MATH 381: Fourier Analysis and Boundary Value Problems

Northwestern University, Lecture Notes

Written by Santiago Cañez

These are notes which provide a basic summary of each lecture for MATH 381, "Fourier Analysis and Boundary Value Problems for ISP", taught by the author at Northwestern University. The book used as a reference is the 11th edition of *Elementary Differential Equations and Boundary Value Problems* by Boyce and DiPrima. Watch out for typos! Comments and suggestions are welcome.

Contents

Lecture 1: Introduction

This is a course about partial differential equations. A *partial differential equation* (PDE for short) is a differential equation that involves an unknown multivariable function and its partial derivatives. The basic goal is understand the solutions of such an equation (or at least their behaviors), or in other words the functions which satisfy it. Partial differential equations model a wide variety of phenoma, and their study constitutes a vast area of mathematics.

Example. Here is an example of a basic PDE, known as the *heat equation*:

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

Here, $u(x,t)$ is a function of two variables, x which represents position and t as time. In the standard interpretation, the value of $u(x,t)$ gives the temperature at position x and time t along a metal rod that has been previously heated. The heat equation characterizes the time evolution of the temperature function $u(x,t)$, and says that at every instant of time and at every position the rate of change of $u(x,t)$ with respect to time must equal the second derivative of $u(x,t)$ with respect to position. (We will see why this is the correct equation needed to model heat diffusion later. Also, there are some constants usually included that I am omitting for now.)

Now, one solution to the heat equation is given by the function

$$u_1(x,t) = e^{-t}\sin x.$$

Indeed, the partial derivative u_t is $-e^{-t}\sin x$, and the second-order partial derivative u_{xx} is $-e^{-t}\sin x$, so this function does satisfy $u_t = u_{xx}$. More generally, one can check that for any n the function

$$u_n(x,t) = e^{-n^2 t}\sin(nx)$$

is also a solution of the heat equation. So, we have many solutions so far, such as

$$u_1(x,t) = e^{-t}\sin x,\ u_2(x,t) = e^{-4t}\sin(2x),\ u_3(x,t) = e^{-9t}\sin(3x)$$

and so on for other positive integers n.

To get more solutions, we can exploit some properties of the heat equation, namely that it is *linear* and *homogeneous*. We will clarify what this means precisely later, but you might recall that in the ODE case it is true that sums of scalar multiplies of solutions of a linear homogeneous ODE are also solutions of the same ODE. This is also true in the PDE case, so something like

$$5e^{-t}\sin x + 7e^{-9t}\sin(3x)$$

is also a solution of the heat equation. More generally, things like

$$c_1 e^{-t}\sin x + \cdots + c_n e^{-n^2 t}\sin(nx)$$

and other "linear combinations" of solutions are also solutions. But why stop with sums of scalar multiples of *finitely* many solutions? Does it make sense to take an *infinite* sum of solutions and still get a solution? The answer is that—once we handle convergence issues appropriately–yes it is true that something like the infinite series

$$\sum_{n=1}^{\infty} c_n e^{-n^2 t}\sin(nx)$$

will also be a solution of the heat equation. Again, there are some subtleties here dealing with convergence (in particular, does this infinite sum even exist as a finite value?), but we will clarify these as we go.

Towards Fourier series. Suppose now we go back and impose an initial condition on the heat equation, by which we mean a function $f(x)$ telling us what the initial temperature (at time 0) should be at any position along the rod:

$$u(x,0) = f(x) \text{ for all } x.$$

We can ask whether it's possible to find a solution of the heat equation of the "infinite sum" form

$$u(x,t) = \sum_{n=1}^{\infty} c_n e^{-n^2 t} \sin(nx)$$

which in particular satisfies this given initial condition. In order for this candidate to satisfy $u(x,0) = f(x)$, we need the following to be true:

$$f(x) = u(x,0) = \sum_{n=1}^{\infty} c_n e^{-n^2 0} \sin(nx) = \sum_{n=1}^{\infty} c_n \sin(nx).$$

The sum at the end is an example of a *Fourier series*, so the conclusion is that in order for

$$u(x,t) = \sum_{n=1}^{\infty} c_n e^{-n^2 t} \sin(nx)$$

to be a solution of the initial value problem

$$\frac{\partial u}{\partial t} = \frac{\partial^2 u}{\partial x^2}, \; u(x,0) = f(x),$$

the coefficients c_n *must* come from a series expansion

$$f(x) = \sum_{n=1}^{\infty} c_n \sin(nx)$$

of the function $f(x)$. If we can figure out what these coefficients c_n should be, we can construct a solution of our initial value problem.

Why Fourier? So, Fourier series will help us construct solutions of PDEs. Now, so far we are only saying that such Fourier series indeed give solutions, but we are not saying anything about whether they give *all* solutions. Perhaps there are other solutions of the same initial value problem above apart from the ones derived from Fourier series. More conceptually, apart from saying that "it works", why should we even anticipate that Fourier series should play *any* role in solving PDEs ahead of time? This is a deep issue, that we will shed a bit of light on later in the quarter.

For now, we can gather *some* sense as to why considering Fourier series when solving PDEs might be reasonable by considering an analogous method for solving ODEs using power series. For example, consider the ODE $y' = y$. The single-variable functions which satisfy this are precisely the scalar multiples of e^x. One way to see this is by *assuming* that the unknown solution we want can be expressed as a power series such as

$$y = \sum_{n=0}^{\infty} c_n x^n.$$

Then we can determine what properties this power series would need to have in order for the function y it defines to satisfy $y' = y$. From this you can derive that in fact the coefficients c_n must be of the form

$$c_n = \frac{C}{n!}$$

for some constant C, so that the solution we want looks like

$$y = \sum_{n=0}^{\infty} \frac{C}{n!} x^n,$$

which is precisely the power series expansion of $y = Ce^x$. So, the idea here is to assume you can write your solution as a power series, and then work out what that power series actually has to look like.

But power series are not the only types of series we can use to express functions in a "nice" way. As we will see, many functions, in particular those which are continuously differentiable, can in fact be expressed in the form

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

for appropriate choices of a_0, a_n, b_n, and L. This right hand side is the form of a general Fourier series, and *if* we assume that our solution of a given differential equation can be expressed in this way, we can then try to work out information about the unknown coefficients in the hopes of determining the solution explicitly. Again, there are deeper reasons for why Fourier series are the right types of series to consider, but we'll come back to this later.

Boundary values. The discussion above explains why the phrase "Fourier analysis" shows up in the title of this course, so now we say something about the remaining phrase: "boundary value problems". Consider the heat equation, now with specified *boundary values*:

$$\frac{\partial u}{\partial t} = \frac{\partial^u}{\partial x^2}, \; u(0,t) = 0 = u(L,t).$$

In the standard interpretation, we are imagining a metal rod of length L, and here we are stating that the temperatures at the ends of the rod—i.e., at the boundaries—should be zero at all times. The question is whether we can find a solution of the heat equation satisfying these boundary conditions. Or, we can consider boundary conditions like

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

which characterizes a rod with *insualted* ends, and other examples. We will see here as well that Fourier series end up being a natural thing to consider. But, unlike, say, an ODE with initial conditions, the question of exitence and uniqueness of solutions is more subtle when imposing boundary conditions. (We will also consider the two-dimensional heat equation

$$\frac{\partial u}{\partial t} = \frac{\partial^u}{\partial x^2} + \frac{\partial^u}{\partial y^2},$$

characterizing temperature functions $u(x, y, t)$ on two-dimensional regions, and corresponding boundary and initial conditions.)

To give a sense of what can happen when imposing boundary conditions, consider the example of the following second-order ODE with boundary conditions:

$$y'' + 4y = 0,\ y(0) = 0,\ y(\pi) = 0.$$

The functions satisfying $y'' + 4y = 0$ are all of the form

$$y = c_1 \cos(2t) + c_2 \sin(2t)$$

where c_1 and c_2 are scalars, as we will briefly review at some point. In order to satisfying the given boundary conditions, we need

$$0 = y(0) = c_1 \cos 0 + c_2 \sin 0 = c_1,$$

so that $y = c_2 \sin(2t)$, and also

$$0 = y(\pi) = c_2 \sin(2\pi),$$

which is satisfied by all c_2. Hence this particular boundary value problem in fact has infintely many solutions, namely all scalar multplies of $\sin(2t)$. But, we can find other boundary conditions for which there will only be one solution, and others for which there will be none. The point is that more analysis is needed when determining the behavior of solutions with boundary conditions, as there is no "general" answer that always applies.

Lecture 2: Fourier Series

Warm-Up. We give examples of boundary conditions on the ODE

$$y'' + 9y = 0$$

which result in zero, exactly one, and infinitely many solutions. First, we note that all functions satisfying $y'' + 9y = 0$ are of the form

$$y = c_1 \cos(3t) + c_2 \sin(3t).$$

Consider the boundary conditinos $y(0) = 0, y(\pi) = 0$. To satisfy the first condition requires that

$$0 = y(0) = c_1 \cos(0) + c_2 \sin(0) = c_1,$$

so that our candidate solution of the form $y = c_2 \sin(3t)$. Any such function automatically satisfies $y(\pi) = 0$ since $\sin(3\pi) = 0$, so these boundary conditions result in infinitely many solutions.

Now, consider $y(0) = 0, y(\pi) = 1$. The first conditions again gives $y = c_2 \sin(3t)$. But now no such function satisfies $y(\pi) = 1$ regardless of c_2, so these boundary condiitons result in no solutions. Finally, for $y(0) = 0, y(\frac{\pi}{2}) = 1$, we first have $y = c_2 \sin(3t)$ from the first condition, and the second requires

$$1 = y(\tfrac{\pi}{2}) = c_2 \sin(\tfrac{3\pi}{2}) = -c_2.$$

Thus $y = -\sin(3t)$ is the only solution of $y'' + 9y = 0$ satisfying the boundary conditions $y(0) = 0$ and $y(\frac{3\pi}{2}) = 1$.

Fourier series. As we saw briefly last time, Fourier series will play a role in solving various PDEs. To be clear, a *Fourier series* is an infinite series of the form

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

(We will see shortly why the constant terms is written as $\frac{a_0}{2}$ instead of simply a_0.) For now we bypass the question of which functions *can* be expressed as such a series, and focus instead on determining, in the case where a function *is* expressible in a such a way, any information we can derive from having such a expression available. That is, if

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

what can we say about f or the required series? As a first observation, we note that the series itself is periodic with period $2L$, since the cosine and sine terms are $2L$-periodic:

$$\cos \frac{n\pi(x+2L)}{L} = \cos\left(\frac{n\pi x}{L} + 2n\pi\right) = \cos \frac{n\pi x}{L}$$

and similarly for $\sin \frac{n\pi x}{L}$. Thus, the function f had better be $2L$-periodic as well:

$$f(x+2L) = f(x).$$

By appropriately choosing the value of L we can in effect set the period to be whatever we need.

So we suppose that f is indeed $2L$ periodic. The next step is to see what we can say about the unknown coefficients a_0, a_n, and b_n in the equation

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

We'll refer to these coefficients as the *Fourier coefficients* of f, and we claim that the value of these can be determined explicitly, so that there is only one possible set of coefficients for which this equality can hold. Compare this to what happens for power series: if f is to be expressible as a power series

$$f(x) = \sum_{n=0}^{\infty} c_n x^n,$$

it turns out that the value of c_n must be given by $\frac{f^{(n)}(0)}{n!}$, where $f^{(n)}$ denotes the n-th order derivative. The same type of conclusion is true of the Fourier coefficients above, only that the explicit formulas will be a little more involved.

Orthogonality relations. Before deriving formulas for the Fourier coefficients, we highlight the following integral identities, called the *orthogonality relations* of sine and cosine. (We will see next time why we use the term "orthogonality" here.) For nonnegative integers n and m, we have:

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

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

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

For the first result, when $m = n = 0$ note simply that we have

$$\int_{-L}^{L} \cos\tfrac{0\pi x}{L} \cos\tfrac{0\pi x}{L}\, dx = \int_{-L}^{L} dx = 2L.$$

When $m = n \neq 0$, we can use the double-angle formula for cosine $\cos^2 A = \frac{1}{2}(1 + \cos 2A)$ to get

$$\int_{-L}^{L} \cos^2 \tfrac{n\pi x}{L}\, dx = \frac{1}{2}\int_{-L}^{L} (1 - \cos\tfrac{2n\pi x}{L})\, dx = \tfrac{1}{2}\left(x - \tfrac{L}{2n\pi}\sin\tfrac{2n\pi x}{L}\right)\Big|_{-L}^{L} = L.$$

When $m \neq n$, we can use the following trig identity:

$$\cos A \cos B = \frac{1}{2}[\cos(A + B) + \cos(A - B)].$$

This gives

$$\int_{-L}^{L} \cos\tfrac{n\pi x}{L} \cos\tfrac{m\pi x}{L}\, dx = \frac{1}{2}\int_{-L}^{L} [\cos\tfrac{(n+m)\pi x}{L} + \cos\tfrac{(n-m)\pi x}{L}]\, dx,$$

which can be computed directly to get the value 0. (Both pieces give terms involving sine which evaluate to zero at both L and $-L$, since sine of an integer multiple of π is zero.)

Similar trig identities for the product of two sine functions, or a cosine function and a sine function can be used to derive the other orthogonality relations.

Deriving Fourier coefficients. We now use the orthogonality relations to derive formulas for the coefficients in

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

The computation we go through might seem to come out of nowhere at first, but next time we will put it in its proper context.

First, for a nonnegative integer m, we multiply both sides of the equality above through by $\cos\frac{m\pi x}{L}$:

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

Now we integrate both sides:

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

There is a subtle point here in that, although it is always true that the integral of a sum of finitely many functions equals the sum of the integrals of each of those individual functions, this is not true for integrals of infinite sums in general. In other words, some care is needed in order to guarantee that we can exchange the integration and summation to get that

$$\int_{-L}^{L} \left[\sum_{n=1}^{\infty} (a_n \cos\tfrac{n\pi x}{L}\cos\tfrac{m\pi x}{L}\, dx + b_n \sin\tfrac{n\pi x}{L}\cos\tfrac{m\pi x}{L})\right] dx$$

equals

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

This is related to the problem of understanding the convergence of Fourier series, but we postpone this discussion for now.

The orthgonality relations tell us what all the integrals in

$$\frac{a_0}{2}\int_{-L}^{L}\cos\tfrac{0\pi x}{L}\cos\tfrac{m\pi x}{L}\,dx+\sum_{n=1}^{\infty}\left(a_n\int_{-L}^{L}\cos\tfrac{n\pi x}{L}\cos\tfrac{m\pi x}{L}\,dx+b_n\int_{-L}^{L}\sin\tfrac{n\pi x}{L}\cos\tfrac{m\pi x}{L}\right)\,dx$$

evaluate to. (Note that we introduced $\cos\frac{0\pi x}{L}$ into the first integral; this of course is simply the constant 1, but we write it in this way so as to make the application of the orthogonality relations to this part clear.) If $m=0$, the only nonzero integral in this entire expression is the first one, in which case the value of the integral is $2L$. Thus we get in this case that

$$\int_{-L}^{L}f(x)\cos\tfrac{0\pi x}{L}\,dx=\frac{a_0}{2}(2L)=a_0L,$$

so

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

When $m>0$, the only nonzero integral on the right side of the equation

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

is the specific one in the infinite series occuring for $n=m$ and involving the two cosine terms. That is, the entire right side above simplifies to $a_m\int_{-L}^{L}\cos\frac{m\pi x}{L}\cos\frac{m\pi x}{L}\,dx$, so we get

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

and thus

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

Here we assumed $m>0$, but note that this same expression also gives the correct value we derived aboe for a_0 when $m=0$. *This* is why we wrote the constant term in our Fourier series as $\frac{a_0}{2}$ instead of as a_0: in the latter case, the formula for the value of a_0 would be $\frac{1}{2L}\int_{-L}^{L}f(x)\,dx$, where the extra 2 comes from the differenence in the orthogonality relations when $m=n\neq 0$ versus $m=n=0$; by incorporating this $\frac{1}{2}$ into the Fourier series form instead, we can use the same formula for a_0 as for the other a_m.

So, we have now derived the explicit formulas for the Fourier coefficients a_n ($n\geq 0$) when we have equality in

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

A similar argument where we instead multiply through by $\sin\frac{m\pi x}{L}$ $(m > 0)$, integrate, and then use the same orthgonality relations produces

$$\int_{-L}^{L} f(x)\sin\tfrac{m\pi x}{L}\,dx = b_m\int_{-L}^{L}\sin\tfrac{m\pi x}{L}\sin\tfrac{m\pi x}{L}\,dx = b_m L,$$

so that $b_m = \frac{1}{L}\int_{-L}^{L} f(x)\sin\frac{m\pi x}{L}\,dx$. The conclusion is thus that, *if* a periodic function f is to be expressible as a Fourier series

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

the values of the Fourier coefficients a_n, b_n *must* be given by

$$a_n = \frac{1}{L}\int_{-L}^{L} f(x)\cos\tfrac{m\pi x}{L}\,dx \text{ for } n \geq 0 \quad \text{and} \quad b_n = \frac{1}{L}\int_{-L}^{L} f(x)\sin\tfrac{n\pi x}{L}\,dx \text{ for } n \geq 1.$$

Thus, determining an explicit Fourier series simply comes down to computing some integrals. (That's not to say that computing these integrals is necessarily easy, as we will see.)

Lecture 3: More on Fourier Series

Warm-Up 1. We compute the Fourier series of the "square wave" function defined by setting

$$f(x) = \begin{cases} 0 & -1 \leq x < 0 \\ 1 & 0 \leq x < 1 \end{cases}$$

and then extending to all of $\mathbb{R}$ to have period 2. Thus, the portion of the graph for $1 \leq x < 2$ looks just like the portion for $-1 \leq x < 0$, the portion for $2 \leq x < 3$ looks like the portion for $0 \leq x < 1$, and so on:

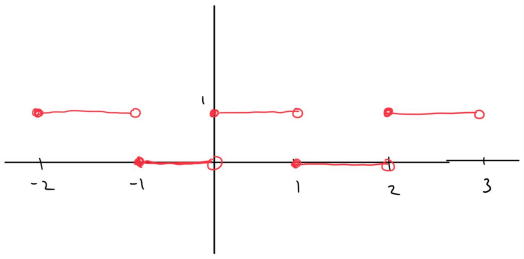

(Hopefully we can see where the term "square wave" comes from!) In this case thus the value of L in the general form of a Fourier series is 1, which is half the period.

First, we compute a_n:

$$a_n = \frac{1}{1}\int_{-1}^{1} f(x)\cos\tfrac{n\pi x}{1}\,dx = \int_{0}^{1}\cos n\pi x\,dx = \begin{cases} 1 & n = 0 \\ 0 & n \neq 0. \end{cases}$$

where we use the fact that $f(x) = 0$ on the interval between -1 and 0, so that the integral over this portion is zero, and that $f(x) = 1$ on the interval between 0 and 1 so that the integral of $f(x)\cos n\pi x$ over this portion is just the integral of $\cos n\pi x$. For b_n, $n \geq 1$ we have:

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

$$= \int_0^1 \sin n\pi x \, dx$$
$$= -\frac{1}{n\pi} \cos n\pi x \Big|_0^1$$
$$= \frac{1}{n\pi}(1 - \cos n\pi).$$

Thus the Fourier series of this function f is

$$\frac{1}{2} + \sum_{n=1}^{\infty} \frac{1}{n\pi}(1 - \cos n\pi) \sin n\pi x.$$

(There are no $\cos n\pi x$ terms for $n \geq 1$ since they all have coefficient $a_n = 0$)

We can rewrite this series in another way using the fact that $\cos n\pi = (-1)^n$, so that

$$1 - \cos n\pi = \begin{cases} 0 & \text{if } n \text{ is even} \\ 2 & \text{if } n \text{ is odd.} \end{cases}$$

So it is only the odd-indexed terms that show up in our series, and if we express these odd indices as $2n + 1$ for $n \geq 0$, we get that the Fourier series of f is

$$\frac{1}{2} + \sum_{n=0}^{\infty} \frac{2}{(2n+1)\pi} \sin[(2n+1)\pi x].$$

As we've said before, the question as to whether this series does in fact *equal* the square wave function f is one we'll come back to.

Warm-Up 2. We compute the Fourier series of the "triangular wave" function defined by setting

$$f(x) = \begin{cases} -x & -L \leq x < 0 \\ x & 0 \leq x < L \end{cases}$$

and extending to have period $2L$. (Draw the graph of this to see why we use the phrase "triangular wave".) First, let us actually determine the value of b_n:

$$b_n = \frac{1}{L} \int_{-L}^{L} f(x) \sin \tfrac{n\pi x}{L} \, dx = \frac{1}{L} \left(\int_{-L}^{0} -x \sin \tfrac{n\pi x}{L} \, dx + \int_0^L x \sin \tfrac{n\pi x}{L} \, dx \right)$$

Now, we can save time by noting that the first integral on the right is precisely the negative of the second one, which we can see by making a change of variables $u = -x$:

$$\int_{-L}^{0} -x \sin \tfrac{n\pi x}{L} \, dx = \int_L^0 u \sin \tfrac{-n\pi u}{L}(-du) = -\int_L^0 -u \sin \tfrac{n\pi u}{L} du = -\int_0^L u \sin \tfrac{n\pi u}{L} du.$$

Thus, the sum of the integral from $-L$ to 0 with the integrl from 0 to L is zero. The *better* explanation for why this happens is to note that since $f(x)$ is an *even* function (meaning $f(-x) = f(x)$ for all x) and $\sin \frac{n\pi x}{L}$ an odd function, the product $f(x) \sin \frac{n\pi x}{L}$ is odd and integrating an odd function over a symmetric interval centered at zero results in the value zero. So $b_n = 0$ for all $n \geq 1$, and there will be no $\sin \frac{n\pi x}{L}$ terms in the Fourier series.

Now, for a_n we have

$$a_n = \frac{1}{L}\int_{-L}^{L} f(x)\cos\tfrac{n\pi x}{L}\,dx = \frac{2}{L}\int_0^L x\cos\tfrac{n\pi x}{L}\,dx,$$

where now we use the fact that $\cos\frac{n\pi x}{L}$ is an even function to say that the product $f(x)\cos\frac{n\pi x}{L}$ is even, and integrating an even function over a symmetric interval gives twice the value of its integral over half of the interval. (Alterantively we can simply break up the integral into one piece over $[-L,0]$ where $f(x)=-x$, and a piece over $[0,L]$ where $f(x)=x$, and then compute both integrals. Using the fact that we are integrating an even function will save some of work.) For $n=0$ we get

$$a_0 = \frac{2}{L}\int_0^L x\,dx = L,$$

and for $n\geq 1$ we get (using integration by parts):

$$\begin{aligned} a_n &= \frac{2}{L}\int_0^L x\cos\tfrac{n\pi x}{L}\,dx \\ &= \frac{2}{L}\left(\frac{L}{n\pi}x\sin\frac{n\pi x}{L}\bigg|_0^L + \frac{L^2}{n^2\pi^2}\cos\frac{n\pi x}{L}\bigg|_0^L\right) \\ &= \frac{2L}{n^2\pi^2}(\cos n\pi - 1). \end{aligned}$$

Thus the Fourier series of the triangular wave function f is

$$\frac{L}{2} + \sum_{n=1}^{\infty}\frac{2L}{n^2\pi^2}(\cos n\pi - 1)\cos\frac{n\pi x}{L}.$$

If we again distinguish between even and odd n, so that $\cos n\pi - 1$ is either -2 or 0, we can write this Fourier series as

$$\frac{L}{2} - \sum_{n=0}^{\infty}\frac{4L}{(2n+1)^2\pi^2}\cos\frac{(2n+1)\pi x}{L}.$$

Here comes the linear algebra. Now we revisit the derivation of the Fourier coefficients we went through last time, and view it from the "proper" perspective. The correct context behind all of this is the linear algebra of *inner products*. In this particular case, consider a "space" V of $2L$-periodic functions. (We will place restrictions on the types of functions we need to consider later.) The sum of periodic functions is still periodic, as is any scalar multiple of a periodic function. This says that V is actually a *vector space*, which, if you recall, is just a set equipped with addition and scalar multiplication operations, satisfying some standard axioms. (We won't have a real need to recall the formal definition of "vector space", but you can look it up elsewhere.)

On this space V we can define the following *inner product*: for $f,g\in V$, we define (f,g) to be

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

Think of this as analogous to the usual dot product of vectors in $\mathbb{R}^n$, where we think of a function f as a "vector" with an infinite number of components indexed by x in $[-L,L]$. The expression $f(x)g(x)$ is the product of the "x-components" of f and g, and taking the integral of $f(x)g(x)$ then "adds" these products together, just as in what happens for $\mathbf{x}\cdot\mathbf{y} = x_1y_1+\cdots+x_ny_n$ when

$\mathbf{x} = (x_1, \ldots, x_n)$ and $\mathbf{y} = (y_1, \ldots, y_n)$. This inner product has the usual properties one would expect of a dot product, like:

$$(f, f) \geq 0, \ (f, g) = (g, f), \text{ and } (f, c_1 g_1 + c_2 g_2) = c_1(f, g_1) + c_2(f, g_2)$$

where in the final one c_1 and c_2 are scalars. After using the definition of this particular inner product, we see that each of these amounts to a basic property of integrals. The point is that anything you did you with dot/inner products before can be carried over to the setting of this particular vector space and inner product. (A vector space equipped with an inner product is close to the definition of what a *Hilbert space* is, although the formal definition of a Hilbert space has one more technical requirement we will mention later.)

Back to orthogonality. With this inner product at hand, we now define what it means for two functions f and g to be orthogonal. Based on what happens for the usual dot product in $\mathbb{R}^2$ or $\mathbb{R}^3$, we say that f and g are *orthogonal* if their inner product is zero:

$$(f, g) = \int_{-L}^{L} f(x)g(x)\, dx = 0.$$

Now we see that the "orthogonality relations" we had last time are precisely the statements that the various sine and cosine functions we are considering in a Fourier series are all orthogonal to one another:

$$(\cos \tfrac{n\pi x}{L}, \cos \tfrac{m\pi x}{L}) = 0 = (\sin \tfrac{n\pi x}{L}, \sin \tfrac{m\pi x}{L})$$

for $m \neq n$ and

$$(\cos \tfrac{n\pi x}{L}, \sin \tfrac{m\pi x}{L}) = 0$$

for all m, n. (This is why we call these the "orthogonality" relations.) Moreover:

$$(\cos \tfrac{0\pi x}{L}, \cos \tfrac{0\pi x}{L}) = 2L \quad \text{and} \quad (\cos \tfrac{n\pi x}{L}, \cos \tfrac{n\pi x}{L}) = L = (\sin \tfrac{n\pi x}{L}, \sin \tfrac{n\pi x}{L})$$

for $n \geq 1$. The point is that these cosine and sine functions make up an "orthogonal basis" of our space V, and the problem of expressing a function as a Fourier series is just the problem of writing a vector as a linear combination of orthogonal basis vectors. (To be clear, so far we only that these cosine and sine functions form an orthogonal set of elements of V; to say that they form an orthogonal *basis* also requires knowing that they "span" V, which means that every element can indeed be written as a linear combination of these. This again boils down to a question about convergence.)

Rederiving the coefficients. For the sake of convenient notation, we will use ϕ_n and ψ_n to denote these cosine and sine functions:

$$\phi_n(x) = \cos \frac{n\pi x}{L} \quad \text{and} \quad \psi_n(x) = \sin \frac{n\pi x}{L}.$$

Then our desired Fourier series is

$$f(x) = \frac{a_0}{2}\phi_0(x) + \sum_{n=1}^{\infty} [a_n \phi_n(x) + b_n \psi_n(x)].$$

(Again, $\psi_0(x) = \cos 0 = 1$, but we write it here explicitly next to $\frac{a_0}{2}$ simply to make this term look like the rest, apart from the extra $\frac{1}{2}$ which comes from the fact that $(\phi_0, \phi_0) = 2L$ instead of $(\phi_n, \phi_n) = L$ for $n \geq 1$.)

