenvalues in higher-dimensional cases lead to solutions containing terms of the form

$$e^{\lambda t}\mathbf{v}_0 + te^{\lambda t}\mathbf{v}_1 + \frac{1}{2}t^2e^{\lambda t}\mathbf{v}_2 + \cdots + \frac{1}{k!}t^ke^{\lambda t}\mathbf{v}_k$$

where $\mathbf{v}_0, \mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_k$ form a *chain* of generalized eigenvectors. Still, even here, the behavior of the exponentials outweigh the behavior of the polynomial factors.

The difference in higher dimensions is that now deficient *imaginary* eigenvalues are possible! This was not the case for 2-dimensional systems: imaginary eigenvalues for $n = 2$ always lead to two linearly independent *real* solutions, coming from the real and imaginary parts of a complex solution. For $n > 2$, it is possible to have a deficient imaginary eigenvalue e^{ibt}, which would require the use of generalized complex eigenvectors when forming a solution:

$$e^{ibt}\mathbf{v}_0 + te^{ibt}\mathbf{v}_1 + \text{higher-order terms in } t.$$

Such expressions will lead to terms of the form

$$t^k \cos bt, t^k \sin bt$$

with $k > 0$ in the explicit solution, and such things will blow up as $t \to \infty$. Thus, the case where A has all imaginary eigenvalues with at least one being deficient leads to similar behavior as the case where it has an eigenvalue with a positive real part, causing unstable behavior. This case is not one we will consider much, but it is worth mentioning in order to get a complete picture.

Where to now? As mentioned previously, the case of $\mathbf{x}' = A\mathbf{x} + \mathbf{F}$ where $\mathbf{F} \neq \mathbf{0}$ is covered by the analysis above after *translating*: the behavior of the equilibrium for $\mathbf{x}' = A\mathbf{x} + \mathbf{F}$ is obtained by translating the behavior of the equilibrium for $\mathbf{x}' = A\mathbf{x}$. So, the case of linear stability covers non-origin equilibrium points as well. The upshot, as we will see over the next two days, is that in many cases, linear stability determines nonlinear stability as well, coming from the fact that can *approximate* nonlinear systems using linear ones. So, it was important to get a good handle on the linear case, since it will tell us most of what we want to know in general.

Lecture 5: Nonlinear Stability

Warm-Up. We determine the nature of the equilibrium at the origin for each of the following systems:

$$\mathbf{x}' = \begin{bmatrix} 3 & -1 \\ 2 & -2 \end{bmatrix} \mathbf{x} \qquad \mathbf{x}' = \begin{bmatrix} -3 & -1 \\ 2 & -2 \end{bmatrix} \mathbf{x} \qquad \mathbf{x}' = \begin{bmatrix} 3 & 1 \\ 2 & 2 \end{bmatrix} \mathbf{x}$$

Now, we approach this in a way which avoids computing (at least initially) eigenvalues directly, but rather using the fact that information about the eigenvalues can be gleamed from the determinant and trace: the determinant of a matrix is the product of its eigenvalues, and the trace (which is the sum of diagonal terms) is the sum of its eigenvalues. This will make some of the analysis computationally simpler.

The matrix defining the first system has negative determinant, so since the determinant is the product of the eigenvalues, it must have one positive and one negative eigenvalue. Hence the origin is an unstable saddle equilibrium point of the first system.

The matrix defining the second system has positive determinant, so one possibility is for it to have real eigenvalues of the same sign. But in fact, positive determinant can also arise with complex eigenvalues: if $a \pm ib$ are nonzero complex eigenvalues, the determinant will be $(a + ib)(a - ib) = a^2 + b^2$, which is positive. Moreover, in this case the trace is $(a + ib) + (a - ib) = 2a$, which could be either negative or positive. Since this matrix has negative trace, it either has two negative eigenvalues or complex eigenvalues with negative real part. Either way, for sure the origin is a stable equilibrium point, but to distinguish between the different types of stable equilibria, we would have to know more about the eigenvalues. The characteristic polynomial is

$$\lambda^2 + 5\lambda + 8,$$

which has complex, non-real roots, and so the origin for this system is a stable spiral node.

The matrix defining the final system has positive determinant and positive trace, so it either has two positive real eigenvalues or complex eigenvalues with positive real part. In either case we get an unstable equilibrium point. Since the characteristic polynomial

$$\lambda^2 - 5\lambda + 4$$

has real distinct roots, the origin is thus a unstable improper node.

General stability. Now we take a step back and give proper definitions of some of the terms we introduced last time in the linear case, in order so that they apply to the nonlinear case as well. Ultimately, we want to capture the idea that "an orbit which is initially close to an equilibrium remains close" or not in a more formal way.

We will phrase this using the notion of a *ball* in $\mathbb{R}^n$: the ball of radius $r > 0$ centered at $\mathbf{p}$ is the set of all points whose distance to $\mathbf{p}$ is less than r. Visually, this looks like a disk in $\mathbb{R}^2$ or a literal ball in $\mathbb{R}^3$, where the usual notions of "radius" and "center" apply. In what follows, $\mathbf{x}(t)$ will denote the orbit with initial point $\mathbf{x}(0)$.

Suppose $\mathbf{p}$ is an equilibrium point of a system $\mathbf{x}' = f(\mathbf{x})$. We say that $\mathbf{p}$ is:

- *stable* if given any ball $B_\epsilon(\mathbf{p})$ around $\mathbf{p}$, there is a another ball $B_\delta(\mathbf{p})$ such that if $\mathbf{x}(0)$ is in $B_\delta(\mathbf{p})$, then $\mathbf{x}(t)$ is in $B_\epsilon(\mathbf{p})$ for all $t \geq 0$;
- *unstable* if it is not stable, meaning that there is some ball $B_\epsilon(\mathbf{p})$ around $\mathbf{p}$ so that no matter how close $\mathbf{x}(0)$ initially is to $\mathbf{p}$, eventually $\mathbf{x}(t)$ will not be in $B_\epsilon(\mathbf{p})$ for some $t \geq 0$.

Let us unpack these definitions. To be stable means that given some measure (as given by the first ball) for how close to $\mathbf{p}$ we want to remain for all time, there is a corresponding measure (given by the second ball) of how close to $\mathbf{p}$ we have to be initially in order to guarantee that happens. In other words, if we are "close" to $\mathbf{p}$ to begin with, we will remain "close" to $\mathbf{p}$ forever:

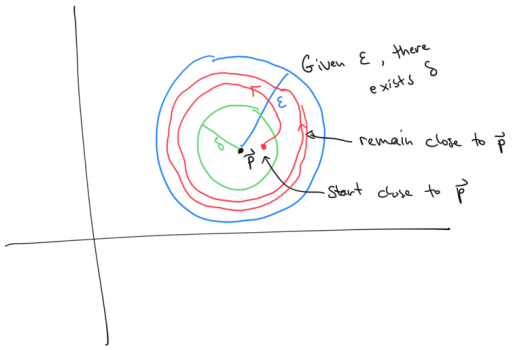

To be unstable means that there is some ball around $\mathbf{p}$ so that no matter who close to $\mathbf{p}$ we are initially, we will eventually jump outside of that ball. Intuitively, we will always move away from $\mathbf{p}$ eventually no matter close we are initially:

So, stable vs unstable has to do with the behavior of orbits near $\mathbf{p}$, in terms of whether they remain close or not. Now, a second consideration is whether orbits *approach* $\mathbf{p}$ in the long-term, which is not quite the same phenomena as being stable. We say that $\mathbf{p}$ is:

- an *attractor* if there is some ball $B_\epsilon(\mathbf{p})$ around $\mathbf{p}$ so that if $\mathbf{x}(0)$ is in that ball, then $\mathbf{x}(t)$ converges (i.e approaches) $\mathbf{p}$ as $t \to \infty$, meaning that $\|\mathbf{x}(t) - \mathbf{p}\| \to 0$ as $t \to \infty$ where $\|\cdot\|$ denotes the length of a vector, so that $\|\mathbf{x}(t) - \mathbf{p}\|$ is the distance between $\mathbf{x}(t)$ and $\mathbf{p}$;
- *asymptotically stable* if it is stable and an attractor;
- *neutrally stable* if it is stable but not an attractor;
- a *repeller* if it is an attractor as $t \to -\infty$ instead of positive ∞, so $\|\mathbf{x}(t) - \mathbf{p}\| \to 0$ as $t \to -\infty$.

Thus, an attractor is an equilibrium point for which orbits which are initially close to $\mathbf{p}$ will in the long-term approach $\mathbf{p}$, and a repeller is one where orbits which are initially close to $\mathbf{p}$ will approach $\mathbf{p}$ if we go *backwards in time*, so that the orbit in a sense emanates *from* $\mathbf{p}$ in this case:

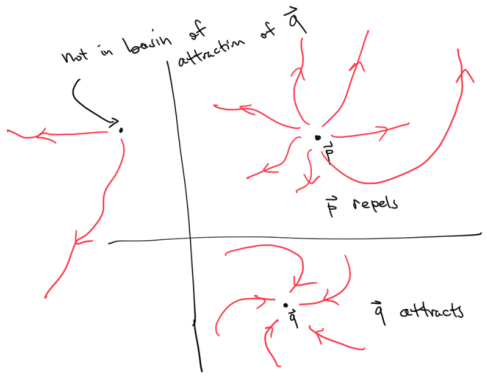

The distinction between asymptotically and neutrally stable depends on what happens long-term: in either case we remain "close" to $\mathbf{p}$, but might or might not actually head *towards* $\mathbf{p}$ as t increases:

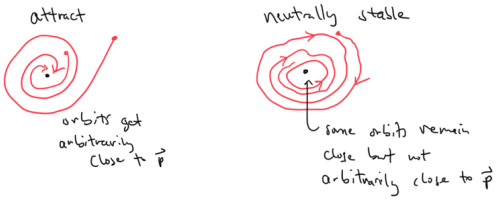

An attractor is only required to attract things which are initially close to it (as measured by the ball in the definition), and might not attract things which are sufficiently far away. The set of initial points $\mathbf{x}(0)$ which are close enough to $\mathbf{p}$ so that they determine orbits which do converge to $\mathbf{p}$ is called the *basin of attraction* of $\mathbf{p}$; if this basin is all of $\mathbb{R}^n$, we say the attraction is *global*, and if not it is only *local*. Similarly, in reverse time, we can speak of the *basin of repulsion* of a repelling equilibrium point, and whether it is a global repeller or only a local repeller.

The linear case. The formal definitions above do capture the expected behavior in the linear case. That is, the terms "stable", "asymptotically stable", "neutrally stable", and "unstable" we used last time there agree with the precise definitions here. We will not give proofs of this fact here, but you can check the book for some of the details. In the end it comes down to the fact that we have explicit solutions $e^{At}\mathbf{x}(0)$ in the linear case on whose norms (i.e. lengths) which we can find bounds of the form $Me^{\lambda t}\|\mathbf{x}(0)\|$ where λ bounds the eigenvalues of A in some way. Special to the linear case is that any attractor must be global, any repeller must be global, and there do not exist unstable attractors. (The existence of unstable attractors in general might seem counterintuitive,

but we'll look at an example in a bit which will make the difference between being "stable" and "attracting" clear.)

But perhaps it is worth clarifying the definition of "neutrally stable" in this setting. Consider a linear autonomous system whose defining matrix has imaginary eigenvalues, in which we get orbits which are shaped like ellipses:

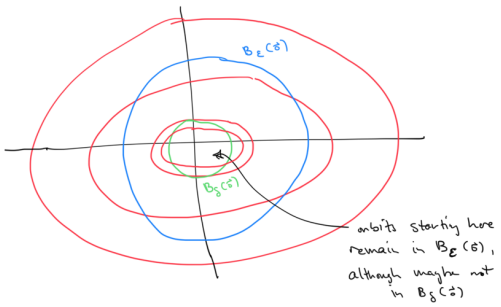

To be stable means that given a ball around $\mathbf{p}$ to start with, we can find another ball around $\mathbf{p}$ characterizing initial points which give rise to orbits which remain inside the first ball for all time. These do not have to be the same ball, as in the picture above, so that orbits which are initially close to $\mathbf{p}$ do not have to remain that *same* measure of closeness for all time, but still close as measured by something different instead. In this case, there is thus no need to have the orbit itself approach $\mathbf{p}$ in the long-term, as long as it does not stray too far away, which is what gives rise to the "neutrally stable" condition.

Unstable attractors. As stated above, a nonlinear system in general can have unstable attractors, which might seem strange at first. Understanding how this is possible helps to further clarify our definitions.

Consider the system

$$x' = x - y - x^3 - xy^2 + \frac{xy}{\sqrt{x^2+y^2}}$$
$$y' = x + y - x^2y - y^3 - \frac{x^2}{\sqrt{x^2+y^2}}$$

As you'll show on the homework, this system has an equilibrium point at $(1,0)$, and then has an orbit which consists of the *rest* of the unit circle oriented counterclockwise. The point here is that moving along this circle as $t \to \infty$ brings us eventually closer and closer to $(1,0)$, even though initially we might move *away* from $(1,0)$ for a short while:

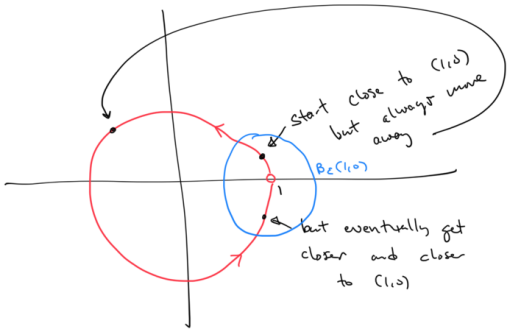

For the ball of initial points which is drawn above, no matter how close to $(1,0)$ we are initially on the unit circle, as time moves forward we will eventually jump outside of this ball, which says that $(1,0)$ is unstable: being close to $(1,0)$ initially does not force us to remain close to $(1,0)$.

But, $(1,0)$ is an attractor: if we continue along the unit circle we will in the long-term loop back around and converge towards $(1,0)$. Now, this alone does not say that $(1,0)$ is an attractor, since for this we have to know about the behavior of other orbits which are initially close to $(1,0)$, not just the one given by the rest of the unit circle, but it will in fact be true that all such orbits which start close $(1,0)$ will converge to $(1,0)$ as $t \to \infty$. Again, the point here is the difference between long-term behavior (attractor) and behavior for all time (stable vs unstable).

Example. Consider the system

$$\begin{aligned} x' &= -x(1+\sqrt{x^2+y^2}) - y \\ y' &= -y(1+\sqrt{x^2+y^2}) + x. \end{aligned}$$

In polar coordinates this becomes:

$$\begin{aligned} r' &= -r - r^2 \\ \theta' &= 1 \end{aligned}$$

Thus first equation gives $r' + r = -r^2 \leq 0$, which after multiplying through by e^t becomes

$$(r(t)e^t)' \leq 0 \text{ for all } t.$$

Hence $r(t)e^t$ is a non-increasing function, so $r(t)e^t \leq r(0)r^0$ for $t \geq 0$, and thus $r(t) \leq r(0)e^{-t}$. As $t \to \infty$, this implies $r(t) \to 0$, which says that the equilibrium at the origin should be an attractor.

In fact, the origin is also stable, so that it is asymptotically stable. Indeed, given any ball $B_\epsilon(\mathbf{0})$ around $\mathbf{0}$, the same ball $B_\delta(\mathbf{0}) = B_\epsilon(\mathbf{0})$ satisfies the necessary requirement on initial conditions given in the definition of stable: if $\mathbf{x}(0)$ is in this ball, so that $r(0) \leq \epsilon$, then for all time $t \geq 0$ we have

$$r(t) \leq r(0)e^{-t} \leq r(0) \leq \epsilon,$$

so that $\mathbf{x}(t)$ remains in $B_\epsilon(\mathbf{0})$ for all time. Moreover, the origin is an asymptotically stable spiral since $\theta(t) = t + \theta_0$ results in counterclockwise motion with decreasing radius $r(t)$ converging to 0.

Another example. For the system

$$\begin{aligned} x' &= x(1+\sqrt{x^2+y^2}) - y \\ y' &= y(1+\sqrt{x^2+y^2}) + x, \end{aligned}$$

the origin is an unstable spiral point. In this case the resulting polar equation for r is $r' = r(1+r)$, which gives

$$r' - r = r^2 \geq 0, \text{ and implies } (re^{-t})' \geq 0 \text{ for } t \geq 0.$$

Thus $r(t)e^{-t}$ is non-decreasing, so $r(t)e^{-t} \geq r(0)e^0 = r(0)$ and thus

$$r(t) \geq r(0)e^t \text{ for } t \geq 0.$$

The right side blows up as $t \to \infty$, which is what causes the unstable behavior. To be more precise, take the ball of radius 1 around $\mathbf{0}$, and any initial radius $r(0)$ however close to 0 you like—then $r(0)e^t$ will be larger than 1 once t is large enough, so that $\mathbf{x}(t)$ jumps outside of $B_1(\mathbf{0})$, which

shows that $\mathbf{0}$ is unstable, and it is a spiral point since $\theta = t + \theta_0$ still results in counterclockwise motion throughout time. (The origin is actually a repeller in this case, but justifying this requires a different approach since $r(0)e^t \to 0$ as $t \to -\infty$ does not say anything about the behavior of the larger value $r(t)$ as $t \to -\infty$.)

Linearization. In the two examples above we were able to determine the nature of the equilibrium at the origin by deriving explicit bounds on $r(t)$ in polar coordinates. But, this will not always be possible, and it might be even more challenging when considering a non-origin equilibrium. So, we need a way to characterize equilibria in a way which does not depend on such explicit bounds, or in other words, which avoids the direct use of the formal definitions of stable, unstable, etc.

The saving grace is an idea we briefly considered last quarter: approximating a nonlinear system by a linear one. In particular, we use the fact from multivariable calculus that for $\mathbf{x}$ near $\mathbf{p}$, the value of $f(\mathbf{x})$ is well-approximated by the linearization $f(\mathbf{p}) + Df(\mathbf{p})(\mathbf{x}-\mathbf{p})$ of f at $\mathbf{p}$, where $Df(\mathbf{p})$ denotes the *Jacobian matrix* of f at $\mathbf{p}$, which is the matrix of partial derivatives of f. If $\mathbf{p}$ is an equilibrium point, $f(\mathbf{p})$ is $\mathbf{0}$ so

$$f(\mathbf{x}) \approx f(\mathbf{p}) + Df(\mathbf{p})(\mathbf{x} - \mathbf{p}) = Df(\mathbf{p})(\mathbf{x} - \mathbf{p})$$

for $\mathbf{x}$ near $\mathbf{p}$. Thus, orbits of the nonlinear system $\mathbf{x}' = f(\mathbf{x})$ can be well-approximated by orbits of the *linear* system $\mathbf{x}' = Df(\mathbf{p})(\mathbf{x}-\mathbf{p})$, which we already know how to study quite well. (The use of $\mathbf{x} - \mathbf{p}$ instead of simply $\mathbf{x}$ in the linear system just has effect of translating the behavior to be centered at $\mathbf{p}$ instead of $\mathbf{0}$. Note that $\mathbf{p}$ is indeed a constant solution of this translated system, but not of $\mathbf{x}' = Df(\mathbf{p})\mathbf{x}$ alone.) The upshot is that in *most* situations, the nature of the equilibrium $\mathbf{p}$ for the linear system is the *same* as that for the nonlinear system, and comes down to looking at the eigenvalues of $Df(\mathbf{p})$. We'll clarify this more next time (in particular what "most" means here), but will finish now with one example.

Final example. Consider the following system, which models cooperation between two species with overcrowding:

$$\begin{aligned} x' &= (1 - x + y)x \\ y' &= (1 + \tfrac{1}{2}x - y)y. \end{aligned}$$

This has an equilibrium at $(4, 3)$—there are others which are not in the first quadrant—and a quick sketch using nullclines suggests that $(4, 3)$ should be an asymptotically stable (i.e. stable attractor) equilibrium. We can now confirm this using a Jacobian matrix.

This system is of the form $\mathbf{x}' = f(\mathbf{x})$ for the function $f : \mathbb{R}^2 \to \mathbb{R}^2$ defined by

$$f(x, y) = (x - x^2 + yx, y + \tfrac{1}{2}xy - y^2).$$

To be clear, the first component $f_1(x, y) = x - x^2 + yx$ comes from the equation for x' and the second component $f_2(x, y) = y + \frac{1}{2}xy - y^2$ comes from y'. The Jacobian matrix of this function at an arbitrary point is formed by taking the partial derivatives of these two compnents:

$$Df(\mathbf{x}) = \begin{bmatrix} \frac{\partial f_1}{\partial x} & \frac{\partial f_1}{\partial y} \\ \frac{\partial f_2}{\partial x} & \frac{\partial f_2}{\partial y} \end{bmatrix} = \begin{bmatrix} 1 - 2x + y & x \\ \frac{1}{2}y & 1 + \frac{1}{2}x - 2y \end{bmatrix}.$$

The Jacobian at the equilibrium of interest is

$$Df(4, 3) = \begin{bmatrix} -4 & 4 \\ \frac{3}{2} & -3 \end{bmatrix}.$$

This has positive determinant and negative trace, so has eigenvalues with negative real part. (This includes the case where the eigenvalues are real, in which case the eigenvalue equals its own real part.) Thus this equilibrium is asymptotically stable for the linear system, and hence for the original nonlinear system as well. The eigenvalues turn out to be real and distinct, so that the equilibrium is actually an improper node, as predicted by the nullclines.

Lecture 6: More on Stability

Warm-Up. Consider the system

$$x' = (1 - x + y)x \qquad y' = (1 + \tfrac{1}{2}x - y)y$$

we finished with last time, which models cooperation between two species with overcrowding. Previously we verified that the equilibrium at $(4, 3)$ was asymptotically stable, and now we consider the behavior at the other equilibria: $(0, 0), (1, 0)$, and $(0, 1)$. Recall that the Jacobian matrix for this system is:

$$Df(\mathbf{x}) = \begin{bmatrix} 1 - 2x + y & x \\ \frac{1}{2}y & 1 + \frac{1}{2}x - 2y \end{bmatrix}.$$

The Jacobian at each equilibrium point is:

$$Df(0,0) = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \qquad Df(1,0) = \begin{bmatrix} -1 & 1 \\ 0 & \frac{3}{2} \end{bmatrix} \qquad Df(0,1) = \begin{bmatrix} 2 & 0 \\ \frac{1}{2} & -1 \end{bmatrix}.$$

The first has positive eigenvalues, so $(0, 0)$ is an unstable improper node (so actually a repeller). The second has eigenvalues of mixed sign, so $(1, 0)$ is an unstable saddle point, and the third also has eigenvalues of mixed sign (negative determinant), so $(0, 1)$ is also an unstable saddle. Note that the existence of other equilibria implies that $(4, 3)$ is only a local attractor, since if it were global then $(0, 0), (1, 0), (0, 1)$ would be in its basin of attraction. Similarly, $(0, 0)$ is only locally repelling.

Now, the linear system which approximates this system near $(4, 3)$ is given by:

$$\mathbf{x}' = Df(4, 3)(\mathbf{x} - [\tfrac{4}{3}]),$$

which results in the equations:

$$x' = -4(x - 4) + 4(y - 3)$$
$$y' = \frac{3}{2}(x - 4) - 3(y - 3).$$

Let us plot some orbits of both this and the original nonlinear system near $(4, 3)$:

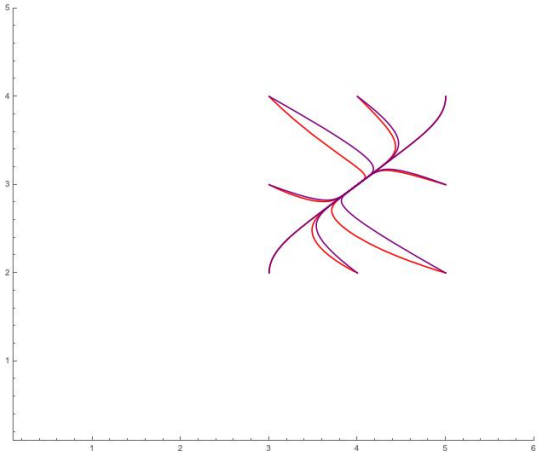

As expected, the nonlinear orbits (in red) do not behave that differently than the linear orbits (in purple). Note, however, that although these orbits are almost indistinguishable when we are very close to $(4, 3)$, they start to differ more so the further away we get. Indeed, once we are outside the basin of attraction of $(4, 3)$, the nonlinear orbits do something different while the linear orbits still approach $(4, 3)$ since this is *globally* attracting for the approximating linear system. So, in general, we can only expect the linear orbits to provide a reasonable approximation to the nonlinear orbits if we are sufficiently close to the equilibrium in question.

The linear system which approximates this system near $(0, 1)$ is given by:

$$x' = 2(x - 0) + 0(y - 1)$$
$$y' = \frac{1}{2}(x - 0) - (y - 1).$$

The same is true here: if we are sufficiently close to $(0, 1)$ the linear (purple) and nonlinear (red) orbits are close to one another, but this breaks down as we move further away:

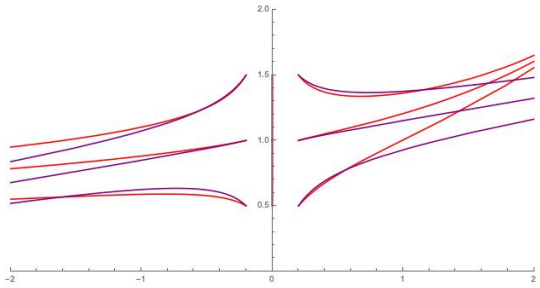

Stability theorem. Let us now give the formal statement of the result which allows us to use the Jacobian matrix to determine stability:

> Suppose $\mathbf{p}$ is a *hyperbolic* equilibrium point of the system $\mathbf{x}' = f(\mathbf{x})$, which means that the Jacobian matrix $Df(\mathbf{p})$ has only eigenvalues with nonzero real part. Then if all eigenvalues of $Df(\mathbf{p})$ have negative real part, $\mathbf{p}$ is asymptotically stable, while if at least one eigenvalue has positive real part, $\mathbf{p}$ is unstable.

The assumption that $\mathbf{p}$ be hyperbolic is only added to simplify the discussion, and technically this is not needed to get a slightly more general version: as long as $Df(\mathbf{p})$ has *some* eigenvalue with positive real part (even if some others have zero real part), $\mathbf{p}$ is unstable. What is definitely true, however, is that if all eigenvalues of $Df(\mathbf{p})$ have zero real part, then we cannot determine stability properties using the Jacobian alone. We'll briefly consider some examples of this later, to get a sense for what can be done in such a scenario.

We will not give a formal proof of this result, but let us give a very rough sketch, at least in the case where we want to establish asymptotic stability to get a sense for how this type of argument works. In order to say that $\mathbf{p}$ is asymptotically stable, we want to say that an orbit $\mathbf{x}(t)$ will approach $\mathbf{p}$ as $t \to \infty$ if $\mathbf{x}(0)$ is sufficiently close to $\mathbf{p}$ to begin with. This comes down to wanting to make the distance $\|\mathbf{x}(t) - \mathbf{p}\|$ between $\mathbf{x}(t)$ and $\mathbf{p}$ small, where $\|\cdot\|$ denotes the length/norm of a vector. Now, denote the orbit of the linear approximation with same initial value $\mathbf{x}(0) = \mathbf{y}(0)$ by $\mathbf{y}(t)$. Since $\mathbf{x}' = Df(\mathbf{x})(\mathbf{x}-\mathbf{p})$ approximates $\mathbf{x}' = f(\mathbf{x})$, we expect $\mathbf{y}(t)$ to be close to $\mathbf{x}(t)$, meaning that we expect $\|\mathbf{x}(t) - \mathbf{y}(t)\|$ to be small. On the other hand, if $Df(\mathbf{p})$ only has eigenvalues with negative real part, then $\mathbf{p}$ should be attracting for the linear system, so we expect $\mathbf{y}(t)$ to be close to $\mathbf{p}$, or in other words $\|\mathbf{y}(t) - \mathbf{p}\|$ should be small.

The *triangle inequality* then allows us to relate these three distances of interest to one another:

$$\|\mathbf{x}(t) - \mathbf{p}\| \leq \|\mathbf{x}(t) - \mathbf{y}(t)\| + \|\mathbf{y}(t) - \mathbf{p}\|.$$

The idea is that since we can make the two terms on the right small, then we can make the term on the left small as well, which implies that $\mathbf{x}(t)$ should converge to $\mathbf{p}$ as required. We can be a little clearer about why it is exactly that we can make the middle term $\|\mathbf{x}(t) - \mathbf{y}(t)\|$ small. We can express the solution $\mathbf{x}(t)$ of $\mathbf{x}' = f(\mathbf{x})$ with initial condition $\mathbf{x}(0)$ via an integral as

$$\mathbf{x}(t) = \mathbf{x}(0) + \int_0^t f(\mathbf{x}(s))\, ds,$$

as we saw last quarter when discussing the proof of the existence and uniqueness theorem. Similarly, we can express the solution $\mathbf{y}(t)$ of $\mathbf{x}' = Df(\mathbf{p})(\mathbf{x} - \mathbf{p})$ with initial condition $\mathbf{y}(0) = \mathbf{x}(0)$ as:

$$\mathbf{y}(t) = \mathbf{x}(0) + \int_0^t Df(\mathbf{p})(\mathbf{y}(s) - \mathbf{p})\, ds.$$

Thus we can write the difference $\mathbf{x}(t) - \mathbf{y}(t)$ as:

$$\mathbf{x}(t) - \mathbf{y}(t) = \int_0^t [f(\mathbf{x}(s)) - Df(\mathbf{p})(\mathbf{y}(s) - \mathbf{p})]\, ds,$$

and we can bound the norm by

$$\|\mathbf{x}(t) - \mathbf{y}(t)\| \leq \int_0^t \|f(\mathbf{x}(s)) - Df(\mathbf{p})(\mathbf{y}(s) - \mathbf{p})\|\, ds.$$

The definition of differentiability for a multivariable function such as f says precisely that we can control how large the norm $\|f(\mathbf{x}(s)) - Df(\mathbf{p})(\mathbf{y}(s) - \mathbf{p})\|$ is (i.e. this is what is meant by saying

that $f(\mathbf{x}) \approx f(\mathbf{p}) + Df(\mathbf{p})(\mathbf{x} - \mathbf{p})$ for $\mathbf{x}$ near $\mathbf{p}$), and being able make this quantity appropriately small is why we are able to make the integral above appropriate small, and hence $\|\mathbf{x}(t) - \mathbf{y}(t)\|$ as well. Of course, there is a lot to fill in in order to make this an actual proof, but the core idea is present.

Another example. Consider the system

$$x' = -xy \qquad y' = x^2 - 1 - y.$$

This has equilibria $(0, -1), (\pm 1, 0)$. The Jacobian matrix at an arbitrary point is

$$Df = \begin{bmatrix} -y & -x \\ 2x & -1 \end{bmatrix},$$

and at the equilibria we get:

$$Df(0,-1) = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} \quad Df(1,0) = \begin{bmatrix} 0 & -1 \\ 2 & -1 \end{bmatrix} \quad Df(-1,0) = \begin{bmatrix} 0 & 1 \\ -2 & -1 \end{bmatrix}.$$

These have respectively eigenvalues of mixed signs, negative real parts, and negative real parts, so $(0, -1)$ is a saddle point and $(\pm 1, 0)$ are both stable spirals:

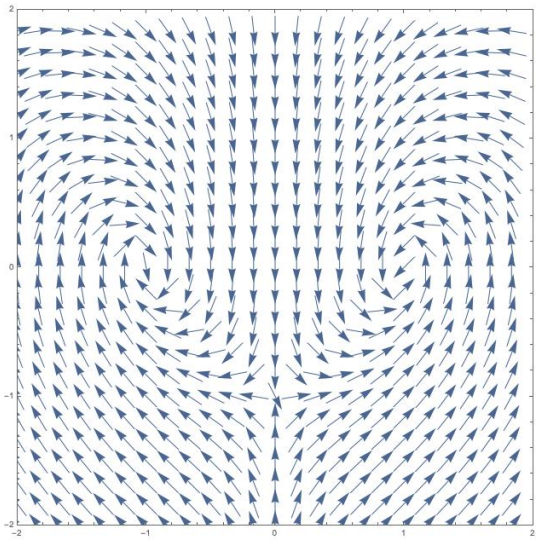

Non-hyperbolic example. ***TO BE FINISHED***

What's the problem? ***TO BE FINISHED***

Lecture 7: Conservative Systems

Warm-Up. We classify the equilibrium points of the system

$$x' = x^2 + y^2 - 1 \qquad y' = x^2 - y^2.$$

TO BE FINISHED

Example. Consider the system

$$x' = y \qquad y' = -x,$$

which is equivalent to the second-order ODE $x'' + x = 0$ which models the harmonic motion of a spring with no damping. Set H to be the function $H(x, y) = \frac{1}{2}y^2 + \frac{1}{2}x^2$. Consider how H changes along the *orbits* of this system: that is, if $(x(t), y(t))$ is an orbit, we consider the derivative of $H(x(t), y(t))$ with respect to t. This can be computed using the multivariable chain rule:

$$\begin{aligned}\frac{dH}{dt} &= \frac{\partial H}{\partial x}\frac{dx}{dt} + \frac{\partial H}{\partial y}\frac{dy}{dt} \\ &= xx' + yy' \\ &= x(y) + y(-x) = 0,\end{aligned}$$

where we use the given system to determine that $x' = y$ and $y' = -x$. This says that the value of H should not change along this orbit, so that H is constant along this orbit. But the orbit was arbitrary, so we conclude that H is constant along any orbit of this system.

In particular, this implies that each orbit of this system has to lie on a single level-set of H, so that if we can determine what these level sets look like, we can get a sense for what the orbits look like as well. The level sets of $H(x, y) = \frac{1}{2}y^2 + \frac{1}{2}x^2$ are all circles:

$$H(x, y) = \text{constant} \rightsquigarrow x^2 + y^2 = \text{constant},$$

so the orbits of this system should lie on these circles. In this case, we know how to solve this system explicitly, and are able to see that the orbits are in fact full circles, but the point is that at least the fact that the orbits lie on circles can be deduced without having the solution available. The function H we used here is called an *integral* of this system (we will define this notion more carefully in a bit), and in this particular instance has a physical interpretation: $\frac{1}{2}y^2$ gives what's called the *kinetic energy* of the spring, and $\frac{1}{2}x^2$ the *potential energy*, so $H = \frac{1}{2}y^2 + \frac{1}{2}x^2$ is the total energy. What we have observed here is that this total energy is *conserved* along the trajectories of the motion of the spring.

Another example. Now consider the system

$$x' = y \qquad y' = -\sin x,$$

which is equivalent to $x'' + \sin x = 0$ and models the motion of a pendulum. We claim that too has an integral, meaning (for the most part) a function which is constant along the orbits. Such an integral H should satisfy:

$$\frac{dH}{dt} = \frac{\partial H}{\partial x}\frac{dx}{dt} + \frac{\partial H}{\partial y}\frac{dy}{dt} = \frac{\partial H}{\partial x}y - \frac{\partial H}{\partial y}\sin x = 0.$$

One way to force this equality to be true is to find a function where

$$\frac{\partial H}{\partial x} = \sin x \quad \text{and} \quad \frac{\partial H}{\partial y} = y,$$

so that $\frac{\partial H}{\partial x}y$ will equal $\frac{\partial H}{\partial y}\sin y$. In this case this is possible to do, essentially because the equations for $\frac{\partial H}{\partial x}$ and $\frac{\partial H}{\partial y}$ are *decoupled* in that each only depends on its own variable; to satisfy $\frac{\partial H}{\partial x} = \sin x$ we need H to have a $k - \cos x$ term where k is constant, and to satisfy $\frac{\partial H}{\partial y} = y$ the function H should have a $\frac{1}{2}y^2$ term, also with possibly a constant added. Let us take the particular choice

$$H = \frac{1}{2}y^2 + (1 - \cos x),$$

which satisfies the requirements. (We took this choice because of its physical significance: $\frac{1}{2}y^2$ is the kinetic energy of the pendulum, and $1-\cos x$ it turns out is its potential energy.)

Since $H = \frac{1}{2}y^2 + (1-\cos x)$ is constant along the orbits (because its derivative along an orbit is zero), we again know that that the orbits of this system lie on the level sets of H. As opposed to the spring, in this case a given level set of H might consist of *multiple* orbits, not just one. Plotting these level sets gives the familiar orbital picture for the pendulum:

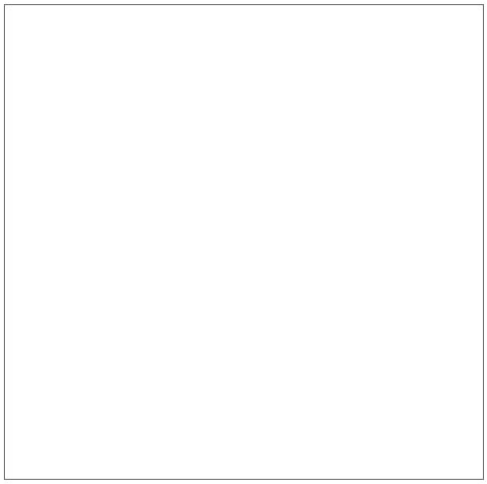

TO BE FINISHED

Conservative systems. Let us now be more precise. An *integral* of a system $\mathbf{x}' = f(\mathbf{x})$ is a differentiable function H such that:

- $H' = \frac{dH}{dt}$ is zero along the orbits of the system, and
- H is nowhere locally constant, meaning that there does not exist any ball on which H is constant.

(The second requirement is there to rule out trivial cases: if H is globally constant, then it satisfies $H' = 0$ for *any* system, and so systems which have integrals would not be different from systems in general. We will see in a bit a better reason why the second requirement is there, and why we also want to rule out being locally constant and not just globally constant.) Note that a system can have multiple integrals (in particular, adding a constant to an integral still gives an integral), so there is no uniqueness.

A system which has an integral is said to be *conservative*. The name comes from the physical examples considered above, where the integral corresponds to a quantity which is conserved, such as energy. But this interpretation is not literally the case in all examples: an integral H is a quantity which is conserved along orbits, but does not necessarily mean that it represents some type of "energy" or some other physical observable. (Soon we will single out one specific type of conservative system, one which is *Hamiltonian*, where H really does mean some type of energy.) For our purposes, you should view H as some auxillary function we introduce—which might not have any obvious relation to any interpretation the system might have—in order to help us study the orbits of the system in question. The key point is that, as we saw before, since H is constant along orbits, each orbit must lie on a level set of H. (In fact, each level set is a union of orbits in general.) So, if we know the level sets, we have a pretty good handle on the orbits themselves.

No attractors/repellers. Here then is one non-obvious feature of a conservatives system: it cannot have any attractors or repellers. Thus, any equilibrium point of a conservative system must be either neutrally stable or something like a saddle; stable/unstable spirals are not allowed, nor are stable/unstable improper nodes. This was indeed the case for the equilibria in the spring and pendulum examples.

Let us prove this fact for the case of attractors. Suppose H is an integral and that $\mathbf{p}$ is an attracting equilibrium. Take any orbit $\mathbf{x}(t)$ whose initial point $\mathbf{x}(0)$ lies within the basin of attraction of $\mathbf{p}$. Since H is an integral, it must be constant along this orbit, so that $H(\mathbf{x}(t)) = c$ for some constant c and all $t \geq 0$. But as $t \to \infty$, $\mathbf{x}(t) \to \mathbf{p}$, and continuity of H implies that

$$H(\mathbf{x}(t)) \to H(\mathbf{p}).$$

But the left side is the constant c, and the limit of a constant is itself, so c must in fact equal $H(\mathbf{p})$. That is to say, the value of H along any point of the orbit $\mathbf{x}(t)$ is $H(\mathbf{p})$. Now if we pick a different initial point within the basin of attraction, the same will be true: the value of H along this new orbit is also $H(\mathbf{p})$. Thus, this implies that the value of H at all points within the basin of attraction is $H(\mathbf{p})$, so that H is constant throughout this basin of attraction, contradicting the nowhere locally constant condition in the definition of an integral. (It is always possible to surround $\mathbf{p}$ by some ball which remains within the basin of attraction; for those of who you know the lingo, the basin of attraction is an *open* subset of $\mathbb{R}^n$.) Thus we conclude that $\mathbf{p}$ could not have been an attractor after all. The proof in the case of repellers is exactly the same, only we take $t \to -\infty$ instead of $t \to \infty$ and use the basin of repulsion instead of the basin of attraction.

Hamiltonian systems. Take $\mathbb{R}^{2n}$ and write its coordinates (of which there are an even number) as $(\mathbf{x}, \mathbf{y}) = (x_1, \dots, x_n, y_1, \dots, y_n)$. A *Hamiltonian system* is a system on $\mathbb{R}^{2n}$ of the form

$$\begin{aligned} x_i' &= \frac{\partial H}{\partial y_i} \\ y_i' &= -\frac{\partial H}{\partial x_i} \end{aligned}$$

where $H(\mathbf{x}, \mathbf{y})$ is a $2n$-variable function. (Note that above there are $2n$ equations in total which come in pairs: two which relate x_1 and y_1, two which relate x_2 and y_2, and so on.) The function H is called the *Hamiltonian* of the system, and is always constant along the orbits:

$$H' = \sum_i \left(\frac{\partial H}{\partial x_i} x_i' + \frac{\partial H}{\partial y_i} y_i' \right) = \sum_i \left(\frac{\partial H}{\partial x_i} \frac{\partial H}{\partial y_i} + \frac{\partial H}{\partial y_i} \frac{\partial H}{\partial x_i} \right) = 0.$$

Thus, any Hamiltonian system is conservative, at least if we assume that H is non-trivial. Hamiltonian systems arise in physics, where H is literally meant to be interpreted as an energy function. In particular, the spring and pendulum systems we first considered are in fact Hamiltonian, where the integrals introduced there are functions which give total energy.

Hamiltonian systems are nice examples of conservatives systems, since it is usually simple to determine if a given system is Hamiltonian, but not so simple to determine if a system is conservative in the non-Hamiltonian cases; in other words, integrals in general are not easy to find, but Hamiltonians are, if they exist. Nevertheless, we know that the orbits of a Hamiltonian system must lie on level sets of the Hamiltonian (conservation of energy), and that a Hamiltonian system has no attractors nor repellers. We will mention one other nice geometric feature of Hamiltonian systems next time, just for fun.

Example. The system

$$x' = 3x - 2y \qquad y' = 5x - 3y$$

is Hamiltonian. Indeed, to be Hamiltonian requires the existence of a function $H(x, y)$ such that

$$x' = \frac{\partial H}{\partial y} = 3x - 2y \quad \text{and} \quad y' = -\frac{\partial H}{\partial x} = 5x - 3y.$$

If we anti-differentiate with respect to y, the first condition says that H must be of the form

$$H = 3xy - y^2 + V(x)$$

where V is some function of x, and then the second condition requires that $\frac{dV}{dx} = -5x$, so we can take $V = -\frac{5}{2}x^2$. Thus $H = 3xy - y^2 - \frac{5}{2}x^2$ is a Hamiltonian function for this sytem.

Lecture 8: Lyapunov Functions

Warm-Up 1. We show that the system

$$x' = -6x + 4y \qquad y' = -10x + 6y$$

is Hamiltonian, and use this to visualize the orbits. ***TO BE FINISHED***

Warm-Up 2. We show that the system

$$x' = x^2 \qquad y' = -2xy$$

is Hamiltonian, and use this to visualize the orbits. ***TO BE FINISHED***

Symplectic geometry. ***TO BE FINISHED***

Example. Consider the system

$$x' = y - x^5 \qquad y' = -x - y^7.$$

Set V to be the function $V(x, y) = x^2 + y^2$. Let us compute the change in V along orbits:

$$\begin{aligned} V' &= \frac{\partial V}{\partial x}x' + \frac{\partial V}{\partial y}y' \\ &= 2x(y - x^5) + 2y(-x - y^7) \\ &= -2x^6 - 2y^8. \end{aligned}$$

We see that V' is zero at the origin, and negative elsewhere. This says that V must be *decreasing* along an orbit, or in other words motion along an orbit must occur towards decreasing values of V. In this case the level sets of V are circles centered at the origin, with smaller values of V corresponding to smaller circles, so orbits of this system must pass through smaller and smaller circles forward in time:

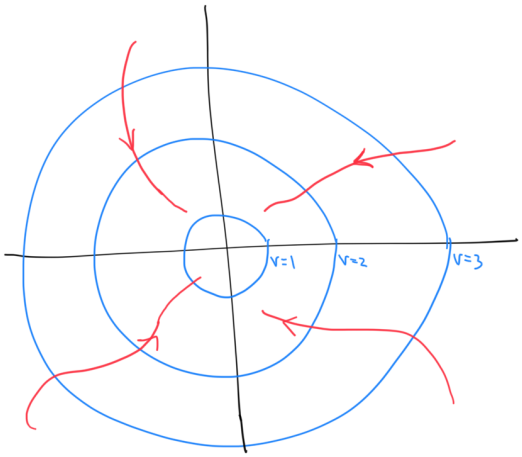

This seems to suggest that the orbits should be approaching the equilibrium at the origin, which is in fact true. The function V here is called a *(strong) Lyapunov function* for this system, and the point is that finding such functions help us determine some stability properties. Think of V as an auxiliary function we introduce to help us control the behavior of orbits, similar to an integral for a conservative system; here, though, the auxiliary Lyapunov function is not constant along the orbits as in the integral case, but rather decreasing so that the direction of the orbits can be determined from the level sets of V.

Definiteness. Here are the key notions we need. Suppose $\mathbf{p}$ is an equilibrium point of some system. We say that a (differentiable) function V is:

- *positive definite* at $\mathbf{p}$ if V is 0 at $\mathbf{p}$ and positive elsewhere. We can distinguish between *local* and *global* positive definiteness depending on whether V is positive only on some ball around $\mathbf{p}$ or everywhere outside $\mathbf{p}$;
- *positive semidefinite* at $\mathbf{p}$ if we take the same definition only relax V being positive elsewhere to only nonnegative, so that it might be zero at some other points as well. Again we have local and global versions of this;
- (locally/globally) *negative definite* at $\mathbf{p}$ if V is 0 at $\mathbf{p}$ and negative elsewhere; and
- (locally/globally) *negative semidefinite* if we only require that V be nonpositive elsewhere.

Using V' (the derivative along orbits) in place V, we can also talk about V' being some version of definite or semidefinite.

One class of functions for which the type of definiteness it might have is simple-ish to determine are those which are *quadratic forms*, which are functions which looks like:

$$V = ax^2 + 2bxy + cy^2 = \begin{bmatrix} x & y \end{bmatrix} \begin{bmatrix} a & b \\ b & c \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}.$$

The final expression is given in terms of matrix multiplication. (There are analogous versions in higher dimensions.) The definiteness of V comes down to the eigenvalues of the defining matrix $\left[\begin{smallmatrix} a & b \\ b & c \end{smallmatrix}\right]$: positive definite corresponds to positive eigenvalues, semidefinite allows one or both to be zero, and similarly for the negative definite conditions. (In general, quadratic forms give "second-order approximations" via Taylor polynomials to general definite functions.)

Lyapunov's Theorem. The precise version of our stability result in terms of Lyapunov functions is the following:

if there exists a function V which is locally/globally positive definite at $\mathbf{p}$ such that V' is locally/globally negative definite at $\mathbf{p}$, then $\mathbf{p}$ is locally/globally asymptotically stable; while if there exists V positive definite but for which V' is only negative semidefinite at $\mathbf{p}$, then $\mathbf{p}$ can only be guaranteed to be stable.

In the first case, we say that V is a *strong* Lyapunov function for the system, while in the second it is a *weak* Lyapunov function. In the weak case, we cannot distinguish using V alone whether $\mathbf{p}$ is asymptotically or neutrally stable. But, the idea is that perhaps there is some *other* Lyapunov function we can use instead which can distinguish between these two. In general, finding a "good" Lyapunov function to use is somewhat of an art form, but in many cases we should start by considering a quadratic form.

Let us give a sense of the proof of Lyapunov's Theorem. The picture we drew before gives the basic idea, but does not along guarantee that an orbit should converge to the equilibrium; in particular, we can imagine an orbit with the following behavior:

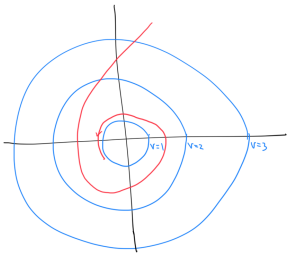

which *does* pass through circles of smaller and smaller radii as time increases, but can get arbitrarily close to the circle at $V = 1$ without ever actually hitting it. It is not clear that this cannot actually happen when $V' < 0$, so we have to be careful about ruling this out. (This can happen when $V' \leq 0$, which is why in the V' being negative semidefinite case we might only get neutral stability.)

TO BE FINISHED

Back to example. Let us return to our first example:

$$x' = y - x^5 \qquad y' = -x - y^7$$

and think about how we can find an appropriate Lyapunov function if we did not know of one beforehand. Since quadratic forms are functions can definiteness can be easily determined, let us first consider a function of the form $V = ax^2 + cy^2$. The question is whether we can find a, c which will make this a Lyapunov function. Of course, we should already have a sense that the origin here is meant to be asymptotically stable, which we can guess by looking at the direction field and determine the direction of orbits using nullclines.

We have:

$$V' = 2ax(y - x^5) + 2cy(-x - y^7) = -2ax^6 - 2cy^8 + (2a - 2c)xy.$$

The terms $-2ax^6$ and $-2cy^8$ will certainly be negative if a, c are chosen to be positive, but the mixed term with xy can take on different signs. One easy way to avoid dealing with this is to make the coefficient $2a - 2c$ equal zero by choosing $a = c$. Thus for $a = c$ positive the function V' will be (globally) negative definite, and these also make V definite. Thus any $V = ax^2 + ay^2$ with $a > 0$ will be global strong Lyapunov function, so the origin is globally asymptotically stable.

Another example. Consider the system

$$x' = y \qquad y' = -2x - 3y.$$

Take a general quadratic form $V = ax^2 + 2bxy + cy^2$. Then:

$$\begin{aligned} V' &= (2ax + 2by)y + (2bx + 2cy)(-2x - 3y) \\ &= (2b - 6c)y^2 - 4bx^2 + (2a - 6b - 4c)xy. \end{aligned}$$

The squared terms will be negative if $2b - 6c < 0$ and $b > 0$. Now, again the mixed xy term might pose a problem, but this can be handled by picking $2a - 6b - 4c$ to be zero. On top of this, since:

$$V = \begin{bmatrix} x & y \end{bmatrix} \begin{bmatrix} a & b \\ b & c \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix},$$

in order for V to be positive definite we need $\left[\begin{smallmatrix} a & b \\ b & c \end{smallmatrix}\right]$ to have positive eigenvalues. (Since this matrix is symmetric, all eigenvalues are real.)