Let us rederive the Fourier coefficients using this new notation, and since we worked it out explicitly for a_n last time, this time we will show the work for b_n. Fix $m \geq 1$ and take the inner product of both sides above with ψ_m:

$$(f, \psi_m) = \frac{a_0}{2}(\phi_0, \psi_m) + \left(\sum_{n=1}^{\infty} [a_n \phi_n + b_n \psi_n], \psi_m \right).$$

By linearity of the inner product, we can distribute the inner product through the series on the right to get:

$$(f, \psi_m) = \frac{a_0}{2}(\phi_0, \psi_m) + \sum_{n=1}^{\infty} [a_n(\phi_n, \psi_m) + b_n(\psi_n, \psi_m)].$$

(Again, some care is needed here: this is the same "exchange integration and infinite summation" issue we mentioned last time. In this case, although we know that something like

$$(c_1 f_1 + \cdots + c_n f_n, g)$$

distributes as $c_1(f_1, g) + \cdots + c_n(f_n, g)$, more work is needed—convergence issues—to know that this is still true if we have an *infinite sum* in the first component of the inner product, as we do above.) By the orthogonality of the functions ϕ_n and ψ_n, we have

$$(\phi_0, \psi_m) = 0, \ (\phi_n, \psi_m) = 0 \text{ for all } m, n, \text{ and } (\phi_n, \psi_m) = 0 \text{ for } n \neq m.$$

Thus the only nonzero term on the right side of

$$(f, \psi_m) = \frac{a_0}{2}(\phi_0, \psi_m) + \sum_{n=1}^{\infty} [a_n(\phi_n, \psi_m) + b_n(\psi_n, \psi_m)]$$

is the one where $n = m$, so we get

$$(f, \psi_m) = b_m(\psi_m, \psi_m).$$

Hence

$$b_m = \frac{(f, \psi_m)}{(\psi_m, \psi_m)}.$$

If we put in what the definition of this inner product actually is, this expression is just

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

as we had before. Similarly, by taking inner products of both sides with ϕ_n (which is just what we did last time only without the language of inner products), we get

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

for $n \neq 0$ and

$$\frac{a_0}{2} = \frac{(f, \phi_0)}{(\phi_0, \phi_0)}.$$

You might recognize the formula for these coefficients as precisely the same one you get when considering orthogonal basis vectors of, say, $\mathbb{R}^3$: if $\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3$ form an orthogonal basis of $\mathbb{R}^3$, then when writing $\mathbf{x}$ as a linear combination of $\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3$, the coefficient of $\mathbf{u}_i$ needed to do so

$$\frac{\mathbf{x} \cdot \mathbf{u}_i}{\mathbf{u}_i \cdot \mathbf{u}_i}$$

where $\cdot$ is the usual dot product. Of course, this is no accident, since is the same computation in that case as the one here, all dependent on having orthogonality.

Orthogonal projections. Hence our desired Fourier series written solely using inner products is

$$f = \frac{(f, \phi_0)}{(\phi_0, \phi_0)}\phi_0 + \sum_{n=1}^{\infty}\left[\frac{(f, \phi_n)}{(\phi_n, \phi_n)}\phi_n + \frac{(f, \psi_n)}{(\psi_n, \psi_n)}\psi_n\right].$$

An individual term of this series, such

$$\frac{(f, \psi_n)}{(\psi_n, \psi_n)}\psi_n$$

for example, is again one you might recognize from linear algebra, say in the form

$$\left(\frac{\mathbf{x} \cdot \mathbf{u}}{\mathbf{u} \cdot \mathbf{u}}\right)\mathbf{u}$$

when dealing with vectors in $\mathbb{R}^n$: this is precisely the *orthogonal projection* of $\mathbf{x}$ onto $\mathbf{u}$! By analogy we then refer to

$$\frac{(f, \psi_n)}{(\psi_n, \psi_n)}\psi_n$$

as being the orthogonal projection of f onto ψ_n, and similarly for the terms in the Fourier series involving ϕ_n.

Thus we can reformulate the problem of finding a Fourier series as one about reconstructing a function f as the (infinite) sum its orthogonal projections onto the orthogonal basis functions ϕ_n and ψ_n. The moral is that the study of Fourier series is essentially the study of a certain orthogonal basis for a certain vector space and the orthogonal projections onto said basis elements. The same idea will make an appearance later in the context of *Sturm-Louiville* theory.

Lecture 4: Complex Fourier Series

Warm-Up. Define the *norm* of a function f defined on the interval $[-\pi, \pi]$ to be $\|f\| = \sqrt{(f, f)}$, where (f, f) is the inner product we defined last time. So, explicitly:

$$\|f\| = \sqrt{\int_{-\pi}^{\pi} f(x)^2\, dx}.$$

(This is called the L^2-*norm* of f, and is an infinite-dimensional version of the usual norm $\|\mathbf{x}\| = \sqrt{x_1^2 + \cdots + x_n^2}$ of vectors $\mathbf{x} = (x_1, \ldots, x_n)$ in $\mathbb{R}^n$.) We find the function in the span of

$$\phi_0(x) = 1,\ \phi_1(x) = \cos x,\ \psi_1(x) = \sin x,\ \phi_2(x) = \cos 2x,\ \psi_2(x) = \sin 2x$$

which minimizes $\left\|x^2 - g(x)\right\|$ among all g in this span. We are using the term "span" here in the same way as in linear algebra: it is the set of all possible linear combinations of the given functions, so in this particular case functions that look like

$$c_0 + c_1 \cos x + c_2 \sin x + c_3 \cos 2x + c_4 \sin 2x$$

where the c_i are scalars.

Now, if you think about the analogous question in the setting of $\mathbb{R}^n$, the answer there is one you would have seen previously: given $\mathbf{x}$ in $\mathbb{R}^n$ and $\mathbf{u}_1, \dots, \mathbf{u}_k$ in $\mathbb{R}^n$, the vector in $\text{span}(\mathbf{u}_1, \dots, \mathbf{u}_k)$ which minimizes $\|\mathbf{x} - \mathbf{y}\|$ among all $\mathbf{y}$ in this span is precisely the orthogonal projection of $\mathbf{x}$ onto this span! Indeed, this is one way of defining what "orthgonal projection" actually means. The same fact is true in more general vector spaces, so this problem is just asking to compute the orthogonal projection of the function x^2 onto the span of the functions $\phi_0, \phi_1, \psi_1, \phi_2, \psi_2$ above. But if you recall the linear-algebra interpretation we gave behind Fourier series last time, this orthoogonal projection is precisely the portion of the full Fourier series of x^2 which only goes up to the $\cos 2x$ and $\sin 2x$ terms:

$$\underbrace{\frac{a_0}{2}}_{\text{proj onto } \phi_0} + \underbrace{a_1 \cos x}_{\text{proj onto } \phi_1} + \underbrace{b_1 \sin x}_{\text{proj onto } \psi_1} + \underbrace{a_2 \cos 2x}_{\text{proj onto } \phi_2} + \underbrace{b_1 \sin 2x}_{\text{proj onto } \psi_2}$$

where for $n \geq 1$:

$$a_n = \frac{(x^2, \phi_n)}{(\phi_n, \phi_n)} = \frac{1}{\pi} \int_{-\pi}^{\pi} x^2 \cos nx \, dx \quad \text{and} \quad b_n = \frac{(x^2, \psi_n)}{(\psi_n, \psi_n)} = \frac{1}{\pi} \int_{-\pi}^{\pi} x^2 \sin nx \, dx.$$

(This uses the fact that the functions $\phi_0, \phi_1, \psi_1, \phi_2, \psi_2$ make up an *orthogonal* basis of this span, so that the orthogonal projection onto this entire span is the sum of the orthogonal projections onto each individual basis vector. For a_0 it is $\frac{a_0}{2}$ which equals the orthogonal projection, not simply a_0 alone, which stems from the difference in values between $(\phi_0, \phi_0) = 2\pi$ and $(\phi_n, \phi_n) = \pi$ for $n \neq 0$. The point of this problem is to emphasize this connection between orthgonal projections and Fourier series; I will not asking something along these lines phrased in terms of linear algebra on an exam, but you should know about it for the sake of context!)

So, we compute some Fourier coefficients. Since x^2 is even on the interval $[-\pi, \pi]$, the coefficients b_n are zero since they involve integrals of the odd functions $x^2 \sin nx$ over symmetric intervals centered at zero. Then:

$$\begin{aligned}
a_0 &= \frac{1}{\pi} \int_{-\pi}^{\pi} x^2 \, dx \\
&= \frac{2}{\pi} \int_{0}^{\pi} x^2 \, dx && \text{since the integrand is even} \\
&= \frac{2\pi^2}{3} \\
a_n &= \frac{1}{\pi} \int_{-\pi}^{\pi} x^2 \cos nx \, dx && (n \geq 1) \\
&= \frac{2}{\pi} \int_{0}^{\pi} x^2 \cos nx \, dx && \text{since the integrand is even} \\
&= \frac{2}{\pi} \left(\frac{1}{n} x^2 \sin nx \Big|_0^{\pi} - \frac{2}{n} \int_0^{\pi} x \sin nx \, dx \right) \\
&= -\frac{4}{\pi n} \left(-\frac{1}{n} x \cos nx \Big|_0^{\pi} + \frac{1}{n^2} \sin nx \Big|_0^{\pi} \right) \\
&= \frac{4 \cos n\pi}{n^2}
\end{aligned}$$

The desired orthogonal projection, and hence function in the span of $\phi_0, \phi_1, \psi_1, \phi_2, \psi_2$ which minimizes $\|x^2 - g\|$ among all g in this span, is

$$\frac{a_0}{2} + a_1 \cos x + b_1 \sin x + a_2 \cos 2x + b_2 \sin 2x = \frac{\pi^2}{3} - 4\cos x + \cos 2x.$$

Orthonormal functions. As you saw in linear algebra, we can also turn an orthogonal basis of a space into an *orthonormal* one by dividing each basis vector by its norm. In our case, for the functions

$$\phi_n = \cos\frac{n\pi x}{L} \quad \text{and} \quad \psi_n = \sin\frac{n\pi x}{L},$$

the orthogonality relatoins say that:

$$\|\phi_0\| = \sqrt{(\phi_0, \phi_0)} = \sqrt{2L},\ \|\phi_n\| = \sqrt{(\phi_n, \phi_n)} = \sqrt{L} = \sqrt{(\psi_n, \psi_n)} = \|\psi_n\|),$$

where $n \geq 1$. Thus, we can produce the following orthonormal basis of our space of periodic functions:

$$\frac{\phi_0}{\|\phi_0\|} = \frac{1}{\sqrt{2L}},\ \frac{\phi_n}{\|\phi_n\|} = \frac{1}{\sqrt{L}}\cos\frac{n\pi x}{L},\ \frac{\psi_n}{\|\psi_n\|} = \frac{1}{\sqrt{L}}\sin\frac{n\pi x}{L}.$$

With these *normalized* functions—call them $\widetilde{\phi_0}$, $\widetilde{\phi_n}$, $\widetilde{\psi_n}$ respectively—the Fourier series looks like:

$$(f, \widetilde{\phi_0})\widetilde{\phi_0} + \sum_{n=1}^{\infty}[(f, \widetilde{\phi_n})\widetilde{\phi_n} + (f, \widetilde{\psi_n})\widetilde{\psi_n}]$$

This perspective will not be crucial for us, but just makes another connection with linear algebra.

Complex orthogonality. Computing a Fourier series explicitly takes quite a bit of work, since there are multiple integrals to compute, some involving cosine and some sine. (Sure, for certian functions, namely those which are even or odd, we can cut down the number of integrals we actually have to compute since some will automatically be zero, but that is not necessarily the case for most functions.) The form of the Fourier series can also be a bit cumbersome to write down, since there's one portion which needs cosines and another which needs sines. It would be nice to have a more compact way of expressing a (real) Fourier series, while at the same reducing the number of integral computations needed.

To see how this can be done, we start by recaling the following identity when dealing with *complex* numbers:

$$e^{i\theta} = \cos\theta + i\sin\theta.$$

(Indeed, this is one possible way of defining what the complex exponential e^{ix} even means.) Then for any n we also have

$$e^{inx} = \cos nx + i\sin nx.$$

The point is that the cosine and sine terms on the right are precisely the ones which show up in a Fourier series (say for a 2π-periodic function), so that it seems plausible that we can rewrite the entire Fourier series using complex exponentials instead.

To set the stage, we consider a vector space of *complex-valued* 2π-periodic functions (still of a real variable), and define take the following as an inner product:

$$(f, g) = \int_{-\pi}^{\pi} f(x)\overline{g(x)}\,dx.$$

The bar on $g(x)$ denotes the *complex conjugate*. (Recall that g is complex-valued, so that even if x is real, $g(x)$ will be a complex number in general.) The need to use a conjugate for this second term can be traced back to the fact that to get the norm of a complex z, we need to take $\sqrt{z\overline{z}}$, which will always be real. (Concretely, when taking the "norm" of f, using $\int_{-\pi}^{\pi} f(x)^2\,dx$ does not work if $f(x)$ is complex since this will not produce a nonnegative real number as a result; we need to use $\int_{-\pi}^{\pi} f(x)\overline{f(x)}\,dx$ instead, which *is* always real and nonnegative, so that then taking the square root makes sense.) Then for the functions e^{inx}, where n now can be any integer (not just nonnegative), we have for $n \neq m$:

$$\begin{aligned}(e^{inx}, e^{imx}) &= \int_{-\pi}^{\pi} e^{inx}\overline{e^{imx}}\,dx \\ &= \int_{-\pi}^{\pi} e^{inx}e^{-imx}\,dx \\ &= \int_{-\pi}^{\pi} e^{i(n-m)x}\,dx \\ &= \frac{1}{i(n-m)}e^{i(n-m)x}\Big|_{-\pi}^{\pi} \\ &= \frac{1}{i(n-m)}[e^{i(n-m)\pi} - e^{-i(n-m)\pi}].\end{aligned}$$

(Note that computing integrals with complex numbers works the same as any other integral computations; simply treat i as a constant like you normally would.) Using $e^{i\theta} = \cos\theta + i\sin\theta$, we can see that:

$$e^{i(n-m)\pi} = \cos([n-m]\pi) + i\sin([n-m]\pi) = \cos([n-m]\pi)$$

and

$$e^{-i(n-m)\pi} = \cos(-[n-m]\pi) + i\sin(-[n-m]\pi) = \cos([n-m]\pi)$$

since sine of an integer multiple of π is zero, and cosine is even. The upshot is that

$$(e^{inx}, e^{imx}) = \frac{1}{i(n-m)}[e^{i(n-m)\pi} - e^{-i(n-m)\pi}] = 0,$$

so the functionx e^{inx} are orthogonal with respect to this complex inner product. Moreover,

$$(e^{inx}, e^{inx}) = \int_{-\pi}^{\pi} e^{inx}e^{-inx}\,dx = \int_{-\pi}^{\pi} dx = 2\pi,$$

so all of these functions have "norm" $\sqrt{2\pi}$.

Complex Fourier series. So, by analogy with the real Fourier series using the orthgonal cosine and sine functions, we can now consider a complex Fourier series using the orthgonal e^{inx} functions instead. Since these functions are orthogonal, *if* a complex-valued function f is expressible as such a Fourier series:

$$f(x) = \sum_{n=-\infty}^{\infty} c_n e^{inx},$$

we get that the coefficients c_n must be given by

$$c_n = \frac{(f, e^{inx})}{(e^{inx}, e^{inx})} = \frac{1}{2\pi}\int_{-\pi}^{\pi} f(x)\overline{e^{inx}}\,dx = \frac{1}{2\pi}\int_{-\pi}^{\pi} f(x)e^{-inx}\,dx.$$

We call this series with these specific coefficients the *complex Fourier series* (also called the *exponential Fourier series*) of f. (Note that the sum here allows for n to be negative, to make sure we cover all of the orthogonal functions e^{inx}. For the cosine/sine Fourier series there was no need to allow n negative, since $\cos(-nx) = \cos(nx)$ and $\sin(-nx) = -\sin(nx)$, so that the terms which would occur for negative n are already accounted for in the terms with positive n. In the complex case, it is not true that e^{inx} is the same as or is the negative of e^{-inx}—rather they are complex conjugates of one another—so each of these must be included in the series.)

The derivation of the formula for c_n comes from the same one we used for the real Fourier coefficients, but just to be clear: take inner product of both sides of

$$f(x) = \sum_{n=-\infty}^{\infty} c_n e^{inx}$$

with e^{imx} to get:

$$(f, e^{imx}) = \left(\sum_{n=-\infty}^{\infty} c_n e^{inx}, e^{imx} \right) = \sum_{n=-\infty}^{\infty} c_n (e^{inx}, e^{imx}),$$

and use orthogonality to see that the only potential nonzero term on the right is the one for $n = m$. Note now that, instead of having to compute a_n and b_n coefficients separately, here we have only one (complex) integral to compute—apart from the $n = 0$ case which often requires its own computation.

The complex Fourier series of f certainly looks cleaner than the previous series using sines and cosines, or at least it is a more compact form to write down. But it is fair to ask just what relation there is between this complex Fourier series and the previous real one? The answer is a very nice one: when f is a real-valued function, the complex Fourier is *literally the same* as the real Fourier series! That is, when f is real, we have equality

$$\frac{a_0}{2} + \sum_{n=1}^{\infty} (a_n \cos nx + b_n \sin nx) = \sum_{n=-\infty}^{\infty} c_n e^{inx}$$

for the values of a_0, a_n, b_n, c_n given before.

So, we are not computing a different type of Fourier series after all, at least for real functions—all we are doing is computing the same Fourier series in an alternative (maybe quicker since it involves fewer integrals?) way. Showing that the complex Fourier series is just a rewritten form of the real Fourier series, for real-valued functions, is a problem on the homework. For the most part this quarter, we will mainly care about the real Fourier series, but having the complex form will make certain things simpler to express from time to time; in particular, the complex form will make the connection between Fourier series and Fourier transforms clearer.

Example. We compute the complex Fourier series of the triangular wave function:

$$f(x) = \begin{cases} -x, & -\pi \le x < 0 \\ x, & 0 \le x < \pi, \end{cases} \qquad f(x + 2\pi) = f(x).$$

Note that since f is real-valued, the complex Fourier series will be exactly the same as the real Fourier series we computed previously, although written to use complex exponentials instead of sines and cosines.

First, we compute:

$$c_0 = \frac{1}{2\pi}\int_{-\pi}^{\pi} f(x)\overbrace{e^{-i0x}}^{1}\,dx = \frac{1}{\pi}\int_0^{\pi} x\,dx = \frac{\pi}{2},$$

where we use the fact that f is even to say that $\int_{-\pi}^{\pi} = 2\int_0^{\pi}$. Next, for $n \neq 0$ we have:

$$c_n = \frac{1}{2\pi}\int_{-\pi}^{\pi} f(x)e^{-inx}\,dx = \frac{1}{2\pi}\left(\int_{-\pi}^{0} -xe^{-inx}\,dx + \int_0^{\pi} xe^{-inx}\,dx\right).$$

Now, in this case we cannot simplify this to twice the integral from 0 to π since, although f is even, $f(x)e^{-inx}$ is *not* even: replacing x by $-x$ gives the complex conjugate of $f(x)e^{-inx}$, but not $f(x)e^{-inx}$ itself. Nonetheless, there is a way to simplify our computation as follows. If we make the change of variables $u = -x$ in the first integral we get:

$$\int_{-\pi}^{0} -xe^{-inx}\,dx = \int_{\pi}^{0} ue^{inu}(-du) = \int_0^{\pi} ue^{inu}\,du,$$

which is the same as $\int_0^{\pi} xe^{inx}\,dx$ after renaming the variable of integration. This is the *conjugate* of the second integral in our expression:

$$\overline{\int_0^{\pi} xe^{-inx}\,dx} = \int_0^{\pi} \overline{xe^{-inx}}\,dx = \int_0^{\pi} xe^{inx}\,dx.$$

Thus, the expression

$$\int_{-\pi}^{0} -xe^{-inx}\,dx + \int_0^{\pi} xe^{-inx}\,dx$$

is the sum of a complex number and its conjugate, so it should equal twice the real part of that complex number:

$$\int_{-\pi}^{0} -xe^{-inx}\,dx + \int_0^{\pi} xe^{-inx}\,dx = 2\left(\text{real part of } \int_0^{\pi} xe^{-inx}\,dx\right).$$

Hence there is only one integral we need to compute, namely $\int_0^{\pi} xe^{-inx}\,dx$.

We have:

$$\begin{aligned}
\int_0^{\pi} xe^{-inx}\,dx &= -\frac{1}{in}xe^{-inx}\Big|_0^{\pi} + \frac{1}{in}\int_0^{\pi} e^{-inx}\,dx \\
&= -\frac{1}{in}\pi e^{-in\pi} - \frac{1}{i^2n^2}e^{-inx}\Big|_0^{\pi} \\
&= -\frac{1}{in}\pi e^{-in\pi} + \frac{1}{n^2}(e^{-in\pi} - 1) \\
&= i\frac{\pi\cos n\pi}{n} + \frac{1}{n^2}(\cos n\pi - 1).
\end{aligned}$$

We only need twice the real part of this, so overall we get

$$c_n = \frac{1}{2\pi}\left(\int_{-\pi}^{0} -xe^{-inx}\,dx + \int_0^{\pi} xe^{-inx}\,dx\right) = \frac{1}{2\pi}\left(\frac{2}{n^2}[\cos n\pi - 1]\right) = \frac{\cos n\pi - 1}{\pi n^2}.$$

The complex Fourier series of f is thus

$$\sum_{n=-\infty}^{\infty} c_n e^{inx} \text{ where } c_0 = \frac{\pi}{2} \text{ and } c_n = \frac{\cos n\pi - 1}{\pi n^2} \text{ for } n \neq 0,$$

which we can also write as

$$\frac{\pi}{2} + \sum_{\substack{n=-\infty \\ n\neq 0}}^{\infty} \frac{\cos n\pi - 1}{\pi n^2} e^{inx},$$

where the $n \neq 0$ in the notation for the sum indicates that the $n = 0$ term is meant to be excluded; the $n = 0$ term is written separately at the start since it has a different form from the other c_n. As alluded to above (based on a homework problem), this complex Fourier series can be shown to be literally the same as the Fourier series

$$\frac{\pi}{2} + \sum_{n=1}^{\infty} \frac{2}{n^2\pi}(\cos n\pi - 1)\cos nx$$

we previously computed in Warm-Up 2 of Lecture 3.

Lecture 5: Fourier Convergence

Warm-Up. We compute the complex Fourier series of $f(x) = x^2, -\pi \leq x \leq \pi$, extended to be 2π-periodic. We have

$$c_0 = \frac{1}{2\pi}\int_{-\pi}^{\pi} x^2\,dx = \frac{\pi^2}{3}$$

and for $n \geq 0$:

$$\begin{aligned}
c_n &= \frac{1}{2\pi}\int_{-\pi}^{\pi} x^2 e^{-inx}\,dx \\
&= \frac{1}{2\pi}\left(-\frac{1}{in}x^2 e^{-inx}\Big|_{-\pi}^{\pi} + \frac{2}{in}\int_{-\pi}^{\pi} xe^{-inx}\,dx\right) \\
&= -\frac{i\pi}{2n}(e^{in\pi} - e^{-in\pi}) + \frac{1}{in\pi}\left(-\frac{1}{in}xe^{-inx}\Big|_{-\pi}^{\pi} - \frac{1}{i^2n^2}e^{-inx}\Big|_{-\pi}^{\pi}\right) \\
&= -\frac{i\pi}{2n}(e^{in\pi} - e^{-in\pi}) + \frac{1}{n^2}e^{-in\pi} + \frac{1}{n^2}e^{in\pi} - \frac{i}{\pi n^3}(e^{in\pi} - e^{-in\pi}),
\end{aligned}$$

which simplies to $\frac{2}{n^2}\cos n\pi$ after using $e^{\pm in\pi} = \cos(n\pi) \pm i\sin(n\pi) = \cos(n\pi)$. The complex Fourier series of f is thus

$$\frac{\pi^2}{3} + \sum_{\substack{n=-\infty \\ n\neq 0}}^{\infty} (-1)^n \frac{2}{n^2} e^{inx}.$$

Note that if we use $e^{inx} = \cos(nx) + i\sin(nx)$, and group the term occuring for $-n$ with the term occuring for n, we get can rewrite this as

$$\frac{\pi^2}{3} + \sum_{\substack{n=-\infty \\ n\neq 0}}^{\infty} (-1)^n \frac{2}{n^2} e^{inx} = \frac{\pi^2}{3} + \sum_{n=1}^{\infty}(-1)^n \frac{2}{n^2}(e^{inx} + e^{-inx}) = \frac{\pi^2}{3} + \sum_{n=1}^{\infty}(-1)^n \frac{4}{n^2}\cos(nx),$$

which is the real Fourier series of f.

Partial sums. We now know how to compute Fourier series, real or complex, but we have not said anything about whether these series actually converge, let alone what they converge to. Ideally we would hope that the Fourier series of a function actually converges to that function itself, since the entire point of this process was in trying to express a given function as a Fourier series. We know that *if* a given function is expressible as a Fourier series, the only possible Fourier series that can work is the one we've defined, but that's different than saying that the series we get in this way does in fact converge to our function.

Before we answer these questions, we first recall what convergence of a series actually means. Given a series $\sum_{n=1}^{\infty} A_n$, we say that it *converges* to S if the sequence of *partial sums*

$$S_N = A_1 + A_2 + \cdots + A_N$$

converges to S as $N \to \infty$. That is, we look at the sequence formed by

$$\begin{aligned} S_1 &= A_1 \\ S_2 &= A_1 + A_2 \\ S_3 &= A_1 + A_2 + A_3 \\ S_4 &= A_1 + A_2 + A_3 + A_4 \\ &\vdots \end{aligned}$$

and ask whether *these* values are converging as we add on more and more terms. In the case of a Fourier series the partial sums looks like

$$\begin{aligned} &\frac{a_0}{2} + \sum_{n=1}^{N} (a_n \cos \tfrac{n\pi x}{L} + b_n \sin \tfrac{n\pi x}{L}) \\ &= \frac{a_0}{2} + a_1 \cos \tfrac{\pi x}{L} + b_1 \sin \tfrac{\pi x}{L} + a_2 \cos \tfrac{2\pi x}{L} + b_2 \sin \tfrac{2\pi x}{L} + \cdots + a_N \cos \tfrac{N\pi x}{L} + b_N \sin \tfrac{N\pi x}{L}. \end{aligned}$$

Specifically, this is called the *N-th order* partial sum of the Fourier series, and is an example of a *trigonometric polynomial*, which in general refers to a linear combination of sines and cosines. If we are looking at the Fourier series of a function f, the question is then whether these trigonometric polynomials converge to f.

Fourier convergence. We will give the answer to the question above, which is easy to state, and next time we'll say something about *why* it's true. The *Fourier convergence theorem* states the following:

> Suppose f is a piecewise C^1 function, which means that it is piecewise continuous with a piecewise continuous derivative. Then the Fourier series of f:
>
> $$\frac{a_0}{2} + \sum_{n=1}^{\infty} (a_n \cos \tfrac{n\pi x}{L} + b_n \sin \tfrac{n\pi x}{L}),$$
>
> where
>
> $$a_n = \frac{1}{L} \int_{-L}^{L} f(x) \cos \tfrac{n\pi x}{L}\, dx \quad \text{and} \quad b_n = \frac{1}{L} \int_{-L}^{L} f(x) \sin \tfrac{n\pi x}{L}\, dx,$$

converges to the function given by

$$\frac{f(x^-) + f(x^+)}{2}$$

where $f(x^-) = \lim_{t \to x^-} f(t)$ is the limit of f as t approaches x from the left and $f(x^+) = \lim_{t \to x^+} f(t)$ the limit as t approaches x from the right.

So, at any x, the value to which the Fourier series of f converges is the average of the left- and right-hand limits of f at x. Note that if f happens to be continuous at x, the left- and right-hand limits above both equal $f(x)$, so this average value is just $f(x)$ in this case. Hence at points x where f is continuous, the equality

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

does hold for a_n, b_n as defined above.

There's more to say about this, and in particular there are some subtleties behind what type of "convergence" we actually mean here, but we'll come back to this. Also, we'll say something about how the bad the "error" can be when approximating a function by its various Fourier partial sums.

Example. Consider the square wave function

$$f(x) = \begin{cases} 0, & -1 \leq x < 0 \\ 1, & 0 \leq x < 1 \end{cases}; \quad f(x+2) = f(x)$$

whose Fourier series we previously found to be

$$\frac{1}{2} + \sum_{n=1}^{\infty} \frac{1 - \cos n\pi}{n\pi} \sin n\pi x.$$

Note that f is piecewise constant, so it is piecewise C^1 as well. Thus the Fourier convergence theorem applies, and says that this particular series converges at a point x to the average of the one-sided limits at x.

For $-1 < x < 0$, f is continuous at x, so at these points this Fourier series converges to $f(x) = 0$. For $1 < x < 0$, f is again continuous at x, so here the series above converges to $f(x) = 1$. Now, f is not continuous at, say $x = 0$, so here the Fourier series converges to the average of 0 (the left-hand liimt at 0) and 1 (the right-hand limit), so it converges to $\frac{1}{2}$. This is also true at $x = 1$,$x = -1$, and in fact at any integer value of x. Thus overall the Fourier series above converges to the function defined by

$$g(x) = \begin{cases} 0, & -1 < x < 0 \\ \frac{1}{2}, & x = -1, 0, 1 \\ 1, & 0 < x < 1 \end{cases}$$

and extended periodically. For example, at $x = \frac{1}{2}$, we thus have the following identity:

$$1 = \frac{1}{2} + \sum_{n=1}^{\infty} \frac{1 - \cos n\pi}{n\pi} \sin \frac{n\pi}{2} = \frac{1}{2} + \sum_{n=0}^{\infty} \frac{2}{(2n+1)\pi} \sin \frac{(2n+1)\pi}{2} = \frac{1}{2} + \sum_{n=0}^{\infty} (-1)^n \frac{2}{(2n+1)\pi}.$$

After rearranging terms, we can turn this into

$$\frac{\pi}{4} = \sum_{n=0}^{\infty} \frac{(-1)^n}{2n+1} = 1 - \frac{1}{3} + \frac{1}{5} - \frac{1}{7} + \cdots,$$

which is a nice identity.

The function to which the series above converges has a graph which looks like:

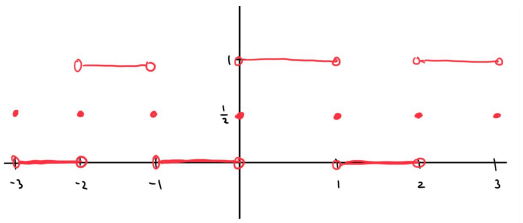

This is exactly the graph of the original square wave function f, except for at the integer points where f was not continuous. Now, we can also plot some of the Fourier partial sums to see the convergence actually happening before our eyes. The 1-st order partial sum is

$$\frac{1}{2} + \frac{2}{\pi} \sin \pi x,$$

whose graph (super imposed on the triangular wave) is:

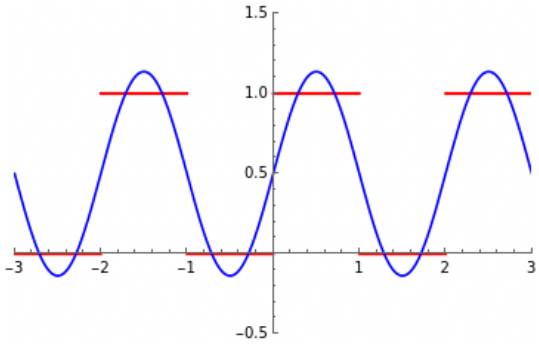

The 3-rd order partial sum

$$\frac{1}{2} + \frac{2}{\pi} \sin \pi x + \frac{2}{3\pi} \sin 3\pi x$$

and 5-th order parital sum

$$\frac{1}{2} + \frac{2}{\pi} \sin \pi x + \frac{2}{3\pi} \sin 3\pi x + \frac{2}{5\pi} \sin 5\pi x$$

have the following graphs, respectively:

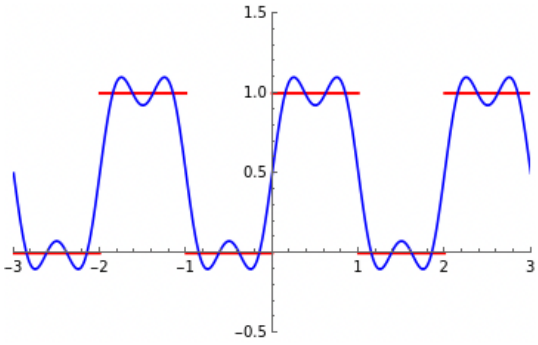

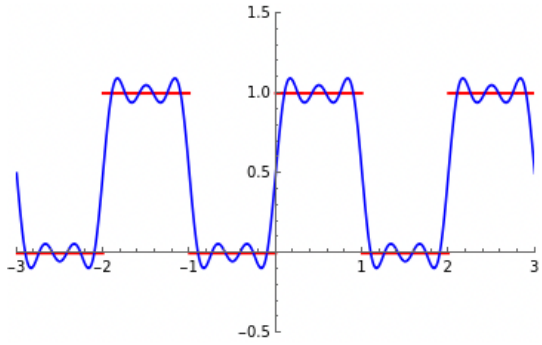

We can literally see the convergence happening: as the order of the partial sum increases, so that we add on more and more terms, we get even better approximations to the square wave:

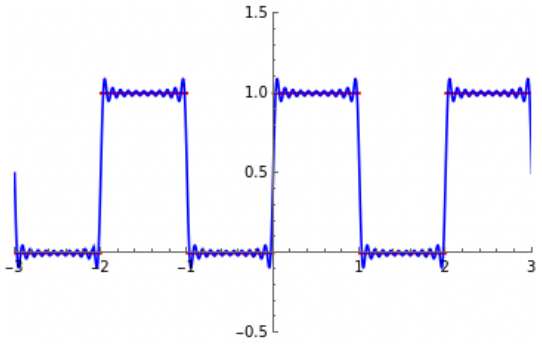

In the limit, the steep jumps at the points at which the triangular wave is discontinuous become the single points at heights $\frac{1}{2}$ in the limit function.

Another example. The triangular wave function

$$f(x) = \begin{cases} -x & -\pi \le x < 0 \\ x & 0 \le x < \pi \end{cases}; \quad f(x+2\pi) = f(x)$$

has Fourier series

$$\frac{\pi}{2} + \sum_{n=1}^{\infty} \frac{2}{n^2\pi}(\cos n\pi - 1)\cos nx = \frac{\pi}{2} - \sum_{n=0}^{\infty} \frac{4}{(2n+1)^2\pi}\cos([2n+1]x).$$

In this case, f is continuous at all points, so the Fourier series converges to the value of f at all points. Thus we have. for example,

$$-x = \frac{\pi}{2} - \sum_{n=0}^{\infty} \frac{4}{(2n+1)^2\pi}\cos([2n+1]x) \quad \text{for } -\pi \le x \le 0$$

and

$$x = \frac{\pi}{2} - \sum_{n=0}^{\infty} \frac{4}{(2n+1)^2\pi}\cos([2n+1]x) \quad \text{for } 0 \le x \le \pi.$$

At $x = \pi$ then, this gives

$$\pi = \frac{\pi}{2} - \sum_{n=0}^{\infty} \frac{4}{(2n+1)^2\pi}\cos([2n+1]\pi) = \frac{\pi}{2} + \sum_{n=0}^{\infty} \frac{4}{(2n+1)^2\pi},$$

which after rearranging gives the following nice identity:

$$\frac{\pi^2}{8} = \sum_{n=0}^{\infty} \frac{1}{(2n+1)^2} = 1 + \frac{1}{3^2} + \frac{1}{5^2} + \frac{1}{7^2} + \cdots.$$

The first few nonconstant Fourier partial sums are

$$\frac{\pi}{2} - \frac{4}{\pi}\cos x, \ \frac{\pi}{2} - \frac{4}{\pi}\cos x - \frac{4}{9\pi}\cos 3x, \ \frac{\pi}{2} - \frac{4}{\pi}\cos x - \frac{4}{9\pi}\cos 3x - \frac{4}{25\pi}\cos 5x,$$

whose graphs (superimposed on the triangular wave graph), respectively, look like;

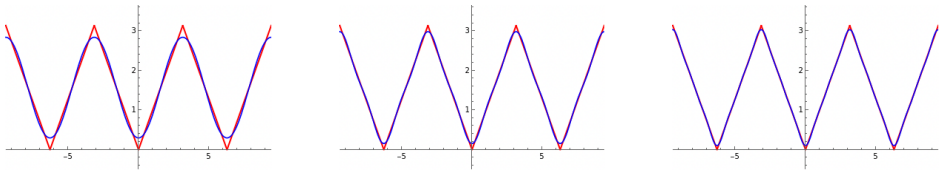

Again, these Fourier approximations just get better and better as we take more and more terms. (This last one in particular is already pretty good!)

Lecture 6: More on Convergence

Warm-Up. The Fourier series of $f(x) = x^2, -\pi \leq x \leq \pi$ (extend periodically) is

$$\frac{\pi^2}{3} + \sum_{n=1}^{\infty} \frac{4\cos n\pi}{n^2} \cos nx.$$

We use this to find the value of

$$\sum_{n=1}^{\infty} \frac{1}{n^2}.$$

(This is the value of what's called the *Riemnan zeta function* at 2. This function is more generally defined by $\zeta(s) = \sum_{n=1}^{\infty} \frac{1}{n^s}$ and is an important function in number theory, complex analysis, and elsewhere.) The key point is that since f is continuous everywehere, this Fourier series converges to $f(x)$ at all x, so that in particular

$$x^2 = \frac{\pi^2}{3} + \sum_{n=1}^{\infty} \frac{4\cos n\pi}{n^2} \cos nx \text{ for } -\pi \leq x \leq \pi.$$

Set $x = \pi$, so that

$$\pi^2 = \frac{\pi^2}{3} + \sum_{n=1}^{\infty} \frac{4\cos n\pi}{n^2} \cos n\pi = \frac{\pi^2}{3} + \sum_{n=1}^{\infty} \frac{4}{n^2}.$$

Rearranging gives

$$\sum_{n=1}^{\infty} \frac{1}{n^2} = \frac{1}{4}\left(\pi^2 - \frac{\pi^2}{3}\right) = \frac{\pi^2}{6}$$

as the desired value.

We can see the convergence of the Fourier series in action by plotting a few Fourier partial sums:

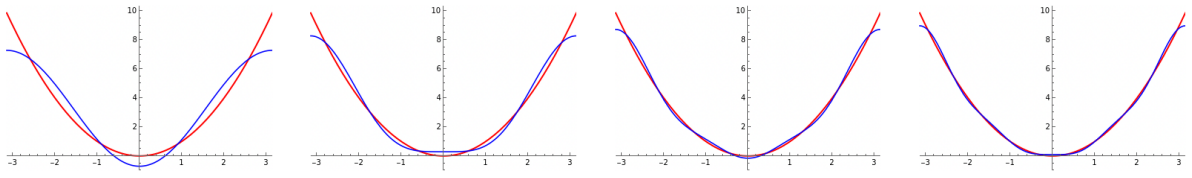

Pointwise vs uniform convergence. Now we clarify something we alluded to last time, namely just what *type* of convergence we're talking about when we say that a Fourier series convergence. If you have not had a course in real analysis (specfically MATH 320-2 or MATH 321-2 at Northwestern), you've probably never even heard that different types of convergence are possible, but

understanding what is actually happenging will be important in order to clarify an earlier computation we performed—that of derviing the formula for the Fourier coefficients. Now, we will not give formal definitions here, since after all having had an analysis course is not a prerequisite for this course, so our aim is simply to give some intuitive sense behind what we are talking about.

We say that a sequence of functions g_n converges *pointwise* to the function g if for any fixed x, the number $g_n(x)$ (obtained by evaluating g_n at the specific input x) converges in the usual numerical sense to the number $g(x)$. Then, a *series* of functions $\sum_{n=1}^{\infty} f_n$ converges pointwise to a function f if the sequence of partial sums

$$f_1 + \cdots + f_N$$

converges pointwise to f as $N \to \infty$. We will look at an example in a bit, but ultimately this just depends on the notion of convergence for numbers we're already used to: plug in x, and look at the behavior of the resulting sequence (or series) of numbers.

One point here is that, when considering pointwise convergence, the *rate* at which the convergence occurs can vary as we move from point to point. We won't give a precise definition of "rate" here, but will use the following intuition: when considering convergence of a sequence of numbers a_n, to say that they converge to the number a means that once we are far enough along the way in our sequence, the numbers a_n are "close enough" to a. We can then intrepret the term at which we do end up "close" enough to a as a measure of how quickly the convergence occurs, so that if we have to go very far in our sequence in order to end up "close enough" to a, we think of the convergence as occuring slowly. So, for a sequence of functions g_n converging pointwise, for a fixed x we have a certain rate of convergence for the numbers $g_n(x)$ to the number $g(x)$, measured by how large n must be to end up "close" to $g(x)$, but this n might change if we change the x we are considering. Maybe the convergence happens more slowly for some x (so we have to go further in the sequence $g_n(x)$) than it does for other x.

We say that the sequence of functions g_n converges *uniformly* to the function g if this rate of convergence can be chosen to be the same for *all* x at once. So, we can find a single n so that $g_n(x)$ ends up "close enough" to $g(x)$ *no matter* what x we plug in. (The name "uniform" comes from the idea that this one n works "uniformly" across all x.) The same then applies to series of functions: $\sum_{n=1}^{\infty} f_n(x)$ converges uniformly to $f(x)$ if the partial sums converge uniformly to f. (Again, there are precise definitions of "pointwise convergence" and "uniform convergence" that we are skipping here.)

Example. To give a basic example of this distinction, consider the sequence of functions $f_n(x) = x^n$ for $n \geq 1$ on the interval $(0, 1)$. For any x in this interval, the numbers x^n get closer and closer to 0 as n increases, so these numbers converge to 0, and hence the sequence of functions f_n converges pointwise to the constant zero function. Here is what this looks like visually:

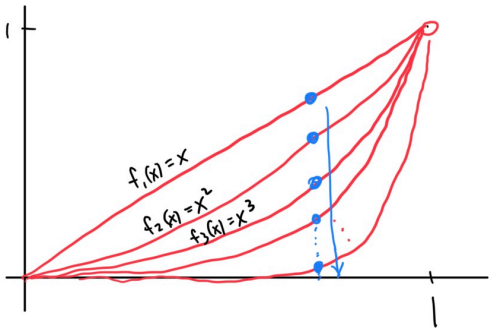

Again, note that for any $0 < x < 1$, the values of x^n (i.e. the heights of the corresponding points on the graph) converge vertically down to zero as n increases.

But, we now ask, for a fixed x and measure of how close we want x^n to actually be to zero, how large must n be in order to make this happen? Consider vertical intervals around 0 at various x:

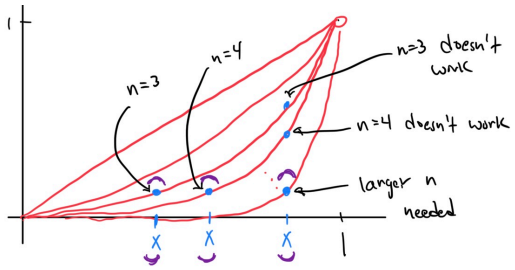

To end up within this distance away from 0 at the first x we have drawn, it appears that $n = 3$ works. But for the next x we have drawn, $n = 3$ does *not* put us within this same distance away from zero, and indeed we need to use $n = 4$ in this case. This is a reflection of the fact that the convergence of x^n to 0 occurs more "slowly" at the second x as opposed to the first. For the third x, we need an even larger value of n to put us within this prescribed distance away from zero, and so on: as x gets closer and closer to 1, the value of n needed to put us within this range away from 0 gets larger and larger, so there will be no one single value of n which puts x^n within the prescribed distance away from 0 for *all* x at once. Thus, the sequence of functions x^n does *not* converge uniformly to the constant zero function.

Imagine now that we do not care about the entire interval $(0, 1)$, and only about x in the interval $(0, \frac{1}{2}]$. The sequence x^n converges pointwise to the constant zero function on this interval, but now we claim that the convergence *is* uniform. The key point is that for x in this particular interval, we have

$$x^n \leq \left(\frac{1}{2}\right)^n = \frac{1}{2^n}.$$

Thus, for a given measure of how close we want to end up to 0, as long as we pick an n large enough to make $\frac{1}{2^n}$ within this measure, we are guaranteed that x^n will also be within this measure *regardless* of x in $(0, \frac{1}{2})$ is. This one n works *uniformly* for all x in this interval:

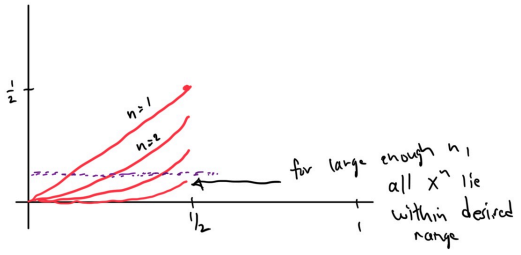

For example, if we want to end up with 0.01 away from 0, we can take $n = 7$ since $\frac{1}{2^7} < 0.01$; and if we want to end up within 0.001 away from zero we can take $n = 10$ since $\frac{1}{2^{10}} < 0.001$. In fact, as long as we fix any $0 < a < 1$ and only consider x in the interval $[0, a]$, we will have that x^n converges uniformly to 0 on $[0, a]$: since $x^n \leq a^n$ for x in this interval, picking n large enough (independent of x) to make a^n small enough will force x^n to be small enough as well. This type of reasoning does not work when we consider all of $(0, 1)$ at once we the best bound we can find is $x^n \leq 1^n = 1$ for x in this interval, but we cannot make 1^n small enough no matter what n is since

1^n is always exactly 1; so for example, we can never make x^n be uniformly within 0.01 away from zero for all x in $(0, 1)$ at once.

The conclusion is that x^n does not converge uniformly to 0 on all of $[0, 1)$, but it does converge uniformly to 0 on any smaller interval $[0, a]$ for a fixed $a < 1$.

Why do we care? The reason for why uniform convergence is important in our case traces back to a subtle point we made when discussing the derivation of the formula for the Fourier coefficients of a function. Recall that we began with

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

Then we multiplied both sides by, say, $\cos \frac{m\pi x}{L}$ and integrated the resulting expression. As part of this, in order to integrate the infinite sum we exchanged summation and integration:

$$\int_{-L}^{L} \left[\sum_{n=1}^{\infty}(a_n \cos \tfrac{n\pi x}{L} + b_n \sin \tfrac{n\pi x}{L})\right] dx = \sum_{n=1}^{\infty} \left[\int_{-L}^{L}(a_n \cos \tfrac{n\pi x}{L} + b_n \sin \tfrac{n\pi x}{L})\, dx\right].$$

This ability to exchange integration and summation is certainly always true for finite sums:

$$\int (f_1(x) + \cdots + f_n(x))\, dx = \int f_1(x)\, dx + \cdots + \int f_n(x)\, dx,$$

but it turns out that in order to guarantee that this works for infinite sums as well actually *requires* that the series question converge *uniformly* and not just pointwise! That is, the equation

$$\int \left[\sum_n f_n(x)\right] dx = \sum_n \left[\int f_n(x)\, dx\right]$$

is *not* automatically true if we only know that the series $\sum_n f_n$ converges pointwise, but it *is* definitely true if we know that the convergence is uniform. Thus, in order for our derivation of the Fourier coefficients to in fact be valid, we have to know that our Fourier series converges uniformly.

(The same type of phenomena is also true for power series, where the method of integrating a power series "term-by-term" depends on *uniform* convergence of a power series. A similar thing is also true for *differentiating* power series and Fourier series, in that something like

$$\left(\sum_n f_n\right)' = \sum_n f_n'$$

for infinite series is only true under a type of uniform convergence assumption.)

Fourier convergence, more precisely. So, let us now state a more precise version of the Fourier convergence theorem, which justifies that our formula for the Fourier coefficients was in fact valid:

> Under the same setup as before, with f a $2L$-periodic piecewise C^1 function, we have that the Fourier series of f converges *pointwise* to the function
>
> $$\frac{f(x^-) + f(x^+)}{2}.$$
>
> Moreover, on any closed interval $[a, b]$ on which f is actually continuous, the convergence is uniform, so that the Fourier series converges to f uniformly on such intervals.

(There is still a bit of subtlety in terms of what happens at the endpoints $[-L, L]$ of the interval of integration in terms of whether f is continuous at them—since we need continuity to guarantee uniform convergence—but suffice it to say that this can be dealt with in a way which guarantees everything works as it should. Huzzah!)

Lecture 7: Even/Odd Extensions

Warm-Up. Consider the functions defined by $f_n(x) = \frac{1}{n}\sin(nx)$. We argue that these functions converge uniformly (as n increases) to the constant zero function. First we note that these function do converge *pointwise* to the constant zero function: for any fixed x, the numbers $\sin(nx)$ are always between -1 and 1 as n varies, so

$$\lim_{n\to\infty} \frac{1}{n}\sin(nx) = 0.$$

Now, to say that the f_n converge *uniformly* to the constant zero function, we need to know that no matter how close we want the value of $\frac{1}{n}\sin(nx)$ to end up to 0, we can find one single n which guarantees we do so regardless of what x actually is. In this case, the point is that since

$$\left|\frac{1}{n}\sin(nx) - 0\right| \leq \frac{1}{n}$$

for all x, we can pick n which makes $\frac{1}{n}$ within whatever "error" we want independent of x. For example, if we want to end up within 0.1 away from the limit 0, we can pick $n = 11$ since $\frac{1}{11} < 0.1$. Then for any x the value of $\frac{1}{n}\sin(nx)$ will be within 0.1 away from 0 for any $n \geq 11$. If we want the value of $f_n(x)$ to be within 0.01 away from 0, then we can take $n = 101$, so the starting with $\frac{1}{101}\sin(101x)$ we will be within 0.01 away from 0 for all values of x. Since this is possible for any desired "error", f_n does converge uniformly to 0.

The reason why this works, as in the x^n on $[0, 1/2]$ example last time, is because we can find a bound on the error which is independent of x—$\frac{1}{n}$ in this case and $\frac{1}{2^n}$ last time—and will approach 0 as n increases. Graphically, if we are given a "tube" around the graph of the constant zero function, to say that $\frac{1}{n}\sin(nx)$ converges uniformly to 0 is to say that once n is large enough, the graph of $\frac{1}{n}\sin(nx)$ will lie completely within this tube:

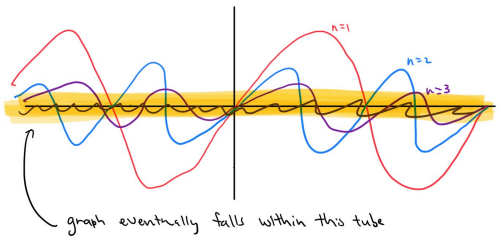

As the "tube" shrinks perhaps we need to take n larger, but still we will eventually find such an n.

Back to partial sums. We can see the uniform convergence properties of Fourier series visually if we consider some graphs. Take the Fourier series

$$\frac{1}{2} + \sum_{n=1}^{\infty} \frac{1 - \cos n\pi}{n\pi} \sin n\pi x$$

of the square wave function

$$f(x) = \begin{cases} 0, & -1 \leq x < 0 \\ 1, & 0 \leq x \leq 1 \end{cases}; \quad f(x+2) = f(x)$$

we've considered previously. Let us focus only the behavior for $0 < x < 1$ to get a clearer picture. For such x the Fourier series converges to $f(x) = 1$ since f is continuous on $0 < x < 1$. Moreover, we should have uniform convergence on any closed interval $[a, b]$ sitting inside $(0, 1)$.

To see this, take an "error tube" around the graph of the square wave:

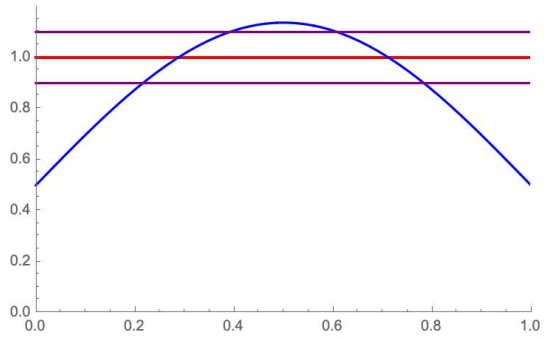

The graph of the 1-st order partial sum which is drawn does not fall within this tube, but if we take a higher-order partial sum it does, at least on some closed interval:

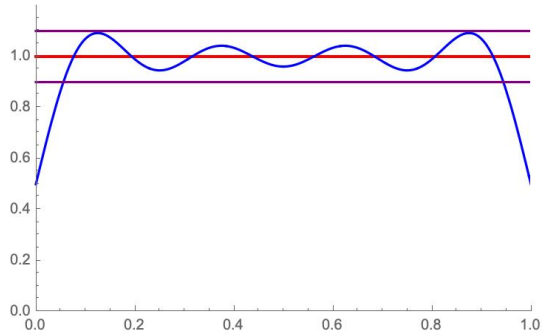

If we take a smaller tube, this specific partial sum no longer works, so we need to take an even higher-order partial sum to find one which does:

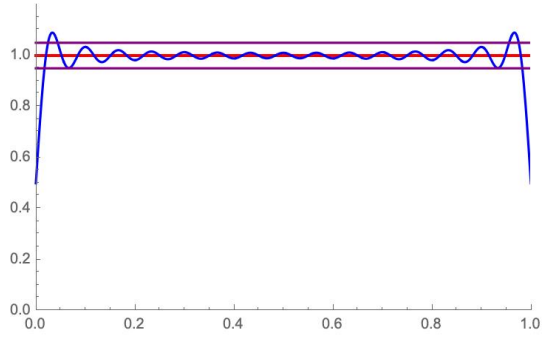

And so on, no matter how small of a tube we start with, we will find that eventually our partial sums will have graphs that fall within this tube, whic his what uniform convergence requires. (The fact that we always have points near 0 or 1 where the partial sum graph "jumps" out of the tube is a reflection of the fact that f is discontinuous at these points, so that uniform convergence does not extend to these.)

Dirichlet kernels. Let us say just a bit more about the errors obtained when trying to approximate a function via its Fourier series. Take the N-th order Fourier partial sum of f:

$$(S_N f)(x) := \frac{a_0}{2} + \sum_{k=1}^{N} (a_k \cos \tfrac{k\pi x}{L} + b_k \sin \tfrac{k\pi x}{L}).$$

To simplify some notation, let us consider only the case of period $2L = 2\pi$. We want to say something about the difference/error

$$f(x) - (S_N f)(x).$$

Now, using the integral formulas for a_0, a_k, b_k, we can write $(S_N f)(x)$ as

$$(S_N f)(x) = \frac{1}{2\pi}\int_{-\pi}^{\pi} f(y)\,dy + \sum_{k=1}^{N}\left(\frac{1}{\pi}\int_{-\pi}^{\pi} f(y)\cos(ky)\,dy\right)\cos(kx) + \sum_{k=1}^{N}\left(\frac{1}{\pi}\int_{-\pi}^{\pi} f(y)\sin(ky)\,dy\right)\sin(kx)$$

Note that we are using y as the variable of integration in the formulas for the coefficients here, so as to not confuse it with x, which is the value at which the partial sum is being evaluated. Since the $\cos(kx)$ and $\sin(kx)$ terms are independent of y, they can each be brought inside the integrals they are being multiplied by, and since all sums here are finite, we can certainly exchange summation and integration. After doing so and combining some things we can put everything under one big integral to obtain:

$$(S_N f)(x) = \int_{-\pi}^{\pi}\left(\frac{1}{2\pi}f(y) + \frac{1}{\pi}\sum_{k=1}^{N}[f(y)\cos(ky)\cos(kx) + f(y)\sin(ky)\sin(kx)]\right)dy.$$

We can manipulate once more by factoring out $\frac{1}{\pi}$ and $f(y)$:

$$(S_N f)(x) = \frac{1}{\pi}\int_{-\pi}^{\pi}\left(\frac{1}{2} + \sum_{k=1}^{N}[\cos(ky)\cos(kx) + \sin(ky)\sin(kx)]\right)f(y)\,dy.$$

Finally, use the trigonometric identity $\cos(a-b) = \cos(a)\cos(b) + \sin(a)\sin(b)$ to rewrite this as

$$(S_N f)(x) = \frac{1}{\pi}\int_{-\pi}^{\pi}\left(\frac{1}{2} + \sum_{k=1}^{N}\cos[k(x-y)]\right)f(y)\,dy.$$

Why did we go through all this trouble? The point is that we have now found a way to express the partial sum $S_N f$ as an integral involving f and the function

$$D_N(x-y) = \frac{1}{2} + \sum_{k=1}^{n}\cos[k(x-y)].$$

It turns out that understanding the convergence of these partial sums comes down to understanding properties of these functions, and *this* is why having such an integral expression is useful. The function $D_N(x)$ is called the *N-th order Dirichlet kernel* and looks like:

$$D_N(x) = \frac{1}{2} + \sum_{k=1}^{N}\cos kx = \frac{1}{2} + \cos x + \cos 2x + \cdots + \cos nx.$$

After making a change of variables $u = x - y$, and using periodicity, it turns out that we can write the expression for $(S_N f)(x)$ as

$$(S_N f)(x) = \frac{1}{\pi}\int_{-\pi}^{\pi} D_N(x-y)\,f(y)\,dy = \frac{1}{\pi}\int_{-\pi}^{\pi} D_N(u)f(x-u)\,du.$$

Coming back to the error, it can be checked the Dirichlet kernels have the property that

$$\frac{1}{\pi}\int_{-\pi}^{\pi} D_N(u)\,du = 1,$$

and using this we can rewrite the value $f(x)$ as

$$f(x) = f(x)\frac{1}{\pi}\int_{-\pi}^{\pi} D_N(u)\,du = \frac{1}{\pi}\int_{-\pi}^{\pi} D_N(u)f(x)\,du.$$

Thus, putting everything together gives the following expression for our error:

$$f(x) - (S_N f)(x) = \frac{1}{\pi}\int_{-\pi}^{\pi} D_N(u)[f(x) - f(x-u)]\,du.$$

This is the Fourier series analog of expressions you might have seen for Taylor remainders for power series previously, and is the key to demonstrating convergence properties of Fourier series. In particular, as mentioned earlier, the fact that this "error" approaches 0 as $N \to \infty$ comes down to properties of the Dirichlet kernel. We will not delve into this more in this course, but it is something you would see more of in a course (second quarter) in real analysis. Our main goal here was to show that there *was* a way to give a concrete expression for the Fourier series error term, albeit one that depends on a new type of function, namely the Dirichlet kernels.