So we start making some choices: for $b > 0$ let us take $b = 1$; then for $2b - 6c = 2 - 6c < 0$ let us take $c = 1$; and then $2a - 6b - 4c = 0$ forces $a = 5$. The resulting matrix $\left[\begin{smallmatrix} 5 & 1 \\ 1 & 1 \end{smallmatrix}\right]$ does have positive eigenvalues, so these choices will indeed give a global strong Lyapunov function $V = 5x^2 + 2xy + y^2$. Thus the origin is globally asymptotically stable.

Lecture 9: Cycles and Limiting Behavior

Warm-Up. We use Lyapunov functions to determine the stability at the origin of

$$x' = (x - y)^2(y - x) \qquad y' = -(x - y)^2(x + y).$$

Let us look for one of the form $V = ax^2 + cy^2$. (Of course, we do not know beforehand that this will work—it will in this case—but it is simpler to try this first before a more general quadratic form.) We have:

$$\begin{aligned} V' &= 2ax(x - y)^2(y - x) - 2cy(x - y)^2(x + y) \\ &= (x - y)^2[-2ax^2 - 2cy^2 + (2a - 2c)xy]. \end{aligned}$$

This will be zero at point along the line $y = x$, so the best we can hope for here is a weak Lyapunov function. But this can be achieved by taking $a = c$ to get rid of the mixed terms and then $a = c > 0$. Any such choice will make V positive definite, so we do get a weak Lyapunov function, say $V = x^2 + y^2$. Thus the origin is stable.

Now, to distinguish between asymptotic and neutral stability, Lyapunov functions are not enough. But, note that any point on the line $y = x$ is an equilibrium, and that such points get arbitrarily close to $(0, 0)$. Thus the origin cannot be an attractor since it it were some of the points on $x = y$ would be within its basin of attraction, which is not possible. Thus the origin is neutrally stable.

Other Lyapunov functions. Other types of Lyapunov functions can be used to indicate other types of behavior. For example, instability can be detected this way (the book gives this result), as can a stronger form of stability known as "exponential stability". Beyond stability, auxiliary Lyapunov functions can be used to prove that orbits are bounded (we'll see an example of this

soon), or unbounded, or whatever. We will not get into this in any more detail here, but suffice it to say that a lot more is known.

Limit cycles. There exist other types of limiting behavior beyond converging to an equilibrium point. We saw an example in the first week in the following system:

$$x' = x(1 - \sqrt{x^2 + y^2}) - y \qquad y' = y(1 - \sqrt{x^2 + y^2}) + x.$$

In polar coordinates this becomes

$$r' = r(1 - r) \qquad \theta' = 1,$$

so the unit circle $r = 1$ forms a single orbit. By tracking the sign of r', we see that r is decreasing above $r = 1$ and increasing below, all nonconstant orbits will approach in the long-term. So, we have in this case not an attracting equilibrium point, but an entire cycle which attracts, which is an example of a "limit cycle":

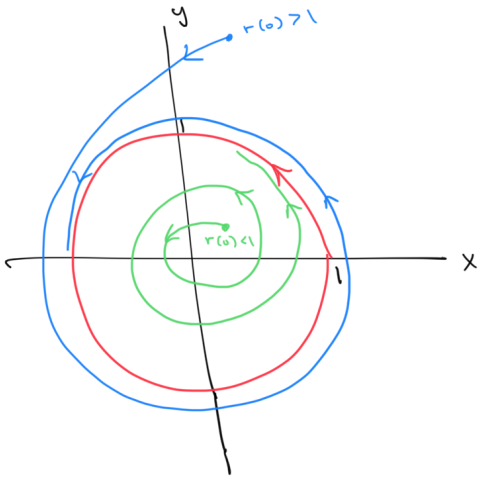

To be clear, a *limit cycle* is a cycle which attracts or repels *some* (at least one) nearby orbit. Then, we say that this an *attracting/repelling* limit cycle if it attracts/repels *all* nearby orbits. In the example above, the unit circle is a globally attracting limit cycle. Note that orbits within the circle are repelled from the equilibrium at the origin towards the attracting circle. Also, this limit cycle contains the equilibrium at the origin in its interior, which is a general phenomenon: any limit cycle must enclose an equilibrium point.

Cycle-graphs. Now consider the system

$$x' = x(1 - x - 3.75y + 2xy + y^2) \qquad y' = y(-1 + y + 3.75x - 2x^2 - xy),$$

which the book considers as well. Here is a picture of an orbit, which moves counterclockwise :

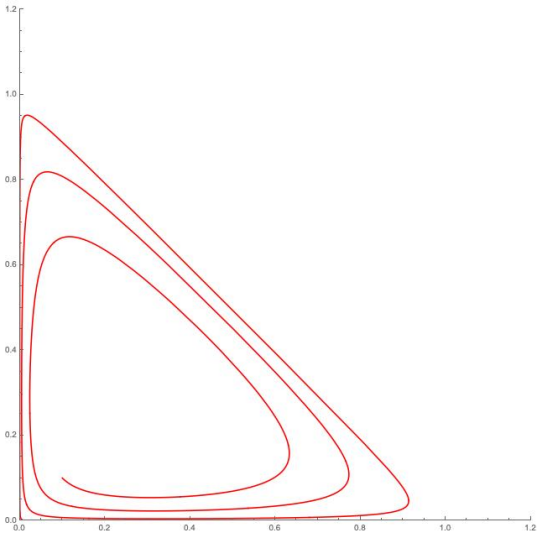

Going backward in time, this orbit emanates from some equilibrium point (it is $(0.25, 025)$) in the first quadrant, but forward in time it approaches a triangle-shaped figure:

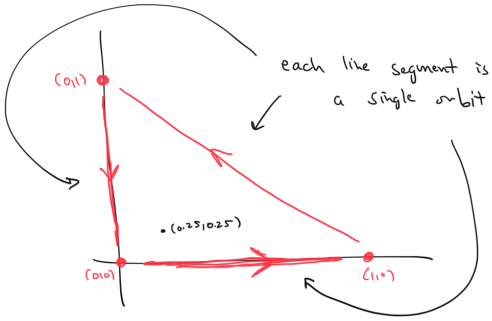

This figure is NOT a cycle, but rather consists of three equilibrium points (the corners) and three orbits (line segments), each of which heads from one equilibrium to another. Such a figure is called a *cycle-graph*, which in general is a graph (a collection of vertices and edges) whose vertices are equilibrium points and whose edges are orbits which are repelled from a vertex as $t \to -\infty$ and are attracted to a vertex as $t \to \infty$. Thus, in addition to equilibrium points and limit cycles, it can also be the case that an orbit approaches a cycle-graph

Poincaré-Bendixon. And that is it as far planar systems goes! The *Poincaré-Bendixon Theorem* states that any *bounded* (as $t \to \infty$ or $t \to -\infty$) orbit of a planar 2-dimensional system has to approach either an equilibrium point, a limit cycle, or a cycle-graph. Thus, (bounded) orbits of planar systems behave in controlled, *predictable* ways in the long-term.

In the example above, the orbit which was drawn originally approaches a cycle-graph (the triangle) as $t \to \infty$ and an equilibrium point as $t \to -\infty$. In the example before (in polar coordiantes), orbits within the unit circle approach a limit cycle as $t \to \infty$ and an equilibrium point as $t \to -\infty$. Orbits which are outside the unit circle approach the same limit cycle as $t \to \infty$, but do not approach anything as $t \to -\infty$: the Poincaré-Bendixon Theorem does not apply in this direction since these orbits are unbounded as $t \to -\infty$. Another useful notion is that of a *limit set*: the *positive limit set* of an orbit is the set of points it comes arbitrarily close to as $t \to \infty$, and the *negative limit set* is the set of points it comes arbitrary close to as $t \to -\infty$. So, these

limit sets consist of equilibrium points, limit cycles, portions of cycle-graphs (not necessarily the full cycle-graph), or are empty in the unbounded case.

Example 1. ***TO BE FINISHED***

Example 2. Consider the system

$$x' = y + x(4 - 2x^2 - 3y^2) \qquad y' = -x + y(4 - 2x^2 - 3y^2)$$

and take $V = x^2 + y^2$, which for now is just some auxiliary "Lyapunov-like" function we will try to use to control the behavior of orbits. We compute:

$$V' = 2x[y + x(4 - 2x^2 - 3y^2)] + 2y[-x + y(4 - 2x^2 - 3y^2)] = (2x^2 + 2y^2)(4 - 2x^2 - 3y^2).$$

Now, for $x^2 + y^2 > 4$, we have

$$V' = (2x^2 + 2y^2)(4 - 2x^2 - 3y^2) < (2x^2 + 2y^2)(4 - x^2 - y^2) < 0$$

since $4 - x^2 - y^2 < 0$, and where the middle step holds since we end up subtracting less away from 4 than before. Thus, V should be decreasing along orbits outside the circle $x^2 + y^2 = 4$, meaning that orbits outside this circle have to come inside of it as $t \to \infty$ since they move towards smaller level sets of V, so towards smaller circles.

On the other hand, for $x^2 + y^2 < 1$, we have

$$V' = (2x^2 + 2y^2)(4 - 2x^2 - 3y^2) > (2x^2 + 2y^2)(4 - 4x^2 - 4y^2) > 0$$

since $4 - 4x^2 - 4y^2 < 0$ and the middle step holds we subtract more away from 4 than before. Thus V should be increasing along orbits within the circle $x^2 + y^2 = 1$, so orbits which are inside this circle move towards it as $t \to \infty$. Thus, we find that orbits in the long-term are trapped within an annulus (region between two circles:

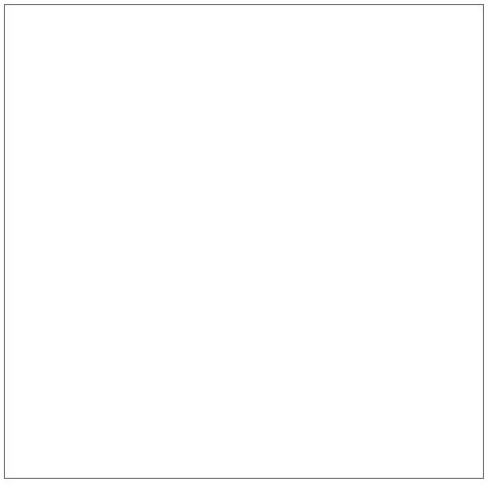

since once an orbit enter this region it can never leave due the behavior of V outside of the annulus. We say that orbits will become "trapped" within this annulus. Such an orbit will thus be bounded as $t \to \infty$, so by Poincaré-Bendixon it must approach either an equilibrium point, a limit cycle, or a cycle-graph within this annulus. The only equilibrium point this system has is $(0, 0)$ (verified

through some tedious algebra), and this is not in the annulus so the orbit cannot approach this. A cycle-graph must have vertices which are equilibria, and again since the origin is the only equilibrium such a cycle-graph cannot be found within the annulus.

Thus we conclude that the orbit has to approach a limit cycle, thereby proving this system indeed has a limit cycle. If you plot orbits more carefully, you can see what the limit cycle should roughly look like:

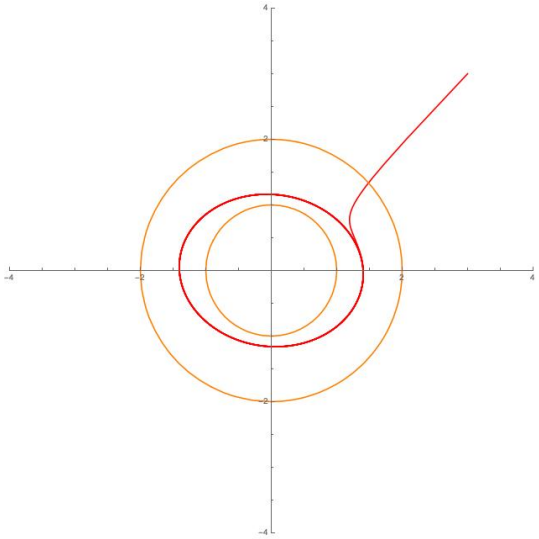

(The annulus is the region between the two orange circles, and the orbit is in red. The limit cycle in this case is the ellipse $2x^2 + 3y^2 = 4$, precisely where $V' = 0$,) In general, trapping orbits within some bounded region, say by using an auxiliary Lyapunov-like function or by studying the direction of orbits using nullclines, will prove the existence of limit cycles when they cannot be found explicitly.

Lecture 10: Hopf Bifurcations

TO BE FINISHED

Lecture 11: Chaotic Systems

TO BE FINISHED

Lecture 12: Partial Differential Equations

More examples of chaos. ***TO BE FINISHED***

End of the line for ODEs. And with that, we have said all about ODEs that we can pretty much say in general. Of course, the various topics we've looked at can be pushed further and branch off into their entire own sub-disciplines—chaos theory, bifurcation theory, Lyapunov stability theory, Hamiltonian mechanics, etc—but in broad strokes there is not much more one can say "in general". The behavior of planar 2-dimensional autonomous systems is essentially fully characterized by the Poincaré-Benidixon Theorem, and studying higher-dimensional systems use tools which are often unique to that one system or ones which are very similar to it. So, we now leave the world of ordinary differential equations, and dive head-first into the world of partial differential equations.

Partial differential equations. A *partial differential equation* (or PDE for short) is a differential equation which characterizes a multivariable function via its partial derivatives. For instance, the equation

$$\frac{\partial u}{\partial t} = \frac{\partial^2 u}{\partial x^2} \text{ (which in alternate notation is } u_t = u_{xx})$$

states that the first-order partial derivative of $u(x,t)$ with respect to t should equal its second-order derivative with respect to x twice. (This particular equation is an instance of what's called the *heat equation*, which we will study soon enough.) As in the ODE case, solving such an equation amounts to finding all such functions, possibly satisfying some given initial or boundary conditions, although what should count as an "initial" condition is still to-be-clarified. One problem we will come across is that in general there is no simple "existence and uniqueness" theorem analogous to what we saw for ODEs, and the situation is much more dependent on the particulars of the specific equation at hand. We characterize equations as first- or second-order depending on the types of partial derivatives used, and also as driven vs undriven, linear vs nonlinear, etc as before.

A first-order linear PDE (in two variables) is thus of the form:

$$f(x,y)\frac{\partial u}{\partial x} + g(x,y)\frac{\partial u}{\partial y} = h(x,y).$$

Here, $u(x,y)$ is the function for which we are looking. To say that this is linear is to say that the defining partial differential operator:

$$L[u] = f(x,y)\frac{\partial u}{\partial x} + g(x,y)\frac{\partial u}{\partial y}$$

is linear in the sense that $L[u_1 + u_2] = L[u_1] + L[u_2]$ and $L[cu_1] = cL[u_1]$. As with ODEs, linearity gives some nice properties, such as the fact that, in the undriven case, solutions form a "vector space", meaning that linear combinations of solutions are still solutions, and that in the driven case all solutions can be expressed as a particular solution plus undriven solutions.

Method of characteristics. Consider a first-order linear PDE in two variables:

$$f(x,y)\frac{\partial u}{\partial x} + g(x,y)\frac{\partial u}{\partial y} = h(x,y).$$

One approach to finding $u(x,y)$ is to describe its *graph* instead, which is the set of points in $\mathbb{R}^3$ satisfying $z = u(x,y)$. Here we will outline what's known as the *method of characteristics* for doing so. We should note that this method is NOT covered in our book at all, and indeed will not really play a role going forward; we are simply mentioning it now to get a sense for the types of techniques which are out there, and in particular as something different than the main methods we will soon look at, which are based on Fourier analysis and related concepts.

After subtracting $h(x,y)$, the equation above can be written as

$$\langle f(x,y), g(x,y), h(x,y)\rangle \cdot \langle u_x, u_y, -1\rangle = 0$$

where the left is the dot product of two vectors. This equality says that at all points, the vector $\langle f(x,y), g(x,y), h(x,y)\rangle$ should be orthogonal to the vector $\langle u_x, u_y, -1\rangle$, which is itself orthogonal to the surface $z = u(x,y)$: the vector $\langle u_x, u_y, -1\rangle$ is the gradient vector of the three-variable function $u(x,y) - z$, and so is orthogonal to any level surface of this function, of which $z = u(x,y)$ is the one at level 0. Thus, $\langle f(x,y), g(x,y), h(x,y)\rangle$ should be *tangent* to the graph of $u(x,y)$ at any point.

Any orbit of the system

$$x' = f(x,y) \qquad y' = g(x,y) \qquad z' = h(x,y),$$

which describes curves which are tangent to the vector field $\langle f(x,y), g(x,y), h(x,y)\rangle$, will hence lie on the graph of $z = u(x,y)$. So, the idea is to *construct* the graph of $u(x,y)$, and hence determine $u(x,y)$ itself, by finding these curves—they will sweep out the graph in question. The resulting curves are called the *characteristic curves* of the given PDE, and the ODEs above are called its *characteristic equations*. In equations for which this method will be useful, we are essentially able to reformulate our PDE as a system of ODEs, which we have a spent a good amount of time studying already. Of course, not all PDEs of interest will be of this type, so we will need to consider other methods later on.

Example. Consider the PDE

$$\frac{\partial u}{\partial x} + a\frac{\partial u}{\partial y} = 0$$

where a is a constant. (This is known as the *transport equation*, and models propagation along a curve.) The corresponding characteristic equations are

$$x' = 1 \qquad y' = a \qquad z' = 0.$$

Thus the characteristic curves (orbits of this system) are given explicitly by

$$x = t + x_0 \qquad y = at + y_0 \qquad z = z_0$$

where x_0, y_0, z_0 are constants. In this case, these curves are lines occurring in horizontal planes, so the graph of $z = u(x,y)$ should have level sets consisting of lines.

From the x and y equations we see that the quantity $y - ax = y_0 - ax_0$ must be constant along any characteristic curve, and then $z = z_0$ says that the value of $z = u(x,y)$ must be the same along any such quantity. Thus, we can express the value $u(x,y)$ as the value of a single-variable function of the input $y - ax$:

$$u(x,y) = f(y - ax) \text{ for some single-variable } f.$$

This f should actually be differentiable in order for the partial derivatives of $u(x,y)$ to actually exist. We can verify using the chain rule that such a function $u(x,y) = f(y-ax)$ does indeed satisfy our original PDE:

$$\frac{\partial u}{\partial x} = f'(y-ax)\frac{\partial (y-ax)}{\partial x} = -af'(y-ax) \quad \text{and} \quad \frac{\partial u}{\partial y} = f'(y-ax)\frac{\partial (y-ax)}{\partial y} = f'(y-ax),$$

so $\frac{\partial u}{\partial x} + a\frac{\partial u}{\partial y} = -af'(y-ax) + af'(y-ax) = 0$ as expected. Hence we conclude that all solutions of $u_x + au_y = 0$ are given by $u(x,y) = f(y-ax)$ where f is a single-variable differentiable function.

One thing to note is that, whereas in the ODE case where we ended up with general solutions expressed using undetermined constants, here we end up with an undetermined *function* in our general solution. This is a common occurrence with PDEs. In particular, this implies that specifying an initial condition such as $u(x_0, y_0) = z_0$, where we give the value of $u(x,y)$ at a single point, is not enough to single out one solution: for this PDE, any f for which $f(y_0 - ax_0) = z_0$ will satisfy this initial condition. Rather, if we impose an initial condition like:

$$u(0,y) = \phi(y)$$

where we give a range of initial values in terms of some single-variable function ϕ, then we do single-out a unique solution: in order for $u(x, y) = f(y - ax)$ to satisfy $u(0, y) = \phi(y)$, we need:

$$u(0, y) = f(y - a0) = f(y) = \phi(y),$$

so that $f = \phi$ and thus we get the unique solution $u(x, y) = \phi(y - ax)$. This need to specify a range of initial values is also a common feature we'll see with PDEs.

Other PDEs. As stated above, some of the features we saw in these first-order linear PDEs will be true of PDEs in general. The first PDE we will study in more detail is known as the *wave equation*, and takes the form:

$$\frac{\partial^2 u}{\partial t^2} = c^2 \frac{\partial^2 u}{\partial x^2}$$

where c^2 is a positive constant. Here, $u(x, t)$ is a function of a position variable x and a time variable t, and indicates the height of a point on a vibrating string at positive x and time t:

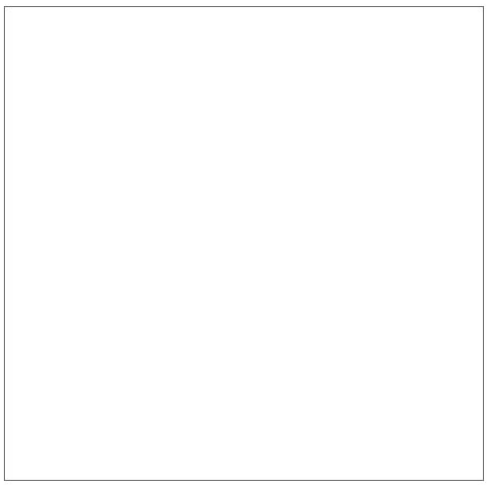

The wave equation thus models the motion of this string. We will look at a derivation of the wave equation next time, and see why it models what we say it does. For now, let us briefly think about what types of conditions we should impose on a solution. First, we can indicate the initial position of the entire string at time $t = 0$, so we impose something like

$$u(x, 0) = f(x)$$

where f is the function given this initial height at every point. If our string has length L, we can require that the endpoints of the string never move, so that they are always held down at the bottom; this gives the boundary condition

$$u(0, t) = 0 \quad \text{and} \quad u(L, t) = 0 \text{ for all } t \geq 0.$$

We will see that if also impose an initial condition for the *velocity* of the motion at time $t = 0$, we will be able to derive a unique solution. Deriving this solution will introduce a fundamental idea we will use for other equations: most immediately the use of *Fourier series* to construct solutions, but more broadly the use of *eigenfunctions* of linear differential operators to construct solutions.

Lecture 13: Wave Equation

Warm-Up. We find the solution of the following initial value PDE using the method of characteristics:

$$-y\frac{\partial u}{\partial x} + x\frac{\partial u}{\partial y} = 0,\ u(x,0) = \sin(x^2).$$

The characteristic equations are:

$$x' = -y \qquad y' = x \qquad z' = 0.$$

The orbits of this system are given by:

$$x = c_1\cos t + c_2 \sin t \qquad y = c_1 \sin t - c_2 \cos t \qquad z = z_0$$

where c_1, c_2, z_0 are arbitrary constants. Visually, these look like circles in horizontal planes. A key observation is that the quantity $x^2 + y^2$ is thus constant along a characteristic, and the value of z only depends on this quantity, so that we can write the value of $u(x,y)$ as the value of a single-variable function evaluated at this quantity:

$$u(x,y) = f(x^2+y^2) \text{ for some single-variable differentiable } f.$$

The chain rule can be used to show that such a function does indeed satisfy our PDE:

$$u_x = f'(x^2+y^2)2x \text{ and } u_y = f'(x^2+y^2)2y,\ \text{so } -yu_x + xu_y = [-2xy + 2xy]f'(x^2+y^2) = 0.$$

In order to satisfy our initial condition $u(x,0) = \sin(x^2)$, we need:

$$u(x,0) = f(x^2+0^2) = f(x^2) = \sin(x^2),$$

so that we can take f to be sine itself, at least for nonnegative inputs like x^2. Thus

$$u(x,y) = \sin(x^2+y^2) \text{ is the solution of } -yu_x + xu_y = 0,\ u(x,0) = \sin(x^2).$$

Note that even though f itself in $u(x,y) = f(x^2+y^2)$ is not uniquely determined, the function $u(x,y)$ *is* uniquely determined: we don't know what something like $f(-1)$ has to be, but such a value is irrelevant in describing $u(x,y)$ since $u(x,y) = f(x^2+y^2)$ only cares about the value of f at nonnegative inputs anyway, and for such inputs $f(x^2+y^2)$ must equal $\sin(x^2+y^2)$.