(The word "kernel" here is not used in the same way as you might have seen in linear algebra, where there the kernel of a linear transformation refers to all input that result in the zero vector upon applying said transformation. There *is* a relation between our usage of the word "kernel" here and linear algebra, but we will save this discussion for when talk about the Fourier transform. In short, the Dirichlet "kernel" can be viewed as an "infinite-dimensional" matrix, in a sense.)

Extensions. In the PDEs we will soon consider, we will only be given the data of a function $f(x)$ for x in some interval $0 < x < L$. It does not make sense to talk about the Fourier series of such a function just yet, since deriving a Fourier series requires integrating over a *symmetric* interval $[-L, L]$. But, we can get around this essentially by just *defining* the value of function to be whatever we want for $-L < x < 0$, and considering the Fourier series of this *extension* of f.

For example, take $f(x) = x^2$ for $0 \le x \le \pi$. To express this as a Fourier series requires knowing values for $-\pi \le x \le 0$, so we can take the so-called *even* extension of f and consider the function on $-\pi \le x \le \pi$ which looks like

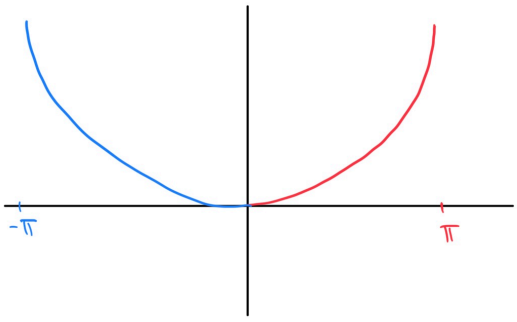

(The even extension is the one defined by enforcing $f(-x) = f(x)$ to hold, so that its graph is symmetric about the y-axis. In this case, all we are doing is simply noting that $f(x) = x^2$ already makes sense for $-\pi \le x \le \pi$ as well, and the even extension is just given by $f(x) = x^2$ itself.) The

Fourier series of this extended function will then converge to $f(x) = x^2$ at least on $0 \leq x \leq -\pi$, since f equals the extension on this interval and is continuous. (Again, in this case the even extension has the same formula x^2 on $-\pi \leq x \leq 0$ too, so the Fourier series also converges to x^2 on this interval. This will not be true for even extensions of arbitrary functions, however.)

Now, the Fourier series of this even extension will not involve any sine terms, since the sine coefficients b_n all end up being zero because they involve the integral of an odd function (sine is odd, and even times odd is odd) over a symmetric integral:

$$b_n = \frac{1}{L} \int_{-L}^{L} \overbrace{\underbrace{f(x)}_{even} \underbrace{\sin \frac{n\pi x}{L}}_{odd}}^{odd} \, dx = 0.$$

To be clear, the point is that if you break this integral up into two pieces:

$$\int_{-L}^{L} f(x) \sin \frac{n\pi x}{L} \, dx = \int_{-L}^{0} f(x) \sin \frac{n\pi x}{L} \, dx + \int_{0}^{L} f(x) \sin \frac{n\pi x}{L} \, dx,$$

the change of variables $x \mapsto -x$ will turn the second integral into the negative of the first, since the integrand is odd. Along similar lines, the a_0, a_n coefficients all involve the integral of an even function (cosine is even, and even times even is even) over a symmetric interval, and hence can be written as twice the integral over only half of that interval:

$$\int_{-L}^{L} (even) \, dx = \int_{-L}^{0} (even) \, dx + \int_{0}^{L} (even) \, dx = 2 \int_{0}^{L} (even) \, dx$$

since the change of variables $x \mapsto -x$ turns the first integral in the middle into the second one in the middle. Thus, the Fourier series of the even extension of f looks like

$$\frac{a_0}{2} + \sum_{n=1}^{\infty} a_n \cos \frac{n\pi x}{L}$$

where

$$a_n = \frac{2}{L} \int_{0}^{L} f(x) \cos \frac{n\pi x}{L} \, dx.$$

This series converges to the original f on $0 < x < L$ (assuming f is continuous here), and is called the *Fourier cosine series* of f. Again, the point is that we are only given the values of f originally on $0 < x < L$, and we obtain a way to express it as a Fourier cosine series on this interval by computing the Fourier series of its even extension. For the example of $f(x) = x^2, 0 \leq x \leq \pi$, the Fourier cosine series is

$$\frac{2\pi^2}{3} + \sum_{n=1}^{\infty} \frac{4 \cos n\pi}{n^2} \cos nx.$$

But of course, the even extension of $f(x) = x^2, 0 \leq x \leq \pi$ is not the only extension we could have chosen. Let us also consider the *odd* extension, which is obtained by enforcing $f(-x) = -f(x)$ to be true, so that the graph is symmetric about the origin:

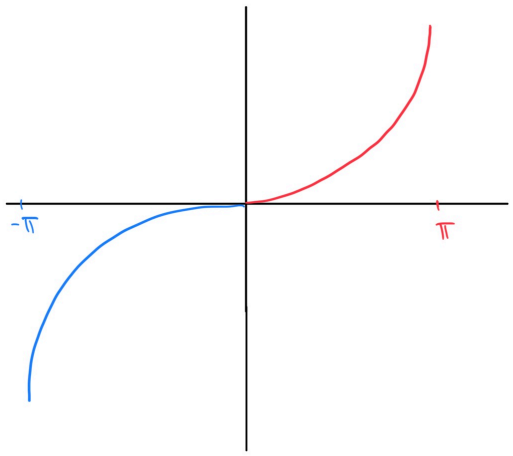

The Fourier series of this odd extension will then converge to the original $f(x) = x^2$ on $0 < x < \pi$ (since f is continuous here), and will give an alternative way of writing x^2 as a Fourier series as compared to the cosine series above. In this case, all a_0, a_n coefficients will be zero since they each involving integrating an odd function (odd times even) over a symmetric interval, and the b_n coefficients (which involve the product of two odd functions, which is itself even) "double-up". The resulting Fourier series thus looks like

$$\sum_{n=1}^{\infty} b_n \sin \frac{n\pi x}{L}$$

where

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

This is called the *Fourier sine series* of f on $0 < x < L$. In the $f(x) = x^2, 0 < x < \pi$ example, the Fourier sine series is

$$\sum_{n=1}^{\infty} \frac{2[(2 - \pi^2 n^2) \cos n\pi - 2]}{\pi n^3} \sin nx.$$

We thus now have two different ways of writing x^2 for $0 < x < \pi$ as a Fourier series:

$$x^2 = \frac{2\pi^2}{3} + \sum_{n=1}^{\infty} \frac{4 \cos n\pi}{n^2} \cos nx = \sum_{n=1}^{\infty} \frac{2[(2 - \pi^2 n^2) \cos n\pi - 2]}{\pi n^3} \sin nx.$$

There are of course other extensions we could consider, say ones that are neither even nor odd:

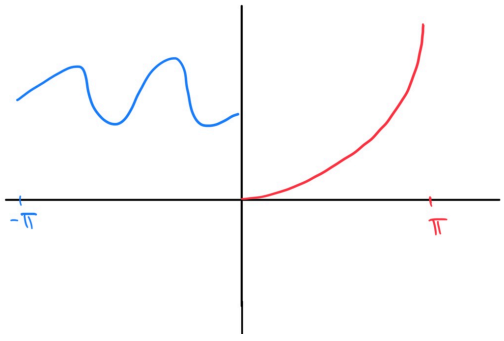

Any such random (let's assume at least continuous) way of extending f to $-L < x < 0$ will result in a Fourier series that converges to f on the original interval $0 < x < L$. So, Fourier series expressions on half-intervals like $(0, L)$ are never unique, and only become unique once we specify what should happen over $(-L, 0)$ as well. The Fourier cosine and sine series, resulting from the even and odd extensions, will be the most useful ones for the purposes of solving PDEs.

Lecture 8: Heat Equation

Warm-Up. We find the Fourier cosine and sine series of $f(x) = 1 - x,\ 0 \le x \le 1$. The Fourier cosine coefficients are:

$$a_0 = \frac{2}{1}\int_0^1 (1-x)\,dx = 1$$
$$a_n = \frac{2}{1}\int_0^1 (1-x)\cos(\tfrac{n\pi x}{1})\,dx = \frac{2(1-\cos n\pi)}{\pi^2 n^2},$$

so the Fourier cosine series of f is

$$\frac{1}{2} + \sum_{n=1}^{\infty} \frac{2(1-\cos n\pi)}{\pi^2 n^2}\cos(n\pi x).$$

The Fourier sine coefficients of f are

$$b_n = \frac{2}{1}\int_0^1 (1-x)\sin(\tfrac{n\pi x}{1})\,dx = \frac{2}{n\pi},$$

so the Fourier sine series of $f(x) = 1 - x$ is

$$\sum_{n=1}^{\infty} \frac{2}{n\pi}\sin(n\pi x).$$

Both of these series converge to $1 - x$ for $0 < x < 1$. (We have to exclude the endpoints in the odd case, since the odd extension of f will not be continuous at $x = 0$.)

But just for fun, we can also take the extension which is still defined by the same formula $f(x) = 1 - x$, now for $-1 < x < 0$. Altogether the three extensions we're using have graphs that look like

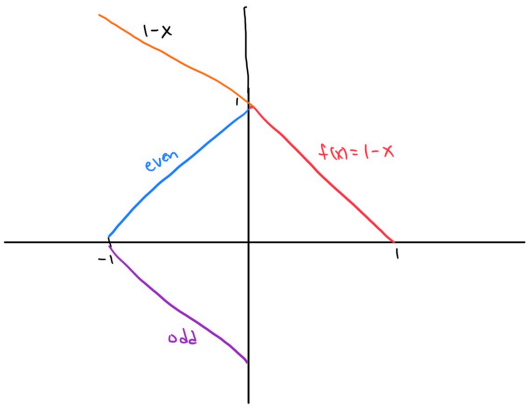

This new extension has Fourier series that looks like

$$1 + \sum_{n=1}^{\infty} \frac{2\cos n\pi}{n\pi} \sin n\pi x,$$

which is not a cosine series because of the sine terms, nor a sine series because of the constant term. This too will converge to $1 - x$ for $0 < x < 1$.

Heat equation. We now come back, finally, to discussing PDEs, beginning with the *heat equation*. Recall from the motivation we gave the first day of class that this is an equation characterizing the function which models the temperature at different points and times in a thin rod. To be precise, we assume our rod is of length L, and that the temperature of the circular cross section at some horizontal position x is constant, so that the temperature $u(x,t)$ only depends on that horizontal position x and time t. We also assume our rod has insulated circular sides, so that that heat can only escape or pass through the ends of the rod:

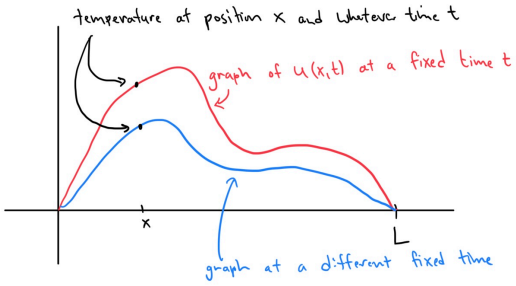

The heat equation states that at any position and time, the rate of change of the temperature with respect to time has to be proportional to the second derivative of the temperature with respect to position, so that $u(x,t)$ satisfies

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

(We will come back to where this equation comes from later.) Here, α^2 is a positive constant that depends on the material the rod is made out of.

Our goal is to derive the explicit temperature function $u(x,t)$ that satisfies this equation, subject to some boundary and initial conditions we will soon impose. But first, in order to understand something about the structure we can expect solutions to have, we note that the heat equation is an example of a *linear, homogeneous* PDE. The correct context from which to view this is linear algebra. Indeed, consider a vector space V of functions (we will restrict the types of functions we consider later), and consider the map $L : V \to V$ defined by

$$L = \frac{\partial}{\partial t} - \alpha^2 \frac{\partial^2}{\partial x^2}, \text{ so that } L(u) = \frac{\partial u}{\partial t} - \alpha^2 \frac{\partial^2 u}{\partial x^2}.$$

(So, we should assume our functions are such that at least the partial derivatives needed on the right exist.) Such an $L = \frac{\partial}{\partial t} - \alpha^2 \frac{\partial^2}{\partial x^2}$ is what's called a *partial differential operator*, since it is a mapping from functions to functions made up out of partial differentiation operations. To say that the heat equation is *linear* means that L is a linear transformation in the sense of linear algebra:

$$L(u_1 + u_2) = L(u_1) + L(u_2) \quad \text{and} \quad L(cu) = cL(u)$$

where u_1, u_2 are functions and c a scalar. (These properties just come from the linearity properties of derivatives.) The heat equation thus takes the form

$$L(u) = 0,$$

and the zero on the right is what makes the heat equation *homogeneous*. This in particular just says that u is in the *kernel* or *null space* of L (in the linear-algebraic sense), so a linear homogeneous PDE in general is one which arises from considering the kernel of a linear partial differential operator.

The fact that the heat equation describes the kernel of a linear map right away tells us that the set of solutions of the heat equation is closed under addition and scalar multiplication, so that it forms a *subspace* of our vector space of functions. (In other words, adding solutions of the heat equation produces other solutions, and scaling solutions produces other ones.) This will be important when we solve the heat equation, as producing new solutions from old ones in this way will be crucial.

Boundary and initial conditions. Now we impose some conditions on the heat solutions we want. First, we will require that the ends of the rod are held at zero temperature for all time:

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

Second, we will specify an initial temperature distribution along the entire rod, by saying that at time $t = 0$ the temperature at x is given by a function $f(x)$:

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

The goal is then to find a solution of the heat equation which satisfies these given boundary and initial conditions.

Will we succeed? In other words, do we have any kind of existence result? What about uniqueness? We will see in our work that, if we impose some mild assumptions on the initial data $f(x)$, say for example that it should be C^1, we will construct a valid solution satisfying the given boundary and initial conditions, which gives existence. In fact, the solution we construct will be the only one that can exist—so we do have uniqueness—but we will not able to prove uniqueness in full in this course. (Later we *will* prove uniqueness for what's called the *Laplace equation*, which might give *some* sense of how *some* uniqueness arguments work.)

Separation of variables. Let's get to work then and start trying to solve the heat equation. We make the assumption (for the time being) that the solution we seek can be written as

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

for some functions X of x and T of t. That is, we assume the the x and t dependences in u can be "separated" from one another. Such a solution is called a *separated solution*, and the method we use here is called *separation of variables*. We do not know at the outset that all solutions of the heat equation arise in this way, so this method will only seemingly produce *some* solutions.

The point of separation is that we can now derive from the heat equation PDE two separate ODEs which X and T must satisfy individually. Indeed, to say that $u(x,t) = X(x)T(t)$ satisfies the heat equation $u_t = \alpha^2 u_{xx}$ is to say that

$$XT' = \alpha^2 X''T$$

holds, which is obtained simply by plugging $u(x,t) = X(x)T(t)$ into the heat equation and computing derivatives. We'll assume that neither X nor T are zero, since otherwise $u(x,t) = X(x)T(t)$ is zero as well, which is not a very interesting solution. Thus we can rearrange the equation above to get

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

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

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

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:

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

These are now the ODEs for X and T we want, which we can write as

$$T' = \lambda \alpha^2 T \quad \text{and} \quad X'' - \lambda X = 0.$$

Solving these ODEs separately will give the possible functions X and T, and from these we can then form a separated solution $u(x,t) = X(x)T(t)$ of our heat equation.

Back to boundary conditions. Let us start with the equation X. Recalling some facts about second-order ODEs shows that the form the solutions of $X'' - \lambda X = 0$ take depends on whether λ is positive, zero, or negative. First we consider $\lambda > 0$. The general solution of $X'' - \lambda X = 0$ is then

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

(We are assuming that we know how to solve such equations already, from a previous course.) But we have more data to work with, namely the boundary conditions

$$u(x,0) = 0 = u(L,t)$$

we're imposing. For our separated solution $u(x,t) = X(x)T(t)$, these boundary conditions become simply boundary conditions $X(0) = 0 = X(L)$ for X, at least if we assume that T is nonzero. Thus what we're really doing is solving the following boundary value problem for X:

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

In the $\lambda > 0$ case above where $X = c_1 e^{\sqrt{\lambda} x} + c_2 e^{-\sqrt{\lambda} x}$, these boundary conditions become the requirements that

$$c_1 + c_2 = 0 \quad \text{and} \quad c_1 \sqrt{\lambda} - c_2 \sqrt{\lambda} = 0,$$

and some algebra shows that c_1 and c_2 must both be zero here. Thus $X = 0$ is the solution we get, which we ignore since we are looking for nonzero solutions.

If $\lambda = 0$, the general solution of $X'' - 0X = 0$ (or simply $X'' = 0$) 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 $X'' - \lambda X = 0$, or equivalently $X'' + k^2 X = 0$, 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 not fall into the case of $\lambda < 0$.) Thus, we get a nontrivial solution to our boundary value problem in the $\lambda < 0$ case when $\lambda = -(\frac{n\pi}{L})^2$, when the solution is

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

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

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

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

We will come back to find T and then our general solution next time, but hopefully you can already start to see why a Fourier series might arise.

Lecture 9: Boundary Conditions

Warm-Up. We derive the ODEs satisfied by the factors of separated solutions of the PDE

$$u_{xx} + u_{xt} + u_t = 0.$$

If $u(x,t) = X(x)T(t)$ satisfies this PDE, then

$$X''T + X'T' + XT' = 0.$$

We can rearrange terms to write this as

$$X''T = -(X' + X)T', \text{ so } \frac{X''}{X' + X} = -\frac{T'}{T}$$

if assume that certain things are nonzero. Since each side here depends on a different variables, both sides be constant (the same constant!), so that

$$\frac{X''}{X' + X} = \lambda = -\frac{T'}{T},$$

which results in the ODEs $X'' = \lambda(X' + X)$ and $T' = -\lambda T$.

Back to the heat equation. Recall that the separated solutions of the heat equation $u_t = \alpha^2 u_{xx}$ satisfy the ODEs

$$X'' - \lambda X = 0 \quad \text{and} \quad T' = \lambda \alpha^2 T.$$

The *homogeneous* boundary conditions $u(x,0) = 0 = u(L,t)$ (we call these homogeneous because they are set to zero) then become the conditions $X(0) = 0 = X(L)$ on X, and we saw last time that the resulting boundary value problem

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

has nontrivial solutions only for $\lambda_n = -\frac{n^2\pi^2}{L^2}$, where n is a positive integer, with a nonzero solution being $X_n = \sin\frac{n\pi x}{L}$. (Other solutions are scalar multiples of these.)

Now, for the λ we have thus determined, we can consider the corresponding ODE for T:

$$T' = \lambda_n \alpha^2 T = -\frac{\alpha^2 n^2 \pi^2}{L^2} T.$$

This has nonzero solution given by

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

which we have again indexed by n. The solutions for X and T we have found are called *fundamental solutions*, and corresponding to each we thus get a "fundamental solution" to the heat equation satisfying the given boundary conditions:

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

Since the heat equation is linear and homogeneous, and since the boundary conditions are homogeneous, linear combinations of solutions still produce solutions. So, we can take an "infinite linear combination" of these fundamental solutions to get a more general solution:

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

There are convergence issues to mention, which we will come back to, so so far this is only a "formal" solution of the heat equation. Note that naively taking a derivative with respect to t and comparing to the second derivative with respect to x gives what we expect: the former gives an extra coefficient of $-\frac{\alpha^2 n^2 \pi^2}{L^2}$, which is precisely α^2 times the coefficient obtained from the latter, exactly as the heat equation specifies.

Initial condition. The formal solution above satisfies the boundary conditions $u(0,t) = 0 = u(L,t)$, so now we consider our initial condition $u(x,0) = f(x)$. In order for this to be satisfied, the following must hold:

$$f(x) = u(x,0) = \sum_{n=1}^{\infty} c_n e^{-\frac{\alpha^2 n^2 \pi^2 0}{L^2}} \sin\frac{n\pi x}{L} = \sum_{n=1}^{\infty} c_n \sin\frac{n\pi x}{L},$$

where the exponential terms are all just $e^0 = 1$. But this equation says precisely that the series on the right should be the Fourier sine series of f (recall that initially f is only defined along the rod for $0 < x < L$), so that the unknown coefficients must be given by

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

Thus, the final conclusion is that

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

where

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

is a (formal) solution of the heat equation $u_t = \alpha^2 u_{xx}$ satisfying the homogeneous boundary conditions $u(0,t) = 0 = u(L,t)$ and the initial condition $u(x,0) = f(x)$. Huzzah!

In order to remove the word "formal" throughout the discussion above requires knowing that the series above defines an *actual* function, meaning that it converges. For the initial condition

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

this requires f to be C^1, or even just piecewise C^1 (where we don't quite get the series converging exactly to $f(x)$ at all points—only those where f is continuous—but this will not really affect the "heat" interpretation anyway, so we ignore such minor details), and in fact if this is true, it turns out that the full solution

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

converges as well. (We will talk about this later, but essentially the exponential term decreases rapidly enough as n increases to force convergence. Also, there are some things to say about whether the convergence is uniform or not, as we'll see.)

Example. Consider the heat equation $u_t = u_{xx}$ (so $\alpha^2 = 1$) with boundary conditions $u(0,t) = 0 = u(1,t)$ and initial condition $u(x,0) = 1 - x$. The solution is

$$u(x,t) = \sum_{n=1}^{\infty} c_n e^{-n^2\pi^2 t} \sin(n\pi x)$$

where the coefficients are the Fourier sine coefficients of $u(x,0) = 1 - x$. We computed these in a previous example to be

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

so the solution of this specific heat problem is

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

Now, consider what happens to this temperature as time goes on. In the limit $t \to \infty$, we have that $e^{-n^2\pi^2 t} \to 0$, and this in fact will force the entire series to approach 0, at any x:

$$\lim_{t\to\infty} u(x,t) = 0.$$

This makes sense physically: as time goes on, heat escapes through the ends of the rod with no new heat being introduced, so the temperature should decrease more and more, eventually "zeroing out" at infinity. We can see this happening if we plot a few instances of $u(x,t)$ at fixed values of t: at $t = 0$ we have our initial temperature $u(x,0) = 1 - x$; at $t = 1$ we have a temperature that has decreased at all x; at $t = 2$ is has decreased more; and so on as we increase t

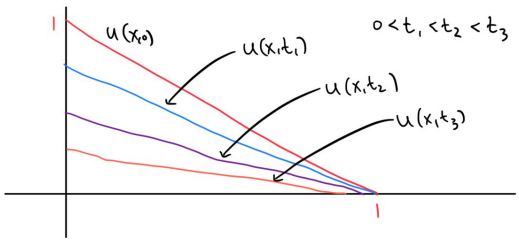

This phenomena occurs in general with homogeneous boundary conditions $u(0,t) = 0 = u(L,t)$, since in general it is true that heat will escape at all times, so that the temperature gets smaller and smaller. Mathematically, this occurs because of the exponential term (with negative exponent) in the general heat solution.

Insulated ends. Let us now consider a different set of boundary conditions, namely boundary conditions describing the case of insulated ends. Here, we imagine that no heat passes through the ends of the rod, which we can describe by the conditions

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

Note that this does not force the temperature at the ends to necessarily be zero, only that the rate at which the temperature changes at the ends with respect to *position* is zero. The analysis here is for the most part the same as for the boundary conditions we considered previously, with the only difference coming in the boundary value problem for X, which now looks like

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

As in the previous case, for $\lambda > 0$ it turns out that there are no nontrivial solutions: we get $X = c_1 e^{\sqrt{\lambda}x} + c_2 e^{-\sqrt{\lambda}x}$, and requiring $X'(0) = 0 = X'(L)$ leads to $c_1 = c_2 = 0$. For $\lambda = 0$, however, our solution (to $X'' = 0$) looks like

$$X = c_1 + c_2 x,$$

and lo-and-behold we *do* get a nonzero solution now: $X'(0) = 0 = X'(L)$ both force $c_2 = 0$, but there are no restrictions on c_1. Thus we can take $X_0 = 1$ to be a constant fundamental solution for $\lambda_0 = 0$. For $\lambda = -k^2 < 0$, we again get

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

Here $X'(0) = 0$ forces $c_2 = 0$—so only the cosine term remains—and $X'(L) = 0$ requires that

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

Hence we get a nonzero solution when $kL = n\pi$ for a positive integer n, which gives $X_n = \cos \frac{n\pi x}{L}$ as the fundamental solution for $\lambda_n = -\frac{n^2\pi^2}{L^2}$.

As for T, $\lambda_0 = 0$ gives $T' = 0T$, so that T is constant and hence we can take $T_0 = e^0 = 1$ to be a fundamental solution. For $\lambda_n = -\frac{n^2\pi^2}{L^2}$ we get the same exponential fundamental solution we had for the previous boundary conditions, so overall we get fundamental solutions of the form

$$u_0(x,t) = T_0(t)X_0(x) = 1 \quad \text{and} \quad u_n(x,t) = T_n(t)X_n(x) = e^{-\frac{n^2\pi^2 t}{L^2}} \cos \frac{n\pi x}{L}.$$

The general solution is thus

$$u(x,t) = \frac{c_0}{2}u_0(x,t) + \sum_{n=1}^{\infty} u_n(x,t) = \frac{c_0}{2} + \sum_{n=1}^{\infty} c_n e^{-\frac{n^2\pi^2 t}{L^2}} \cos \frac{n\pi x}{L}.$$

(I took the constant term to be $\frac{c_0}{2}$ instead of just c_0 to match because of the what happens when we impose our initial condition.) Imposing an initial condition $u(x,0) = f(x)$ (let's assume that f is piecewise C^1) gives

$$f(x) = u(x,0) = \frac{c_0}{2} + \sum_{n=1}^{\infty} c_n \cos\frac{n\pi x}{L},$$

so that now the coefficients c_n must be the Fourier cosine coefficients of f. Thus, the solution to the heat equation $u_t = \alpha^2 u_{xx}$ with insulated boundary conditions $u_x(0,t) = 0 = u_x(L,t)$ and piecewise C^1 initial condition $u(x,0) = f(x)$ is

$$u(x,t) = \frac{c_0}{2} + \sum_{n=1}^{\infty} c_n e^{-\frac{n^2\pi^2 t}{L^2}} \cos\tfrac{n\pi x}{L}$$

where

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

The upshot is that the form the solution takes is determined by the boundary conditions.

Steady-states. In the insulated ends case, we get a different limit as $t \to \infty$ than in the original case:

$$\lim_{t\to\infty} u(x,t) = \frac{c_0}{2}.$$

(Again, the exponential terms force everything else to approach 0.) Thus, in the limit the temperature "evens out" to a nonzero constant. This makes sense physically: if no heat is escaping our rod, then the total heat in the rod remains the same throughout but it it bounces around from pieces of high temperature to pieces of low temperature until it settles down. (This constant can also be interpreted as the average temperature throughout the rod at any time.)

Moreover, if our initial temperature had been this exact constant $u(x,0) = \frac{c_0}{2}$, then we get $u(x,t) = \frac{c_0}{2}$, so that the temperature never changes. Thus, this constant temperature is a "steady-state" solution of the heat equation, which means a solution that does not change in time. Hence, steady-state solutions are ones which satisfy $u_t = 0$, so that the heat equation becomes simply $u_{xx} = 0$, which means that the potential steady-state solutions are of the form $u(x,t) = A + Bx$ (recall no dependence on t). In the case of insulated ends, the boundary conditions $u_x(0,t) = 0 = u_x(L,t)$ force $B = 0$, so the steady-state solution $u(x,t) = A$ is indeed constant. We can interpret the general solution

$$u(x,t) = \frac{c_0}{2} + \sum_{n=1}^{\infty} c_n e^{-\frac{n^2\pi^2 t}{L^2}} \cos\tfrac{n\pi x}{L}$$

as the sum of the steady-state one and a "transient" term that characterizes the dependence on t.

Lecture 10: More on Heat Solutions

Warm-Up. Consider the heat equation with boundary conditions

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

so that the left end of the rod is held at temperature zero while the right end is insulated. We determine the form a series solution will take. As with previous examples, the key point is the resulting boundary value problem for X:

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

For $\lambda > 0$, we have $X = c_1 e^{\sqrt{\lambda}x} + c_2 e^{-\sqrt{\lambda}x}$, and the boundary conditions force $c_1 = c_2 = 0$, so that's no good. For $\lambda = 0$ we have $X = c_1 + c_2 x$, and again the boundary conditions force $c_1 = c_2 = 0$.

For $\lambda = -k^2 < 0$, we have $X = c_1 \cos kx + c_2 \sin kx$. The first boundary condition gives $c_1 = 0$, and the second then gives

$$c_2 k \cos(kL) = 0.$$

We thus get a nonzero solution when kL is an odd multiple of $\frac{\pi}{2}$, so that $k = \frac{(2n-1)\pi}{2L}$ for a positive integer n. Hence the given boundary value problem for X has a nonzero solution

$$X_n = \sin \frac{(2n-1)\pi x}{2L} \text{ for } \lambda_n = -\frac{(2n-1)^2\pi^2}{4L^2}.$$

The corresponding solutions for T are $T_n = e^{-\frac{(2n-1)^2\pi^2 t}{4L^2}}$, so the fundamental solutions for u are

$$u_n(x,t) = e^{-\frac{(2n-1)^2\pi^2 t}{4L^2}} \sin \tfrac{(2n-1)\pi x}{2L}.$$

Hence the general solution is

$$u(x,t) = \sum_{n=1}^{\infty} c_n e^{-\frac{(2n-1)^2\pi^2 t}{4L^2}} \sin \tfrac{(2n-1)\pi x}{2L}.$$

If we then impose an initial condition $u(x,0) = f(x)$, we get the requirement that

$$f(x) = \sum_{n=1}^{\infty} c_n \sin \tfrac{(2n-1)\pi x}{2L}.$$

This type of Fourier series showed up on the homework, where the coefficients c_n are given by

$$c_n = \frac{2}{L} \int_0^L f(x) \sin \tfrac{(2n-1)\pi x}{2L}\, dx.$$

It arises from extending f into $L < x < 2L$ for requiring $f(2L - x) = f(x)$, and then extending as an odd function into $-2L < x < 0$. The moral is that these different types of specific forms of Fourier series arise from various types of boundary considerations.

Other boundary conditions. There are plenty of other types of boundary conditions we can impose. For example, we can consider non-homogeneous boundary conditions like

$$u(0,t) = T_1 \quad \text{and} \quad u(L,t) = T_2,$$

where the ends of the rod are held at nonzero temperature. Here a solution can be found which is of the form

$$u(x,t) = v(x) + w(x,t)$$

where $v(x)$ is steady-state solution and $w(x,t)$ a solution satisfying the homogeneous boundary conditions $u(0,t) = 0 = u(L,t)$. (Essentially, the non-homogeneous boundary conditions can "absorbed" by the stead-state term by requiring $V(0) = T_1, V(L) = T_2$.) We will look at an example along these lines next time, and you can also check the book for the general approach.

More general boundary conditions are found by assuming that the rate at which the temperature changes with respect to position at the end should be proportional to the temperature itself:

$$u_x(0,t) = (\text{scalar})u(0,t) \quad \text{and} \quad u_x(L,t) = (\text{scalar})u(L,t).$$

We will consider such boundary conditions later on from the perspective of Sturm-Liouville theory. We can also consider the effect of an external heat or cooling source (so, not technically a boundary condition), which affects the value of u_t:

$$u_t = \alpha^2 u_{xx} + (\text{something else}).$$

We will look at an example like this next time.

Behavior of solutions. Let us now say a bit more about the heat solutions we are finding. Much of the behavior and structure of solutions comes from the fact that the heat equation is an example of a *parabolic* PDE. (In fact, the heat equation is the prototypical model of a parabolic PDE.) Here's the definition: a second-order linear homogeneous PDE with constant coefficients

$$Au_{xx} + Bu_{xt} + Cu_{tt} + Du_x + Eu_t = 0$$

is *parabolic* if $B^2 - 4AC = 0$. The name comes from the fact that this property characterizes when the corresponding quadratic equation

$$Ax^2 + Bxy + Cy^2 + Dx + Ey = 0$$

defines a parabola in $\mathbb{R}^2$.

Parabolic PDEs in general have the following two key properties:

- there are basic existence and uniqueness results under basic boundary and initial condition assumptions, and
- solutions exhibit a type of "smoothing" behavior.

The second property essentially means that although the initial condition $u(x,0) = f(x)$ might only be piecewise C^1 with possibly many discontinuities, these discontinuities disappear in time, in fact in a short amount of time. The flow of heat from high to low temperature areas thus causes the overall heat distribution to "smooth out".

For the first property, let us now come back to elaborate on our method for finding solutions via separation. As we pointed out at the start of our whole heat discussion, we do not know at the outset that separation of variables will find all solutions. But in fact after everything we've done we now know that this method does in fact find all solutions subject to homogeneous boundary conditions and piecewise C^1 initial data: in this setting, our construction for sure produces *a* solution, and so if we take the uniqueness property of parabolic PDEs for granted, this solution we found must be the only one. (We will not be able to prove this uniqueness result here in full generality, but will note that among solutions obtained via operation, uniqueness follows from the uniqueness of Fourier series coefficients.)

But there is still one lingering issue, which is how we know that the solution we've constructed is in fact a solution; that is, why does

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

actually satisfy the heat equation $u_t = \alpha^2 u_{xx}$? Naively we can simply try to compute u_t and u_{xx}:

$$u_t = \sum_{n=1}^{\infty} -\frac{\alpha^2 n^2 \pi^2}{L^2} c_n e^{-\frac{\alpha^2 n^2 \pi^2 t}{L^2}} \sin \frac{n\pi x}{L} \text{ and } u_{xx} = \sum_{n=1}^{\infty} -\frac{n^2 \pi^2}{L^2} c_n e^{-\frac{\alpha^2 n^2 \pi^2 t}{L^2}} \sin \frac{n\pi x}{L}$$

and plug them in. However, and this is the subtle point, this computation depends on being able to compute the derivative of a series by exchanging differentiation and summation; for example,

$$\frac{\partial}{\partial x}\left(\sum_{n=1}^{\infty} c_n e^{-\frac{\alpha^2 n^2 \pi^2 t}{L^2}} \sin\tfrac{n\pi x}{L}\right) = \sum_{n=1}^{\infty} c_n e^{-\frac{\alpha^2 n^2 \pi^2 t}{L^2}} \frac{\partial}{\partial x}\left(\sin\tfrac{n\pi x}{L}\right).$$

We have previously spoken about how this requires uniform convergence of our original series, but actually a bit more is required as well: we need to know that the series obtained after differentiating in this way converges uniformly as well! The same applies to the series obtained by differentiating with respect to x once more. In all of these cases, it is the presence of the exponential term that guarantees everything converges uniformly and so everything works out nicely. Yes it is then possible to verify that our proposed solution is in fact a solution (so no longer merely a "formal" one) by differentiating term-by-term and seeing that the heat equation is satisfied. (This will not be the case with the wave equation, however, and we will need a different approach in order to verify that our proposed solution there is actually a solution.)

Heat equation elsewhere. The heat equation has applications beyond simply modeling temperature distributions. Here are two examples of other ways in which it shows up. First, in finance the *Black-Scholes equation* models the pricing of what are called options over time. If V denotes the pricing function, the Black-Scholes equation states that

$$\frac{\partial V}{\partial t} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 V}{\partial S^2} + rS\frac{\partial V}{\partial S} - rV = 0.$$

(I'll ignore here what all the various symbols mean.) Solving the Black-Scholes equation then gives the pricing function. The point is that, after some suitable (multiple) changes of variables, the Black-Scholes equation becomes precisely the heat equation! If we accept that options pricing should be a "diffusive" process, then it makes some sense that the behavior of such prices should exhibit some "heat-like" qualities.

Second, the notion of *Ricci flow* in geometry (and possibly general relativity) also depends on a version of the heat equation. Here we consider a geometric object which is initially curved in a possibly random way:

The Ricci flow equation describes the evolution of the curvature of this surface, and says that in time the curvature smooths out until we are left with a much simpler smoother surface:

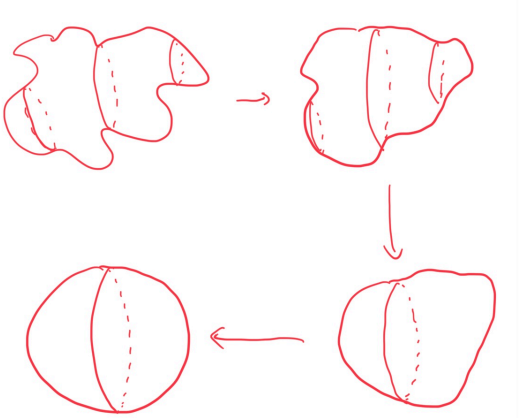

The point is that although the original surface might have been difficult to identity directly, after it "smooths out" we'll have an easier time of doing so. (The notion Ricci flow was used heavily in the proof of what's called the *Poincaré conjecture*, which gives a classification of certain 3-dimensional geometric objects. It's proof back in the mid 2000's was a big deal at the time.)

In general, given any type of "diffusive" phenomena, where extreme behaviors in one location are "smoothed out" towards less extreme behaviors elsewhere, there is a decent chance that there is a version of the heat equation underlying those phenomena.

Two dimensions. The heat equation also shows up in higher-dimensions, say in describing the distribution of heat in a two-dimensional region, or in a three-dimensional region. The two-dimensional heat equation is

$$\frac{\partial u}{\partial t} = \alpha^2 \left(\frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2} \right),$$

where $u(x, y, t)$ now depends on two space parameters. (In higher dimensions we simply add another second-order derivative term on the right for each new spatial variable.) We will not solve this equation in full (although you'll see it a bit in discussion), but we will solve a special case of it soon enough when we consider the Laplace equation.

The basic idea is the same: assume your solution is separated, and derive some ODEs. Here we have $u(x, y, t) = X(x)Y(y)T(t)$, so we get

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

Separating gives

$$\frac{T'}{\alpha^2 T} = \frac{X''Y + XY''}{XY},$$

so that

$$T' = \lambda\alpha^2 T \quad \text{and} \quad X''Y + XY'' = \lambda XY$$

for some constant λ. We can separate again in the second equation to get

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

so that

$$X'' = \mu X \quad \text{and} \quad Y'' = (\lambda - \mu)Y$$

for some other constant μ. Then we can solve after imposing some boundary conditions, such as boundary conditions on a rectangle for example. Or we can consider boundary conditions on, say,

a disk, but it turns out here that the separated form above is not so useful, and that it is better to consider the heat equation in terms of *polar* coordinates and then use separation there. We'll see glimpses of these things in discussion and later. Good stuff!

Lecture 11: Wave Equation

Warm-Up. We solve the non-homogeneous heat equation

$$u_t = u_{xx} + 2,$$

with boundary conditions $u(0,t) = 1, u(1,t) = 2$ and initial condition $u(x,0) = x, 0 < x \leq 1$. This equation characterizes a rod with an external heat source applied, increasing the temperature by a constant rate of 2 at every point and time. Let us suppose that the solution we seek can be written as

$$u(x,t) = v(x) + w(x,t),$$

where $v(x)$ is a steady-state solution which is independent of time (so that $\frac{\partial v}{\partial t} = 0$ and hence $0 = v'' + 2$) and satisfying the boundary conditions $v(0) = 1, v(1) = 2$. Note that then we have

$$w_t = u_t - 0 = u_{xx} + 2 = (v'' + w_{xx}) + 2 = w_{xx}$$

and

$$w(0,t) = u(0,t) - v(0) = 1 - 1 = 0 \quad \text{and} \quad w(1,t) = u(1,t) - v(1) = 2 - 2 = 0,$$

so that $w(x,t)$ satisfies our usual homogeneous heat equation withe usual homogeneous boundary conditions. Essentially, we are absorbing the external heat term and the non-homogeneous boundary conditions $u(0,t) = 1, u(1,t) = 2$ into the steady-state term.

The heat equation for the steady-state solution is $0 = v'' + 2$, so v looks like

$$v = -x^2 + c_2 x + c_1.$$

The boundary conditions on v then give

$$1 = c_1 \quad \text{and} \quad 2 = -1^2 + c_2 + c_1,$$

so that $c_2 = 2$. Thus the steady-state solution is

$$v = -x^2 + 2x + 1.$$

The remaining term $w(x,t)$ has the form

$$w(x,t) = \sum_{n=1}^{\infty} a_n e^{-n^2\pi^2 t} \sin(n\pi x)$$

since w satisfies the homogeneous heat equation with usual homogeneous boundary conditions. Since $w(x,t) = u(x,t) - v(x)$, the initial condition for w is

$$w(x,0) = u(x,0) - v(x) = x - (-x^2 + 2x + 1) = x^2 - x - 1.$$

Thus the coefficients a_n are the Fourier sine coefficients of $x^2 - x - 1$:

$$a_n = 2\int_0^1 (x^2 - x - 1)\sin(n\pi x)\, dx = \frac{(2\pi^2 n^2 + 4)(\cos \pi n - 1)}{\pi^3 n^3}.$$

The full solution to the given non-homogeneous heat boundary value problem is thus

$$\begin{aligned} u(x,t) &= v(x) + w(x,t) \\ &= -x^2 + 2x + 1 + \sum_{n=1}^{\infty} \frac{(2\pi^2 n^2 + 4)(\cos \pi n - 1)}{\pi^3 n^3} e^{-n^2\pi^2 t} \sin(n\pi x). \end{aligned}$$

Derivation of heat equation. ***TO BE FINISHED***

Wave equation. Consider an elastic vibrating string. We assume that each point on the string can only move vertically, so that the vertical displacement (i.e., height) of a point on the string is given by a function $u(x,t)$ depending only on time t and the horizontal position x of the point:

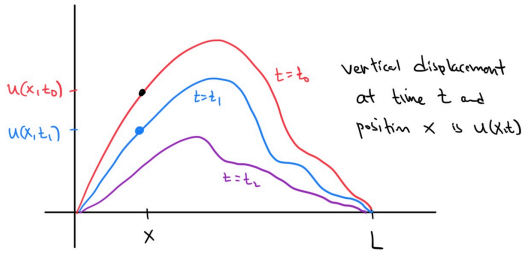

The *wave equation* says that, under only the action of tension with no other external force, the displacement function $u(x,t)$ evolves according to

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

where $a^2 > 0$ is a constant depending only on the string and what it is made of. We will talk about the derivation of this equation from Newton's laws next time. Like the heat equation, the wave equation is an example of a linear homogeneous PDE: here the relevant linear partial differential operator is

$$L = \frac{\partial^2}{\partial t^2} - a^2 \frac{\partial^2}{\partial x^2},$$

and the wave equation says that $L(u) = 0$. Thus, being the kernel of a linear operator, we know that the space of solutions is closed under addition and scalar multiplication.

As one might expect after having seen the heat equation, we can find solutions to the wave equation via Fourier series methods. The presence of such series comes, as usual, from boundary conditions. In this case, we impose the following basic boundary conditions:

$$u(0,t) = 0 = u(L,t),$$

which say that the ends of the string are tied down throughout all time. Later we will also consider boundary conditions which say that the partial derivative u_x should be 0 at an end, which allow for one or both ends to move freely.

We will also impose initial conditions, where the first $u(x,0) = f(x)$ gives the starting position of string. Since the wave equation is second-order in time (meaning it involves a second derivative with respect to time), to single-out a solution requires that we also specify the first derivative with respect to time, so we impose $u_t(x,0) = g(x)$, which gives the starting velocity at any point along the string. The goal is then to find the function $u(x,t)$ satisfying the wave equation and the given boundary and initial conditions.

Fundamental and general solutions. Using separation of variables, for a separated solution $u(x,t) = X(x)T(t)$, we get the following ODEs:

$$T'' - \lambda a^2 T = 0, \qquad X'' - \lambda X = 0.$$

The boundary conditions $u(0,t) = 0 = u(L,t)$ become $X(0) = 0 = X(L)$, and the resulting boundary value problem for X

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

is precisely the same one we had in the heat equation. Thus we get the same fundamental solutions as before:

$$X_n = \sin\left(\frac{n\pi x}{L}\right) \quad \text{for} \quad \lambda_n = -\frac{n^2\pi^2}{L}$$

where n is a positive integer.

The ODE for T then becomes

$$T'' + \frac{a^2 n^2 \pi^2}{L} T = 0.$$

Rather than the exponential solutions we got for the heat equation, here we get trigonometric solutions:

$$T_n = c_n \cos\left(\frac{an\pi t}{L}\right) + d_n \sin\left(\frac{an\pi t}{L}\right)$$

where c_n, d_n are scalars. After multiplying this by X_n, we see that we get fundamental solutions for $u_n(x,t)$ which consist of terms involving

$$\cos\left(\frac{an\pi t}{L}\right)\sin\left(\frac{n\pi x}{L}\right) \quad \text{and} \quad \sin\left(\frac{an\pi t}{L}\right)\sin\left(\frac{n\pi x}{L}\right),$$

and so the general (formal) solution is

$$u(x,t) = \sum_{n=1}^{\infty} \left[c_n \cos\left(\frac{an\pi t}{L}\right) + d_n \sin\left(\frac{an\pi t}{L}\right)\right] \sin\left(\frac{n\pi x}{L}\right).$$

Now we impose initial conditions. For $u(x,0) = f(x)$ we get the requirement that

$$f(x) = u(x,0) = \sum_{n=1}^{\infty} c_n \sin\left(\frac{n\pi x}{L}\right),$$

so the c_n should be the Fourier sine coefficients of f. For $u_t(x,0) = g(x)$, we get

$$g(x) = u_t(x,0) = \sum_{n=1}^{\infty} d_n \frac{an\pi}{L} \sin\left(\frac{n\pi x}{L}\right),$$

so $d_n \frac{an\pi}{L}$ should give the Fourier sine coefficients of the initial velocity $g(x)$, meaning that d_n should be this Fourier sine coefficient *times* $\frac{L}{an\pi}$. Thus, the solution of the wave equation

$$u_{tt} = a^2 u_{xx}$$

satisfying

$$u(0,t) = 0 = u(L,t), \ u(x,0) = f(x), \text{ and } u_t(x,0) = g(x)$$

is

$$u(x,t) = \sum_{n=1}^{\infty} \left[c_n \cos\left(\frac{an\pi t}{L}\right) + d_n \sin\left(\frac{an\pi t}{L}\right) \right] \sin\left(\frac{n\pi x}{L}\right)$$

where

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

and

$$d_n = \frac{L}{an\pi} \cdot \frac{2}{L} \int_0^L g(x) \sin\left(\frac{n\pi x}{L}\right)\, dx = \frac{2}{an\pi} \int_0^L g(x) \sin\left(\frac{n\pi x}{L}\right)\, dx$$

Example. Consider the string on $[0, \pi]$ with initial position

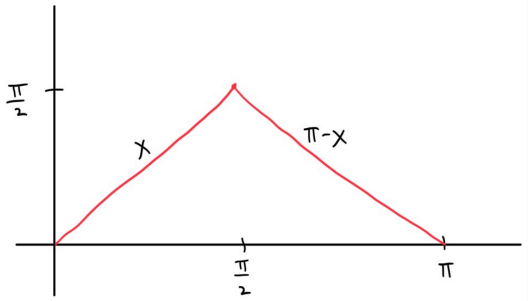

so $f(x) = x$ for $0 \le x < \frac{\pi}{2}$ and $f(x) = \pi - x$ for $\frac{\pi}{2} \le x \le \pi$, and with initial velocity zero, evolving according to $u_{tt} = u_{xx}$. (So $a^2 = 1$ in this case.) The Fourier sine coefficients of f are

$$c_n = \frac{2}{\pi} \int_0^{\pi} f(x) \sin(nx)\, dx = \frac{2}{\pi} \left[\int_0^{\pi/2} x \sin(nx)\, dx + \int_{\pi/2}^{\pi} (\pi - x) \sin(nx)\, dx \right] = \frac{4 \sin(\frac{n\pi}{2})}{\pi n^2},$$

and the Fourier sine coefficients $u_t(x, 0) = 0$ are all zero. Hence the displacement of the string is given by

$$u(x,t) = \sum_{n=1}^{\infty} \frac{4 \sin(\frac{n\pi}{2})}{\pi n^2} \cos(nt) \sin(nx).$$

Note what happens as t increases—here are some plots for $t = 0, 0.5, 1, 1.5, 2$:

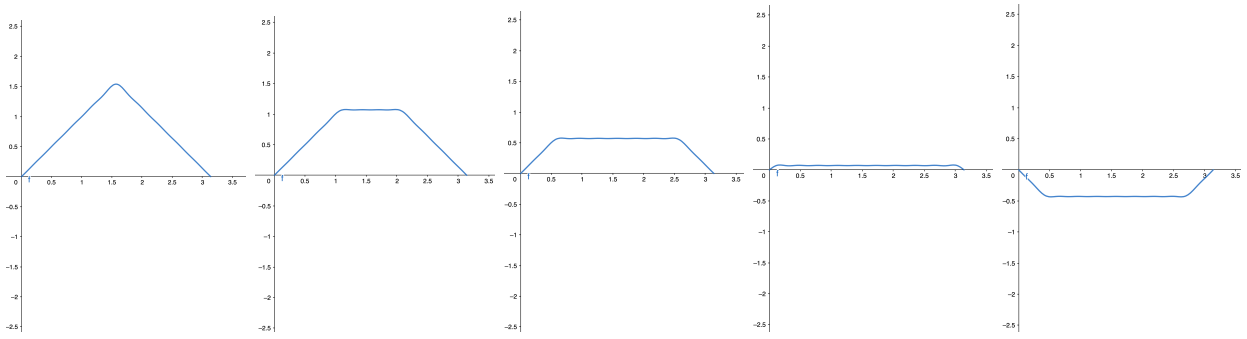

This exhibits the "wave-like" behavior we expect: the initial graph begins to move down as t increases, eventually dipping below the x-axis. It will begin to do so until it reaches its minimal (i.e. most negative) displacement at $t = \pi$ (the graph will be the mirror image of the initial one), after which the graphs begins to move back up, reaching the initial displacement again at $t = 2\pi$. The point is that the wave equation truly does describe wave-like behavior!

Lecture 12: More on Waves

Warm-Up. We solve the wave equation $u_t = 4u_{xx}$ with boundary conditions

$$u(0,t) = 0 = u(1,t)$$

and initial conditions $u(x,0) = 10,\ u_t(x,0) = x^2$. The solution is

$$u(x,t) = \sum_{n=1}^{\infty} \left[c_n \cos(2n\pi t) + d_n \sin(2n\pi t)\right] \sin(n\pi x)$$

where the c_n are the Fourier sine coefficients of $u(x,0) = 10$ and $2n\pi d_n$ are the Fourier sine coefficients of $u_t(x,0) = x^2$. Thus we have

$$c_n = 2\int_0^1 10\sin(n\pi x)\,dx = \frac{20(1-\cos n\pi)}{n\pi}$$
$$d_n = \frac{1}{n\pi}\int_0^1 x^2\sin(n\pi x)\,dx = \frac{(2-\pi^2n^2)\cos n\pi - 2}{\pi^4 n^4}.$$

The solution we want is

$$u(x,t) = \sum_{n=1}^{\infty}\left[\frac{20(1-\cos n\pi)}{n\pi}\cos(2n\pi t) + \frac{(2-\pi^2n^2)\cos n\pi - 2}{\pi^4 n^4}\sin(2n\pi t)\right]\sin(n\pi x).$$

But note that care must be taken in order to know that this solution is indeed a solution, in the sense that it gives a *well-defined* function. In other words, how do we know that this series in fact converges? If it doesn't converge it would make no sense to try to construct a candidate solution from it, so this is an important issue. In particular, note that taking two partial derivatives with respect to x (as required in the wave equation) will produce an n^2 term, and it is not obvious that a series with terms like

$$\sum_{n=1}^{\infty} n\frac{20[1-\cos n\pi]}{\pi}\cos(2n\pi t)\sin(n\pi x)$$

(I'm omitting the second piece for now) will converge. This was not really an issue with the heat equation because there we had an exponential term in our solution that decayed rapidly enough to balance out any other term like "n" that grows, but there is no such exponential term here. We will come back to this issue in a bit.

Visualizing velocity. ***TO BE FINISHED***

Derivation of the wave equation. ***TO BE FINISHED***

Convergence issues. ***TO BE FINISHED***

Hyperbolic PDEs. ***TO BE FINISHED***

Other boundary conditions. Other typical boundary conditions come from allowing one or both ends of the string to move freely. For example, with

$$u(0,t) = 0,\ u_x(L,t) = 0,$$

we hold the end on the left always at zero but the right end can move. (This is a wave analog of an "insulated end".) This gives boundary conditions $X(0) = 0 = X'(L)$ for a separated solution, which results in fundamental solutions

$$X_n = \sin\left(\frac{(2n-1)\pi x}{2L}\right) \quad \text{for} \quad \lambda_n = -\frac{(2n-1)^2\pi^2}{4L^2}.$$

The case where both ends can move freely:

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

gives $X'(0) = 0 = X'(L)$, which results in fundamental solutions

$$X_0 = 1 \text{ for } \lambda_0 = 0, \text{ and } X_n = \cos\left(\frac{n\pi x}{L}\right) \text{ for } \lambda = -\frac{n^2\pi^2}{L^2}.$$

In this case initial conditions $u(x,0) = f(x),\ u_t(x,0) = g(x)$ require the use of Fourier cosine series when deriving the unknown coefficients.

2D wave equation. ***TO BE FINISHED***

Lecture 13: Laplace Equation

Warm-Up. We show that the function

$$u(x,t) = \sum_{n=1}^{\infty} a_n \cos\left(\frac{an\pi t}{L}\right)\sin\left(\frac{n\pi x}{L}\right),$$

which describes the solution to the wave equation $u_{tt} = a^2 u_{xx}$ with boundary conditions $u(0,t) = 0 = u(L,t)$ and initial conditions $u(x,0) = f(x)$, $u_t(x,0) = 0$, is $\frac{1}{2}(h(x-at) + h(x+at))$ where h is the odd extension of the initial condition $u(x,0) = f(x)$. This reflects a general fact about the structure of wave solutions which we'll clarify afterwards.

The odd extension h has Fourier series

$$h(x) = \sum_{n=1}^{\infty} a_n \sin\left(\frac{n\pi x}{L}\right)$$

where the a_n are the Fourier sine coefficients of f. Using the angle-addition formula for sine:

$$\sin(A+B) = \cos A \sin B + \cos B \sin A,$$

we get that

$$\begin{aligned} h(x-at) &= \sum_{n=1}^{\infty} a_n \sin\left(\frac{n\pi(x-at)}{L}\right) \\ &= \sum_{n=1}^{\infty} a_n \cos\left(\frac{n\pi x}{L}\right)\sin\left(-\frac{\pi nat}{L}\right) + \sum_{n=1}^{\infty} a_n \cos\left(-\frac{n\pi at}{L}\right)\sin\left(\frac{n\pi x}{L}\right) \end{aligned}$$

and

$$h(x+at) = \sum_{n=1}^{\infty} a_n \sin\left(\frac{n\pi(x+at)}{L}\right)$$

$$= \sum_{n=1}^{\infty} a_n \cos\left(\frac{n\pi x}{L}\right) \sin\left(\frac{\pi nat}{L}\right) + \sum_{n=1}^{\infty} a_n \cos\left(\frac{n\pi at}{L}\right) \sin\left(\frac{n\pi x}{L}\right).$$

Since cosine is even and sine is odd, when adding $h(x-at)$ and $h(x+at)$ we see that the terms with $\sin\left(\frac{\pi nat}{L}\right)$ cancel out, so we are left with

$$h(x-at) + h(x+at) = 2\sum_{n=1}^{\infty} a_n \cos\left(\frac{n\pi at}{L}\right) \sin\left(\frac{n\pi x}{L}\right).$$

Thus after dividing by 2 we get the desired claim:

$$u(x,t) = \frac{1}{2}(h(x-at) + h(x+at)).$$

Superpositions and standing waves. Now, we can interpret this result as follows. First, note that $h(x-at)$ has the same graph as $h(x)$, only *shifted* to the right, with the "shift" increasing as t does. Similarly, $h(x+at)$ is the graph of h shifted to the left, more and more as t increase. Thus, if we interpret $h(x-at)$ as describing a *rightward* moving wave and $h(x+at)$ a *leftward* moving wave, the result

$$u(x,t) = \frac{1}{2}(h(x-at) + h(x-at))$$

in the case of no initial velocity says that any wave is a superposition of a rightward moving wave with a leftward moving wave. The constant a then describes the speed of the moving wave.