The wave equation. ***TO BE FINISHED***

Separation of variables. To solve the wave equation more explicitly we make use of the *method of separation of variables*, which is similar-in-spirit but different in practice to the method we look at last quarter which went by this name. Concretely, we look for solutions of the wave equation which are of the form

$$u(x,t) = X(x)T(t)$$

for single-variable functions X and T. (Thus, solutions were the variables are "separated" from one another.) In order for this function to satisfy $u_{tt} = c^2 u_{xx}$, we must have:

$$XT'' = c^2 X''T.$$

To be clear, the derivatives here are actually taken with respect to different variables—x in X'' and t in T''—but we use the same prime notation either way. Now, we ideally want non-trivial

solutions since having $u(x,t)$ be identically zero is not so interesting. Thus, we will assume that neither X nor T are ever zero. In this case, the equation above can be rewritten as

$$\frac{T''}{c^2T} = \frac{X''}{X}.$$

The punchline now is that *both* sides here must actually be *constant*! Indeed, if we put variables back then, this equation really says that

$$\frac{T''(t)}{c^2T(t)} = \frac{X''(x)}{X(x)}.$$

If we hold t fixed and vary x, this says that the right side must be a constant function, and if we fix x and vary t, so must the left side. (This comes down to the fact that the two sides depend on different variables.) Thus, we get that for some constant λ:

$$\frac{T''}{c^2T} = \lambda \quad \text{and} \quad \frac{X''}{X} = \lambda.$$

Separated solutions. This thus gives rise to a pair of *ordinary* differential equations, one for T and one for X:

$$T'' = c^2\lambda T \quad \text{and} \quad X'' = \lambda X,$$

which we know how to solve! This is the entire point of the separation of variables method: reduce a PDE into a collection of ODEs. Now we impose our boundary conditions: $u(0,t) = 0 = u(L,t)$. These turn into $X(0)T(t) = 0 = X(L)T(t)$, which since we are assuming T is nonzero, results in a pair of boundary conditions for X alone:

$$X(0) = 0 = X(L).$$

Hence we end up with the following ODEs:

$$T'' = c^2\lambda T \quad \text{and} \quad X'' = \lambda X,\ X(0) = 0 = X(L).$$

First we consider the equation for X. If $\lambda > 0$, the general solution of $X'' = \lambda X$ is

$$X = c_1 e^{\sqrt{\lambda}x} + c_2 e^{-\sqrt{\lambda}x}.$$

But the boundary conditions then turn into the requirements that:

$$X(0) = c_1 + c_2 = 0 \quad \text{and} \quad X(L) = c_1 e^{\sqrt{\lambda}L} + c_2 e^{-\sqrt{\lambda}L} = 0.$$

Some algebra shows that this is only satisfied when $c_1 = c_2 = 0$, so in this case $X = 0$, which is not the type of solution we are after. If $\lambda = 0$, the general solution of $X'' = 0X = 0$ is

$$X = c_1 + c_2 x,$$

but again the boundary conditions $X(0) = 0 = X(L)$ force $c_1 = c_2$, so that $X = 0$.

Thus, if we are looking for nontrivial solutions, we need only consider $\lambda < 0$. To make the notation simpler, we set $\lambda = -k^2 < 0$ for some positive number k. The general solution of $X'' = \lambda X = -k^2 X$ (equivalently the simple harmonic oscillator equation $X'' + k^2 X = 0$) here is:

$$X = c_1 \cos kx + c_2 \sin kx.$$

The boundary condition $X(0) = 0$ forces $c_1 = 0$, but $X(L) = 0$ forces

$$c_2 \sin kL = 0,$$

which can be satisfied non-trivially! For $c_2 \neq 0$, this equality holds as long as $\sin kL = 0$, so for

$$kL = n\pi, \text{ or equivalently } k = \frac{n\pi}{L}$$

where n is a positive integer. (Allowing n to be negative gives the same value of $\lambda = -k^2 = -(\frac{n\pi}{L})^2$ as for n positive, so we do not get any other new solutions this way. Also, $n = 0$ gives $\lambda = 0$, which does fall into the case of $\lambda < 0$.) Thus, we get a nontrivial solution to our boundary value problem when $\lambda = -(\frac{n\pi}{L})^2$, namely:

$$X = 0 \cos \frac{n\pi x}{L} + c_2 \sin \frac{n\pi x}{L} = c_2 \sin \frac{n\pi x}{L}$$

for $c_2 \neq 0$. We only need one such solution, so will simply take $c_2 = 1$. To summarize, we have found that the nontrivial solutions of

$$X'' = \lambda X, \ X(0) = 0 = X(L)$$

are $X_n = \sin \frac{n\pi x}{L}$ for $\lambda_n = -(\frac{n\pi}{L})^2$ where n is a positive integer.

And now finally we come back to T, which must satisfy $T'' = -c^2\lambda T$. Using the allowable values of λ we already determined turns this into

$$T'' = -c^2(\tfrac{n\pi}{L})^2 T.$$

This has general solution

$$T_n = A_n \cos(\tfrac{cn\pi t}{L}) + B_n \sin(\tfrac{cn\pi t}{L}).$$

So, we find that the separated solutions—indexed by positive integers n—of our wave equation with boundary conditions:

$$u_{tt} = c^2 u_{xx}, \ u(0,t) = 0 = u(L,t)$$

are given by

$$u_n(x,t) = X_n(x)T_n(t) = \sin(\tfrac{n\pi x}{L})[A_n \cos(\tfrac{cn\pi t}{L}) + B_n \sin(\tfrac{cn\pi t}{L})].$$

Formal solutions. For previous linear differential equations we have seen, once we had *some* solutions we were able to form more by taking arbitrary linear combinations. We will do the same thing here, were the main difference is that now we are taking a linear combination of an *infinite* number of solutions $u_n(x,t)$, so that the sum we obtain is actually an infinite series:

$$u(x,t) = \sum_{n=1}^{\infty} d_n u_n(x,t).$$

The coefficient d_n can be absorbed into the scalars A_n, B_n which show up in the expression for $u_n(x,t)$ we derived above, so we may just as well take $d_n = 1$, so that we get:

$$u(x,t) = \sum_{n=1}^{\infty} u_n(x,t) = \sum_{n=1}^{\infty} \sin(\tfrac{n\pi x}{L})[A_n \cos(\tfrac{cn\pi t}{L}) + B_n \sin(\tfrac{cn\pi t}{L})].$$

Of course, whenever we are dealing with an infinite series, there are questions of convergence to address (i.e. does the infinite sum actually make sense as a finite number?), but for now we will ignore this and refer to this series expression as given the *formal solution* of our wave equation with boundary conditions. You can verify that for this function $u(x,t)$, the derivative with respect to t twice will indeed equal c^2 times the derivative with respect to x twice.

Lecture 14: More on Wave Solutions

Warm-Up. Suppose a function f is given by the following infinite series on the interval $[0, L]$:

$$f(x) = \sum_{n=1}^{\infty} a_n \sin \tfrac{n\pi x}{L}.$$

We derive an explicit formula for the coefficient a_n. First, multiply this equality through by $\sin \frac{m\pi x}{L}$:

$$f(x) \sin \tfrac{m\pi x}{L} = \sum_{n=1}^{\infty} a_n \sin \tfrac{n\pi x}{L} \sin \tfrac{m\pi x}{L}.$$

Next, integrate both sides:

$$\int_0^L f(x) \sin \tfrac{m\pi x}{L}\, dx = \int_0^L \left(\sum_{n=1}^{\infty} a_n \sin \tfrac{n\pi x}{L} \sin \tfrac{m\pi x}{L} \right) dx = \sum_{n=1}^{\infty} a_n \int_0^L \sin \tfrac{n\pi x}{L} \sin \tfrac{m\pi x}{L}\, dx.$$

(If you have seen some real analysis before, you might ask why it is that we can change integration and summation in this way—we will discuss convergence of such series later, but the answer is that such series will converge *uniformly* if f is continuously differentiable, which is the property we need in order for the manipulation above to be valid.)

Now we make use of the following fact: the value of

$$\int_0^L \sin \tfrac{n\pi x}{L} \sin \tfrac{m\pi x}{L}\, dx$$

is 0 whenever m and n are different, and $\frac{L}{2}$ when $m = n$. This can be verified through a direct computation which makes use of some trigonometric identities, but you can check the book for the details. (Later we will see what this equality "really" says about the functions $\sin \frac{n\pi x}{L}$.) Thus, in the series equality above, all terms in the summation are zero except for the one where $n = m$, in which case we get:

$$\int_0^L f(x) \sin \tfrac{m\pi x}{L}\, dx = \sum_{n=1}^{\infty} a_n \int_0^L \sin \tfrac{n\pi x}{L} \sin \tfrac{m\pi x}{L}\, dx = a_m \int_0^L \sin \tfrac{m\pi x}{L} \sin \tfrac{m\pi x}{L}\, dx = a_m \frac{L}{2}.$$

Thus, after dividing by $\frac{L}{2}$, we get that a_n must be given by:

$$a_n = \frac{2}{L} \int_0^L f(x) \sin \frac{n\pi x}{L}\, dx.$$

Initial conditions. Recall that so far we have the following formal solution to our wave equation with boundary conditions:

$$u(x,t) = \sum_{n=1}^{\infty} u_n(x,t) = \sum_{n=1}^{\infty} \sin(\tfrac{n\pi x}{L})[A_n \cos(\tfrac{cn\pi t}{L}) + B_n \sin(\tfrac{cn\pi t}{L})].$$

Now we impose our initial conditions, which give the initial position and speed of any point on the string at time $t = 0$:

$$u(x,0) = f(x),\ u_t(x,0) = g(x).$$

Using the formal solution above, the first condition turns into:

$$f(x) = u(x,0) = \sum_{n=1}^{\infty} \sin(\tfrac{n\pi x}{L})[A_n \cos(\tfrac{cn\pi 0}{L}) + B_n \sin(\tfrac{cn\pi 0}{L})] = \sum_{n=1}^{\infty} A_n \sin(\tfrac{n\pi x}{L})$$

and the second turns into:

$$g(x) = u_t(x,0) = \sum_{n=1}^{\infty} \sin(\tfrac{n\pi x}{L})[-A_n \tfrac{cn\pi}{L} \sin(\tfrac{cn\pi 0}{L}) + B_n \tfrac{cn\pi}{L} \cos(\tfrac{cn\pi 0}{L})] = \sum_{n=1}^{\infty} B_n \tfrac{cn\pi}{L} \sin(\tfrac{n\pi x}{L}).$$

Thus, if our initial conditions are to be satisfied, the functions f and g must be expressible as these infinite series expressions.

But from the Warm-Up we then know what the so-far-undetermined coefficients A_n, B_n have to be: A_n must be equal to

$$A_n = \frac{2}{L} \int_0^L f(x) \sin \frac{n\pi x}{L}\, dx$$

and $B_n \frac{cn\pi}{L}$ (the coefficient of $\sin \frac{n\pi x}{L}$) must be equal to

$$B_n \frac{cn\pi}{L} = \frac{2}{L} \int_0^L g(x) \sin \frac{n\pi x}{L}\, dx, \text{ and hence } B_n = \frac{2}{cn\pi} \int_0^L g(x) \sin \frac{n\pi x}{L}\, dx.$$

Hence, with these specific coefficients, we have found the solution to our initial value wave equation with boundary conditions.

Example. ***TO BE FINISHED***

Eigenvalue problems. Before moving on, it will be useful to recast one thing we did in deriving the formal solution to the wave equation in linear-algebraic terms. Consider the equation we obtained for X, together with its boundary conditions:

$$X'' = \lambda X,\ X(0) = 0 = X(L).$$

Letting $D = \frac{d}{dx}$ denote the differentiation operator, this can be written as

$$D^2 X = \lambda X, X(0) = 0 = X(L).$$

But $D^2 X = \lambda X$ says precisely that X should be an "eigenvector" of the linear operator D^2 with eigenvalue λ. In this context, we call X an *eigenfunction* of D^2, since it is actually a function. So, the point is that solving for such X can be recast as the problem of finding (nontrivial) eigenfunctions and eigenvalues of this linear differential operator, taking the boundary conditions into account as well. We will refer to this setup as an *eigenvalue problem*, and we will consider more problems of this type as we go. This particular eigenvalue problem, as we derived last time, has solutions (i.e. eigenvalues and eigenfunctions) given by:

$$\lambda_n = -(\tfrac{n\pi}{L})^2 \quad \text{and} \quad X_n = \sin \tfrac{n\pi x}{L}$$

Towards Fourier series. So, we have seen that solutions of the wave equation can be expressed as infinite series of various sine and cosine functions, and in particular the initial conditions amount

to expressing single-variable functions $f(x), g(x)$ as infinite series of sine functions alone. Such a series is called a *Fourier series*, which in the most general form looks like:

$$A_0 + \sum_{n=0}^{\infty} (A_n \cos \tfrac{n\pi x}{L} + B_n \sin \tfrac{n\pi x}{L}).$$

When all B_n are zero, we call this a Fourier *cosine* series, and when all A_n are zero we call this a Fourier *sine* series. Just as with the Fourier sine series which show up in the initial conditions, it turns out that the coefficients A_n, B_n can all be determined explicitly. We will come back to this next time, and say a bit about the *convergence* of such series.

Lecture 15: Fourier series

Warm-Up. ***TO BE FINISHED***

Fourier series. A *Fourier series* is an infinite series of the form

$$\frac{a_0}{2} + \sum_{n=1}^{\infty} (a_n \cos \tfrac{n\pi x}{L} + b_n \sin \tfrac{n\pi x}{L}).$$

If we want to express a function $f(x)$ in this form:

$$f(x) = \frac{a_0}{2} + \sum_{n=1}^{\infty} (a_n \cos \tfrac{n\pi x}{L} + b_n \sin \tfrac{n\pi x}{L})$$

it turns out that the coefficients a_n, b_n needed must be given by the following integrals:

$$a_n = \frac{1}{L} \int_{-L}^{L} f(x) \cos \tfrac{n\pi x}{L}\, dx \qquad b_n = \frac{1}{L} \int_{-L}^{L} f(x) \sin \tfrac{n\pi x}{L}, dx.$$

This comes from the same computation used in the Warm-Up last time to determine the coefficients needed when solving the wave equation with initial conditions. We will do this exact computation in a much more general way as a Warm-Up next time, but let us go through it in this specific case here anyway, at least for finding a_n.

Multiplying the required Fourier series equality for $f(x)$ through by $\cos \frac{m\pi x}{L}$ gives:

$$f(x) \cos \tfrac{m\pi x}{L} = \tfrac{a_0}{2} \cos \tfrac{m\pi x}{L} + \sum_{n=1}^{\infty} (a_n \cos \tfrac{n\pi x}{L} \cos \tfrac{m\pi x}{L} + b_n \sin \tfrac{n\pi x}{L} \cos \tfrac{m\pi x}{L}).$$

Now we integrate both sides:

$$\int_{-L}^{L} f(x) \cos \tfrac{m\pi x}{L}\, dx = \tfrac{a_0}{2} \int_{-L}^{L} \cos \tfrac{m\pi x}{L}\, dx + \sum_{n=1}^{\infty} (a_n \int_{-L}^{L} \cos \tfrac{n\pi x}{L} \cos \tfrac{m\pi x}{L}\, dx + b_n \int_{-L}^{L} \sin \tfrac{n\pi x}{L} \cos \tfrac{m\pi x}{L})\, dx.$$

The following identities, called the *orthogonality relations*, come from some clever use of trig identities, but we will omit the computations here:

$$\int_{-L}^{L} \cos \tfrac{n\pi x}{L} \cos \tfrac{m\pi x}{L}\, dx = \begin{cases} 0 & m \neq n \\ L & m = n \neq 0 \\ 2L & m = n = 0 \end{cases}$$

$$\int_{-L}^{L} \cos\tfrac{n\pi x}{L} \sin\tfrac{m\pi x}{L}\, dx = 0$$

$$\int_{-L}^{L} \sin\tfrac{n\pi x}{L} \sin\tfrac{m\pi x}{L}\, dx = \begin{cases} 0 & m \neq n \text{ or } m = n = 0 \\ L & m = n \end{cases}$$

Thus, if $m \neq 0$ the only nonzero integral on the right-side of the expression derived above is the one with the product of two cosine terms when $n = m$ (thinking of the integrand in the first integral as $\cos\frac{0\pi x}{L}\cos\frac{m\pi x}{L}$ shows that the orthogonality relations apply to this integral as well), so the entire expression above simplifies to

$$\int_{-L}^{L} f(x)\cos\tfrac{m\pi x}{L}\, dx = a_m \int_{-L}^{L} \cos\tfrac{m\pi x}{L}\cos\tfrac{m\pi x}{L}\, dx = a_m L.$$

Hence we have

$$a_m = \frac{1}{L}\int_{-L}^{L} f(x)\cos\tfrac{m\pi x}{L}\, dx$$

as claimed. If $m = 0$ in the series expression above, then it is the first integral on the right-side which is nonzero with all others being zero, so we get

$$\int_{-L}^{L} f(x)\cos\tfrac{0\pi x}{L}\, dx = \frac{a_0}{2}\int_{-L}^{L} \cos\tfrac{0\pi x}{L}\cos\tfrac{0\pi x}{L}\, dx = \frac{a_0}{2}2L = a_0 L$$

and thus

$$a_0 = \frac{1}{L}\int_{-L}^{L} f(x)\cos\tfrac{0\pi x}{L}\, dx = \frac{1}{L}\int_{-L}^{L} f(x)\, dx$$

as well. (The fact that integral with $\cos\frac{0\pi x}{L}\cos\frac{0\pi x}{L}$ is $2L$ instead of simply L is the reason for writing the constant term in the Fourier series as $\frac{a_0}{2}$: if we had used a_0 alone as the constant term, the integral formula for this coefficient would have to be modified to have a denominator of $2L$ instead of L, so that it would not be of the same form as the formula for other a_n. In order to have a consistent formula which applies to all a_n we thus instead incorporate this extra $\frac{1}{2}$ into the Fourier series expression itself.) The same argument works for the sine coefficients, along the lines of the Warm-Up from last time.

Example. ***TO BE FINISHED***

Convergence. Now we should say something about the convergence of a Fourier series, since, after all, trying to construct solutions of PDEs as such infinite series only makes literal sense when these series actually produce finite results. Here is the basic result:

> If f is a piecewise continuously differentiable (also called piecewise C^1) function on $[-L, L]$, the Fourier series of f converges to $f(x)$ itself at any point x where f is continuous, and to the average $\frac{1}{2}(f(x^-) + f(x^+))$ of the left- and right-hand limits of f at x at points where f is not continuous.

Proving this is beyond the scope of our course (indeed, it is a difficult thing to prove in this generality), but nonetheless the convergence is easy to determine. As long as we only deal with piecewise continuously differentiable functions—say in specifying initial conditions—we are guaranteed that our solutions formulated as Fourier-type series will indeed converge.

TO BE FINISHED

Inner products. Finally, we provide a more coherent understanding of what is truly underlying the concept of a Fourier series. The key notion is that of an *inner product*, which is meant to generalize the dot product used in linear algebra. In that case, the dot product of two vectors $\mathbf{x}, \mathbf{y}$ in $\mathbb{R}^n$ is defined to be:

$$\mathbf{x} \cdot \mathbf{y} = x_1 y_1 + \cdots + x_n y_n$$

where the x_i and y_i are the components of $\mathbf{x}$ and $\mathbf{y}$ respectively. One basic fact is $\mathbf{x} \cdot \mathbf{y} = 0$ if and only if $\mathbf{x}$ and $\mathbf{y}$ are *orthogonal* to one another.

Now, given functions f and g, we define their *inner product* on an interval $[a, b]$ to be

$$\langle f, g \rangle = \int_a^b f(x) g(x)\, dx.$$

The point is that this is an infinite-dimensional version of the usual dot product: we multiply corresponding "components" $f(x)$ and $g(x)$ of f and g and then "add" (i.e. integrate) the results together. (The term "inner product" is used for such more general versions of the dot product. We will not give the precise meaning of this term here, but will point out at least that inner products always satisfy a linearity property: $\langle c_1 f_1 + c_2 f_2, g \rangle = c_1 \langle f_1, g \rangle + c_2 \langle f_2, g \rangle$ and similarly if we instead take a linear combination in the second argument of the inner product. General inner products are covered more fully in a higher-level linear algebra course, such as Math 334.)

Following the terminology used in the case of $\mathbb{R}^n$, we say that f and g are *orthogonal* if their inner product is zero: $\langle f, g \rangle = 0$. We will see more clearly what key properties orthogonality gives us next time, but for now we highlight the fact that the "orthogonality relations" given before say precisely that the functions $\cos \frac{n\pi x}{L}$ and $\sin \frac{m\pi x}{L}$ are all orthogonal to each other on the interval $[-L, L]$: cosines and sines are always orthogonal to each other, and two cosines are orthogonal when they occur for different n, and similarly for two sines. This is the absolutely crucial property which makes the theory of Fourier series work. The upshot is that these sine and cosines functions form an *orthogonal basis* for an appropriate space of functions, where we use the term "basis" in the linear-algebraic sense: they are "linearly independent" (this just comes from being orthogonal) and any function can be written as a linear combination (infinite in this case) of them, so that these "span" the space of all (piecewise continuous) functions.

Lecture 16: More on Fourier and Wave

Warm-Up. Suppose e_n give an orthogonal basis of some space with respect to some inner product. Then any x can be expressed in terms of this basis as

$$x = \sum c_n e_n$$

and we determine the values of c_n explicitly. (We are not making any assumptions on how large this basis is, so this sum could in fact be an infinite series.) Taking the inner products of both sides with some e_m gives:

$$\langle x, e_m \rangle = \langle c_1 e_1 + c_2 e_2 + \cdots, e_m \rangle = c_1 \langle e_1, e_m \rangle + c_2 \langle e_2, e_m \rangle + \cdots,$$

where we use the linearity products of the inner product to break up the right side: break up into individual sums, and then pull the scalar out of each. Since the e_n are orthogonal, the only nonzero inner product occurring on the right is $\langle e_m, e_m \rangle$, so the above expression simplifies to

$$\langle x, e_m \rangle = c_m \langle e_m, e_m \rangle.$$

Thus, $c_m = \frac{\langle x, e_m \rangle}{\langle e_m, e_m \rangle}$, so we conclude that

$$x = \sum \frac{\langle x, e_m \rangle}{\langle e_m, e_m \rangle} e_m$$

is how we can explicitly write x in terms of our orthogonal basis. (In linear algebra, the $\frac{\langle x, e_m \rangle}{\langle e_m, e_m \rangle} e_m$ terms is called the *orthogonal projection* of x onto e_m, so the result is that x can be written as a sum of these orthogonal projections.)

This explains the formulas for Fourier coefficients we saw before: if we use the inner product $\langle f, g \rangle = \int_{-L}^{L} f(x)g(x)\,dx$ on functions, then the sine and cosine functions used are orthogonal by the orthogonality relations, and the coefficient of $\sin \frac{n\pi x}{L}$ in the Fourier series of $f(x)$ is

$$\frac{\left\langle f(x), \sin \frac{n\pi x}{L} \right\rangle}{\left\langle \sin \frac{n\pi x}{L}, \sin \frac{n\pi x}{L} \right\rangle} = \frac{\int_{-L}^{L} f(x) \sin \frac{n\pi x}{L}}{\int_{-L}^{L} \sin \frac{n\pi x}{L} \sin \frac{n\pi x}{L}} = \frac{1}{L} \int_{-L}^{L} f(x) \sin \frac{n\pi x}{L}\,dx$$

and similarly for the cosine coefficients. With the inner product $\langle f, g \rangle = \int_0^L f(x)g(x)\,dx$ instead, we get the coefficients of the Fourier sine series of $f(x)$

$$\frac{\left\langle f(x), \sin \frac{n\pi x}{L} \right\rangle}{\left\langle \sin \frac{n\pi x}{L}, \sin \frac{n\pi x}{L} \right\rangle} = \frac{\int_0^L f(x) \sin \frac{n\pi x}{L}}{\int_0^L \sin \frac{n\pi x}{L} \sin \frac{n\pi x}{L}} = \frac{1}{\frac{L}{2}} \int_{-L}^{L} f(x) \sin \frac{n\pi x}{L}\,dx = \frac{2}{L} \int_0^L f(x) \sin \frac{n\pi x}{L}\,dx$$

and similarly for the Fourier cosine series of $f(x)$. The study of Fourier series thus essentially *is* simply the study of expansions with respect to particular orthogonal bases of functions.

Eigeninterpretation. We mentioned before that the functions $\sin \frac{n\pi x}{L}$ used in the formal solution of the wave equation come from solving an *eigenvalue problem*:

$$X'' = \lambda X,\ X(0) = 0 = X(L)$$

and in particular are eigenfunctions of the linear differential operator D^2. If we ignore the boundary conditions, the functions $\cos \frac{n\pi x}{L}$ used in the full-blown Fourier series are also eigenfunctions of D^2, so a Fourier series is nothing but an expansion in terms of orthogonal eigenfunctions of D^2. The question is: what about D^2 suggests that such orthogonal eigenfunctions will exist, and what makes it possible to express an arbitrary function in terms of these eigenfunctions? We will see soon that the answer is that D^2 is an example of a *symmetric* operator, which are infinite-dimensional analogs of symmetric matrices. *This* is the underlying reason why we should expect such Fourier-series type solutions of the wave equation, and, as we'll see, other PDEs as well.

Complex Fourier series.

Back to wave. ***TO BE FINISHED***

Where are we? Let us recap what we've done in our approach to solving the wave equation, using the language and ideas we have now further developed. We start with the wave equation:

$$u_{tt} = c^2 u_{xx}.$$

We look for separated solutions $u(x,t) = X(x)T(t)$, and see that X and T individually must satisfy some ODEs:

$$X'' = \lambda X \qquad T'' = c^2 \lambda T.$$

Next, we impose boundary conditions $u(0,t) = 0 = u(L,t)$, which turn into boundary conditions on X and give rise the an eigenvalue problem:

$$X'' = \lambda X,\ X(0) = 0 = X(L).$$

The fact that the linear differential operator D^2 which describes the left-hand side is symmetric with respect to the inner product $\langle f, g\rangle = \int_0^L f(x)g(x)\,dx$ (keeping the boundary conditions in mind) implies (as we will discuss) that the solutions $X_n = \sin\frac{n\pi x}{L}$ of this eigenvalue problem will form an orthogonal eigenfunction basis for our space of functions, so that will be able to use these to express our desired solution.

Then the ODE for T yields solutions $T_n = c_1\cos\frac{cn\pi t}{L} + c_2\sin\frac{cn\pi t}{L}$, and thus we have the separated solutions $u_n(x,t) = X_n(x)T_n(t)$. The general solution is then a linear combination of all of these, and the initial conditions allow us to determine any unknown coefficients via a Fourier coefficient integral computation. This outline is the same approach we will take towards the heat equation next time, and eventually towards other possible PDEs. The role of orthogonal eigenfunctions is absolutely crucial!