If we go back to the form

$$u(x,t) = \sum_{n=1}^{\infty} a_n \cos\left(\frac{an\pi t}{L}\right) \sin\left(\frac{n\pi x}{L}\right),$$

we can also give an interpretation to each summand. The function $\sin(\frac{n\pi x}{L})$ describes a *standing wave*, since it does not depend on t. The $\frac{n\pi}{L}$ constant is then the *frequency* of this standing wave and is inversely proportional to the period $2L$. (The frequency tells you how many "peaks" there are within a given interval.) The coefficient $\cos(\frac{n\pi at}{L})$ describes an *amplitude* of the wave, so that as time increases our standing wave remains "standing" only with varying amplitude:

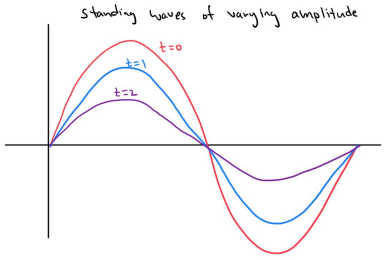

The general solution

$$u(x,t) = \sum_{n=1}^{\infty} a_n \cos\left(\frac{an\pi t}{L}\right) \sin\left(\frac{n\pi x}{L}\right)$$

in the zero initial velocity case thus says that an arbitrary wave is a superposition of standing waves, each corresponding to a certain frequency. (We will generalize this "frequency" perspective a bit later to the case of a *continuous* range of frequencies when we discuss Fourier transforms.)

Justifying wave solutions. Moreover, the form $u(x,t) = \frac{1}{2}(h(x-at) + h(x+at))$ allows us to formally justify the fact that the series above does give an actual solution to the wave equation. Recall the issue we mentioned last time that even though

$$u(x,t) = \sum_{n=1}^{\infty} a_n \cos\left(\frac{an\pi t}{L}\right) \sin\left(\frac{n\pi x}{L}\right)$$

might converge, the series obtained by differentiating might not, and so it is not clear that we can verify $u_{tt} = a^2 u_{xx}$ directly in such a situation. But now we have a way out: the function $u(x,t) = \frac{1}{2}(h(x-at) + h(x+at))$ definitely exists and is definitely differentiable as much times as we want (assuming h is), so we *can* verify by direct substitution that it satisfies the wave equation:

$$u_{tt} = \frac{1}{2}(a^2 h''(x-at) + a^2 h''(x+at)) \quad \text{and} \quad u_{xx} = \frac{1}{2}(h''(x-at) + h''(x+at),$$

so $u_{tt} = a^2 u_{xx}$ is true. The upshot is that, although it may be the case that "term-by-term" differentiation in

$$u(x,t) = \sum_{n=1}^{\infty} a_n \cos\left(\frac{an\pi t}{L}\right) \sin\left(\frac{n\pi x}{L}\right)$$

does not actually work, the function defined by this series satisfies the wave equation nonetheless via alternative means. A similar thing is true when we include a nonzero initial velocity.

Laplace equation. The final PDE we consider is the *Laplace equation*, which in the two-dimensional case is

$$\frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2} = 0.$$

(This is the prototypical example of an *elliptic* partial differential equation.) This is a linear homogeneous PDE, corresponding to the linear partial differential operator given by

$$\Delta = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}.$$

This is called the *Laplacian* operator, and with this notation the Laplace equation is $\Delta u = 0$.

Solutions of the Laplace equation are known as *harmonic* functions, and show up in various contexts. Perhaps the main one is in characterizing various types of *potential* functions in physics. For example, the standard gravitational and electric potential and functions are harmonic. That this might be a desirable thing to have comes from recognizing that being harmonic says that the gradient of u is divergence-free:

$$\nabla u = (u_x, u_y) \rightsquigarrow \operatorname{div}(\nabla u) = \frac{\partial (u_x)}{\partial x} + \frac{\partial (u_y)}{\partial y} = u_{xx} + u_{yy}.$$

Harmonic functions also show up in *complex analysis* as the real and imaginary parts of *complex-differentiable* functions. That is, if $f : \mathbb{C} \to \mathbb{C}$ has a complex derivative (you'll learn what this means in MATH 382) and we write f as $f = u + iv$ where u, v are both real-valued, it is a basic fact that u and v must both be harmonic. (This follows from what are called the *Cauchy-Riemann equations*, which is a system of PDEs you'll see in a complex analysis course.)

But most relevant to things we've seen in this course is that the fact that the Laplace equation describes steady-state solutions of the two-dimensional heat equation:

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

For a steady-state solution we require that $u_t = 0$, which turns the steady-state heat equation into precisely the Laplace equation.

Rectangular regions. The form which solutions of the Laplace equation takes depends heavily on the region being considered. We will mainly be concerned with rectangular regions and, later, disks. We begin by considering the Laplace equation on a rectangle

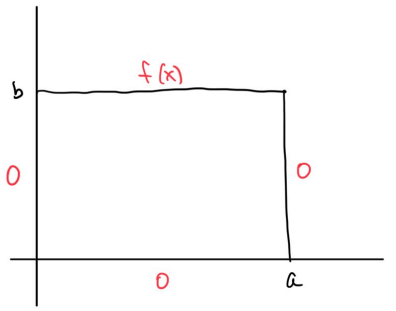

with boundary conditions

$$u(0,y) = 0 = u(a,y),\ u(x,0) = 0,\ u(x,b) = f(x).$$

(There is no initial condition since there is no dependence on t.) Now, as usual, we separate our solution and get some ODEs:

$$u(x,y) = X(x)Y(y) \rightsquigarrow X'' - \lambda X = 0,\ Y'' + \lambda Y = 0.$$

The boundary conditions $u(0,y) = 0 = u(a,y)$ on the left and right edges give $X(0) = 0 = X(a)$. The equation $X'' - \lambda X = 0$ with these boundary conditions is one we've solved already, where we get nonzero solutions

$$X_n = \sin\left(\frac{n\pi x}{a}\right) \quad \text{for} \quad \lambda_n = -\frac{n^2\pi^2}{a^2}.$$

For these λ_n, the ODE for Y is $Y'' - \frac{n^2\pi^2}{a^2}Y = 0$, which has solution

$$Y = c_1 e^{n\pi y/a} + c_2 e^{-n\pi y/a}.$$

Now, the boundary condition $u(x,0) = 0$ along the bottom edge gives $Y(0) = 0$, so $c_1 + c_2 = 0$. Thus $c_2 = -c_1$, so

$$Y = c_1(e^{n\pi y/a} - e^{-n\pi y/a}).$$

After setting c_1 we take

$$Y_n = e^{n\pi y/a} - e^{-n\pi y/a}$$

as a fundamental solution. In fact, this expression is precisely 2 times the *hyperbolic sine* function: $Y_n = 2\sinh(\frac{n\pi y}{a})$. So then we can actually simply take $Y_n = \sinh(\frac{n\pi y}{a})$ as a fundamental solution.

The general solution of the Laplace equation with these boundary conditions is then

$$u(x,y) = \sum_{n=1}^{\infty} c_n \sinh\left(\frac{n\pi y}{a}\right)\sin\left(\frac{n\pi x}{a}\right).$$

The remaining boundary condition $u(x,b) = f(x)$ along the top edge then gives

$$f(x) = u(x,b) = \sum_{n=1}^{\infty} c_n \sinh\left(\frac{n\pi b}{a}\right) \sin\left(\frac{n\pi x}{a}\right),$$

which says that the $c_n \sinh(\frac{n\pi b}{a})$ should be the Fourier sine coefficients of $f(x)$. Then our solution is

$$u(x,y) = \sum_{n=1}^{\infty} c_n \sinh\left(\frac{n\pi y}{a}\right) \sin\left(\frac{n\pi x}{a}\right)$$

where

$$c_n = \frac{2}{\sinh(\frac{n\pi b}{a})} \int_0^a f(x) \sin\left(\frac{n\pi x}{a}\right)\, dx.$$

Lecture 14: More on Laplace

Warm-Up. We solve the Laplace equation $\Delta u = 0$ on the rectangle $0 \le x \le a,\ 0 \le y \le b$ with boundary conditions

$$u(0,y) = 0 = u(1,y),\ u(x,0) = 0,\ u(x,2) = 1 - x.$$

As we derived last time, the solution is

$$u(x,y) = \sum_{n=1}^{\infty} c_n \sinh(n\pi y) \sin(n\pi x)$$

where $b_n \sinh(\frac{n\pi b}{a})$ are the Fourier sine coefficients of $u(x,0)$. In this specific case we have

$$c_n = \frac{2}{\sinh(2n\pi)} \int_0^1 (1-x) \sin(n\pi x)\, dx = \frac{2}{n\pi \sinh(2n\pi)}.$$

Hence our solution is

$$u(x,y) = \sum_{n=1}^{\infty} \frac{2}{n\pi \sinh(2n\pi)} \sinh(n\pi y) \sin(n\pi x).$$

Here is a plot of the graph of $u(x,y)$:

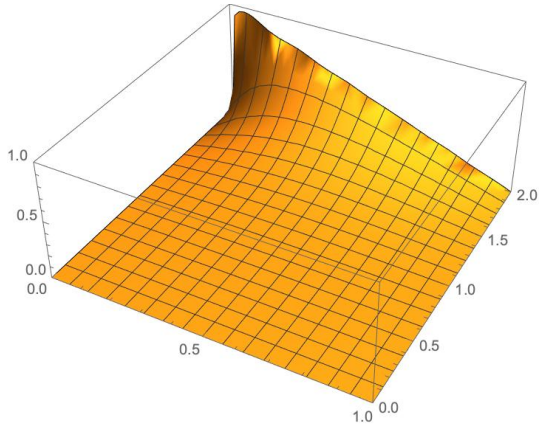

If we interpret this as a steady-state solution of the heat equation, we see that, as expected from the boundary conditions, the temperature of $u(x,2) = 1 - x$ along the top boundary "dissuses" away and decreases as we move away from this boundary, reaching a temperature of zero along the other pieces of the boundary. The diffusion occurs quite rapidly, meaning we don't have to move away from the top boundary by too much before we find that the temperature has fallen dramatically.

Other rectangular boundaries. Let us now consider the Laplace equation on the same rectangle $0 \le x \le a, 0 \le y \le b$ as before, but now with nonzero boundary condition on the left side and zero on the other pieces of the boundary:

$$u(0,y) = g(y),\ u(a,y) = 0,\ u(x,0) = 0 = u(x,b).$$

Here we end up with the following boundary value problems for Y:

$$Y'' + \lambda Y = 0,\ Y(0) = 0 = Y(b).$$

This has fundamental solutions

$$Y_n = \sin\left(\frac{n\pi y}{b}\right) \quad \text{for} \quad \lambda_n = \frac{n^2\pi^2}{b^2}.$$

For X we then get (imposing only $u(a,y) = 0$ for now)

$$X'' - \frac{n^2\pi^2}{b^2}X = 0,\ X(a) = 0.$$

The solution thus looks like

$$X = c_1 e^{n\pi x/b} + c_2 e^{-n\pi x/b},$$

and $X(a) = 0$ forces $c_2 = -c_1 e^{2n\pi a/b}$. Hence

$$X = c_1 e^{n\pi x/b} - c_1 e^{2n\pi a/b} e^{-n\pi x/b}.$$

To find a "nice" fundamental solution, we can rewrite this by factoring out a factor of $e^{n\pi a/b}$:

$$X = c_1 e^{n\pi a/b}(e^{n\pi(x-a)/b} - e^{-n\pi(x-a)/b}).$$

The expression in parentheses is $2\sinh(\frac{n\pi(x-a)}{b})$, so we can take

$$X_n = \sinh\left(\frac{n\pi(x-a)}{b}\right)$$

as a fundamental solution.

Thus our solution looks like

$$u(x,y) = \sum_{n=1}^{\infty} c_n \sinh\left(\frac{n\pi(x-a)}{b}\right)\sin\left(\frac{n\pi y}{b}\right).$$

The remaining boundary condition $u(0,y) = g(y)$ gives

$$g(y) = \sum_{n=1}^{\infty} -c_n \sinh\left(\frac{n\pi a}{b}\right)\sin\left(\frac{n\pi y}{b}\right),$$

so that $-c_n \sinh\left(\frac{n\pi a}{b}\right)$ must be the Fourier sine coefficients of $g(y)$.

Now, for the problem where we have both a nonzero boundary condition on the top and left sides of the rectangle, with zero boundary conditions on the other two sides, we can simply add together the solutions we've found so far. That is, if $u_1(x, y)$ is the solution we had last time for the case of a nonzero boundary condition $f(x)$ on the to edge, and $u_2(x, y)$ is the solution above with a nonzero boundary condition $g(y)$ on the left edge, then

$$u(x, y) = u_1(x, y) + u_2(x, y)$$

still satisfies the Laplace equation (the Laplace equation is linear) and satisfies the boundary conditions

$$u(0, y) = u_1(0, y) + u_2(0, y) = 0 + g(y) = g(y),\ u(x, b) = u_1(x, b) + u_2(x, b) = f(x) + 0 = f(x),$$

and zeroes on the other two edges. More generally, for a rectangle with nonzero boundary conditions on *all* sides, we can first find the solutions with one nonzero boundary conditions at-a-time, and then add these together to get the full solution. Thus, we can solve the Laplace equation on any rectangular region.

Dirichlet vs Neumann boundary conditions. The boundary conditions give above, where we specify the values of our solution on the boundary, are known as *Dirichlet* boundary conditions. Alternatively, *Neumann* boundary conditions are ones where we specify instead the *rate of change* in the "normal" direction of the solution along boundary. In the case of a rectangle, Neumann boundary conditions thus looks like

$$u_x(0, y) = f(y),\ u_x(a, y) = g(y),\ u_y(x, 0) = h(x),\ u_y(x, b) = \ell(x).$$

The process of solving Neumann boundary value problems is the same as Dirichlet problems (indeed, we already saw some Neumann conditions with the heat equation, say with insulated ends), but the structure of the solutions can be a bit different. You'll see some examples on the homework.

Uniqueness. We now prove uniqueness of solutions to the Laplace equation on *any* region with Dirichlet boundary conditions. (At least, any region nice enough to which Green's theorem is applicable!) This is one of the only uniqueness results we can definitely justify in this class, since it just depends on Green's theorem (more precisely, the divergence form of Green's theorem) from multivariable calculus. The upshot is that the behavior of harmonic functions is fairly restrictive: the behavior along the boundary of a region completely determines the behavior through the entire region, since there is only one possible function which those boundary values can give rise to. (You will see a similar phenomenon show in complex analysis, which is no surprise given the connection between harmonic functions and complex differentiable functions we mentioned last time.)

Suppose D is a nice two-dimensional region with boundary ∂D. (For example, "nice" could mean that the boundary ∂D consists of piecewise smooth curves.) The claim is that if f and g are two functions satisfying the Laplace equation on D which have the same values along the boundary of D, then $f = g$ on all of D. Note that $f = g$ on D if and only if $f - g$ is zero on all of D, and $f = g$ on ∂D if and only if $f - g = 0$ on ∂D, so it is enough to know that a harmonic function which is zero along ∂D must then be zero on all of D.

So, suppose $\Delta u = 0$ on D and $u = 0$ on ∂D. Consider the vector field $u\nabla u$. The divergence form of Green's theorem states that

$$\int_{\partial D} u\nabla u \cdot \mathbf{n}\, ds = \iint_D \operatorname{div}(u\nabla u)\, dA$$

when ∂D has the proper orientation. (The orientation won't matter in this application. Also, the left side is a line integral where we measure the extent to which the field $u\nabla u$ flows "through" the boundary, as indicated by the dot product with the unit normal vector $\mathbf{n}$ along the boundary. It's not important to go back and recall what this line integral actually means, since it will be zero in our case.) Since $u = 0$ on ∂D, the vector field $u\nabla u$ is zero on the boundary as well, so the left side of the equation above is zero.

For the right side we first compute:

$$\operatorname{div}(u\nabla u) = \operatorname{div}(uu_x, uu_y) = \frac{\partial(uu_x)}{\partial x} + \frac{\partial(uu_y)}{\partial y} = u_x u_x + uu_{xx} + u_y u_y + uu_{yy}.$$

Since $\Delta u = 0$ on D, we have that $uu_{xx} + uu_{yy} = u(u_{xx} + u_{yy}) = 0$ on D, so the integral on the right simplifies and we are left with

$$0 = \iint_D [(u_x)^2 + (u_y)^2]\, dA.$$

But the integrand $(u_x)^2 + (u_y)^2$ is never negative, so in order for this integral to be zero we must have $(u_x)^2 + (u_y)^2 = 0$ at all points of D. This forces u_x and u_y to be zero on all of D, which implies that u must be constant on D. Since the value of u on ∂D is zero, this "constant" must be zero, so $u = 0$ on all of D as desired. Thus, solutions of the Laplace equation on a region with specified Dirichlet boundary conditions are unique.

Lecture 15: Polar Laplace

Warm-Up. We solve the Laplace equation on a rectangle with some *Neumann* boundary conditions, say

$$u_x(0, y) = f(y),\ u_x(a, y) = 0,\ u_y(x, 0) = 0 = u_y(x, b).$$

As usual, we separate variables and for Y in particular we get

$$Y'' + \lambda Y = 0, \quad Y'(0) = 0 = Y'(b).$$

This is the boundary value problem we saw previously when considering the heat equation with insulated ends, and the fundamental solutions are

$$Y_n = \cos\left(\frac{n\pi y}{b}\right) \text{ for } \lambda_n = \frac{n^2\pi^2}{b^2}, \text{ and } Y_0 = 1 \text{ for } \lambda_0 = 0.$$

Then for X and λ_n $(n \neq 0)$ we get

$$X'' - \frac{n^2\pi^2}{b^2}X = 0,\ X'(a) = 0,$$

which has solution

$$X = c_1 e^{n\pi x/b} + c_1 e^{2n\pi a/b} e^{-n\pi x/b}.$$

We can express in terms of hyperbolic cosine via:

$$X = c_1 e^{n\pi a/b}(\underbrace{e^{n\pi(x-a)/b} + e^{-n\pi x - a/b}}_{2\cosh(n\pi(x-a)/b}),$$

so we take

$$X_n = \cosh\left(\frac{n\pi(x-a)}{b}\right)$$

as the fundamental solution. For $\lambda_0 = 0$, we have $X = c_1 + c_2 x$, and $X'(a) = 0$ gives $X_0 = 1$.

Thus, our solution is

$$u(x,y) = \frac{c_0}{2} + \sum_{n=1}^{\infty} c_n \cosh\left(\frac{n\pi(x-a)}{b}\right)\cos\left(\frac{n\pi y}{b}\right),$$

and the remaining boundary condition $u_x(0,y) = f(y)$ requires

$$f(y) = u_x(0,y) = \sum_{n=1}^{\infty} -\frac{n\pi c_n}{b}\sinh\left(\frac{n\pi a}{b}\right)\cos\left(\frac{n\pi y}{b}\right),$$

so that $-\frac{n\pi c_n}{b}\sinh\left(\frac{n\pi a}{b}\right)$ should be the Fourier *cosine* coefficients of $f(y)$. (In particular, the constant Fourier cosine coefficient c_0 should be zero since it does not appear in the series above, so f should satisfy $\int_0^b f(y)\,dy = 0$ in order for there to be a solution at all.)

Laplace on a disk. The Laplace equation on other regions sides rectangles is not so easy to solve, but at least we can also work out a general solution in the case of a disk. This will take quite a bit more work than the case of a rectangle, and depends more heavily on some ODE material, but we get a nice explicit form in the end.

The key point is that we can do all the required computations in polar coordinates. The Laplace equation in polar coordinates is

$$u_{rr} + \frac{1}{r}u_r + \frac{1}{r^2}u_{\theta\theta} = 0.$$

(We briefly saw the left side a bit when considering the heat equation in polar coordinates.) We consider a disk centered at the origin of radius a, with boundary condition

$$u(a,\theta) = f(\theta).$$

Since $0 = 2\pi$ in terms of values for θ on the disk, we must require that f be 2π-periodic. Moreover, we will also assume that the solution $u(r,\theta)$ we seek is bounded on the disk. All of these (mild) assumptions are necessary in order to derive an explicit solution.

After separating variables $u(r,\theta) = R(r)\Theta(\theta)$, we get

$$\Theta'' - \lambda\Theta = 0 \text{ and } r^2R'' + rR' + \lambda R = 0.$$

Now, for $\lambda > 0$ we get

$$\Theta = c_1 e^{\sqrt{\lambda}\theta} + c_2 e^{-\sqrt{\lambda}\theta}.$$

But such a thing cannot be periodic on the boundary unless c_1 and c_2 are both zero, so we get no nonzero solution here. For $\lambda = 0$, we get $\Theta = c_1 + c_2\theta$, and the periodic assumption forces $c_2 = 0$, so that Θ is constant. The ODE for R in this case becomes

$$r^2R'' = -rR'.$$

This can be solved by rewriting it as $\frac{R''}{R'} = -\frac{1}{r}$, integrating once to obtain

$$\ln|R'| = -\ln r + C,$$

exponentiating to get

$$R' = c_2 e^{-\ln r} = \frac{c_2}{r} \text{ for some constant } c_2,$$

and finally integrating once more to get $R = c_1 + c_2 \ln r$. (I told you the case of a disk was more involved than that of a rectangle!) But our solution is meant to be bounded, and $\ln r$ is unbounded as $r \to 0$ in the disk, so we must have $c_2 = 0$ and thus $R = c_1$ is constant as well. Thus, in the $\lambda = 0$ case we get a constant solution for u, so $u_0 = 1$ is our fundamental solution.

Finally, suppose $\lambda = -k^2 < 0$. For Θ we get

$$\Theta = c_1 \cos(k\theta) + c_2 \sin(k\theta),$$

and the 2π-periodic requirement forces k to be an integer. The equation for R is then

$$r^2 R'' + rR' - k^2 R = 0.$$

This is an example of an *Euler equation*, which you perhaps learned how to solve in a previous course. The strategy is to search for solution of the form $R = r^n$. If this is to be solution, then we must have

$$\underbrace{n(n-1)r^n}_{r^2R''} + \underbrace{nr^n}_{rR'} - k^2 r^n = 0,$$

which requires that $n(n-1) + n - k^2 = 0$. Solving for n gives $n = \pm k$ as the possible values, so r^k and r^{-k} are two solutions of this Euler equation, and we can write the general solution as

$$R = c_1 r^k + c_2 r^{-k}.$$

The requirement that $u(r, \theta)$ be bounded forces $c_2 = 0$ since r^{-k} blows up to infinite as $r \to 0$, so $R_n = r^n$ is our fundamental solution for R when $\lambda_n = -n^2$.

After putting everything together, we thus get that our general solution is

$$u(r, \theta) = \frac{c_0}{2} + \sum_{n=1}^{\infty} (c_n r^n \cos(n\theta) + b_n r^n \sin(n\theta)).$$

The boundary condition $u(a, \theta) = f(\theta)$ gives

$$f(\theta) = u(a, \theta) = \frac{c_0}{2} + \sum_{n=1}^{\infty} (c_n a^n \cos(n\theta) + b_n a^n \sin(n\theta)),$$

so the $c_n a^n$ and $b_n a^n$ must be the Fourier coefficients (in the very first sense we considered in the first week of class) of $f(\theta)$.

Lecture 16: Fourier Transform

Warm-Up. We solve the Laplace equation on the region $3 \leq r < \infty$ *outside* the disk $r < 3$, with boundary values

$$u(3, \theta) = \begin{cases} 0, & 0 \leq \theta < \pi \\ 2\pi - \theta, & \pi \leq \theta < 2\pi. \end{cases}$$

Just like last time when considering the Laplace equation *inside* the disk, we look here for a bounded solution $u(r, \theta)$ which is periodic on the boundary circle. Following the same analysis as last time, the first minor difference comes in the $\lambda = 0$ case, where the solution of R is

$$R = c_1 + c_2 \ln r.$$

Previously we used the fact that $\ln r$ is unbounded as $r \to 0$ to force $c_2 = 0$, and here it the fact that $\ln r$ is unbounded as $r \to \infty$ that again forces $c_2 = 0$. Thus we get a constant solution for $\lambda = 0$, just as before.

The only other difference comes in the $\lambda = -k^2 < 0$ case, where we again get

$$R = c_1 r^k + c_2 r^{-k}.$$

However, now it is the c_1 coefficient which must be zero since $r^k \to \infty$ as $r \to \infty$. Thus here we get $R_n = r^{-k}$ as the fundamental solution, so the general solution is

$$u(r,\theta) = \frac{c_0}{2} + \sum_{n=1}^{\infty} [c_n r^{-n} \cos(n\theta) + d_n r^{-n} \sin(n\theta)].$$

The boundary condition $u(3,\theta) = f(\theta)$ implies that $3^{-n}c_n$ and $3^{-n}d_n$ are the Fourier coefficients of $f(\theta)$, so we get;

$$\begin{aligned}
c_0 &= \frac{1}{\pi}\int_0^{2\pi} f(\theta)\, d\theta = \frac{1}{\pi}\int_\pi^{2\pi} (2\pi - \theta)\, d\theta = \frac{\pi}{2} \\
c_n &= \frac{3^n}{\pi}\int_0^{2\pi} f(\theta)\cos(n\theta)\, d\theta = \frac{1}{\pi}\int_\pi^{2\pi} (2\pi - \theta)\cos(n\theta)\, d\theta = \frac{3^n(\cos \pi n - 1)}{\pi n^2} \\
d_n &= \frac{3^n}{\pi}\int_0^{2\pi} f(\theta)\sin(n\theta)\, d\theta = \frac{1}{\pi}\int_\pi^{2\pi} (2\pi - \theta)\sin(n\theta)\, d\theta = \frac{3^n \cos \pi n}{n}.
\end{aligned}$$

Thus our solution is

$$u(r,\theta) = \frac{\pi}{4} + \sum_{n=1}^{\infty} \left[\frac{3^n(\cos \pi n - 1)}{\pi n^2} r^{-n} \cos(n\theta) + \frac{3^n \cos \pi n}{n} r^{-n} \sin(n\theta) \right]$$

Spherical Laplace. ***TO BE FINISHED***

Back to the heat equation. We now move to a new topic, that of the *Fourier transform*, but will motivate its definition by returning to the heat equation. We seek to find a solution of the heat equation

$$u_t = \alpha^2 u_{xx}, \quad u(x,0) = f(x)$$

but now with no boundary conditions and modeling the temperature in a rod of *infinite* length. In other words, we now want $u(x,t)$ to be a function where x can take values on all of $\mathbb{R}$, as opposed to the type of bounded interval $[0, L]$ we considered before. We will make one assumption, however, namely that $u(x,t)$ should be bounded as x varies in $\mathbb{R}$.

Following the same method as before, we come to consider the ODE

$$X'' - \lambda X = 0,$$

with no boundary conditions. For $\lambda > 0$ we have

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

which, if nonzero, does not remain bounded as $x \to \pm\infty$. If $\lambda = 0$, we get $X = c_1 + c_2 x$, and the bounded condition forces $c_2 = 0$. So we get a nonzero solution $X_0 = 1$ for $\lambda_0 = 0$. Then we are left with the $\lambda = -k^2 < 0$ case, where now we write our solution in complex form as

$$X = c_1 e^{ikx} + c_2 e^{-ikx}.$$

(Recall that expressing functions in terms of these complex exponentials is equivalent to expressing in terms of sine and cosine.) This is always bounded since $|e^{ikx}| = 1$, so we have no further restrictions, and in particular there is no restriction in this case that k must be an integer, so that k is now any nonzero real number! (The integer restriction previously came from the boundary conditions on $[0, L]$.) So, we get a whole bunch of fundamental solutions e^{ikx} for nonzero k, and if we allow $k = 0$ we get the one $e^0 = 1$ we had for $\lambda = 0$.

From discrete to continuous. Thus, by analogy with the bounded string we had before, we would expect (or, perhaps more precisely, hope) that our solution in this case should be end up being a "sum" of scalar multiples of e^{ikx}, where k ranges among all real numbers:

$$\sum_{k \text{ is real}} c_k e^{ikx}.$$

But of course, such a sum does not quite make sense since we cannot actually "add"—in the normal sense—such quantities together if there are "$\mathbb{R}$-many" of them. But, the way out of this is to interpret the "sum" as an integral over $-\infty < k < \infty$ instead, so that we want something like

$$\int_{-\infty}^{\infty} c(k)\, e^{ikx}\, dk.$$

(The point is that an integral should be interpreted as a "continuous sum", versus an infinite series, which is still a "discrete sum". We write the coefficient as $c(k)$ instead of c_k to make it look more like a function since we are getting such a value for each real "input" k, but the idea is the same: for each k, $c(k)$ is a scalar coefficient.) As for the "coefficient function" $c(k)$, again if we go purely by analogy with the coefficients obtained in the "discrete" case before:

$$c_k = \frac{1}{2\pi} \int_{-\pi}^{\pi} f(x) e^{-ikx}\, dx,$$

we might expect to get something like

$$c(k) = \frac{1}{2\pi} \int_{-\infty}^{\infty} f(x) e^{-ikx}\, dx$$

The function (of k) defined by this integral $\int_{-\infty}^{\infty} f(x) e^{-ikx}\, dx$ (we'll forget the constant $\frac{1}{2\pi}$ in front for now) is thus a continuous analog of a Fourier coefficient, and will eventually give us a way to express f as a continuous analog of a Fourier series.