Lecture 17: Heat Equation

Warm-Up. We derive ODEs satisfied by the components of a separated solution to the 2-dimensional wave equation

$$u_{tt} = c^2(u_{xx} + u_{yy}).$$

Here, the three-variable function $u(x,y,t)$ describes the height at the point (x,y) and time t of a two-dimensional wave, say for instance caused by a vibrating rubber membrane. (Higher-dimensional wave equations simply add more second-order terms on the right.) If $u(x,y,t) = X(x)Y(y)T(t)$ is to be a solution, we must have

$$XYT'' = c^2(X''YT + XY''T).$$

Assuming X, Y, T are all nonzero, this gives

$$\frac{T''}{c^2T} = \frac{X''Y + XY''}{XY}.$$

The left side depends only on t, and the right only on x and y, so both sides must be constant:

$$\frac{T''}{c^2T} = \lambda = \frac{X''Y + XY''}{XY}.$$

Then the second equality can be written as

$$\frac{X''}{X} = \lambda - \frac{Y''}{Y},$$

so as before both sides must be constant since they depend on different variables:

$$\frac{X''}{X} = \mu = \lambda - \frac{Y''}{Y}.$$

Thus we find that X, Y, T satisfy:

$$T'' = \lambda c^2 T \quad X'' = \mu X \quad Y'' = (\lambda - \mu)Y$$

where λ, μ are constants. With appropriate boundary conditions, solving these will lead to a two-variable Fourier series-like expression.

The heat equation. The next equation we will consider is known as the *heat equation*, since it describes the distribution of heat in a metal rod. If $u(x,t)$ denotes the temperature at position x along the rod and time t, the heat equation says that $u(x,t)$ must satisfy:

$$u_t = \alpha^2 u_{xx}.$$

This is almost the same as the wave equation, except that it is only first-order in t. The positive constant α^2 depends on the material the rod is made of and its "thermal conductivity". We will not derive this equation here (you can check the book), but we point out that it comes from Newton's Second Law of Cooling, which says that the change in temperature in a body is proportional to the difference between its current temperature and the ambient temperature of the region it sits within. (Going from this to the actual heat equation still takes some work.)

Separated and formal solutions. Let us impose the boundary conditions $u(0,t) = 0 = u(L,t)$, which say that the temperature at the endpoints of the rod is always zero. Then we can look for separated solutions $u(x,t) = X(x)T(t)$, where we see that:

$$XT' = \alpha^2 X''T \rightsquigarrow \frac{T'}{\alpha^2 T} = \frac{X''}{X}.$$

As with the wave equation, both sides must be constant, and so we get the pair of ODEs:

$$X'' = \lambda X \quad \text{and} \quad T' = \lambda \alpha^2 T.$$

The boundary conditions turns into $X(0) = 0 = X(L)$, so we get precisely the same eigenvalue problem for X we had for the wave equation:

$$X'' = \lambda X,\ X(0) = 0 = X(L).$$

Borrowing the work we did earlier, this has nontrivial solutions $X_n = \sin \frac{n\pi x}{L}$ with corresponding eigenvalues $\lambda_n = -(\frac{n\pi}{L})^2$.

The equation for T is now first-order:

$$T' = -\alpha^2 (\tfrac{n\pi}{L})^2 T,$$

which has an exponential solution:

$$T_n = e^{-(\frac{\alpha n\pi}{L})^2 t}.$$

Thus we get a formal solution which looks like:

$$u(x,t) = \sum_{n=1}^{\infty} X_n(x) T_n(t) = \sum_{n=1}^{\infty} A_n e^{-(\frac{\alpha n\pi}{L})^2 t} \sin \tfrac{n\pi x}{L}.$$

With an initial condition $u(x,0) = f(x)$ describing the initial temperature at every point, we get:

$$f(x) = \sum_{n=1} A_n \sin \tfrac{n\pi x}{L},$$

so that the A_n are the Fourier sine coefficients of $f(x)$:

$$A_n = \frac{2}{L} \int_0^L f(x) \sin \frac{n\pi x}{L}\, dx.$$

Since the heat equation is first-order in t, we only need specify one initial condition, as opposed to the wave equation where we also needed $u_t(x,0)$.

Example. ***TO BE FINISHED***

Other boundary conditions. We can impose other types of boundary conditions instead of $u(0,t) = 0 = u(L,t)$. For instance, we can ask that the ends of the rod be *insulated*, in the sense that heat does not pass through them with respect to position:

$$u_x(0,t) = 0 = u_x(L,t).$$

Here we end up with the following eigenvalue problem for X:

$$X'' = \lambda X,\ X'(0) = 0 = X'(L).$$

As with the version where the ends of the rod where held to be at zero temperature throughout all time, we can compute that for $\lambda > 0$ there are no nontrivial solutions: $X = c_1 e^{\sqrt{\lambda}x} + c_2 e^{-\sqrt{\lambda}x}$ satisfies $X'(0) = 0 = X'(L)$ only when $c_1 = c_2 = 0$. Now, for $\lambda_0 = 0$ we have the ODE $X'' = 0$, which has general solution

$$X = c_1 + c_2 x.$$

In this case the boundary conditions only require that $c_2 = 0$ but place no restriction on c_1! Thus $\lambda = 0$ *is* a valid eigenvalue here, with corresponding eigenfunction $X_0 = 1$.

For $\lambda = -k^2 < 0$, we have the general solution

$$X = c_1 \cos kx + c_2 \sin kx.$$

The boundary condition $X'(0) = 0$ forces $c_2 = 0$, and then $X'(L) = 0$ requires that

$$-c_1 k \sin kL = 0.$$

Thus we get $k = \frac{n\pi}{L}$, so that $\lambda_n = -(\frac{n\pi}{L})^2$ is an eigenvalue, with the difference between this and the previous boundary conditions being that now the eigenfunction comes from the cosine term: $X_n = \cos\frac{n\pi x}{L}$. The equation for T is the same as before ($T' = -\alpha^2(\frac{n\pi}{L})^2 T$), so we get separated solutions of the form

$$u_n(x,t) = X_n(x)T_n(t) = e^{-(\frac{\alpha n\pi}{L})^2 t}\cos\tfrac{n\pi x}{L}.$$

(Note this applies even for $n = 0$: in that case the equation for T is $T' = 0$, which has constant functions as solutions, and thus gives $T_0 = e^0 = 1$, so that $u_0 = X_0 T_0 = 1$ as well.) The formal solution is then

$$u(x,t) = \frac{A_0}{2}u_0(x,t) + \sum_{n=1}^{\infty} A_n u_n(x,t) = \frac{A_0}{2} + \sum_{n=1}^{\infty} A_n e^{-(\frac{\alpha n\pi}{L})^2 t}\cos\tfrac{n\pi x}{L},$$

and with initial condition $u(x,0) = f(x)$, we get

$$f(x) = \frac{A_0}{2} + \sum_{n=1}^{\infty} A_n \cos\tfrac{n\pi x}{L},$$

so that the A_n are the Fourier *cosine* coefficients of f:

$$A_n = \frac{2}{L}\int_0^L f(x)\cos\frac{n\pi x}{L}\,dx.$$

Wanting this specific integral to give the correct value for all n, including $n = 0$, is, as usual, why we took the constant term in our cosine series to be $\frac{A_0}{2}$ instead of just A_0.

Properties of heat solutions. ***TO BE FINISHED***

Lecture 18: Sturm-Liouville Theory

Warm-Up. We solve the heat equation $u_t = u_{xx}$ with boundary conditions $u(0,t) = 0 = u_x(\pi, t)$ (so one end is held at constant zero temperature and the other is insulated) and initial condition $u(x,0) = 2\sin\frac{3x}{2} + 5\sin\frac{7x}{2}$. The boundary conditions give the following eigenvalue problem for X:

$$X'' = \lambda X,\ X(0) = 0,\ X'(\pi) = 0.$$

For $\lambda > 0$, one can check that there are no nontrivial solutions: in this case we get the general solution $X = c_1 e^{\sqrt{\lambda}x} + c_2 e^{-\sqrt{\lambda}x}$, and the boundary conditions only held for $c_1 = c_2 = 0$. For $\lambda = 0$, $X = c_1 + c_2 x$ and again this satisfies the boundary conditions only for $c_1 = c_2 = 0$.

Thus we are left considering $\lambda = -k^2 < 0$. The general solution for X is

$$X = c_1 \cos kx + c_2 \sin kx.$$

The boundary condition $X(0) = 0$ forces $c_1 = 0$, and the second $X'(\pi) = 0$ then requires:

$$c_2 k \cos k\pi = 0.$$

Since $\cos k\pi = 0$ only when $k = \frac{1}{2}(2n+1)$ is half of an odd integer, we get that this eigenvalue problem has eigenvalues $\lambda_n = -(\frac{2n+1}{2})^2$ and eigenfunctions $X_n = \sin\frac{(2n+1)x}{2}$. Then from

$$T' = \lambda T = -(\tfrac{2n+1}{2})^2 T$$

we get $T_n = e^{-(\frac{2n+1}{2})^2 t}$. Thus the separated solutions are

$$u_n(x,t) = X_n(x)T_n(t) = e^{-(\frac{2n+1}{2})^2 t}\sin\tfrac{(2n+1)x}{2},$$

and the formal solution is

$$u(x,t) = \sum_{n=1}^{\infty} A_n e^{-(\frac{2n+1}{2})^2 t}\sin\tfrac{(2n+1)x}{2}.$$

The given initial condition gives:

$$2\sin\tfrac{3x}{2} + 5\sin\tfrac{7x}{2} = \sum_{n=1}^{\infty} A_n \sin\tfrac{(2n+1)x}{2},$$

so on the one hand we have:

$$A_n = \frac{2}{\pi}\int_0^{\pi} (2\sin\tfrac{3x}{2} + 5\sin\tfrac{7x}{2})\sin\tfrac{(2n+1)x}{2}\,dx$$

which can be computed using the orthogonality relations. On the other hand, we can note that the function $2\sin\frac{3x}{2} + 5\sin\frac{7x}{2}$ is already in the form of a Fourier sine series, so in order to have coefficients on both sides of

$$2\sin\tfrac{3x}{2} + 5\sin\tfrac{7x}{2} = \sum_{n=1}^{\infty} A_n \sin\tfrac{(2n+1)x}{2},$$

match up, we need: $A_1 = 2, A_3 = 5$ and $A_n = 0$ for $n \neq 1, 3$. So, the solution of our initial value heat problem with given boundary conditions is:

$$u(x,t) = 2e^{-(\frac{3}{2})^2 t}\sin\tfrac{3x}{2} + 5e^{-(\frac{7}{2})^2 t}\sin\tfrac{7x}{2}$$

The Laplace equation. Before moving on let us briefly mention one more well-known PDE, the *Laplace equation*: $u_{xx} + u_{yy} = 0$ for a function $u(x, y)$ of two variables. Functions satisfying this equation are called *harmonic*, and are important in numerous areas of mathematics and elsewhere. Some sample solutions can be found using the same variable-separation method we've seen. (Note that if we rewrite the Laplace equation as $u_{xx} = -u_{yy}$, then it looks pretty similar to the wave equation, except for the negative sign.) But here is one interpretation which we can make sense of based on what we've done so far: if we consider the 2-dimensional heat equation

$$u_t = \alpha^2(u_{xx} + u_{yy})$$

we see that $u_{xx}+u_{yy} = 0$ describes solutions which satisfy $u_t = 0$ and hence are constant with respect to time. Thus, one interpretation is that the Laplace equation describes *steady state* solutions of the heat equation. This is not the only interpretation, but is all we will say here.

There is, however, one more reason as to why the Laplace equation is worth mentioning. Last time we briefly introduced the terms *parabolic* and *hyperbolic* when describing the types of PDEs of which the heat and wave equations were prototypical examples. The Laplace equation is the main example of an *elliptic* PDE. In some sense, after making an appropriate "change of variables" it turns out that all elliptic equations can essentially be turned into the Laplace equation, and in fact all hyperbolic and parabolic equations can be turned into the wave and heat equations respectively. This is why these three particular equations are of so much interest: they essentially cover (almost) the full range of nice PDEs in general.

Symmetric operators. We have mentioned before the idea that much of what (conceptually) underlies our derivation of solutions to the wave and heat equations comes from properties of the $D^2 = \frac{d^2}{dx^2}$ operator, namely the fact that it gives rise to an orthogonal basis of eigenfunctions for the space of (ususally piecewise continuous) functions satisfying boundary conditions in which we are interested. We now aim to make this precise.

Given an inner product $\langle \cdot, \cdot \rangle$ on a space, we say that a linear operator L is *symmetric* (also called *self-adjoint*) if it has the following property for all u and v:

$$\langle Lu, v\rangle = \langle u, Lv\rangle .$$

(In general, the *adjoint* or *transpose* of L is the linear operation L^* which satisfies $\langle Lu, v\rangle = \langle u, L^*v\rangle$, so operators which are symmetric are ones which equal their own adjoint.) Being symmetric is what underlies the orthogonality and existence of eigenvectors, as we will see.

In the case of the usual dot product on $\mathbb{R}^n$, where $L = A$ is an $n \times n$ matrix, it is the ordinary matrix transpose A^t which satisfies $\langle Au, v\rangle = \langle u, A^t v\rangle$, so A being symmetric in this case means the usual notion of symmetric $A^t = A$. In a linear algebra course you would have seen that for symmetric matrices, there is an orthogonal basis of $\mathbb{R}^n$ consisting of eigenvectors of that matrix, which is precisely the result we want to generalize. This type of result goes by the name of *Spectral Theorem* in both the finite- and infinite-dimensional setting.

Now take the space of piecewise continuous functions on $[0, L]$ with the inner product

$$\langle f, g\rangle = \int_0^L f(x)g(x)\,dx.$$

For the differentiation operator $D = \frac{d}{dx}$, we have the following via integration by parts:

$$\langle Df, g\rangle = \int_0^L f'(x)g(x)\,dx$$

$$= f(x)g(x)\Big|_0^L - \int_0^L f(x)g'(x)\,dx$$
$$= f(x)g(x)\Big|_0^L + \langle f, -Dg\rangle\,.$$

If we only consider those functions which satisfy an appropriate boundary condition—in this case being 0 at $x = 0$ and $x = L$—we get $\langle Df, g\rangle = \langle f, -Dg\rangle$, so that the adjoint of D is $-D$. Thus, D is not symmetric, but the point is that D^2 then *is* symmetric:

$$\left\langle D^2 f, g\right\rangle = \langle Df, -Dg\rangle = \left\langle f, (-D)^2 g\right\rangle = \left\langle f, D^2 g\right\rangle.$$

A version of the Spectral Theorem then says that we can expect to find orthogonal eigenfunctions for D which give a basis for our space of functions satisfying boundary conditions, so for instance the $X_n = \sin\frac{n\pi x}{L}$ functions we used in both the wave and heat equations.

But there are other types of inner products we can consider, in the hopes of allowing for more symmetric operators. Even in the $\mathbb{R}^n$ case, say $\mathbb{R}^3$ to be specific, we can consider a "weighted" version of the standard dot product:

$$\langle \mathbf{x}, \mathbf{y}\rangle = 2x_1y_1 + 5x_2y_2 + x_3y_3,$$

where we possibly want to put more emphasis on some specific coordinates. Or, we can consider weighted inner products on functions like:

$$\langle f, g\rangle = \int_{-L}^{L} w(x)f(x)g(x)\,dx$$

for a positive continuous function $w(x)$. A linear operator might not be symmetric with respect to one inner product, but it might be with respect to another, so that some version of the Spectral Theorem will still apply. The only thing which will change is the notion of "orthogonality", but not the existence of a basis of eigenvectors.

Sturm-Liouville operators. More generally, we can consider other second-order linear differential operators beyond D^2:

$$L = a_2(x)D^2 + a_1(x)D + a_0(x).$$

This acts on a function X to give:

$$L[X] = a_1(x)X'' + a_1(x)X' + a_0(x)X.$$

It turns out that in nice situations, such an expression can be rewritten as:

$$L[X] = \frac{(p(x)X')' + q(x)X}{\rho(x)}$$

for appropriate choices of $p(x), q(x)$, and $\rho(x)$. An operator L of this form is called a *Sturm-Liouville* operator, and gives a main example of an operator which is symmetric with respect to the correctly-chosen inner product. The corresponding eigenvalue problem:

$$L[X] = \lambda X \iff \frac{(p(x)X')' + q(x)X}{\rho(x)} = \lambda X \iff (p(x)X')' + q(x)X = \lambda\rho(x)X$$

with boundary conditions is called a *Sturm-Liouville problem*. The case of $L = D^2$ which shows up in the wave and heat equations corresponds to the choice $\rho(x) = 1, p(x) = 1, q(x) = 0$.

Example. Consider for instance the second order operator $L = D^2 + D - 2$ defined by

$$L[X] = X'' + X' - 2X.$$

Take $\rho(x) = e^x$. In order to get L into the Sturm-Liouville form, we need:

$$X'' + X' - 2X = \frac{p'(x)X' + p(x)X'' + q(x)X}{e^x},$$

so $\frac{p'(x)}{e^x} = 1$, $\frac{p(x)}{e^x} = 1$, and $\frac{q(x)}{e^x} = -2$. Thus taking $p(x) = e^x$ and $q(x) = -2e^x$ works, so L is a Sturm-Liouville operator with:

$$L[X] = \frac{(e^x X')' - 2e^x X}{e^x}.$$

The book has exact expressions for $p(x), q(x), \rho(x)$ written out in the book, but the point here is that usually these can be determined on-the-fly by comparing $p'(x)X' + p(x)X'' + q(x)X$ to the equation at hand. In general, $\rho(x)$ should be taken to be some type of exponential, since we will want to guarantee that it is always positive, for reasons which will soon be clear.

Symmetry/self-adjointness. And now the point, that Sturm-Liouville operators are symmetric. Recall that we are considering an operator of the form

$$L[X] = \frac{(p(x)X')' + q(x)X}{\rho(x)},$$

where now we ask that $\rho(x)$ be positive. We then consider the following *weighted* inner product

$$\langle f, g\rangle = \int_0^L \rho(x) f(x) g(x)\, dx.$$

Let us compute $\langle L[f], g\rangle$ and $\langle f, L[g]\rangle$, where we note that the weighting term $\rho(x)$ goes away due to the same thing occurring in the denominator of L, and use integration by parts to "move" derivatives from one function over to the other:

$$\begin{aligned}
\langle L[f], g\rangle &= \int_0^L [(p(x)f'(x))' + q(x)f(x)]g(x)\, dx \\
&= \int_0^L (p(x)f'(x))'g(x)\, dx + \int_0^L q(x)f(x)g(x)\, dx \\
&= p(x)f'(x)g(x)\Big|_0^L - \int_0^L p(x)f'(x)g'(x)\, dx + \int_0^L q(x)f(x)g(x)\, dx \\
\langle f, L[g]\rangle &= \int_0^L f(x)[(p(x)g'(x))' + q(x)g(x)]\, dx \\
&= \int_0^L f(x)(p(x)g'(x))'\, dx + \int_0^L f(x)q(x)g(x)\, dx \\
&= f(x)p(x)g'(x)\Big|_0^L - \int_0^L f'(x)p(x)g'(x)\, dx + \int_) f(x)q(x)g(x)\, dx.
\end{aligned}$$

Comparing these, we see that the only difference comes in the boundary terms:

$$p(x)f'(x)g(x)\Big|_0^L \quad \text{vs} \quad p(x)f(x)g'(x)\Big|_0^L.$$

Thus, if we impose boundary conditions which imply that these are both zero (by either requiring that functions or their derivatives are zero at both 0 and L, or by requiring that functions be periodic so that their values at 0 and L are the same), we see that $\langle L[f], g\rangle = \langle f, L[g]\rangle$ on such a space. So, a Sturm-Liouville operator will be symmetric with respect to this weighted inner product on a space of functions satisfying appropriate boundary conditions. As we'll see next time, this will give us the existence of all the orthogonal eigenfunctions we'll need.

Lecture 19: More on Sturm-Liouville

Warm-Up. We verify that the operator $L = D^2 - 5D + 6$ is a Sturm-Liouville operator, and argue that with boundary conditions $X(0) = 0 = X(1)$ it has no eigenvalues larger than or equal to $-\frac{1}{4}$. First, we have:

$$L[X] = X'' - 5X' + 6X = \frac{(e^{-5x}X')' + 6e^{-5x}X}{e^{-5x}},$$

so that L is in Sturm-Liouville form. Concretely, we use e^{-5x} to account for the coefficient -5 which appears in L: in required $(p(x)X')'$ the term, the coefficient of X' is $p'(x)$, so we need this term to produce -5 in order to match up with the given L, and $p(x) = e^{-5x}$ does the trick. This choice of $\rho(x) = e^{-5x}$ in the denominator is also good since it is positive, meaning that it does serve as an appropriate "weight" in the corresponding inner product:

$$\langle f, g\rangle = \int_0^1 e^{-5x} f(x)g(x)\,dx.$$

The operator $L = D^2 - 5D + 6$ is then symmetric with respect to this inner product on the space of (say continuous) functions satisfying the boundary conditions $X(0) = 0 = X(1)$.

Now we look for eigenvalues of L. The eigenvalue equation $L[X] = \lambda X$ with boundary conditions, of which we want nonzero solutions, in this case is:

$$X'' - 5X' + 6X = \lambda X,\ X(0) = 0 = X(1).$$

This seocnd-order ODE can be rewritten as

$$X'' - 5X' + (6 - \lambda)X = 0,$$

which we know how to solve: the characteristic equation is $r^2 - 5r + (6 - \lambda) = 0$, and so the roots which will determine the form of the solution are given by:

$$r = \frac{5 \pm \sqrt{5^2 - 4(6 - \lambda)}}{2} = \frac{5 \pm \sqrt{1 + 4\lambda}}{2}.$$

For $\lambda > -\frac{1}{4}$ this gives two distinct real roots λ_1, λ_2, so the solution is

$$X = c_1 e^{\lambda_1 x} + c_2 e^{\lambda_2 x}.$$

But here the boundary conditions will force $c_1 = c_2 = 0$:

$$\begin{cases} X(0) = c_1 + c_2 = 0 \\ X(1) = c_1 e^{\lambda_1} + c_2 e^{\lambda_2} = 0 \end{cases} \quad \rightsquigarrow \quad \begin{bmatrix} 1 & 1 \\ e^{\lambda_1} & e^{\lambda_2} \end{bmatrix} \begin{bmatrix} c_1 \\ c_2 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix} \quad \rightsquigarrow \quad \begin{bmatrix} c_1 \\ c_2 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$$

where the matrix in the middle is invertible since $\lambda_1 \neq \lambda_2$. So, there are no nonzero solutions here and thus no eigenvalues larger than $-\frac{1}{4}$. For $\lambda = -\frac{1}{4}$, the characteristic equation has a repeated root λ, so the solution looks like

$$X = c_1 e^{\lambda x} + c_2 x e^{\lambda x}.$$

But again the boundary conditions force this to be trivial: $X(0) = 0$ becomes $c_1 = 0$, and then $X(1) = 0$ implies $c_2 = 0$. So $\lambda = -\frac{1}{4}$ is also not an eigenvalue. The upshot is that all eigenvalues of this specific operator (with the given boundary conditions) must be less than $-\frac{1}{4}$. We will finish computing these in a bit.

Eigenproperties. Let us clarify some important properties related to the eigenspaces of a Sturm-Liouville operator L. First, we claim that the eigenspaces corresponding to different eigenvalues are orthogonal to one another: if f, g are eigenfunctions with eigenvalues $\lambda \neq \mu$ respectively, then f and g are orthogonal. (This is what underlies the orthogonality relations we used previously in deriving the coefficients of a Fourier series! We'll highlight this in a bit.) This actually has nothing to do with the specifics of Sturm-Liouville theory at all, and is just a general fact about the linear algebra of symmetric operators. Indeed, if $Lf = \lambda f$ and $Lg = \mu g$, with $\lambda \neq \mu$, then

$$\langle Lf, g\rangle = \langle \lambda f, g\rangle = \lambda \langle f, g\rangle \text{ and } \langle f, Lg\rangle = \langle f, \mu g\rangle = \mu \langle f, g\rangle .$$

But if L is symmetric then these two expressions should be equal, which since λ and μ are different requires that $\langle f, g\rangle = 0$, so f and g are orthogonal as claimed.

(If you have studied complex inner products before, you might point out that technically the second computation above should really result in $\langle f, \mu g\rangle = \overline{\mu} \langle f, g\rangle$ based on argument, so that the reasoning above only works if $\mu = \overline{\mu}$ is real. But, this is really no restriction at all: a symmetric operator can only have real eigenvalues in general. We will not prove this here, but it is a standard computation in a linear algebra course which deals with abstract inner products.)

A second key property of L *does* depend on the fact that we are working within the context of Sturm-Liouville theory: the fact that each eigenspace is only 1-dimensional, and hence is spanned by any one choice of an eigenfunction. (This is not true of symmetric operators in general—for instance, it is not true of symmetric matrices.) This is why we get eigenfunctions like $\sin \frac{n\pi x}{L}$, *one for each* n, when solving the wave and heat equations. This all comes from the fact that we are working with second-order linear ODEs:

$$\frac{(p(x)X')' + q(x)X}{\rho(x)} = \lambda X \iff p(x)X'' + p'(x)X + q(x)X = \lambda \rho(x) X.$$

We know from last quarter than the space of solutions to such a (non-driven) second-order linear ODE is 2-dimensional, where any solution can be written as a linear combination of two linearly independent solutions. If we impose *separated* boundary conditions (i.e. boundary conditions which depend on each endpoint one at a time, but do not relate their simultaneous behavior to one another), we get an eigenspace which, if it were 2-dimensional as well, would imply that *all* solutions of the second-order ODE above would necessarily have to satisfy the boundary conditions too, which is not the case. In other words, the space of solutions satisfying boundary conditions is at least 1-dimensional and sits inside the space of all solutions of the second-order ODE, but cannot be equal to this entire space of solutions since not all solutions satisfy the boundary conditions, so it must have dimension strictly less than 2. This fact makes it possible to easily describe *all* eigenfunctions for each eigenvalue, once we have at least one.

Example. Let us return to the Sturm-Liouville operator $L = D^2 - 5D + 6$. We saw before that the eigenvalue problem

$$X'' - 5X' + 6X = \lambda X,\ X(0) = 0 = X(1)$$

can only possibly have a nonzero solution for $\lambda < -\frac{1}{4}$. Let us now determine the values in this range which are *actually* eigenvalues. The roots of the corresponding characteristic equation in this case are:

$$\frac{5 \pm \sqrt{1+4\lambda}}{2} = \frac{5 \pm i\sqrt{|1+4\lambda|}}{2},$$

where we note that $1 + 4\lambda < 0$ so that we need to take its absolute value in the square root. Thus, the general solution to the second-order ODE above is:

$$X = c_1 e^{\frac{5x}{2}} \cos(\tfrac{x\sqrt{|1+4\lambda|}}{2}) + c_2 e^{\frac{5x}{2}} \sin(\tfrac{x\sqrt{|1+4\lambda|}}{2}).$$

The boundary condition $X(0) = 0$ forces $c_1 = 0$, and the boundary condition $X(1) = 0$ requires that:

$$c_2 e^{\frac{5}{2}} \sin(\frac{\sqrt{|1+4\lambda|}}{2}) = 0.$$

Thus we get a nonzero solution only when

$$\frac{\sqrt{|1+4\lambda|}}{2} = n\pi$$

for a nonzero integer n. (Taking $n = 0$ requires $\lambda = -\frac{1}{4}$, which we already ruled out.) Solving this for λ yields eigenvalues of the form

$$\lambda_n = \frac{-4n^2\pi^2 - 1}{4}.$$

(The negative signs show up when getting rid of the absolute value.) Note that all of these, for $n \neq 0$, are indeed strictly smaller than $-\frac{1}{4}$, as expected. Corresponding eigenfunctions are:

$$X_n = e^{\frac{5x}{2}} \sin(\tfrac{x\sqrt{|1+4\lambda_n|}}{2}) = e^{\frac{5x}{2}} \sin(\tfrac{x\sqrt{4n^2\pi^2}}{2}) = e^{\frac{5x}{2}} \sin(n\pi x),$$

and any other eigenfunction must be a scalar multiple of one of these.

Moreover, these eigenfunctions are orthogonal for different n, with respect to the weighted inner product which corresponds to this Sturm-Liouville operator:

$$\langle X_n, X_m \rangle = \int_0^1 e^{-5x} X_n(x) X_m(x)\, dx = \int_0^1 e^{-5x} e^{\frac{5x}{2}} \sin(n\pi x) e^{\frac{5x}{2}} \sin(m\pi x)\, dx = 0$$

for $m \neq n$. This can be a tedious think to check directly using trig identities (note that we did not actually verify the orthogonality relations used in Fourier series directly), but *must* be true from the general eigenproperties of these operators we gave before.

Sturm-Liouville (Spectral) Theorem. Finally, we state the main result about the operators we are considering. As usual, we are omitting the specific assumptions (continuity or otherwise) needed to make this all work out, but you can check the details in the book. These assumptions will hold for all the scenarios we actually care about, for instance with the wave and heat equations. The key facts are:

- the eigenvalues λ_n form a discrete sequence: $\lambda_1, \lambda_2, \lambda_3, \ldots$ of numbers, and

- if we pick an eigenfunction X_n for each, the collection $\{X_n\}$ forms an orthogonal basis for the space of (continuous) functions satisfying the appropriate boundary conditions.

The second property is what guarantees that we will be able to express arbitrary functions satisfying the boundary conditions in terms of these eigenfunctions, with the first property guaranteeing that such an expression will in fact be an infinite series indexed by n, in addition to guaranteeing implicitly that the eigenvalues exist. As mentioned last time, such results are known as *Spectral Theorems*; the term "spectrum" is used to denote (roughly) the set of eigenvalues of an operator, which is where the "spectral" in Spectral Theorem comes from.

Now, finding these eigenvalues and eigenvectors explicitly in the end does come down to solving some ODE, as we went through in the example of $L = D^2 - 5D + 6$. So, this theorem does not give us a way to get around this brute-force computations. Rather, once we have done the required computation, the machinery of this theorem is what tells us that we will be able to express *other* functions in terms of the eigenfunctions we found.

Examples we have seen. And now let us finish by recasting numerous examples we have seen in the past few weeks as instances of this theorem in action. First, we have just worked out here that for the Sturm-Liouville operator $L = D^2 - 5D + 6$ with boundary conditions $X(0) = 0 = X(1)$, the eigenvalues and orthogonal basis eigenfunctions guaranteed to exist by this theorem are:

$$\lambda_n = \frac{-4n^2\pi^2 - 1}{4},\ X_n = e^{\frac{5x}{2}} \sin(n\pi x).$$

In this case we have orthogonality with respect to $\langle f, g\rangle = \int_0^1 e^{-5x} f(x)g(x)\,dx$. Expressing a function in terms in these eigenfunctions yields a series made up of these exponential-sine terms.

The eigenvalue problem we first considered with the wave and heat equations was

$$X'' = \lambda X,\ X(0) = 0 = X(L).$$

The eigenvalues and eigenfunctions we get from this theorem for the Sturm-Liouville operator D^2 are the ones we computed previously:

$$\lambda_n = -(\tfrac{n\pi}{L})^2,\ X_n = \sin\tfrac{n\pi x}{L},$$

where we have orthogonality with respect to $\langle f, g\rangle = \int_0^1 f(x)g(x)\,dx$. The resulting type of series these give rise to is a Fourier sine series.

Another example we considered previously was the heat equation with insulated boundary conditions, which translated to $X'(0) = 0 = X'(L)$. (Here we are talking about the same Sturm-Liouville operator D^2 as before, only with different boundary conditions.) The eigenvalues and eigenfunctions this theorem yields were found to be:

$$\lambda_0 = 0,\ X_0 = 1 \text{ and } \lambda_n = -(\tfrac{n\pi}{L})^2,\ X_n = \cos(\tfrac{n\pi x}{L}).$$

Again we have orthogonality with respect to $\langle f, g\rangle = \int_0^L f(x)g(x)\,dx$, and the type of series we get is a Fourier cosine series.

Finally, consider again the operator D^2, but now with periodic boundary conditions:

$$X(-L) = X(L),\ X'(-L) = X'(L).$$

These are not "separated" since they mix up the behavior at $-L$ with the behavior at L, but are good enough to still imply that D^2 is symmetric, in this case with respect to $\langle f, g\rangle = \int_{-L}^L f(x)g(x)\,dx$. The eigenvalues and eigenfunctions get here are

$$\lambda_0 = 0,\ X_0 = 1 \text{ and } \lambda_n = -(\tfrac{n\pi}{L})^2,\ X_n = \cos\tfrac{n\pi x}{L},\ Y_n = \sin\tfrac{n\pi x}{L}.$$

(So, in this case we actually get two independent eigenfunctions for each eigenvalue, and together the cosine and sine eigenfunctions are needed to obtain all eigenfunctions. In other words, in this case the eigenspace is 2-dimensional; a 1-dimensional eigenspace is only guaranteed when we have *separated* instead of periodic boundary conditions.) The resulting type of series is a full-blown Fourier series with both sine and cosine terms. Note again, that the orthogonality relations for these sine and cosine functions we mentioned previously but avoided verifying directly now come for free simply from the fact that these are eigenfunctions for a symmetric linear operator.

Lecture 20: Eigenfunction Expansions

Warm-Up. We determine the eigenvalues and eigenfunctions of the operator $L = D^2 + 1$ subject to the boundary conditions $X(0) + X'(0) = 0$ and $X(\pi) + X'(\pi) = 0$. Note that this is indeed a Sturm-Liouville operator:

$$L[X] = X'' + X = \frac{(1X')' + 1X}{1},$$

where the point is that since there is no X' term, a more elaborate $p(x)$ showing up in $(p(x)X')'$ in the numerator is not needed. This operator is then symmetric with respect to the inner product

$$\langle f, g\rangle = \int_0^\pi f(x)g(x)\,dx$$

with no additional weighting. Now, the eigenvalue problem to solve is

$$X'' + X = \lambda X,\ X(0) + X'(0) = 0 = X(\pi) + X'(\pi).$$

This second-order ODE can be written as $X'' + (1 - \lambda)X = 0$, and in the cases where $1 - \lambda \leq 0$, it turns out that the only solution which satisfies the given boundary conditions is the zero function, so we move on. (You'll verify this on the homework: write down what the solutions look like in these cases and use the boundary conditions to show that the unknown c_1, c_2 which appear must both be zero.)

So, we assume $1 - \lambda > 0$. Write $1 - \lambda = k^2$ for a positive number k, so that the general solution of $X'' + kX = 0$ is

$$X = c_1 \cos kx + c_2 \sin kx.$$

Now, the boundary condition $X(0) + X'(0) = 0$ gives

$$c_1 + kc_2 = 0,$$

and the boundary condition $X(\pi) + X'(\pi) = 0$ gives

$$c_1 \cos k\pi + c_2 \sin k\pi - kc_1 \sin k\pi + kc_2 \cos k\pi = 0.$$

(Careful: we do not yet know that k is an integer, so we cannot yet say that for instance $\sin k\pi = 0$.) If we write this second condition as

$$(c_1 + kc_2)\cos k\pi + (c_2 - kc_1)\sin k\pi = 0,$$

we see that the first boundary requirement $c + kc_2 = 0$ forces the first term above to be zero, so that we are left with

$$(c_2 - kc_1)\sin k\pi = 0.$$

Thus either $c_2 - kc_1 = 0$ or $\sin k\pi = 0$, but in the first case we thus have the pair of requirements

$$c_1 + kc_2 \quad \text{and} \quad c_2 - kc_1 = 0,$$

which are only satisfied by $c_1 = c_2 = 0$. This is no good, so we must instead have $\sin k\pi = 0$, which says that k must indeed be a positive integer. Thus $1 - \lambda = n^2$ for some positive integer, and thus the eigenvalues are $\lambda_n = 1 - n^2$. The solution of our ODE in this case is

$$X = c_1 \cos nx + c_2 \sin nx,$$

where the first boundary requirement still forces the condition that $c_1 + nc_2 = 0$. Taking $c_2 = 1$ in order to get a nontrivial solution then gives $c_1 = -n$, so the eigenfunctions are given by

$$X_n = -n \cos nx + \sin nx.$$

By the Sturm-Liouville (Spectral) Theorem, these eigenfunctions give an orthogonal basis for the space of (piecewise) continuous functions on $[0, \pi]$. Here, orthogonality means that

$$\langle X_n, X_m \rangle = \int_0^\pi X_n(x) X_m(x)\, dx = \int_0^\pi (-n \cos nx + \sin nx)(-m \cos mx + \sin mx)\, dx = 0$$

when $m \neq n$. This can be verified directly via a tedious integral computation, but of course we know it must be true solely from the fact that the X_n are eigenfunctions of the symmetric operator $L = D^2 + 1$ which corresponding to different eigenvalues.

Eigenfunction expansions. So, we now finally come to the point where we will see why having such eigenfunctions available is good. The idea is that ODEs or PDEs we might want to solve will become simpler to study once we express any function involved in terms of these bases of eigenfunctions. This is analogous to the idea in finite-dimensional linear algebra that matrix equations can be simpler to study if we express everything in terms of coordinates with respect to a well-chosen basis, and we will elaborate on this point of view in a bit.

So, suppose we consider a *driven* ODE of the form $L[X] = f(x)$, where L is a Sturm-Liouville operator, and were we impose some appropriate boundary conditions. The Sturm-Liuoville Theorem gives a basis of eigenfunctions X_n, with corresponding eigenvalues λ_n, which are orthogonal with respect to some appropriate (possibly weighted) inner product. Our goal is to find the solution X to $L[X] = f(x)$. First, we can expand this to-be-determined solution in terms of our eigenbasis:

$$X = \sum_{n=1}^{\infty} c_n X_n.$$

This is something like a Fourier series, only a "Fourier series" *relative* to the specific eigenfunctions X_n. (What we usually refer to as a "Fourier series" is the situation where the eigenfunctions are given by sines and cosines, but really they could be something else in general.) The goal is to find the coefficients c_n needed to make this satisfy our driven equation. Similarly, we can expand the driving term $f(x)$ in terms of this basis:

$$f(x) = \sum_{n=1}^{\infty} d_n X_n.$$

But in fact, we know precisely what the coefficients d_n have to be, due to the orthogonality of the basis functions X_n: we showed at some point last week that these coefficients must be given by

$$d_n = \frac{\langle f(x), X_n \rangle}{\langle X_n, X_n \rangle},$$

so that the $d_n X_n$ term is something like an "orthogonal projection" of f onto X_n.

Now, we can also determine $L[X]$ in terms of this basis. Using linearity of L, we have:

$$L[X] = L\left[\sum_{n=1}^{\infty} c_n X_n\right] = \sum_{n=1}^{\infty} L[c_n X_n] = \sum_{n=1}^{\infty} c_n L[X_n].$$

Actually, there is one subtlety here: linearity for sure allows to break up *finite* linear combinations in this way: $L[c_1 X_1 + \cdots + c_k X_k] = c_1 L[X_1] + \cdots + c_k L[X_k]$, but here we are trying to this to an *infinite* linear combination. This does in fact work, but depends on the fact that L is what's called a *continuous* linear differential operator. We will not go into what this means here, but is worth pointing out nonetheless; continuity—whatever that means in this context—is what allows us to make the jump from finite sums to infinite sums. Going back to the work above, *now* we use the fact that the X_n are eigenfucntions of L to say that $L[X_n] = \lambda_n X_n$, so that the expression above becomes

$$L[X] = \sum_{n=1}^{\infty} c_n \lambda_n X_n.$$

Indeed, this is the entire point of having a basis of *eigenfunctions*: it is easy to describe how a linear operator will behave in general since it will simply scale each eigenfunction by the appropriate eigenvalue. (The point of having our eigenfunctions be *orthogonal* was to determine the coefficients needed in our sum explicitly, as fractions of inner products as written above.) Thus, if we want X to satisfy $L[X] = f(x)$, after having expressed $X, L[X]$, and $f(x)$ in terms of our basis X_n, we are left with the requirement that

$$\sum_{n=1}^{\infty} c_n \lambda_n X_n = \sum_{n=1}^{\infty} d_n X_n, \text{ where } d_n = \frac{\langle f(x), X_n \rangle}{\langle X_n, X_n \rangle}.$$

But in order for this to be true, the coefficients of X_n on both sides must match up (this comes from some type of "linear independence" property), so

$$c_n \lambda_n = d_n, \text{ and thus } c_n = \frac{d_n}{\lambda_n}$$

and we have achieved our goal: determine what the coefficients c_n in our to-be-determined solution $X = \sum c_n X_n$ actually have to be.

The conclusion thus is that the solution of the driven ODE $L[X] = f(x)$, satisfying boundary conditions, is given by

$$X = \sum_{n=1}^{\infty} \frac{d_n}{\lambda_n} X_n, \text{ where } d_n = \frac{\langle f(x), X_n \rangle}{\langle X_n, X_n \rangle}.$$

So, if we can find the eigenvalues λ_n and eigenfunctions X_n, we can solve any such driven equation. (Actually, such a solution will exist only for certain $f(x)$; what we are really saying is that *if* a solution exists, this is what it has to look like, but it turns out that not all choices of driving term $f(x)$ will actually yield an honest solution, due to the nature of working with boundary conditions. You'll see an example of this on the homework.)

Diagonalization. Let us recast what we have just done in the language of linear algebra, to make a connection with a concept you have no doubt seen before. After expanding X in terms of eigenfunctions:

$$X = \sum_{n=1}^{\infty} c_n X_n,$$

the point is that the c_n should be viewed as the "coordinates" of X with respect to the basis X_n, and we can encode these in an "infinite-dimensional" vector:

$$\begin{bmatrix} c_1 \\ c_2 \\ c_3 \\ \vdots \end{bmatrix}.$$

The fact that the X_n are eigenvectors so that

$$L[X] = \sum_{n=1}^{\infty} c_n \lambda_n X_n$$

says that the coordinates of $L[X]$ are $\lambda_n c_n$, so that L has the following effect on "coordinate" vectors, scaling each coordinate by the appropriate eigenvalue:

$$\begin{bmatrix} c_1 \\ c_2 \\ c_3 \\ \vdots \end{bmatrix} \rightsquigarrow \begin{bmatrix} \lambda_1 c_1 \\ \lambda_2 c_2 \\ \lambda_3 c_3 \\ \vdots \end{bmatrix}.$$

But this final expression can be written in terms of multiplication by an infinite-dimensional matrix:

$$\begin{bmatrix} \lambda_1 c_1 \\ \lambda_2 c_2 \\ \lambda_3 c_3 \\ \vdots \end{bmatrix} = \begin{bmatrix} \lambda_1 & & & \\ & \lambda_2 & & \\ & & \lambda_3 & \\ & & & \ddots \end{bmatrix} \begin{bmatrix} c_1 \\ c_2 \\ c_3 \\ \vdots \end{bmatrix},$$

so the upshot is that the effect the operator L has on "coordinates" is given by this matrix, so that L is *represented* by this infinite-dimensional matrix with respect to this particular eigenbasis.

And now the point: this is *precisely* what happens for a *diagonalizable* matrix! To say that a matrix A is diagonalizable is to say that there is a basis of $\mathbb{R}^n$ consisting of eigenvectors of A, or equivalently that $A = SDS^{-1}$ for some diagonal matrix D and invertible matrix S, whose columns give the eigenvector basis. The diagonal matrix D, whose diagonal entries are the eigenvalues, then describes the effect which A has on coordinates relative to the eigenbasis, so that the effect is simply to scale each coordinate by the appropriate eigenvalue. Thus, what we have done here now is to essentially *diagonalize* (or better yet: *orthogonally* diagonalize) the infinite-dimensional Sturm-Liouville operator L! From this perspective, solving $L[X] = f(x)$ amounts to solving for the c_n in the infinite-dimensional matrix equation

$$\begin{bmatrix} \lambda_1 & & & \\ & \lambda_2 & & \\ & & \lambda_3 & \\ & & & \ddots \end{bmatrix} \begin{bmatrix} c_1 \\ c_2 \\ c_3 \\ \vdots \end{bmatrix} = \begin{bmatrix} d_1 \\ d_2 \\ d_3 \\ \vdots \end{bmatrix}$$

where the d_n are the "coordinates" of $f(x)$ relative to the basis of eigenfunctions. But this matrix is invertible when all the eigenvalues are nonzero, and its inverse is still diagonal with diagonal entries being the reciprocals of the λ_n, so we the solution we want is

$$\begin{bmatrix} c_1 \\ c_2 \\ c_3 \\ \vdots \end{bmatrix} = \begin{bmatrix} \lambda_1 & & & \\ & \lambda_2 & & \\ & & \lambda_3 & \\ & & & \ddots \end{bmatrix}^{-1} \begin{bmatrix} d_1 \\ d_2 \\ d_3 \\ \vdots \end{bmatrix} = \begin{bmatrix} \frac{1}{\lambda_1} & & & \\ & \frac{1}{\lambda_2} & & \\ & & \frac{1}{\lambda_3} & \\ & & & \ddots \end{bmatrix} \begin{bmatrix} d_1 \\ d_2 \\ d_3 \\ \vdots \end{bmatrix}.$$

This explains the values $c_n = \frac{d_n}{\lambda_n}$ we derived before, so being able to "diagonalize" L has now allowed us to easily "invert" it in order to solve $L[X] = f(x) \rightsquigarrow X = L^{-1}[f(x)]$.

Example. Consider the second-order driven ODE $X'' + X = \sin x + \cos x$ with boundary conditions $X(0) + X'(0) = 0 = X(\pi) + X'(\pi)$. We determined the eigenvalues and orthogonal basis of eigenfunctions in the Warm-Up:

$$\lambda_n = 1 - n^2 \qquad X_n = -n\cos nx + \sin nx.$$

so we can use the method of eigenfunction expansion to solve this. Now, this is not necessary in this case, since this it the type of ODE we considered last quarter: the general solution of the non-driven equation is $X_h = c_1 \cos x + c_2 \sin x$, and we can find a particular solution of the form $X_d = Ax\cos x + Bx\sin x$, from which we can form the general solution of this driven equation and then work out the c_1, c_2 needed to satisfy the boundary conditions. Or, we can use the method of variation of parameters or equivalently Green's functions to find a particular solution. Nonetheless, let us work this out using eigenfunction expansions to see how this method plays out in practice.

We can expand our driving term as:

$$\sin x + \cos x = \sum_{n=1}^{\infty} d_n(-n\cos nx + \sin nx)$$

where

$$d_n = \frac{\langle \sin x + \cos x, -n\cos nx + \sin nx\rangle}{\langle -n\cos nx + \sin nx, -n\cos nx + \sin nx\rangle} = \frac{\int_0^\pi (\sin x + \cos x)(-n\cos nx + \sin nx)\,dx}{\frac{1}{2}\pi(n^2+1)}.$$

The denominator comes from computing $\int_0^\pi (-n\cos nx + \sin nx)^2\,dx$, and if we compute the numerator (use a computer) we get:

$$\int_0^\pi (\sin x + \cos x)(-n\cos nx + \sin nx)\,dx = \frac{2n(\cos n\pi + 1)}{n^2 - 1},$$

at least when $n \neq 1$; when $n = 1$ this integral is zero, so $d_1 = 0$. Thus

$$d_n = \frac{4n(\cos n\pi + 1)}{\pi(n^2+1)(n^2-1)} = \frac{4n(\cos n\pi + 1)}{\pi(n^4-1)}, \quad n > 1$$

and hence the solution of our ODE is

$$X = \sum_{n=2}^{\infty} \frac{d_n}{\lambda_n} X_n = \sum_{n=2}^{\infty} \frac{4n(\cos n\pi + 1)}{\pi(n^4-1)(1-n^2)}(-n\cos nx + \sin nx).$$

(Note that we exclude the term with $n = 1$ from this sum, since as derived above we have $d_1 = 0$, which is good since we would not be able to divide by λ_1 in this case because $\lambda_1 = 1 - 1^2$ in fact equals 0. This, as you'll see on the homework, is what places a restriction on the types of driving terms we can consider in $L[X] = f(x)$; for this particular L, these driving terms must be ones which make the d_1 coefficient zero, as in this example.) After using some infinite series identities, this will actually end up giving the same solution as the other methods described above.

Driven PDEs. Let us now apply the method of eigenfunction expansion to the driven wave equation, which we otherwise would not know how to solve:

$$u_{tt} - c^2 u_{xx} = F(x,t)$$

with boundary conditions $u(0,t) = 0 = u(L,t)$ and initial conditions $u(x,0) = f(x), u_t(x,0) = g(x)$. (Next time we will consider an example of the driven heat equation.) In the non-driven case we've seen that the formal solution can be constructed using the eigenfunctions $X_n = \sin\frac{n\pi x}{L}$, which form a orthogonal basis for the space of (pieceise) continuous functions on the interval $[0, L]$ with respect to the inner product $\langle f, g\rangle = \int_0^L f(x)g(x)\,dx$. We will now expand the to-be-determined solution $u(x,t)$ of our driven equation in terms of the same eigenfunctions.

To be precise, if we fix t we can consider $u(x,t)$ as a single-variable function $u(x,t)$ of x alone, and this single-variable function can expanded in terms of the eigenfunction basis above:

$$u(x,t_0) = \sum_{n=1}^{\infty} U_n \sin\frac{n\pi x}{L}.$$

If we pick a different t, only the coefficients U_n change, so we can view these coefficients as varying as t does:

$$u(x,t) = \sum_{n=1}^{\infty} U_n(t) \sin\frac{n\pi x}{L}.$$

For instance, in the case of the undriven wave equation with $F(x,t) = 0$, these functions are the ones we saw before in the undriven formal solution:

$$U_n(t) = A_n \cos\tfrac{cn\pi t}{L} + B_n \sin\tfrac{cn\pi t}{L}.$$

Thus, the problem boils down to determining the coefficient functions $U_n(t)$ needed in order for this $u(x,t)$ to indeed satisfy our driven wave equation. Plugging $u(x,t)$ into our PDE gives:

$$\underbrace{\sum_{n=1}^{\infty} U_n''(t)\sin\tfrac{n\pi x}{L}}_{u_{tt}} - c^2 \underbrace{\sum_{n=1}^{\infty} U_n(t)(-\tfrac{n^2\pi^2}{L^2})\sin\tfrac{n\pi x}{L}}_{u_{xx}} = \underbrace{\sum_{n=1}^{\infty} C_n(t)\sin\tfrac{n\pi x}{L}}_{F(x,t)}$$

where we are using the eigenfunction expansion, with respect to x alone so that the resulting coefficients still depend on t, of the driving term on the right:

$$F(x,t) = \sum_{n=1}^{\infty} C_n(t)\sin\tfrac{n\pi x}{L},\ C_n(t) = \tfrac{2}{L}\int_0^L F(x,t)\sin\tfrac{n\pi x}{L}\,dx.$$

The coefficients of the eigenfunctions on both sides have to match up (due to a linear independence property), so we must have:

$$U_n''(t) + (\tfrac{cn\pi}{L})^2 U_n(t) = C_n(t).$$

Thus, the to-be-determined coefficient functions $U_n(t)$ must be ones which satisfy this driven second-order ODE. We already know methods for solving this, making it feasible to find $U_n(t)$.