Fourier transform. Motivated by the discussion above, we thus make the following definition: the *Fourier transform* of a function f is the function $\hat{f}$ defined by

$$\hat{f}(s) = \int_{-\infty}^{\infty} f(x) e^{-isx}\, dx.$$

We finish with an example, and compute the Fourier transform of the "pulse" function

$$f(x) = \begin{cases} 1 & -T \leq x \leq T \\ 0 & \text{otherwise} \end{cases}$$

where $T > 0$ is a constant. We have:

$$\hat{f}(s) = \int_{-\infty}^{\infty} f(x) e^{-isx}\, dx$$

$$
\begin{aligned}
&= \int_{-\infty}^{-T} 0 \cdot e^{-isx}\,dx + \int_{-T}^{T} 1 \cdot e^{-isx}\,dx + \int_{T}^{\infty} 0 \cdot e^{-isx}\,dx \\
&= \int_{-T}^{T} e^{-isx}\,dx \\
&= -\frac{1}{is} e^{-isx}\Big|_{-T}^{T} \\
&= -\frac{1}{is}(e^{-isT} - e^{isT}) \\
&= -\frac{1}{is}(-2i\sin(sT)) \\
&= \frac{2\sin(sT)}{s}.
\end{aligned}
$$

Note that although this resulting function appears at first to be undefined at $s = 0$, since

$$\lim_{s\to 0} \frac{2\sin(sT)}{s} = 2T$$

we can set its value at $s = 0$ to be $2T$ and still get a continuous (in fact C^1) function on $\mathbb{R}$. This is reflective of a general idea we will highlight next time.

Lecture 17: More on Transforms

Warm-Up. We compute the Fourier transforms of the functions

$$f(t) = \begin{cases} 0 & t < 0 \\ e^{-at} & t \geq 0 \end{cases}$$

and $g(t) = e^{-a|t|}$, where $a > 0$ is a constant. For f we have:

$$
\begin{aligned}
\hat{f}(s) &= \int_{-\infty}^{\infty} f(t) e^{-ist}\,dt \\
&= \int_{0}^{\infty} e^{-at} e^{-ist}\,dt \\
&= \int_{0}^{\infty} e^{-(a+is)t}\,dt \\
&= -\frac{1}{a+is} e^{-(a+is)t}\Big|_{0}^{\infty} \\
&= \frac{1}{a+is}.
\end{aligned}
$$

(Recall that to "evaluate" an improper integral at an infinite bound really means to take a limit as the variable in question approaches that infinite bound, so that evaluating at the upper bound of ∞ in this case really means to compute

$$\lim_{t\to\infty} -\frac{1}{a+is} e^{-(a+is)t} = \lim_{t\to\infty} -\frac{1}{a+is} e^{-at} e^{-ist}.$$

Here, the e^{-ist} term is bounded and the e^{-at} decays to 0 as $t \to \infty$, so the entire limit, and hence the "value" at ∞, is zero.) Thus the Fourier transform of f is $\hat{f}(s) = \frac{1}{a+is}$. Note that even though f is real-valued here, its Fourier transform is complex-valued.

For g we have:

$$\begin{aligned}\hat{g}(s) &= \int_{-\infty}^{\infty} g(x)e^{-isx}\,dx \\ &= \int_{-\infty}^{0} e^{-a(-t)}e^{-ist}\,dt + \int_{0}^{\infty} e^{-at}e^{-ist}\,dt,\end{aligned}$$

where in the first integral we use the fact that $|t| = -t$ for $t < 0$. Now, we can compute each of these in a similar way as we did for f above, but in fact we can directly use the result for f above here. In particular, note that the second integral above is precisely the integral which gave $\hat{f}(s)$. For the first integral we can make a change of variables $t \mapsto -t$ to get:

$$\int_{-\infty}^{0} e^{-a(-t)}e^{-ist}\,dt = \int_{0}^{\infty} e^{-at}e^{ist}\,dt,$$

which is precisely the value of $\hat{f}$ at the input $-s$ instead of s. (We can also characterize this integral as the conjugate of $\hat{f}(s)$.) Thus we have

$$\hat{g}(s) = \hat{f}(-s) + \hat{f}(s) = \frac{1}{a - is} + \frac{1}{a + is} = \frac{2a}{a^2 + s^2}$$

as the Fourier transform of g.

Functions vs their transforms. Let us graph the function g above (in green) and its Fourier transform (in blue) for $a = 1, 2, 3$:

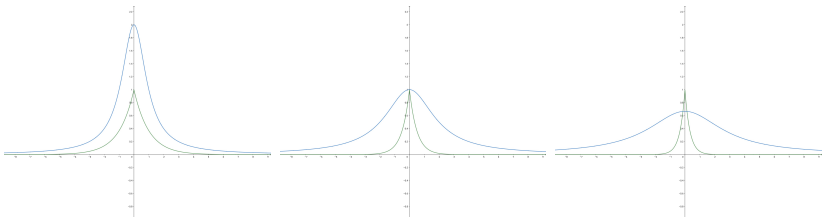

The key observation here is that g fails to be differentiable at 0 (notice the "cusp"), but $\hat{g}$ is perfectly nice and smooth at 0. The cusp gets worse as a increases, and in turn $\hat{g}$ behaves in a more controlled way. The idea is that taking the Fourier transform of a function can "smooth out" any singularities it might have, making the Fourier transformed function a possibly nicer function to work with. (Below we argue that we can always recover g from $\hat{g}$, so that we lose no information when working with the Fourier transform instead.)

We see the same phenomenon with the example of the "pulse" function from last time, which is 1 for $-T \leq x \leq T$ and zero elsewhere:

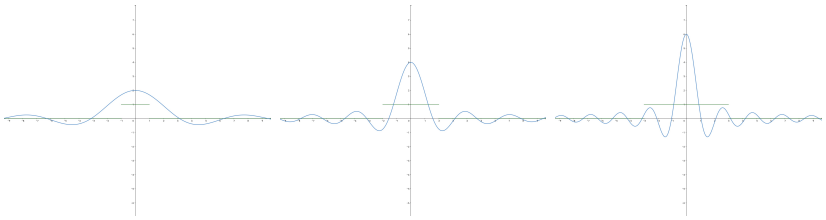

(Here we take $T = 1, 2, 3$.) Now the original function has discontinuities, which disappear after computing the Fourier transform, which is even differentiable!

Fourier transforms and coefficients. We previously alluded to the idea of using Fourier transforms to give "continuous" versions of Fourier coefficients, and now we make this precise. Indeed, we claim that the process of computing the Fourier coefficients of a function is *itself* a type of Fourier transform, only in a discrete setting. Given a 2π-periodic (possibly complex-valued) function f, we can compute its (complex) Fourier coefficients

$$c_n(f) = \frac{1}{2\pi}\int_{-\pi}^{\pi} f(x)e^{-inx}\,dx.$$

(We use $c_n(f)$ here in order to emphasize the coefficients depend on the function f.) These coefficients form a sequence $(c_n(f))$ indexed by integers $n \in \mathbb{Z}$, and we can view such a sequence as giving a function $\mathbb{Z} \to \mathbb{C}$ by sending $n \in \mathbb{Z}$ to $c_n(f) \in \mathbb{C}$. Thus, we can interpret the operation $\mathcal{F}$ of computing Fourier coefficients as one which maps periodic functions on $\mathbb{R}$ to complex-valued functions on $\mathbb{Z}$:

$$\mathcal{F} : \{2\pi\text{-periodic functions } \mathbb{R} \to \mathbb{C}\} \to \{\text{functions } \mathbb{Z} \to \mathbb{C}\}, \quad \mathcal{F}(f) = (c_n(f))$$

The point is that the Fourier transform does the exact same thing, only where we now drop the requirement of periodicity. Now we have the operation $\mathcal{F}$, which sends a (complex-valued) function on $\mathbb{R}$ to its Fourier transform, mapping the space of functions on $\mathbb{R}$ to itself:

$$\mathcal{F} : \{\text{functions } \mathbb{R} \to \mathbb{C}\} \to \{\text{functions } \mathbb{R} \to \mathbb{C}\}, \quad \mathcal{F}(f) = \hat{f}.$$

In other words, the Fourier transform is literally just the "continuous" analog of the "discrete" Fourier coefficient construction above. It is thus no accident that the integral defining the Fourier transform looks very similar to the one defining Fourier coefficients:

$$\int_{-\infty}^{\infty} f(x)e^{-isx}\,dx \quad \text{vs} \quad \frac{1}{2\pi}\int_{-\pi}^{\pi} f(x)e^{-inx}\,dx,$$

since, after all, they are meant to be essentially the "same" construction. (The factor of $\frac{1}{2\pi}$ on the right missing in the definition of the Fourier transform shows up elsewhere, as we will soon clarify.)

Fourier inversion. From the perspective above, we can interpret the construction of a Fourier series as an "inverse" operation. Namely, consider the "inverse" transform $\mathcal{F}^{-1}$ in the discrete case:

$$\mathcal{F}^{-1} : \{\text{functions } \mathbb{Z} \to \mathbb{C}\} \to \{2\pi\text{-periodic functions } \mathbb{R} \to \mathbb{C}\}$$

defined by sending the sequence of coefficients (c_n) to the series

$$\sum_{n=-\infty}^{\infty} c_n e^{inx}.$$

The Fourier convergence theorem can then be interpreted as saying that this $\mathcal{F}^{-1}$ is indeed the inverse of $\mathcal{F}$, at least for C^1 functions:

$$f = \mathcal{F}^{-1}(\mathcal{F}(f)) \quad \text{means} \quad f(x) = \sum_{n=-\infty}^{\infty} c_n(f)e^{inx} \text{ where } c_n(f) = \frac{1}{2\pi}\int_{-\pi}^{\pi} f(x)e^{-inx}\,dx.$$

That is, from f, compute its Fourier coefficients, then form the Fourier series using these coefficients, and you get back f. (If f is only piecewise C^1 and not C^1 everywhere, then $\mathcal{F}^{-1}(\mathcal{F}(f))$ produces the function $\frac{1}{2}(f(x^-) + f(x^+))$ of averages of left- and right-hand limits.)

The analog of this in the continuous setting for

$$\mathcal{F} : \{\text{functions } \mathbb{R} \to \mathbb{C}\} \to \{\text{functions } \mathbb{R} \to \mathbb{C}\}, \quad \mathcal{F}(f) = \hat{f}$$

is called the *Fourier inversion theorem*, or *Fourier inversion formula*. Define the *inverse Fourier transform* $\mathcal{F}^{-1}(g)$ of a function g to be

$$(\mathcal{F}^{-1}(g))(s) = \frac{1}{2\pi} \int_{-\infty}^{\infty} g(x) e^{isx}\, dx.$$

(There's the 2π showing up in the definition of complex Fourier coefficients but missing in the definition of the Fourier transform!) The Fourier inversion formula says that this is indeed in the inverse of the Fourier transform operation $\mathcal{F}$ in the sense that

$$f(x) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \hat{f}(s) e^{isx}\, ds, \text{ or in other notation } f = \mathcal{F}^{-1}(\mathcal{F}(f)),$$

at least when f is C^1. (For f only piecewise C^1, you get $\frac{1}{2}(f(s^-) + f(s^+))$ on the left side of this formula.) Thus, the inverse Fourier transform of the Fourier transform of f gives back f, so that f can be recovered from its Fourier transform. We will give a rough justification of this inversion formula later, but the upshot is that the function $\hat{f}$ contains the exact same information as f does, since we can recover f from $\hat{f}$. (In the discrete case, the coefficients $c_n(f)$ contain the same information as f, since we can recover f from these coefficients via a Fourier series.)

Just as the Fourier transform $\hat{f}$ can be viewed as the continuous version of the sequence of Fourier coefficients $(c_n(f))$, the expression given by the inverse Fourier transform is thus the continuous analog of a Fourier series:

$$\frac{1}{2\pi} \int_{-\infty}^{\infty} \hat{f}(s) e^{isx}\, ds \text{ is the continuous version of } \sum_{n=-\infty}^{\infty} c_n e^{inx},$$

and this is why we expect such a thing to be useful in solving the heat equation on $\mathbb{R}$. Of course, there is still that extra $\frac{1}{2\pi}$ we have to deal with in order to make this analogy complete. In the discrete Fourier series case, the $\frac{1}{2\pi}$ shows up in the definition of the coefficients, while here we are using it in the definition of the inverse transform. Ultimately it is just a matter of convention as to where we put this factor. Many sources define the Fourier transform by

$$\hat{f}(s) = \frac{1}{2\pi} \int_{-\infty}^{\infty} f(x) e^{-ixs}\, dx,$$

so that it truly mimics the definition of the coefficients $c_n(f)$, and in this case the inverse Fourier transform would be

$$(\mathcal{F}^{-1}(\hat{f}))(x) = \int_{-\infty}^{\infty} \hat{f}(s) e^{isx}\, ds,$$

which now closely mimics the Fourier series expression. Other sources "split the difference" and put a factor of $\frac{1}{\sqrt{2\pi}}$ in *both* the definitions of the Fourier transform and its inverse:

$$\hat{f}(s) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} f(x) e^{-ixs}\, dx, \quad (\mathcal{F}^{-1}(\hat{f}))(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \hat{f}(s) e^{ixs}\, ds,$$

so that a $\frac{1}{2\pi}$ appears when composing $\mathcal{F}^{-1}(\mathcal{F}(f))$. It is also common to use a different exponential in the definition of the Fourier and inverse Fourier transform:

$$\hat{f}(s) = \int_{-\infty}^{\infty} f(x) e^{-i2\pi xs}\, dx, \quad (\mathcal{F}^{-1}(\hat{f}))(x) = \int_{-\infty}^{\infty} \hat{f}(s) e^{i2\pi xs}\, ds,$$

which eliminates the need to use any factor of $\frac{1}{2\pi}$ at all. We will stick with the convention we began with, where we don't use a $\frac{1}{2\pi}$ in the Fourier transform and instead use it in the inverse transform.

Frequency/momentum domain. We now give a "physical" interpretation of the Fourier transform of f. To motivate this, again consider the discrete case of a Fourier series:

$$f \rightsquigarrow c_n(f) \rightsquigarrow \sum_{n=-\infty}^{\infty} c_n(f)e^{inx}.$$

Since $e^{inx} = \cos(nx) + i\sin(nx)$, we see that n here is the *frequency* of the corresponding wave, which is the number of "peaks" that occur within a unit time interval, assuming that x represents time. The coefficients $c_n(f)$ can be viewed as a defining a function $\mathbb{Z} \to \mathbb{C}$, which is thus to be interpreted as a function of the frequency variable n and outputs the amplitude $c_n(f)$ of the portion of the wave described the "frequency n" component e^{inx}. So, if f is a periodic function of time, the "transformed" function $\mathbb{Z} \to \mathbb{C}$ describes the same phenomena as f, but as a function of the frequency instead. If x represents position, then it turns out that the frequency can also essentially be interpreted as the momentum of the wave, so $c_n(f)$ is a function of the momentum variable n.

The same interpretation holds for the Fourier transform: if f is a function of time, then $\hat{f}$ is a function of frequency since

$$f(t) = \frac{1}{2\pi}\int_{-\infty}^{\infty} \hat{f}(s)e^{ist}\,ds$$

says that $\hat{f}(s)$ gives the amplitude of the portion e^{ist} of the wave (formed as a "superposition" of a continuous range of simpler waves) occurring at frequency s. We often say that the Fourier transform shifts from the description of the phenomena being studied from the time domain to the frequency domain. If x represents position, then the Fourier transform of f can also be thought of as a function of momentum, so the transform moves us from the position domain to the momentum domain. The overarching point is that a function and its Fourier transform in the end capture the same information, only from different perspectives, and that shifting perspectives might make certain things easier to study, as we'll see.

Uncertainty priciple. Let us make one more observation from the graphs we had previously:

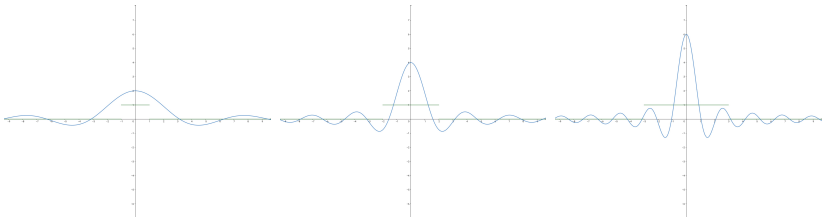

If we interpret the original function as a function of position (or time) and the Fourier transformed function as a function of momentum (or frequency), we see that the better we can isolate the position the worse we can isolate momentum, and vice-versa. In other words, in the first graph we have better knowledge of where position is (between -1 and 1) as compared to the other two, but the momentum is more "spread out" than the other two since the graph of the transform "wider"; in the third graph, now the position is less known (between -3 and 3) than in the other two, but in turn the momentum is more "localized" around the y-axis, so we can better determine its value here (since more of the graph occurs near 0) than we can in the first graph.

The same is apparent in the other graphs we had above, and the upshot is that this is a reflection of what is called the *uncertainty principle* in physics: the closer we are to determining the value

of position (or time), the further we are from determining the value of momentum (or frequency), and vice-versa. There is a restriction on how well we can know both position and momentum simultaneously, and in general we can only pick or the other to compute experimentally.

Transforms and derivatives. Our original motivation in introducing the Fourier transform was in wanting to solve the heat equation over all of $\mathbb{R}$ as opposed to an interval $[0, L]$. We will go ahead and do this in detail next time, but the key tool we need to do so is the relation between Fourier transforms and derivatives. That is, given a (say C^1) function f, we want to know how the Fourier transform of f' relates to the Fourier transform of f. It is this relation which allow us to express the heat equation in the "frequency/momentum" domain, where it will be simpler to solve. Fourier inversion will then allow us to recover our desired solution from the "frequency" solution.

We compute the Fourier transform of f' using integration by parts (!) as follows:

$$\begin{aligned}\widehat{f'}(s) &= \int_{-\infty}^{\infty} f'(t)e^{-ist}\,dt \\ &= f(t)e^{-ist}\Big|_{-\infty}^{\infty} - \int_{-\infty}^{\infty} -isf(t)e^{-ist}\,dt.\end{aligned}$$

The second term is precisely is times the Fourier transform of f! To deal with the first term, we use a fact we have glossed over so far, namely that in order for the improper integral defining the Fourier transform to actually exist as a finite value requires that f have certain properties; that is, it is not true that $\int_{-\infty}^{\infty} f(t)e^{-ist}\,dt$ exists for all functions f, but only for a certain class of functions. To guarantee that this integral exists we will actually assume that f is *absolutely integrable*, which means that

$$\int_{-\infty}^{\infty} |f(t)|\,dt < \infty.$$

We use the absolute value here instead of just f in order to guarantee that the integral is finite, not because of some coincidental "cancellation" that occurs between positive and negative parts, but rather as a result of the overall behavior of the function itself. The main practical consequence for the computation at hand is that if a function is indeed absolutely integrable in this way, it must be true that it decays to zero in both the positive and negative infinite directions:

$$\int_{-\infty}^{\infty} |f(t)|\,dt < \infty \implies \lim_{t\to\pm\infty} f(t) = 0.$$

(Otherwise the "area" under the graph of $|f|$ would have no hope of being finite.) This then implies that the "boundary" term $f(t)e^{-ist}\Big|_{-\infty}^{\infty}$ in the computation above is zero, since "evaluating" at both the upper ∞ and lower $-\infty$ limits give zero. Thus we are left with

$$\widehat{f'}(s) = is\int_{-\infty}^{\infty} f(t)e^{-ist}\,dt = is\hat{f}(s).$$

The conclusion is thus that the Fourier transform of f' is the Fourier transform of f times i times the variable of the transformed function. We summarize this by saying that Fourier transformation turns differentiation into multiplication by (i times) the transform variable. As we will see with the heat equation next time, this gives a way to essentially turns PDEs into ODEs, which will make them easier to solve.

Lecture 18: Yet More Fourier

Warm-Up. We determine the relation between the derivative of $\hat{f}$ (assuming it exists) and $\hat{f}$ itself, and then use this to compute the Fourier transform of the *Gaussian* function $f(x) = e^{-ax^2/2}$, where $a > 0$. (Gaussians are ubiquitous in mathematics, mainly stemming from the fact that they model "normal distributions" and "bell curves".) The Fourier transform of f is

$$\hat{f}(s) = \int_{-\infty}^{\infty} f(x)e^{-ixs}\,dx.$$

To differentiate this respect to s we can "differentiate under the integral sign":

$$(\hat{f})'(s) = \frac{d}{ds}\int_{-\infty}^{\infty} f(x)e^{-ixs}\,dx = \int_{-\infty}^{\infty} \frac{d}{ds}[f(x)e^{-ixs}]\,dx = \int_{-\infty}^{\infty} -ixf(x)e^{-ixs}\,dx.$$

But the right side is $-i$ times the Fourier transform of $xf(x)$, so

$$(\mathcal{F}f)'(s) = -i\int_{-\infty}^{\infty} xf(x)e^{-ixs}\,dx = -i\mathcal{F}(xf(x)).$$

Thus the conclusion is that the derivative of a Fourier transform is $-i$ times the Fourier transform of the original function multiplied by the variable in question. Said another way, this gives

$$i(\mathcal{F}f)'(s) = \mathcal{F}(xf(x)),$$

so Fourier transforms send "multiplication by x" to i times differentiation. This is a counterpart to the relation $(\mathcal{F}f')(s) = is(\mathcal{F}f)(s)$ we derived last time.

Now, we use this to compute the Fourier transform of $f(x) = e^{-ax^2/2}$, where $a > 0$. This Fourier transform is defined by

$$(\mathcal{F}(e^{-ax^2/2}))(s) = \int_{-\infty}^{\infty} e^{-ax^2/2}e^{-ixs}\,dx = \int_{-\infty}^{\infty} e^{-(ax^2/2+isx)}\,dx,$$

and we can ask why we can't just compute this integral directly? The answer is that we can, but doing so requires some *complex analysis*, which in this course we do not have access to. We will come back to this type of integral later to give *some* sense as to what's required, and to give a hint of some of the things you'll see in MATH 382, the complex analysis ISP course, but in this course we must find another way to compute this transform.

The key is that $f(x) = e^{-ax^2/2}$ satisfies the following ODE:

$$f'(x) = -axf(x),$$

which you can check by differentiating $f(x)$. We now take the Fourier transform of both sides of *this* ODE, and use the relations between Fourier transforms and derivatives determined above and last time:

$$(\mathcal{F}f') = -a(\mathcal{F}(xf(x))) \rightsquigarrow is\hat{f}(s) = -ia(\hat{f})'(s).$$

(We used the fact the Fourier transform operator $\mathcal{F}$ is in fact *linear*, so that in particular it preserves scalar multiplication; this was needed in order to say that $\mathcal{F}(axf(x))$ is the same as $a\mathcal{F}(xf(x))$. This can be seen from the integral definition directly, since integration is itself a linear operation.) This says that the Fourier transform $\hat{f}$ we want satisfies the ODE

$$(\hat{f})'(s) = -\frac{s}{a}\hat{f}(s),$$

and solving gives

$$\hat{f}(s) = Ae^{-s^2/2a}$$

for some constant A. This constant is the value of the Fourier transform at $s = 0$, which we can compute directly:

$$\hat{f}(0) = \int_{-\infty}^{\infty} f(x)e^{i0x}\,dx = \int_{-\infty}^{\infty} e^{-ax^2/2}\,dx = \sqrt{\frac{2}{a}}\int_{-\infty}^{\infty} e^{-u^2}\,du,$$

where in the last step we made the change of variables $u = \sqrt{\frac{a}{2}}x$. This final integral is possibly one you saw in a previous course as an application of using polar coordinates in double integrals (ask me if you've never seen this computation before!), and its value is $\sqrt{\pi}$. Thus we get

$$\hat{f}(0) = \sqrt{\frac{2}{a}}\int_{-\infty}^{\infty} e^{-u^2}\,du = \sqrt{\frac{2\pi}{a}}.$$

Hence the Fourier transform of $f(x) = e^{-ax^2/2}$ is $\hat{f}(s) = \sqrt{\frac{2\pi}{a}}e^{-s^2/2a}$. The Fourier transform of a Gaussian is another Gaussian!

Back to linear algebra. Note what happens above in the particular case where $a = 1$: we get that the Fourier transform of $e^{-x^2/2}$ is

$$\mathcal{F}(e^{-x^2/2}) = \sqrt{2\pi}e^{-x^2/2}.$$

(I'm using x as the transform variable on the right instead of s to make the relation between the two sides clearer.) This says that the Fourier transform of $e^{-x^2/2}$ is a scalar multiple of $e^{-x^2/2}$, so that $e^{-x^2/2}$ is in fact an *eigenvector* (!!!) for the Fourier transform operator! In this context, $e^{-x^2/2}$ is called an *eigenfunction* of $\mathcal{F}$, which means an eigenvector in a vector space of functions.

Let us now provide the correct linear-algebraic context from which to view $\mathcal{F}$. For a space X, we define $L^1(X)$ to be the space of *absolutely integrable* complex-valued functions on X, meaning functions $f : X \to \mathbb{C}$ for which the integral of $|f|$ exists and is finite:

$$L^1(X) = \left\{ f : X \to \mathbb{C} \;\middle|\; \int_X |f(x)|\,dx < \infty \right\}.$$

(The superscript ^{1} relates to the fact that we are integrating the first power of $|f(x)|$. Soon we will introduce $L^2(X)$—an example of a *Hilbert space*—which consists of functions for which $f(x)^2$ gives a finite integral value.) The space $L^1(X)$ is a vector space, since adding or scalar multiplying absolutely integrable functions produces an absolutely integrable function. With this notation, the operation sending an L^1 function on $\mathbb{R}$ to its Fourier transform is then a linear transformation

$$\mathcal{F} : L^1(\mathbb{R}) \to L^1(\mathbb{R}),$$

and it is then this linear operator of which $e^{-x^2/2}$ is an eigenfunction. (Actually, this is not quite correct as stated. The issue is that it is not true that the Fourier transform of an L^1 function is necessarily L^1 itself, so $\mathcal{F}$ does not actually have codomain $L^1(\mathbb{R})$ has written above. Either we have to enlarge the codomain to allow for non-absolutely integrable functions, *or* we have to shrink the domain and consider only those functions f for which $\hat{f}$ is indeed L^1. We will sweep this all under the rug in this course, but a graduate course in functional analysis like MATH 410-2 would clear this all up if you were so inclined!)

Even better, the discrete version of the Fourier transform we defined last time, sending a periodic function to its sequence of complex Fourier coefficients, fits into this same exact framework! A key observation is that a periodic function can be viewed as a function on a circle: for f defined on $[-L, L]$ with $f(-L) = f(L)$, "gluing" the two endpoints of $[-L, L]$ produces a "circle", and $f(-L) = f(L)$ guarantees that f is perfectly well-defined on this circle. Thus, the Fourier transform operator in this case looks like

$$\mathcal{F} : L^1(\text{circle}) \to L^1(\mathbb{Z}), \text{ sending } f \text{ to } (c_n(f)).$$

(Again, here we view the sequence $(c_n(f))$ as defining a function $\mathbb{Z} \to \mathbb{C}$.) As we highlighted last time, Fourier series coefficients and Fourier transforms play the exact same role in all of this, where the only difference is the type of space—$\mathbb{R}$ or a circle—we use. (One can ask why it is that for $\mathbb{R}$ we use both $L^1(\mathbb{R})$ as the domain and codomain of the Fourier transform operator, whereas for the circle we use domain $L^1(\text{circle})$ and codomain $L^1(\mathbb{Z})$. The answer has to do with the concept of what's called *Pontryagin duality*, which is unfortunately way beyond the scope of this course. In short, there is a notion of "Pontryagin dual" for certain types of spaces, and $\mathbb{R}$ is its own Pontryagin dual where as the Pontryagin dual of a circle is $\mathbb{Z}$. The concept of a Fourier transform then generalizes to the setting of such duals. This all belongs to the subject of "harmonic analysis on locally compact groups", whatever that means.)