Even better: the initial conditions $u(x,0) = f(x), u_t(x,0) = g(x)$ can be turned into initial conditions on the $U_n(t)$. The first becomes:

$$f(x) = u(x,0) = \sum_{n=1}^{\infty} U_n(0)\sin\tfrac{n\pi x}{L},$$

which says that the $U_n(0)$ must be the Fourier sine coefficients of $f(x)$, and the second initial condition becomes

$$g(x) = u_t(x,0) = \sum_{n=1}^{\infty} U_n'(0)\sin\tfrac{n\pi x}{L},$$

which says that the $U_n'(0)$ must be the Fourier sine coefficients of $g(x)$. Thus we end up with $U_n(t)$ satisfying the following second-order IVP

$$U_n''(t) + (\tfrac{cn\pi}{L})^2 U_n(t) = C_n(t),\ U_n(0) = \frac{2}{L}\int_0^L f(x)\sin\tfrac{n\pi x}{L}\,dx,\ U_n'(0) = \frac{2}{L}\int_0^L f(x)\cos\tfrac{n\pi x}{L}\,dx,$$

which has a unique solution. Depending on the form which $C_n(t)$ actually takes, this can be solved using either the method of undetermined coefficients or more generally the method of variation of parameters (i.e. Green's functions). For an undriven equation with $F(x,t) = 0$, the functions $C_n(t)$ are all zero, and the solutions of $U_n'' + (\frac{cn\pi}{L})^2 U_n = 0$ produce the $A_n \cos\frac{cn\pi t}{L} + B_n \sin\frac{cn\pi t}{L}$ in the previous formal solution we derived.

To summarize, the solution of the driven wave equation $u_{tt} - c^2 u_{xx} = F(x,t)$ with boundary and initial conditions

$$u(0,t) = 0 = u(L,t),\ u(x,0) = f(x),\ u_t(x,0) = g(x)$$

is given by

$$u(x,t) = \sum_{n=1}^{\infty} U_n(t)\sin\tfrac{n\pi x}{L}$$

where $U_n(t)$ is the solution of the second-order ODE

$$U_n''(t) + (\tfrac{cn\pi}{L})^2 U_n(t) = \frac{2}{L}\int_0^L F(x,t)\sin\tfrac{n\pi x}{L}\,dx$$

with initial conditions

$$U_n(0) = \frac{2}{L}\int_0^L f(x)\sin\tfrac{n\pi x}{L}\,dx \quad \text{and} \quad U_n'(0) = \frac{2}{L}\int_0^L g(x)\sin\tfrac{n\pi x}{L}\,dx.$$

Determining $U_n(t)$ explicitly still involves some work, but the point is that the problem of solving a driven PDE (in this case the wave equation) is reduced to a whole bunch of integral computations: find the value of $U_n(0)$, the value of $U_n'(0)$, the driving term of the ODE for $U_n(t)$, and finally the value of $U_n(t)$ using variation of parameters (which is also an integral computation) in the most general case. We'll work through an explicit example of this next time for a driven heat equation, where the work is a bit simpler since we will end up with only a first-order ODE for $U_n(t)$ as opposed to a second-order ODE as above.

Lecture 21: Green's Functions

Warm-Up. We find a solution of the driven ODE

$$X'' - 5X' + 6X = e^{\frac{5x}{2}}\sin 3\pi x - 4e^{\frac{5x}{2}}\sin 7\pi x$$

which satisfies the boundary conditions $X(0) = 0 = X(1)$. Now, of course this is something we could have done last quarter using either the method of undetermined coefficients or variation of parameters, but here we carry this out using the method of eigenfunction expansion. The hard work here is in determining the eigenvalues and eigenfunctions of the appropriate linear differential operator, but once we have these, finding the desired solution the driven ODE is simpler to carry out algebraically than prior methods were.

In fact, we found the eigenvalues and eigenfunctions of $L = D^2 - 5D + 6$ with boundary conditions $X(0) = 0 = X(1)$ as an example in a previous lecture, where we got:

$$\lambda_n = \frac{-4n^2\pi^2 - 1}{4} \quad \text{and} \quad X_n = e^{\frac{5x}{2}} \sin n\pi x.$$

The inner product with respect to which L was symmetric, and with respect to which these eigenfunctions are orthogonal, was:

$$\langle f, g \rangle = \int_0^1 e^{-5x} f(x) g(x)\, dx.$$

So, the point now is to expand all functions involved—both the driving term and the solution we are after—in terms of these eigenfunctions. But the particular form of the driving term we are considering makes this easy: if we want

$$e^{\frac{5x}{2}} \sin 3\pi x - 4e^{\frac{5x}{2}} \sin 7\pi x = \sum_{n=1}^{\infty} d_n e^{\frac{5x}{2}} \sin n\pi x,$$

we must have $d_3 = 1, d_7 = -4$, and all other $d_n = 0$ since the left side is already in the correct form, so we need the coefficients of $e^{\frac{5x}{2}} \sin n\pi x$ on both sides to match up. Alternatively, we can use the inner product to determine these coefficients:

$$d_n = \frac{\langle f(x), X_n \rangle}{\langle X_n, X_n \rangle} = \frac{\int_0^1 e^{-5x}[e^{\frac{5x}{2}} \sin 3\pi x - 4e^{\frac{5x}{2}} \sin 7\pi x] e^{\frac{5x}{2}} \sin n\pi x\, dx}{\int_0^1 e^{-5x}[e^{\frac{5x}{2}} \sin n\pi x][e^{\frac{5x}{2}} \sin n\pi x]\, dx},$$

where the orthogonality relations for sine makes these quick to compute: we only get a nonzero integral in the numerator for $n = 3$ and $n = 7$, and we get the same values of d_n as stated above. (Note that the $e^{\frac{5x}{2}}$ terms all go away in the integrals after multiplying by the weight e^{-5x}.) Again, in this case these integral computations were unnecessary given our specific driving term, but for more general driving terms these integral computations are the way to go.

Having expanded our driving term in terms of these eigenfunctions, the solution can now be found by dividing each coefficient by the appropriate eigenvalue:

$$L[X] = \sum d_n X_n \rightsquigarrow X = \sum \frac{d_n}{\lambda_n} X_n.$$

Note in our case that all eigenvalues are nonzero, so we do not have exclude any terms from this sum. Thus, using $d_3 = 1, d_7 = -4$ and other $d_n = 0$, we get that our solution is:

$$X = \frac{d_3}{\lambda_3} X_3 + \frac{d_7}{\lambda_7} X_7 = \frac{e^{\frac{5x}{2}} \sin 3\pi x}{\frac{-36\pi^2 - 1}{4}} - 4\frac{e^{\frac{5x}{2}} \sin 7\pi x}{\frac{-196\pi^2 - 1}{4}} = -\frac{4e^{\frac{5x}{2}} \sin 3\pi x}{36\pi^2 + 1} + \frac{e^{\frac{5x}{2}} \sin 7\pi x}{196\pi^2 + 1}.$$

Again, this works since applying L to this has the effect of scaling each eigenfunction by the corresponding eigenvalue, which will "undo" the eigenvalues in the denominator.

Driven heat example. Consider now the driven heat equation

$$u_t - K u_{xx} = F(x, t)$$

subject to boundary conditions $u(0, t) = 0 = u(L, t)$ and initial condition $u(x, 0) = f(x)$. As with the wave example last time, we seek to solve this by expanding everything in terms of eigenfunctions.

The eigenfunctions in this case are $X_n = \sin\frac{n\pi x}{L}$ with eigenvalues $\lambda_n = -(\frac{n\pi}{L})^2$, so, we write our to-be-determined solution as

$$u(x,t) = \sum_{n=1}^{\infty} U_n(t) \sin\frac{n\pi x}{L}$$

for some unknown coefficient functions $U_n(t)$. (To be clear, fix $t = t_0$ and expand the single-variable function $u(x,t_0)$ in terms of x using our eigenfunctions, and then vary t.) In order to satisfy our driven PDE, we require that

$$\underbrace{\sum_{n=1}^{\infty} U_n'(t) \sin\frac{n\pi x}{L}}_{u_t} - K \underbrace{\sum_{n=1}^{\infty} U_n(t)(-\tfrac{n^2\pi^2}{L^2}) \sin\frac{n\pi x}{L}}_{u_{xx}} = \underbrace{\sum_{n=1}^{\infty} C_n(t) \sin\frac{n\pi x}{L}}_{F(x,t)}$$

where

$$C_n(t) = \frac{2}{L}\int_0^L F(x,t) \sin\frac{n\pi x}{L}\, dx$$

are the Fourier sine coefficients (with t fixed, x varying) of $F(x,t)$. After comparing coefficients, this says that $U_n(t)$ must satisfy the following first-order ODE:

$$U_n'(t) + K(\tfrac{n\pi}{L})^2 U_n(t) = C_n(t).$$

(Note in the undriven case $F(x,t) = 0$, we have $C_n(t) = 0$ so that this equation has exponential solutions $U_n(t) = A_n e^{-K(\frac{n\pi}{L})^2 t}$, which gives the formal solution $u(x,t) = \sum_{n=1}^{\infty} A_n e^{-K(\frac{n\pi}{L})^2 t} \sin\frac{n\pi x}{L}$ we derived previously for the undriven heat equation.) If we now impose our initial condition $u(x,0) = f(x)$, this translates into

$$u(x,0) = f(x) = \sum_{n=1}^{\infty} U_n(0) \sin\frac{n\pi x}{L},$$

so that the $U_n(0)$ must be the Fourier since coefficients of f:

$$U_n(0) = \frac{2}{L}\int_0^L f(x) \sin\frac{n\pi x}{L}\, dx.$$

Thus we end up with a first-order IVP for $U_n(t)$:

$$U_n'(t) + K(\tfrac{n\pi}{L})^2 U_n(t) = C_n(t),\ U_n(0) = \frac{2}{L}\int_0^L f(x) \sin\frac{n\pi x}{L}\, dx,$$

which after solving will yield the solution $u(x,t) = \sum_{n=1}^{\infty} U_n(t) \sin\frac{n\pi x}{L}$ to our driven heat equation.

Let us work all this out in detail for the following specific example:

$$u_t - u_{xx} = xt,\ u(0,t) = 0 = u(\pi,t),\ u(x,0) = 0.$$

First, our solution will look like

$$u(x,t) = \sum_{n=1}^{\infty} U_n(t) \sin nx$$

for functions $U_n(t)$ satisfying some first-order IVP. To write down this IVP, we first need to find (use a computer) the Fourier since coefficients of $F(x,t) = xt$:

$$C_n(t) = \frac{2}{\pi}\int_0^{\pi} xt \sin nx\, dx = -\frac{2t\cos n\pi}{n}.$$

The Fourier sine coefficients of the initial condition $u(x,0) = 0$ are all zero, so the IVP we want is:

$$U_n'(t) + n^2 U_n(t) = -\frac{2t\cos n\pi}{n},\ U_n(0) = 0.$$

Solving this using an integrating factor (or undetermined coefficients) gives:

$$U_n(t) = -\frac{2\cos n\pi}{n^5}e^{-n^2 t} + \frac{2\cos n\pi}{n^5} - \frac{2\cos n\pi}{n^3}t,$$

and thus the solution to our heat equation with driving term $F(x,t) = xt$, boundary conditions $u(0,t) = 0 = u(\pi,t)$, and initial condition $u(x,0) = 0$ is

$$u(x,t) = \sum_{n=1}^{\infty}\left(-\frac{2\cos n\pi}{n^5}e^{-n^2 t} + \frac{2\cos n\pi}{n^5} - \frac{2\cos n\pi}{n^3}t\right)\sin nx.$$

Green's functions. ***TO BE FINISHED***

Fundamental solutions. ***TO BE FINISHED***

Lecture 22: Series Solutions

Heat-like equations. ***TO BE FINISHED***

Fourier transforms. ***TO BE FINISHED***

Series expansions. For our final topic, we return to second-order linear ODEs, and seek to develop a more general approach to finding solutions. In particular, we will mainly be interested in second-order equations with variable coefficients:

$$a_2(x)y'' + a_1(x)y' + a_0(x)y = 0.$$

We have previously fully solved the case of constant coefficients, and we also know that in this more general setting all we need to do is find two linearly independent solutions, since all solutions can be expressed as a linear combination of these. But what remains is a procedure for finding these independent solutions in the first place.

We now consider, not an eigenfunction expansion, but a *power series* expansion of a solution:

$$y = \sum_{n=0}^{\infty} c_n (x - x_0)^n.$$

The strategy is then the same as what we did with eigenfunction expansions in the PDE case: plug this proposed solution into our ODE, and use the series form to deduce information about what the unknown coefficients c_n must look like. Once we have these coefficients, we have our solution, at least once we clear up issues regarding convergence. So, this is not literally the same we did in the PDE case since this is not an *eigenfunction* expansion, but it is "similar-in-spirit" in the sense that we are still performing all computations in terms of a specific type of "expansion"; later we will make some more direct connections with what we did for PDEs.

Example. Consider the ODE $y'' + y = 0$. Now, we already know how to solve this, but let us nonetheless work out the solution via the power series approach. Suppose we can expand y as a power series centered at 0:

$$y = \sum_{n=0}^{\infty} c_n x^n.$$

By the general theory of power series, we can then compute derivatives term-by-term:

$$y' = \sum_{n=0}^{\infty} (c_n x^n)' = \sum_{n=0}^{\infty} n c_n x^{n-1} = \sum_{n=1}^{\infty} n c_n x^{n-1} \quad \text{and} \quad y'' = \sum_{n=2}^{\infty} n(n-1) c_n x^{n-2}.$$

Thus, in order for the proposed y to satisfy the given ODE, we need:

$$\underbrace{\sum_{n=2}^{\infty} n(n-1) c_n x^{n-2}}_{y''} + \underbrace{\sum_{n=0}^{\infty} c_n x^n}_{y} = 0.$$

We seek to determine the overall coefficient of each x^n on the left side, and for this it is helpful to reindex the first sum to start at $n = 0$ instead:

$$\sum_{n=0}^{\infty} (n+2)(n+1) c_{n+2} x^n + \sum_{n=0}^{\infty} c_n x^n = 0.$$

The basic fact is that in order for this equality to hold, the coefficient of x^n on the left has to match up with the coefficient of x^n on the right (think of 0 as $\sum 0x^n$), so we require that

$$(n+2)(n+1) c_{n+2} + c_n = 0 \text{ for all } n \geq 0.$$

This then gives a *recursive* description of the coefficients we are after, which fully characterizes what these coefficients (for $n \geq 2$) should actually be, so we would say that the solution to our ODE is given by $y = \sum_{n=0}^{\infty} c_n x^n$ where the c_n are given by the recursion above, with c_0, c_1 (the two terms not specified by the recursion) arbitrary constants.

In some nice cases, such as this example, we can go further and use the recursion to determine the coefficients, and hence series solution, explicitly. (This will not always be possible, which does not cause an issue practically since numerically the recursion gives enough information to approximate solutions to any degree of accuracy. We will see an example of this later.) To start, taking $n = 0$ in the recursion above gives:

$$2c_2 + c_0 = 0, \text{ so } c_2 = -\frac{1}{2} c_0.$$

Taking $n = 1$ gives

$$3 \cdot 2 c_3 + c_1 = 0, \text{ so } c_3 = -\frac{1}{3 \cdot 2} c_1.$$

For $n = 2$ and $n = 3$ we get:

$$c_4 = -\frac{1}{4 \cdot 3} c_2 = \frac{1}{4 \cdot 3 \cdot 2} c_0 \qquad c_5 = -\frac{1}{5 \cdot 4} c_3 = \frac{1}{5 \cdot 4 \cdot 3 \cdot 2} c_1,$$

and so on; in general, all even-indexed coefficients can be expressed solely in terms of c_0 and all odd-indexed coefficients in terms of c_1, with c_0 and c_1 being arbitrary:

$$c_{2n} = \frac{(-1)^n}{(2n)!} c_0 \qquad c_{2n+1} = \frac{(-1)^n}{(2n+1)!} c_1.$$

Our solution thus looks lke:

$$\begin{aligned} y &= \sum_{n=0}^{\infty} c_n x^n \\ &= \sum_{n=0}^{\infty} c_{2n} x^{2n} + \sum_{n=0}^{\infty} c_{2n+1} x^{2n+1} \\ &= c_0 \sum_{n=0}^{\infty} \frac{(-1)^n}{(2n)!} x^{2n} + c_1 \sum_{n=0}^{\infty} \frac{(-1)^n}{(2n+1)!} x^{2n+1}, \end{aligned}$$

where in the second line we have split up the sum into portion which consists only of even powers of x and the portion consisting of odd powers, since this split is what the recursion naturally produced. In this case we can go further still: the resulting power series are precisely the standard expansions of $\cos x$ and $\sin x$ respectively, so our solution is

$$y = c_0 \cos x + c_1 \sin x,$$

precisely as we would have expected using roots of the characteristic equation. Again, for this particular ODE, the power series approach is not strictly necessary, but rather provides an easy first example of the method in action.

Another example. Consider now the ODE $x^2y'' - 2xy' + 2y = 0$. We saw this type of thing last quarter as an example of what's called an *Euler equation*, which are equations we will study more carefully in the coming days. Back then, we looked for solutions of the form x^r, and determined the exponents r which would turn this into an actual solution. But let us forget this for now and derive the solution via power series.

Suppose we can expand the solution y as a power series $y = \sum_{n=0}^{\infty} c_n x^n$. Using the same derivatives we wrote down in the prior example, in order for this to satisfy this ODE requires that:

$$x^2 \underbrace{\sum_{n=2}^{\infty} n(n-1)c_n x^{n-2}}_{y''} - 2x \underbrace{\sum_{n=1}^{\infty} n c_n x^{n-1}}_{y'} + 2 \underbrace{\sum_{n=0}^{\infty} c_n x^n}_{y} = 0.$$

After bringing in the coefficients, this becomes:

$$\sum_{n=2}^{\infty} n(n-1)c_n x^n - 2\sum_{n=1}^{\infty} n c_n x^n + 2\sum_{n=0}^{\infty} c_n x^n = 0.$$

The coefficient of x^n on the left must be zero, and so we get the requirement

$$n(n-1)c_n - 2nc_n + 2c_n = 0 \text{ for } n \geq 0.$$

(Note that when $n = 0$, the first two terms are zero, which makes sense since neither of the sums above starting at $n = 2$ or $n = 1$ contribute anything to x^0, and when $n = 1$ the first term is zero, which again makes sense since the sum starting at $n = 2$ contributes nothing to x^1.) The resulting requirement can be written as

$$[n(n-1) - 2n + 2]c_n = [n^2 - 3n + 2]c_n = 0,$$

so we get that either $n^2 - 3n + 2 = 0$ for $c_n = 0$ for each $n \geq 0$. The first holds for $n = 1, 2$, so we conclude that for $n \neq 1, 2$, c_n must be zero, while there are no restrictions on c_1 and c_2.

Thus, our solution looks like:

$$y = \sum_{n=0}^{\infty} c_n x^n = 0 + c_1 x + c_2 x^2 + 0x^3 + 0x^4 + \cdots = c_1 x + c_2 x^2.$$

In particular, x and x^2 are indeed the only powers of x for which x^r is a solution of this Euler equation, so we have reproduced the fact that all solutions are linear combinations of these two independent solutions. So, again, this was not a case where the power series approach was strictly necessary, but rather a nice way to illustrate the general strategy.

Lecture 23: More on Series Solutions

Warm-Up. We solve the ODE

$$y'' - 2xy' + 4y = 0$$

using the power series approach. Suppose we can expand y as

$$y = \sum_{n=0}^{\infty} c_n x^n.$$

Then the ODE becomes:

$$\sum_{n=2}^{\infty} n(n-1)c_n x^{n-2} - 2x \sum_{n=1}^{\infty} n c_n x^{n-1} + 4 \sum_{n=0}^{\infty} c_n x^n = 0,$$

or equivalently

$$\sum_{n=0}^{\infty} (n+2)(n+1)c_{n+2} x^n - 2 \sum_{n=1}^{\infty} n c_n x^n + 4 \sum_{n=0}^{\infty} c_n x^n = 0.$$

Comparing constant terms (coefficients of x^0) throughout gives

$$2c_2 + 4c_0 = 0, \text{ so } c_2 = -2c_0,$$

and comparing coefficients of x^n for $n \geq 1$ gives:

$$(n+2)(n+1)c_{n+2} - 2nc_n + 4c_n = 0, \text{ so } c_{n+2} = \frac{2n-4}{(n+2)(n+1)} c_n.$$

The values of c_0, c_1 are thus arbitrary, and this recursive formula then characterizes our solution.

But we can go a bit further in this example. Working out a few more coefficients gives:

$$\begin{aligned}
c_3 &= \frac{-2}{3 \cdot 2} c_1 \\
c_4 &= \frac{0}{4 \cdot 3} c_2 = 0 \\
c_5 &= \frac{2}{5 \cdot 4} c_3 = \frac{-2 \cdot 2}{5!} c_1 \\
c_6 &= \frac{4}{6 \cdot 5} c_4 = 0 \\
c_7 &= \frac{6}{7 \cdot 6} c_5 = \frac{-2 \cdot 2 \cdot 6}{7!} c_1.
\end{aligned}$$

Hence, in general the even-indexed terms beyond c_2 are all zero, since all of these are multiples of $c_4 = 0$, and the odd-indexed terms look like:

$$c_{2n+1} = \frac{-2 \cdot 2 \cdot 6 \cdots [2(2n+1)-4]}{(2n+1)!} c_1 = \frac{-2 \cdot 2 \cdot 6 \cdots (4n-2)}{(2n+1)!} c_1.$$

The general solution to the given ODE is thus:

$$\begin{aligned} y &= \sum_{n=0}^{\infty} c_n x^n \\ &= \sum_{n=0}^{\infty} c_{2n} x^{2n} + \sum_{n=0}^{\infty} c_{2n+1} x^{2n+1} \\ &= (c_0 - 2c_0 x^2) + \sum_{n=0}^{\infty} \frac{-2 \cdot 2 \cdot 6 \cdots (4n-2)}{(2n+1)!} c_1 x^{2n+1} \\ &= c_0(1 - 2x^2) + c_1 \sum_{n=0}^{\infty} \frac{-2 \cdot 2 \cdot 6 \cdots (4n-2)}{(2n+1)!} x^{2n+1}. \end{aligned}$$

In particular, taking one of c_0, c_1 to be 1 and the other 0 gives two independent solutions

$$y_1 = 1 - 2x^2 \qquad y_2 = \sum_{n=0}^{\infty} \frac{-2 \cdot 2 \cdot 6 \cdots (4n-2)}{(2n+1)!} x^{2n+1}.$$

(So this ODE has multiples of $1-2x^2$ as its only polynomial solutions, where it was not at all obvious at the start that any polynomial solutions exist at all.) The first function, being a polynomial, is well-defined, but the second requires some care since it is defined as the value of an infinite series: this will only make sense as an honest function if this series *converges*. In a bit we will talk about some general considerations which guarantee that such concerns will not cause a problem for the types of equations we'll focus on, but let us verify this formally in this specific example. To show that the series defining y_2 converges we make use of the *ratio test* from calculus, and compute:

$$\lim_{n \to \infty} \frac{\left| \frac{-2 \cdot 2 \cdot 6 \cdots (4[n+1]-2)}{(2[n+1]+1)!} x^{2[n+1]+1} \right|}{\left| \frac{-2 \cdot 2 \cdot 6 \cdots (4n-2)}{(2n+1)!} x^{2n+1} \right|} = \lim_{n \to \infty} \frac{4n+2}{(2n+3)(2n+2)} |x|^2 = 0$$

for any x. Since this limit is smaller than 1, the ratio test guarantees that the series in question converges for all x, so that the expression given above for y_2 defines a valid function. As opposed to the series expressions for, say, $\sin x$ or $\cos x$, in this case there is no simpler standard notation we can use to express this resulting function, so the explicit series is the best we can do.

Eigeninterpretation. ***TO BE FINISHED***

Analyticity. Now that we are working with functions defined via convergent power series, it is useful to introduce a term for the types of functions which arise in this way. A function f is said to be *analytic* at x_0 if it can be expressed as a convergent power series centered at x_0 on some interval around x_0: there exists $\sum c_n (x - x_0)^n$ and an interval I containing x_0 such that

$$\sum c_n (x - x_0)^n \text{ converges to } f(x) \text{ on } J.$$

Most functions we know and love are indeed analytic, in fact on all of $\mathbb{R}$: polynomials, exponentials $e^x = \sum \frac{x^n}{n!}$, $\sin x$, and $\cos x$ to name a few. It can be shown that the coefficients needed in order for this to work out must be given by $c_n = f^{(n)}(x_0)/n!$, so that the power series required in the definition turns out to be the Taylor series of f centered at x_0:

$$\sum_{n=0}^{\infty} \frac{f^{(n)}(x_0)}{n!}(x - x_0)^n.$$

But, even though the ODEs we will focus on going forward are ones where the functions involved will be analytic, it is worthwhile to give an example an a function which is not analytic as a point of reference. Certainly, non-continuous or non-differentiable functions will not be analytic (since power series are always infinitely-differentiable wherever they converge), but the surprising fact is that not all infinitely-differentiable functions are analytic. The standard example is

$$f(x) = \begin{cases} e^{-1/x} & x > 0 \\ 0 & x \leq 0. \end{cases}$$

This is in fact infinitely-differentiable, including at 0 where $f^{(n)}(0) = 0$ for all n. If this function were to be analytic at 0, it would have to equal its own Taylor series centered at 0 on some interval around 0, but this Taylor series is

$$\sum_{n=0}^{\infty} \frac{f^{(n)}(x_0)}{n!}(x - x_0)^n = \sum_{n=0}^{\infty} \frac{0}{n!}x^n = 0,$$

which converges to 0 everywhere and hence not to $f(x) = e^{-\frac{1}{x}}$ for $x > 0$. Hence f is not analytic at 0. (In fact, there is a sense in which "most" infinitely-differentiable functions are non-analytic.)

Theorem. The key fact for our purposes is that analyticity of the coefficients in a second-order linear ODE implies analyticity of the solution. To be precise, suppose

$$y'' + p(x)y' + q(x)y = 0,\ y(x_0) = y_0,\ y'(x_0) = y_0'$$

is a second-order linear IVP. The result is that if $p(x)$ and $q(x)$ are analytic at x_0, then so is the solution. Moreover, if the power series defining $p(x)$ and $q(x)$ converge on a common interval I, so then power series defining the solution also converges on I. Thus, as long as we deal with coefficients which are expressible as power series, we are guaranteed that the solution too is expressible as a power series, so our series solution method will actually work.

Aging spring example. Consider the ODE $y'' + e^{-2t}y = 0$ with initial values $y(0) = 2, y'(0) = 1$. The functions 0 (the coefficient of y') and e^{-2t} are analytic at 0, so the solution is indeed expressible as a power series centered at 0:

$$y = \sum_{n=0}^{\infty} c_n t^n.$$

Moreover, the series expansions of 0 and e^{-2t} converge on all of $\mathbb{R}$, so this series solution should converge on all of $\mathbb{R}$ as well. We will derive a recursive formula for the coefficients.

To give some context, this equation is similar to the harmonic oscillator equation $y'' + k^2y = 0$, only that here the "spring constant" e^{-2t} is not constant but varies with respect to time. As t increases, this "spring constant" decreases, so we can view this equation as modeling a spring where the strength/tension of the spring decreases over time, as would be the case for an "aging spring". In our case, we pull the spring to a starting position of $y(0) = 2$, and then give it an initial velocity of $y'(0) = 1$ as we let go:

This will initially move the spring a bit to the right, but the tension then pulls it back towards the left, with this "pulling" action getting weaker as time progresses. We thus expect the solution to look something like the following, where the upward direction corresponds to rightward motion:

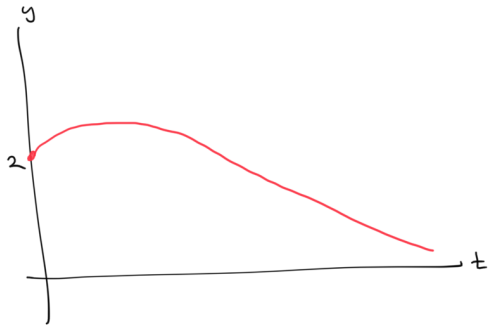

The series defining our solution must thus satisfy:

$$\sum_{n=0}^{\infty}(n+2)(n+1)c_{n+2}t^n + e^{-2t}\sum_{n=0}^{\infty}c_n t^n = 0.$$

Now, the function e^{-2t} can be expanded as:

$$e^{-2t} = \sum_{n=0}^{\infty}\frac{(-2t)^n}{n!} = \sum_{n=0}^{\infty}\frac{(-2)^n}{n!}t^n.$$

We can multiply this series with $\sum_{n=0}^{\infty}c_n t^n$ in order to obtain:

$$e^{-2t}\sum_{n=0}^{\infty}c_n t^n = \left(\sum_{n=0}^{\infty}\frac{(-2)^n}{n!}t^n\right)\left(\sum_{n=0}^{\infty}c_n t^n\right) = \sum_{n=0}^{\infty}\left(\sum_{k=0}^{n}\frac{(-2)^k}{k!}c_{n-k}\right)t^n.$$

To be clear, the expression for the result of multiplying two power series together comes from grouping together like terms in:

$$(a_0 + a_1 t + a_2 t^2 + \cdots)(b_0 + b_1 t + \cdots) = a_0 b_0 + (a_0 b_1 + a_1 b_0)t + (a_2 b_0 + a_1 b_1 + a_0 b_2)t^2 + \cdots$$

where the coefficient of t^n becomes $a_0 b_n + a_1 b_{n-1} + \cdots + a_n b_0 = \sum_{k=0}^{n} a_k b_{n-k}$. In the product we are considering, we take $a_n = (-2)^n/n!$ and $b_n = c_n$.

Thus our ODE in series form becomes:

$$\sum_{n=0}^{\infty}(n+2)(n+1)c_{n+2}t^n + \sum_{n=0}^{\infty}\left(\sum_{k=0}^{n}\frac{(-2)^k}{k!}c_{n-k}\right)t^n = 0.$$

Comparing coefficients then gives the requirement that

$$(n+2)(n+1)c_{n+2} + \sum_{k=0}^{n}\frac{(-2)^k}{k!}c_{n-k} = 0, \text{ so } c_{n+2} = -\frac{1}{(n+2)(n+1)}\sum_{k=0}^{n}\frac{(-2)^k}{k!}c_{n-k}.$$

This is the desired recursive relation, which thus determines the solution to our aging problem. This recursion will only kick once we reach the coefficient c_2, but the values of c_0 and c_1 can be determined from the initial data: if $y = c_0 + c_1 t + c_2 t^2 + \cdots$, then

$$y(0) = c_0 \text{ and } y'(0) = c_1,$$

so in our case $c_0 = 2$ and $c_1 = 1$. We will work out some more coefficients next time, and plot approximations to the solution to see if they make sense based on the physical interpretation above.

Lecture 24: Euler Equations

Warm-Up 1. We return to the aging spring example from last time:

$$y'' + e^{-2t} y = 0,\ y(0) = 2,\ y'(0) = 1$$

and work out some explicit coefficients in the power series solution. We saw last time that the solution can be expressed as a convergent power series on all of $\mathbb{R}$:

$$y = \sum_{n=0}^{\infty} c_n x^n$$

and that the coefficients in this series were $c_0 = 2, c_1 = 1$, and then recursively determined by

$$c_{n+2} = -\frac{1}{(n+2)(n+1)} \sum_{k=0}^{n} \frac{(-2)^k}{k!} c_{n-k} \text{ for } n \geq 0.$$

Thus, for $n = 0$ we get:

$$c_2 = -\frac{1}{2 \cdot 1} \sum_{k=0}^{0} \frac{(-2)^k}{k!} c_{0-k} = -\frac{1}{2}(c_0) = -1.$$

For $n = 1$ we get:

$$c_3 = -\frac{1}{3 \cdot 2}(c_1 - 2c_0) = \frac{1}{2},$$

and so on for a few more coefficients:

$$\begin{aligned}
c_4 &= -\frac{1}{4 \cdot 3}(c_2 - 2c_1 + 2c_0) = -\frac{1}{12} \\
c_5 &= -\frac{1}{5 \cdot 4}\left(c_3 - 2c_2 + 2c_1 - \frac{8}{6}c_0\right) = -\frac{1}{120} \\
c_6 &= -\frac{1}{6 \cdot 5}\left(c_4 - 2c_3 + 2c_2 - \frac{8}{6}c_1 + \frac{16}{24}c_0\right) = \frac{37}{360} \\
c_7 &= -\frac{1}{7 \cdot 6}\left(c_5 - 2c_4 + 2c_3 - \frac{8}{6}c_2 + \frac{16}{24}c_1 - \frac{32}{120}c_0\right) = -\frac{61}{1008}.
\end{aligned}$$

Hence the first few terms in the series solution look like:

$$y = 2 + t - t^2 + \frac{1}{2}t^3 - \frac{1}{12}t^4 - \frac{1}{120}t^5 + \frac{37}{360}t^6 - \frac{61}{1008}t^7 + \cdots.$$

There is no simple explicit formula for all the coefficients in this case, nor is there a simpler of expressing the solution y, but we can *approximate* the solution by taking *partial sums* of this series:

$$2 + t - t^2 \quad \text{vs} \quad 2 + t - t^2 + \frac{1}{2}t^3 \quad \text{vs} \quad 2 + t - t^2 + \frac{1}{2}t^3 - \frac{1}{12}t^4 \quad \text{etc.}$$

Plotting the partial sums of 3-rd, 5-th, and 7-th order gives:

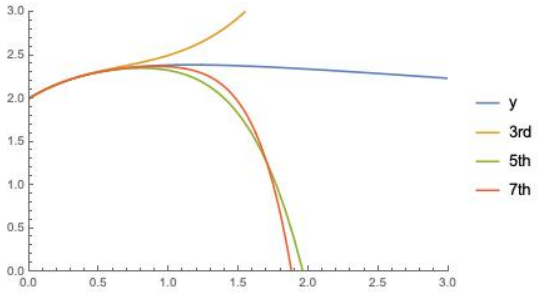

The actual solution (found numerically) is in blue, and has the behavior we expect: jumps up above $y = 2$ initially and then becomes to die off when the spring begins to pull back towards the left, but ever so slowly. Near the initial time $t = 0$, the partial sum do provide a good approximation to this solution, with the approximation getting better and better as we take higher-order partial sums. Now, far enough away from $t = 0$ the partial sum is no longer a good approximation, but since this series solution will converge everywhere we know that there is *some* "high-enough-order" partial sum which will provide a good approximation even further way; for instance, the 7-th partial sum above is not a good approximation at $t = 2.5$, but as we take more and more partial sums we will eventually find one (maybe of order 100) which *does* provide a good approximation even near $t = 2.5$. Power series solutions in general are best suited for analyzing behavior near the initial point, so if we really wanted to study the behavior near $t = 2.5$ we should probably instead use a power series centered at 2.5 instead of at 0.

Warm-Up 2. We derive a recursive expression for the coefficients in a series solution of

$$(1 - x^2)y'' - 3xy' - y = 0.$$

Writing this in standard linear form (where the coefficient of y'' is 1) results in:

$$y'' - \frac{3x}{1 - x^2}y' - \frac{1}{1 - x^2}y = 0.$$

Since the coefficient functions $-\frac{3x}{1-x^2}$ and $-\frac{1}{1-x^2}$ are each analytic at 0 (polynomials are analytic, and the quotient of analytic functions is still analytic wherever the denominator is nonzero), we know that it is possible to express the solution as a convergent power series centered at 0. (Note that these functions are NOT analytic at ± 1, so it will not be possible to express the solution as a convergent power series centered at these points instead. We refer to such points as *singular* points. In a bit we will describe one method for dealing with a certain type of such singularities.)

So, expand y as a power series centered at 0:

$$y = \sum_{n=0}^{\infty} c_n x^n.$$

The given ODE then becomes:

$$(1 - x^2)\sum_{n=2}^{\infty} n(n-1)c_n x^{n-2} - 3x\sum_{n=1}^{\infty} nc_n x^{n-1} - \sum_{n=0}^{\infty} c_n x^n = 0,$$

which after multiplying through by coefficients and reindexing can be written as:

$$\sum_{n=0}^{\infty}(n+2)(n+1)c_{n+2}x^n - \sum_{n=2}^{\infty} n(n-1)c_n x^n - \sum_{n=1}^{\infty} 3nc_n x^n - \sum_{n=0}^{\infty} c_n x^n = 0.$$

(The goal here was to get everything expressed in terms of x^n.) Thus, by equating the coefficient of x^n with zero, for $n = 0$ we get

$$2c_2 - c_0 = 0;$$

for $n = 1$ we get

$$3 \cdot 2c_3 - 3c_1 - c_1 = 0;$$

and for $n \geq 2$ we get

$$(n+2)(n+1)c_{n+2} - n(n-1)c_n - 3nc_n - c_n = 0.$$

Hence:

$$c_2 = \frac{1}{2}c_0 \qquad c_3 = \frac{2}{3}c_1 \qquad c_{n+2} = \frac{n(n-1)+3n+1}{(n+2)(n+1)}c_n = \frac{n+1}{n+2}c_n.$$

Note that this final expression does happen to give the correct value for $n = 0$ and $n = 1$ as well, so this is the recursive relation we want, with c_0 and c_1 arbitrary.

This is all that was required, but let us push further here and derive the explicit series solution. Computing a few more terms explicitly (in terms of c_0, c_1) gives:

$$c_4 = \frac{3}{4}c_2 = \frac{1 \cdot 3}{2 \cdot 4}c_0 \qquad c_5 = \frac{4}{5}c_3 = \frac{2 \cdot 4}{3 \cdot 5}c_1 \qquad c_6 = \frac{5}{6}c_4 = \frac{1 \cdot 3 \cdot 5}{2 \cdot 4 \cdot 6}c_0.$$

The pattern continues, so that in general c_{2n} will have a numerator consisting of products of odd integers and a denominator consisting of products of even integers, and this is flipped for c_{2n+1}:

$$c_{2n} = \frac{1 \cdot 3 \cdot 5 \cdots (2n-1)}{2 \cdot 4 \cdot 6 \cdots 2n}c_0 \qquad c_{2n+1} = \frac{2 \cdot 4 \cdot 6 \cdots 2n}{1 \cdot 3 \cdot 5 \cdots (2n+1)}c_1.$$

Thus the series solution explicitly looks like:

$$\begin{aligned} y &= \sum_{n=0}^{\infty} c_{2n}x^{2n} + \sum_{n=0}^{\infty} c_{2n+1}x^{2n+1} \\ &= c_0 \sum_{n=0}^{\infty} \frac{1 \cdot 3 \cdot 5 \cdots (2n-1)}{2 \cdot 4 \cdot 6 \cdots 2n}x^{2n} + c_1 \sum_{n=0}^{\infty} \frac{2 \cdot 4 \cdot 6 \cdots 2n}{1 \cdot 3 \cdot 5 \cdots (2n+1)}x^{2n+1}, \end{aligned}$$

where the two series individually define linearly independent solutions.

Regular singularities. ***TO BE FINISHED***

Euler equations. ***TO BE FINISHED***

Lecture 25: Frobenius Series

Warm-Up. ***TO BE FINISHED***

General 2nd-order linear ODEs. We now consider more general "Euler-like" ODEs, where we allow the constants p_0, q_0 used in the Euler case to vary:

$$x^2y'' + xp(x)y' + q(x)y = 0.$$

We will assume that $p(x)$ and $q(x)$ are analytic at 0, which guarantees that 0 is a regular singularity of the equation written in normal form after dividing through by x^2:

$$y'' + \frac{p(x)}{x} + \frac{q(x)}{x^2} = 0.$$

Even though this seems like a specialized type of equation, in fact *any* second-order linear ODE can be written in this way:

$$a_2(x)y'' + a_1(x)y' + a_0(x)y = 0 \rightsquigarrow x^2y'' + x\underbrace{\left(x\frac{a_1(x)}{a_2(x)}\right)}_{p(x)}y' + \underbrace{\left(x^2\frac{a_0(x)}{a_2(x)}\right)}_{q(x)}y = 0$$

with appropriate assumptions on $a_0(x), a_1(x)$, and $a_2(x)$—namely ones which say that 0 is a regular singularity of the rewritten form. The point of writing a second-order linear ODE in this particular way is to make use of what we know about Euler equations in order to derive a solution.

Here is the basic idea. Each of the functions $p(x), q(x)$ can be expanded as power series:

$$x^2y'' + x(p(0) + c_1x + c_2x^2 + \cdots)y' + (q(0) + d_1x + d_2x^2 + \cdots)y = 0$$

where we use the fact that the constant term in the expansion of $p(x) = \sum_{n=0}^{\infty} c_nx^n$ is precisely $p(0)$, and similar for $q(x)$. (We won't care about what the other coefficients explicitly look like.) The equation above can then be turned into something like:

$$\underbrace{x^2y'' + xp(0)y' + q(0)y}_{\text{zeroth-order}} + (\text{something with higher order terms}) = 0$$

where we group together the terms which involve the "zeroth-order" constants $p(0)$ and $q(0)$, and put everything else into the "something with higher order terms" expression. Setting this "zeroth-order" part equal to zero then gives an Euler equation:

$$x^2y'' + xp(0)y' + q(0)y = 0,$$

which we can view as a "zeroth-order" approximation to our original ODE. Thus, we can try to build a solution to our original ODE by taking a solution to this Euler approximation, and modifying it by adding "higher-order" terms as needed in order to get a solution of our original equation:

$$(\text{Euler solution}) + (\text{higher-order terms}).$$

If these "higher-order" terms are chosen in just the right way, we will in fact end up with an honest solution of $x^2y'' + xp(x)y + q(x)y = 0$.

The Euler equation $x^2y'' + xp(0)y' + q(0)y = 0$ has indicial equation

$$r^2 + (p(0) - 1)r + q(0) = 0,$$

and we will refer to this as being the *indicial equation* of our original Euler-like equation as well. The Euler solutions can then be determined by finding the indicial roots, as described previously.

Frobenius method. In the case where we we have real and distinct indicial roots, we get Euler solutions of the form x^r, and we then try to construct a solution to our original ODE by adding higher-order terms to this:

$$x^r + (\text{higher-order terms}).$$

To be precise, we look for a solution of the form:

$$x^r + c_1 x^{r+1} + c_2 x^{r+2} + c_3 x^{r+3} + \cdots,$$

which after factoring out a common x^r can be written as x^r times a power series:

$$x^r(1 + c_1 x + c_2 x^2 + \cdots) = x^r \sum_{n=0}^{\infty} c_n x^n.$$

(In fact, we will allow for more general constant terms c_0 apart from $c_0 = 1$ in this power series, since if x^r is a solution to our approximating Euler equation, then so is $c_0 x^r$.) The goal is then to determine the coefficients c_n needed in order for this proposed solution to be an actual solution to our ODE. Usually, we will only be able to determine these coefficients recursively and not necessarily explicitly, but that is good enough for most purposes.

The expression $x^r \sum_{n=0}^{\infty} c_n x^n$ is not necessarily itself a power series, since r need not be a nonnegative integer. For instance, for $r = \frac{1}{2}$ our proposed solution looks like

$$\sqrt{x} \sum_{n=0}^{\infty} c_n x^n = c_0\sqrt{x} + c_2 x\sqrt{x} + c_3 x^2\sqrt{x} + \cdots,$$

which is not a power series due to the fractional powers. Such a series is called a *Frobenius series*, and the method we will use to construct such solutions is called the *Frobenius method*.

In the case where the indicial equation has a repeated root, we know that to get a second Euler solution we need to use something like $x^r \ln x$, so in this case we should be looking to construct a solution to our ODE using something like

$$x^r \ln x + (\text{higher-order terms}) = x^r \ln x \sum_{n=0}^{\infty} c_n x^n,$$

and in the complex root case we should use something like

$$x^\alpha \cos(\beta \ln x) + (\text{higher-order terms}) = x^\alpha \cos(\beta \ln x) \sum_{n=0}^{\infty} c_n x^n \quad \text{or} \quad x^\alpha \sin(\beta \ln x) \sum_{n=0}^{\infty} c_n x^n.$$

These are NOT "Frobenius series", and we will not discuss these types of solutions any further. But, essentially the idea for finding solutions of this form is the same as the Frobenius method we will describe: plug these proposed solutions into the ODE, and then determine the coefficients needed recursively to extent possible.

Example. It is easiest to describe the workings of the Frobenius method via an example, so here we go. We seek to solve the following second-order linear ODE:

$$xy'' + (2 - x)y' - y = 0.$$

First, we write this in the required Euler-like form by multiplying through by x:

$$x^2 y'' + x(2 - x)y' - xy = 0.$$

(This was needed in order to end up with the correct indicial equation.) Then $p(x) = 2 - x$ and $q(x) = -x$, so the indicial equation is

$$r^2 + (p(0) - 1)r + q(0) = r^2 + r = 0,$$

and the indicial roots are $r = -1, 0$. Thus we will build solutions to our ODE by adding higher-order terms to $x^{-1} = \frac{1}{x}$ and $x^0 = 1$.

Suppose our solution takes the form of a Frobenius series:

$$y = x^r \sum_{n=0}^{\infty} c_n x^n = \sum_{n=0}^{\infty} c_n x^{n+r}$$

for some to-be-determined coefficients c_n. Plugging this into our ODE gives the following:

$$\underbrace{x^2 \sum_{n=0}^{\infty} (n+r)(n+r-1)c_n x^{n+r-2}}_{y''} + x(2-x) \underbrace{\sum_{n=0}^{\infty} (n+r)c_n x^{n+r-1}}_{y'} - x \underbrace{\sum_{n=0}^{\infty} c_n x^{n+r}}_{y} = 0.$$

As with prior series solutions, the goal now is to express everything in terms of a common power of x so that the coefficients will be more to easy compare. After bringing the x^2 into the first sum we end up with x^{n+r}, so let us express everything in terms of this. The middle term splits into two pieces after disbributing, so altogether we get:

$$\sum_{n=0}^{\infty} (n+r)(n+r-1)c_n x^{n+r} + 2\sum_{n=0}^{\infty} (n+r)c_n x^{n+r} - \sum_{n=0}^{\infty} (n+r)c_n x^{n+r+1} - \sum_{n=0}^{\infty} c_n x^{n+r+1} = 0.$$

We can rewrite the third and fourth sums to use x^{n+r} instead of x^{n+r+1} by reindexing these to begin at $n = 1$ instead of $n = 0$, which has the effect of decreasing every instance of n by 1:

$$\sum_{n=0}^{\infty} (n+r)(n+r-1)c_n x^{n+r} + 2\sum_{n=0}^{\infty} (n+r)c_n x^{n+r} - \sum_{n=1}^{\infty} (n-1+r)c_{n-1} x^{n+r} - \sum_{n=1}^{\infty} c_{n-1} x^{n+r} = 0.$$

Now we can equate coefficients on both sides, which in particular means that the coefficient of x^{n+r} on the left should be zero for all $n \geq 0$. To start with, the coefficient of x^r when $n = 0$ is

$$(0+r)(0+r-1)c_0 + 2(0+r)c_0$$

which only comes from the first two sums since the third and fourth have x^{r+1} as the smallest power of x which appears. Thus we get the requirement that

$$r(r-1)c_0 + 2rc_0 = 0, \text{ which becomes } (r^2+r)c_0 = 0.$$

In the end, the recursion we will derive for all other coefficients c_n will relate them to the value of c_0, so if $c_0 = 0$ then all remaining coefficients will be zero as well and we will not get a nontrivial solution. Thus in order to get a nontrivial solution, this first condition imposes the requirement that $r^2 + r = 0$, which is precisely the indicial equation of this ODE! This is the point: the inidicial equation $r^2 + (p(0) - 1)r + q(0) = 0$ is precisely the condition needed in order to get nontrivial solutions via the Frobenius method, which is why we assume that r in $x^r \sum c_n x^n$ is in fact an indicial root. For each choice of root (recall we are in the distinct real root case), we will get a Frobenius solution determined recursively.

With r being an indicial root, the value of c_0 is then arbitrary. Next, the coefficient of x^{n+r} for $n \geq 1$ in the series equality above is:

$$(n+r)(n+r-1)c_n + 2(n+r)c_n - (n-1+r)c_{n-1} - c_{n-1} = 0.$$

After simplifying, this turns into $(n+r)(n+r+1)c_n - (n+r)c_{n-1} = 0$, so we get the recursion

$$c_n = \frac{c_{n-1}}{n+r+1} \text{ for } n \geq 1.$$

Thus for the indicial root $r = -1$, we get that

$$y_1 = x^{-1} \sum_{n=0}^{\infty} c_n x^n \text{ for } c_0 \text{ arbitrary and } c_n = \frac{c_{n-1}}{n}, n \geq 1$$

is a solution to our ODE, and for $r = 0$ we get

$$y_2 = x^0 \sum_{n=0}^{\infty} c_n x^n \text{ for } c_0 \text{ arbitrary and } c_n = \frac{c_{n-1}}{n+1}, n \geq 1$$

as a second solution. In general this is about as far as we will be able to get, but in this case we can be more explicit. The recursions give the explicit values

$$c_n = \frac{c_0}{n!} \text{ in the } r = -1 \text{ case, and } c_n = \frac{c_0}{(n+1)!} \text{ in the } r = 0 \text{ case.}$$

Hence, explicitly (let us take $c_0 = 1$ in both cases) we have

$$y_1 = \frac{1}{x} \sum_{n=0}^{\infty} \frac{x^n}{n!} = \frac{e^x}{x} \quad \text{and} \quad y_2 = 1 \sum_{n=0}^{\infty} \frac{x^n}{(n+1)!} = \frac{e^x - 1}{x}$$

as solutions of our ODE. These are in fact linearly independent, and so the general solution to the second-order ODE $xy'' + (2-x)y' - y = 0$ (for $x \neq 0$) is

$$y = c_1 \frac{e^x}{x} + c_2 \frac{e^x - 1}{x}.$$

Note in particular that taking $c_1 = 1, c_2 = -1$ gives $y = \frac{1}{x}$ as a solution, which was not obvious at the start.