Integral kernels. As a brief aside, let us say one more thing in terms of context. The linear transformation defined by the Fourier transform operator

$$\mathcal{F} : L^1(\mathbb{R}) \to L^1(\mathbb{R})$$

(with the caveat we mentioned previously that technically the codomain is not quite $L^1(\mathbb{R})$) is an example of an *integral transform*, which refers to transforms (i.e. linear transformation) defined via integration. One way to produce examples of integral transforms is to start with a function $K(x, y)$ of two variables, and define a transform $T : L^1(\mathbb{R}) \to L^1(\mathbb{R})$ via

$$f \mapsto Tf \text{ given by } (Tf)(x) = \int_{-\infty}^{\infty} K(x, y) f(y)\, dy.$$

In this setting $K(x, y)$ is called the *kernel* of the integral transform T. (Note that this usage of the word "kernel" is different from the usage in "kernel of a linear transformation" as things which the transformations sends to zero.) In fact, the Fourier transform arises in this way, with kernel function $K(s, x) = e^{-isx}$:

$$\hat{f}(s) = \int_{-\infty}^{\infty} e^{-isx} f(x)\, dx.$$

We previously saw another example of an integral transform, namely the transform that sends a 2π-periodic function f to its N-th order Fourier partial sum:

$$f(x) \mapsto S_N f(x) = \frac{a_0}{2} + \sum_{n=1}^{N} (a_n \cos nx + b_n \sin nx).$$

Here, the kernel function was what we called the N-th order *Dirichlet kernel* D_N, which satisfied

$$(S_N f)(x) = \int_{-\pi}^{\pi} \underbrace{D_N(x - y)}_{K(x,y)} f(y)\, dy.$$

(You can back to a previous lecture to see what this Dirichlet kernel looked like, but the explicit expression is not important here.) In fact, it is (essentially) true that all integral transforms arise in this way, so that every integral transform has an associated kernel. (This is not quite literally true as stated, but can be made true if we are more careful about what type of object we allow to serve as a "kernel"; concretely, we should allow kernels which are *distributions*, which is a concept we will briefly mention in the next few days.)

Now, how does this all relate to the finite-dimensional linear algebra you're already used to? I claim that a kernel $K(x, y)$ in this sense is nothing but the "continuous" analog of a matrix! Indeed, consider ordinary matrix multiplication $A\mathbf{x}$:

$$\begin{bmatrix} a(1,1) & \cdots & a(1,n) \\ \vdots & \ddots & \vdots \\ a(n,1) & \cdots & a(n,n) \end{bmatrix} \begin{bmatrix} x(1) \\ \vdots \\ x(n) \end{bmatrix}.$$

Here we write the entries of the matrix as $a(i, j)$ instead of a_{ij} and the entries of the vector as $x(j)$ instead of x_j to make the relation to $K(x, y)$ and $f(x)$ clearer. The i-th component of the resulting product $A\mathbf{x}$ is

$$(A\mathbf{x})(i) = \sum_{j=1}^{n} a(i, j)x(j).$$

In the case of an integral kernel, we instead have a "continuous" range of values of i and j—which we denote by variables x and y instead—and we can interpret the integral transform definition

$$(Tf)(x) = \int_{-\infty}^{\infty} K(x, y)f(y)\, dy$$

as saying that the "x-th coordinate" of the result Tf is the "sum" of $K(x, y)f(y)$ as y ranges over all possible values, just as what happens in $A\mathbf{x}$ above. So, an integral transform is just the continuous analog of a matrix multiplication, and the kernel function is the continuous analog of a matrix as claimed. The fact that every integral transform has an integral kernel is the continuous analog of the fact that any linear transformation $\mathbb{R}^n \to \mathbb{R}^n$ is induced by a matrix.

Back to heat. After going off on that linear-algebraic tangent, we now return to something more concrete. Our original motivation behind introducing the Fourier transform came from wanting to study the heat equation on $\mathbb{R}$, and so that is what we now begin to do. Our goal is to solve

$$u_t = \alpha^2 u_{xx}, \text{ with initial condition } u(x, 0) = f(x)$$

where now $-\infty < x < \infty$ is no longer constrained. The idea is to take the Fourier transform of both sides *with respect to* x (so, leave t alone), which will end up eliminating the need to use derivatives with respect to x.

Consider the Fourier transform of $u(x, t)$ with respect to the variable x:

$$\hat{u}(y, t) = \int_{-\infty}^{\infty} u(x, t)e^{-iyx}\, dx.$$

According the relation between Fourier transforms and derivatives we have seen before, the Fourier transform of u_x is iy times the Fourier transform of u:

$$\widehat{u_x}(y, t) = iy\hat{u}(y, t).$$

Doing this once more gives that the Fourier transform of u_{xx} picks up another factor of iy:

$$\widehat{u_{xx}}(y,t) = (iy)(iy)u(y,t) = -y^2\hat{u}(y,t).$$

Thus, after taking Fourier transforms, the right side of the heat equation becomes

$$\alpha^2 u_{xx} \rightsquigarrow -\alpha^2 y^2 \hat{u}.$$

As for the left side, we use differentiation under the integral sign to get

$$\widehat{u_t}(y,t) = \int_{-\infty}^{\infty} u_t(x,t)e^{-iyx}\,dx = \frac{\partial}{\partial t}\int_{-\infty}^{\infty} u(x,t)e^{-iyx}\,dx,$$

which is just the partial derivative of $\hat{u}$ with respect to t, so that:

$$\widehat{u_t} = \hat{u}_t.$$

(In other words, since the Fourier transform is being taken with respect to the x variable alone, taking derivatives with respect to t doesn't affect anything, so the Fourier transform of u_t is the t-derivative of the Fourier transform of u.)

Thus altogether, taking the Fourier transform of both sides of the heat equation gives

$$\hat{u}_t = -\alpha^2 y^2 \hat{u}.$$

What we have gained is that the x-derivatives are gone, so that what remains is simply an ODE (not a PDE) for $\hat{u}$! ODEs are simpler to solve than PDEs, so the strategy is to use this ODE to determine $\hat{u}$, which is not hard to do, and then to use Fourier inversion to recover u from $\hat{u}$, which is somewhat hard to do, but we'll go ahead and do it in detail next time. Stay tuned!

Lecture 19: The Heat Kernel

Warm-Up. We find the Fourier transform of a solution $u(x,t)$ of the wave equation $u_{tt} = a^2 u_{xx}$ on $\mathbb{R}$. (Imagine an infinitely long string.) As with the heat equation last time, we take the Fourier transform of $u(x,t)$ with respect to x:

$$\hat{u}(y,t) = \int_{-\infty}^{\infty} u(x,t)e^{-ixy}\,dx.$$

Taking the Fourier transform of both sides of the wave equation gives

$$\hat{u}_{tt} = -a^2 y^2 \hat{u},$$

where for the right side we used the same computation as we did last time for the heat equation:

$$\widehat{u_{xx}} = \widehat{(u_x)_x} = iy\,\widehat{u_x} = (iy)(iy)\hat{u} = -y^2\hat{u}.$$

The resulting ODE $\hat{u}_{tt} + a^2y^2\hat{u} = 0$ for $\hat{u}$ has solution

$$\hat{u}(y,t) = c_1(y)\cos(ayt) + c_2(y)\sin(ayt),$$

where we allow the coefficients to depend on y since when solving an ODE with respect to t alone, "arbitrary constants" can depend on the other variable y. The explicit forms of $c_1(y), c_2(y)$ then

come from imposing initial conditions, which we will do later. Alternatively, we can express the solutions of this ODE in complex form:

$$\hat{u}(y,t) = d_1(y)e^{iayt} + d_2(y)e^{-iayt},$$

where again $d_1(y), d_2(y)$ are arbitrary functions of y which can be determined from initial conditions.

Solving the heat equation. We now return to the heat equation on $\mathbb{R}$, and work through the computation which determines the solution. Recall that we are considering

$$u_t = \alpha^2 u_{xx}, \quad u(x,0) = f(x)$$

where $-\infty < x < \infty$. We saw last time that taking the Fourier transform of both sides with respect to x gives

$$\hat{u}_t = -\alpha^2 y^2\, \hat{u}$$

where y is the transform variable. This is a first-order ODE for $\hat{u}$, whose solution is

$$\hat{u}(y,t) = A(y)e^{-\alpha^2 y^2 t}$$

for some "constant with respect to t" $A(y)$. Now, the coefficient function $A(y)$ should be the value of the transform of u at $t = 0$:

$$\hat{u}(y,0) = A(y).$$

But $\hat{u}(y,0)$ is obtained by taking the Fourier transform of $u(x,0)$, so our given initial condition $u(x,0) = f(x)$ says that that $A(y) = \hat{f}(x)$, so that

$$\hat{u}(y,t) = \hat{f}(y)e^{-\alpha^2 y^2 t}$$

where f is our initial condition.

At this point we have completed the first step of our derivation: we have found an explicit expression for the Fourier transform of the solution $u(x,t)$ we want. The next step is to use an inverse Fourier transform to recover u from $\hat{u}$:

$$u(x,t) = \frac{1}{2\pi}\int_{-\infty}^{\infty} \hat{u}(y,t)\, e^{iyx}\, dy.$$

After substituting the expression we found for $\hat{u}(y,t)$, we see that

$$u(x,t) = \frac{1}{2\pi}\int_{-\infty}^{\infty} \hat{f}(y)e^{-\alpha^2 y^2 t}e^{iyx}\, dy.$$

Our goal now is to express this solution in terms of the original f instead of $\hat{f}$.

Comparison with series methods. Let us take a brief aside to note the relation between the solution we have found so far for the heat equation on $\mathbb{R}$ and the solution we found previously for the heat equation on a bounded interval $[0,\pi]$ with boundary conditions $u(0,t) = 0 = u(\pi,t)$. In this latter case, we saw that the solution is given by

$$u(x,t) = \sum_{n=1}^{\infty} c_n(f)\, e^{-\alpha^2 n^2 t}\sin(nx)$$

where $c_n(f)$ are the Fourier sine coefficients of $u(x,0) = f(x)$. Be expressing

$$\sin(nx) = -\frac{i}{2}(e^{inx} - e^{-inx})$$

in terms of complex exponentials, we see that we can write $u(x,t)$ in the form

$$u(x,t) = \sum_{n=-\infty}^{\infty} d_n(f)\, e^{-\alpha^2 n^2 t} e^{inx}.$$

The point is that this the *same* form that our solution on $\mathbb{R}$ takes:

$$u(x,t) = \frac{1}{2\pi}\int_{-\infty}^{\infty} \hat{f}(y) e^{-\alpha^2 y^2 t} e^{iyx}\, dy,$$

only that the series is the discrete version; the sum over n in $\mathbb{Z}$ becomes an integral over y in $\mathbb{R}$, and the Fourier coefficients $d_n(f)$ become the Fourier transform $\hat{f}(y)$. (The $\frac{1}{2\pi}$ in our solution on $\mathbb{R}$ shows up in the formula for the Fourier coefficients $d_n(f)$ in the discrete case instead.) Thus, it is no accident that we get something like

$$u(x,t) = \frac{1}{2\pi}\int_{-\infty}^{\infty} \hat{f}(y) e^{-\alpha^2 y^2 t} e^{iyx}\, dy,$$

since this is just the continuous version of what we had before.

Computing the inverse transform. We now compute

$$u(x,t) = \frac{1}{2\pi}\int_{-\infty}^{\infty} \hat{f}(y) e^{-\alpha^2 y^2 t} e^{iyx}\, dy$$

more explicitly. First, we can put in the definition of $\hat{f}(y)$:

$$u(x,t) = \frac{1}{2\pi}\int_{-\infty}^{\infty} \left[\int_{-\infty}^{\infty} f(s) e^{-isx}\, ds\right] e^{-\alpha^2 y^2 t} e^{iyx}\, dy.$$

(Note we are taking care when introducing variables of integration: we use s for the integral defining $\hat{f}$ since x is already used elsewhere.) The resulting expression is an iterated double integral, and we can switch the order of integration to write it as

$$\begin{aligned} u(x,t) &= \frac{1}{2\pi}\int_{-\infty}^{\infty} \left[\int_{-\infty}^{\infty} f(s) e^{-isx}\, ds\right] e^{-\alpha^2 y^2 t} e^{iyx}\, dy \\ &= \frac{1}{2\pi}\int_{-\infty}^{\infty} f(s) \left[\int_{-\infty}^{\infty} e^{-isx} e^{-\alpha^2 y^2 t} e^{iyx}\, dy\right] ds, \end{aligned}$$

where we have pulled $f(s)$ out of the y-integral since $f(s)$ does not depend on y. After combining some exponentials we get

$$u(x,t) = \frac{1}{2\pi}\int_{-\infty}^{\infty} f(s) \left[\int_{-\infty}^{\infty} e^{-\alpha^2 y^2 t} e^{-i(s-x)y}\, dy\right] ds$$

But what is the expression

$$\int_{-\infty}^{\infty} e^{-\alpha^2 y^2 t} e^{-i(s-x)y}\, dy$$

in brackets? It is nothing but the Fourier transform of $e^{-\alpha^2 y^2 t}$ only *evaluated* at $s - x$! (Recall $\hat{g}(z) = \int_{-\infty}^{\infty} g(y) e^{-izy}\, dy$, so here we have $g(y) = e^{-\alpha^2 y^2 t}$ and $z = s - x$.) The function $e^{-\alpha^2 t y^2}$ is a Gaussian, so using the general computation we did last time of the transform of a Gaussian:

$$\mathcal{F}(e^{-ay^2/2})(z) = \sqrt{\frac{2\pi}{a}} e^{-z^2/2a}$$

we get (with $a = 2\alpha^2 t$ in our case)

$$\mathcal{F}(e^{-\alpha^2 t y^2})(z) = \sqrt{\frac{2\pi}{2\alpha^2 t}}\, e^{-z^2/2(2\alpha^2 t)} = \sqrt{\frac{\pi}{\alpha^2 t}}\, e^{-z^2/4\alpha^2 t}.$$

Thus, finally, we find that the solution $u(x,t)$ we want is

$$\begin{aligned} u(x,t) &= \frac{1}{2\pi} \int_{-\infty}^{\infty} f(s) \left[\int_{-\infty}^{\infty} e^{-\alpha^2 y^2 t} e^{-i(s-x)y}\, dy \right] ds \\ &= \frac{1}{2\pi} \int_{-\infty}^{\infty} f(s) \sqrt{\frac{\pi}{\alpha^2 t}}\, e^{-(s-x)^2/4\alpha^2 t}\, ds \\ &= \int_{-\infty}^{\infty} f(s) \frac{1}{\sqrt{4\pi\alpha^2 t}}\, e^{-(s-x)^2/4\alpha^2 t}\, ds. \end{aligned}$$

(An absolutely epic computation!)

The heat kernel. Our conclusion is thus that the solution of the heat equation $u_t = \alpha^2 u_{xx}$ on $\mathbb{R}$ with initial condition $u(x,0) = f(x)$ is

$$u(x,t) = \int_{-\infty}^{\infty} f(s) \frac{1}{\sqrt{4\pi\alpha^2 t}}\, e^{-(s-x)^2/4\alpha^2 t}\, ds.$$

(This final form looks quite different than the types of series solutions we had before, but there is a way to directly relate the two, which we will look at later.) Note here that the only dependence on the initial condition in this solution comes from the $f(s)$ factor, and that the exponential factor does not depend on f at all.

The upshot is that there is one single function

$$\frac{1}{\sqrt{4\pi\alpha^2 t}}\, e^{-(s-x)^2/4\alpha^2 t}$$

we can use to construct the solution of the heat equation on $\mathbb{R}$ with any initial condition: simply multiply your initial condition by this particular function and integrate. This function is called the *heat kernel* (more precisely, the one-dimensional heat kernel), and is of fundamental importance in the study of the heat equation. The word "kernel" used here refers to the context we gave previously regarding *integral transforms*, where the heat kernel is something like an infinite-dimensional "continuous" matrix. The linear transform defined by the integral transform corresponding to this kernel is hence the one that sends an initial condition function f to the solution of the heat equation with initial condition f.

If we introduce the function

$$h(z,t) = \frac{1}{\sqrt{4\pi\alpha^2 t}}\, e^{-z^2/4\alpha^2 t},$$

then the heat kernel is $h(x-s,t)$. In this notation, the solution of the heat equation is

$$u(x,t) = \int_{-\infty}^{\infty} f(s)h(x-s,t)\,ds.$$

Integral expressions of this type were introduced on the last discussion problem set, and is called the *convolution* of f with h (in the first coordinate, ignoring t). Thus, solutions of the heat equation are produced by convolving (not sure that this an actual word, but I'm going with it) with h. This $h(z,t)$ is called the *fundamental solution* of the heat equation, since it is the solution from which all others can be obtained, which is similar-in-spirit to how we used "fundamental solution" with the Fourier series approach we considered previously.

But, if $h(z,t)$ is in fact to be a "fundamental solution" of the heat equation, then what initial condition does it actually satisfy? It is straightforward enough (but a bit tedious) to verify that

$$h(z,t) = \frac{1}{\sqrt{4\pi\alpha^2 t}}\, e^{-z^2/4\alpha^2 t}$$

does satisfy $u_t = \alpha^2 u_{zz}$, but an attempt to find the appropriate initial condition $u(x,0)$ by plugging in $t=0$ fails since $h(z,t)$ is undefined at $t=0$! Thus, we have to be more careful about what we actually mean by "initial condition" in this case; the answer will come from the notion of a *delta function*, which we'll talk about next time.

Lecture 20: Delta Functions

Warm-Up. We find the solution of the PDE

$$u_t = u_{xx} - u_x,\ u(x,0) = f(x)$$

on $\mathbb{R}$. This is known as a *diffusion-convection* equation, where the u_x term indicates a contribution to the change in u coming from convection. (Think the heat in some liquid moving around, or a chemical which "diffuses" when placed in a liquid.) In general we can allow coefficients in front of the convection term, and the heat equation is the special case where this coefficient is zero. As in our solution of the heat equation, we can take the Fourier transform of both sides with respect to x to get

$$\hat{u}_t = -y^2\hat{u} - iy\hat{u},\ \hat{u}(y,0) = \hat{f}(y).$$

The ODE $\hat{u}_t = (-y^2 - iy)\hat{u}$ has solution

$$\hat{u}(y,t) = A(y)e^{(-y^2-iy)t},$$

and $\hat{u}(y,0) = \hat{f}(y)$ gives $A(y) = \hat{f}(y)$, so

$$\hat{u}(y,t) = \hat{f}(y)e^{(-y^2-iy)t}.$$

Now, to determine $u(x,t)$ we can do what we did for the heat equation last time: use an inverse transform to express $u(x,t)$ in terms of an integral involving $\hat{u}(y,t)$, and then compute the resulting integral. However, we can save some work here by using the exact result of the computation we did last time. Note that we can rewrite $\hat{u}(y,t)$ above as

$$\hat{u}(y,t) = \hat{f}(y)e^{-iyt}e^{-y^2t}.$$

Here, the factor $\hat{f}(y)e^{-iyt}$ is precisely the Fourier transform of the "shift" $f(x-t)$, as was computed on a recent homework problem. That is, if we denote this shift by $g(x) = f(x-t)$, then

$$\hat{g}(y) = e^{-iyt}\hat{f}(y),$$

so that

$$\hat{u}(y,t) = \hat{g}(y)e^{-y^2t}.$$

The point is that *this* expression is precisely the one we would have obtained by considering the heat equation $u_t = u_{xx}$ with initial condition $u(x,0) = g(x)$, so the explicit form of $u(x,t)$ is exactly the one given by the heat kernel integral and initial function $g(x)$:

$$u(x,t) = \int_{-\infty}^{\infty} g(s)\frac{1}{\sqrt{4\pi t}}e^{-(s-x)^2/4t}\,ds = \int_{-\infty}^{\infty} f(s-t)\frac{1}{\sqrt{4\pi t}}e^{-(s-x)^2/4t}\,ds.$$

Thus, by exploiting the heat kernel computation from last time, we can save a bunch of work this time. (This idea works for more general diffusion-convection equations as well.)

Finally, we make a change of variables to replace the integral involving the shifted f with one involving the unshifted f. For $p = s - t$, we get

$$u(x,t) = \int_{-\infty}^{\infty} f(s-t)\frac{1}{\sqrt{4\pi t}}e^{-(s-x)^2/4t}\,ds = \int_{-\infty}^{\infty} f(p)\frac{1}{\sqrt{4\pi t}}e^{-(p+t-x)^2/4t}\,dp$$

as our final solution. (The particular PDE $u_t = u_{xx} - u_x$ thus has "fundamental solution" $h(x,t) = \frac{1}{\sqrt{4\pi t}}e^{-(t-x)^2/4t}$, from which all other solutions can be derived by convolution.)

The heat kernel at time 0. We now come back to determining the initial condition satisfied by the heat equation's fundamental solution

$$h(z,t) = \frac{1}{\sqrt{4\pi\alpha^2 t}}\,e^{-z^2/4\alpha^2 t}.$$

As pointed out last time, evaluating this function at $t = 0$ gives an undefined expression, so we can't do that. (Or can we?) To get a handle on what is happening, consider instead the limit of $h(z,t)$ as $t \to 0$. Here are some plots for values of t approaching 0:

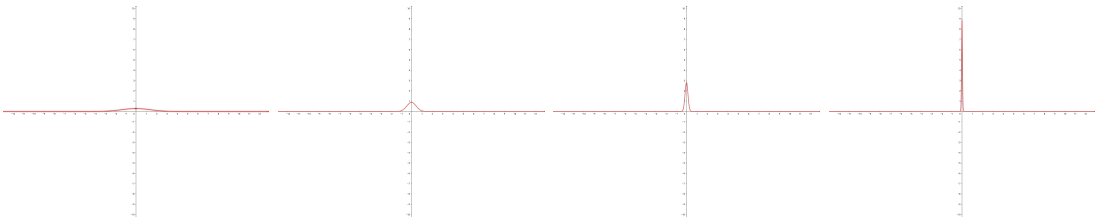

(Note that we if view these pictures in reverse, so as t increases, we get the diffusive type of behavior we expect of a solution of the heat equation: the bump in the middle gets smoothed out as time increases, due to the high temperature in the middle moving towards areas of low temperature.)

The key observation is that as $t \to 0$, the bump gets concentrated more and more near $x = 0$, and grows in magnitude more and more, and regions away from 0 get flattened out towards zero more and more. In some sense, the limit as $t \to 0$ should be "function" which is zero for $x \neq 0$ and infinitely large at $x = 0$. Of course, it doesn't make sense to say that the value of the limit at

$x = 0$ should be ∞ (or does it?), but the situation is even stranger: it turns out that for any $t > 0$, the integral of function

$$\frac{1}{\sqrt{4\pi\alpha^2 t}}\, e^{-z^2/4\alpha^2 t}$$

over the entire real number line is actually 1! The limit of the constant 1 as $t \to 0$ is still 1, so we might expect that the "function" we get by taking $t \to 0$ in $h(z,t)$ should be one whose integral over $\mathbb{R}$ is in fact 1. Thus, altogether, the initial condition we want (obtained by taking $t \to 0$ in $h(z,t)$) should be something which is zero for $x \neq 0$, ∞ for $x = 0$, and yet has integral 1 over $\mathbb{R}$. There is absolutely no actual function that has all these properties, so it seems like we are just talking nonsense.

Fourier transform of a constant. But we are not talking nonsense after all. The non-existent "function" with the properties above *does* exist, and is called the *Dirac delta function*. Now, the first thing you learn about the Dirac delta function is that it is *not* actually a function, since, after all, as we said above no literal function having the desired properties exists. So, if the Dirac delta function is not a function, then what is it?

Before answering this, we will look at a different problem first, namely the problem of determining the Fourier transform of the constant function 1. By definition, this Fourier transform is defined as

$$\hat{1}(s) = \int_{-\infty}^{\infty} 1 \cdot e^{-isx}\, dx = \int_{-\infty}^{\infty} e^{-isx}\, dx.$$

The trouble is that this integral does not exist! As x ranges over all of $\mathbb{R}$, the complex numbers e^{-isx} move around the unit circle in the complex plane, and no "convergence" occurs. Or, said another way, if we write this integral in terms of real and imaginary parts:

$$\int_{-\infty}^{\infty} e^{-isx}\, dx = \int_{-\infty}^{\infty} \cos(sx)\, dx + i \int_{-\infty}^{\infty} \sin(sx)\, dx,$$

we get two divergent integrals. (Note that the constant function 1 is not in $L^1(\mathbb{R})$.) So, the constant function of 1 does not have a Fourier transform in the usual sense.

But, we can start trying to give some meaning to $\hat{1}$ regardless by considering 1 to be a limit of certain functions we considered previously: the "rectangular pulse" functions

$$f_T(x) = \begin{cases} 1 & -T \leq x \leq T \\ 0 & \text{otherwise.} \end{cases}$$

As $T \to \infty$, in some sense the functions $f_T(x)$ do "approach" the constant function 1, since the interval over which the value of $f_T(x)$ is 1 gets larger and larger, eventually becoming all of $\mathbb{R}$ in the limit. We computed the Fourier transform of this pulse function earlier:

$$\hat{f_T}(s) = \frac{2\sin sT}{s},$$

so we can try to essentially *define* the Fourier transform of the constant 1 to be the limit of *these* transformed pulse functions as $T \to \infty$:

$$\hat{1}(s) \stackrel{?}{=} \lim_{T\to\infty} \hat{f_T}(s) = \lim_{T\to\infty} \frac{2\sin sT}{s}.$$

The same problem from before still remains however: the limit on the right does not exist.

Nevertheless, we can give this limit a precise meaning as follows. First we note the following fact: for any $T > 0$ we have

$$\int_{-\infty}^{\infty} \frac{2\sin sT}{s}\, ds = 2\pi,$$

which comes from making a change of variables $x = sT$ in the integral

$$\int_{-\infty}^{\infty} \frac{\sin x}{x}\, dx = \pi.$$

(The function $\frac{\sin x}{x}$ is called the *sinc* function, and this particular integral can be computed using methods from complex analysis. We'll say something about this in a bit to give a flavor of the types of things you would see in a complex analysis course like MATH 382.) Now, for a function $f(s)$, consider the limit

$$\lim_{T\to\infty} \int_{-\infty}^{\infty} f(s)\, \frac{2\sin(sT)}{s}\, dx.$$

It turns out that this limit *does* exist, and it has the value $2\pi f(0)$! Proving this formally is beyond the scope of this course, but we can give some pretty good intuition as follows. As T increases, the functions $\frac{2\sin(sT)}{s}$ look like:

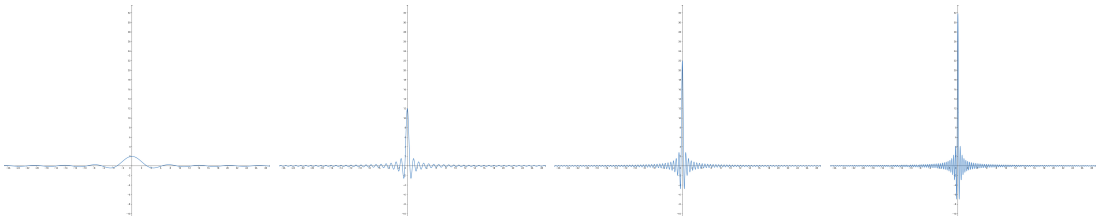

The point is that the values are becoming concentrated more and more near $s = 0$, with points away from zero becoming more and more negligible. (There are still "peaks" away from zero, but these get thinner and thinner, hence why they become more and more negligible too.) The net areas under all of these graphs are 2π as stated above, so no area disappears as T increases, it just gets more and more concentrated. Thus the product $f(s)\,\frac{2\sin(sT)}{s}$ also gets more and more concentrated near $s = 0$ as T increases, with parts of f away from zero being "zeroed" out, so that in the limit the *only* contribution to

$$\lim_{T\to\infty} \int_{-\infty}^{\infty} f(s)\, \frac{2\sin(sT)}{s}\, dx$$

comes from the value of $f(0)$, with the net area of 2π under the graph of $\frac{2\sin sT}{s}$ preserved throughout the limit. Again, this is not a proof, but should hopefully make it plausible that

$$\lim_{T\to\infty} \int_{-\infty}^{\infty} f(s)\, \frac{2\sin(sT)}{s}\, dx = 2\pi f(0).$$

Dirac delta function. Now working even more informally, we exchange limit and integration above to get

$$\int_{-\infty}^{\infty} f(x) \left(\lim_{T\to\infty} \frac{2\sin sT}{s} \right) ds = 2\pi f(0).$$

As stated before, $\lim_{T\to\infty} \frac{2\sin sT}{s}$ does not exist as a function, but if it did, this equation says that it should be a "function" $2\pi\delta(x)$ with the property that

$$\int_{-\infty}^{\infty} f(x)[2\pi\delta(x)]\, dx = 2\pi f(0).$$

The notation $\delta(x)$ denotes what's called the *Dirac delta function*, and this equation is precisely how we define it: $\delta(x)$ is the object defined by requiring that

$$\int_{-\infty}^{\infty} f(x)\delta(x)\,dx = f(0).$$

(We incorporated 2π into $2\pi\delta(x)$ when describing the limit of $\frac{2\sin sT}{s}$ above precisely so that it would not show up in this defining expression.)

It is crucial to note, once again, that $\delta(x)$ is not an actual function here, but rather is merely a symbol which acquires meaning only via the requirement that integrating $f(x)\delta(x)$ produces $f(0)$. But, it is natural to ask: if $\delta(x)$ is not a function, what actually *is* it? The answer is that it is what's called a *distribution*. (Distributions are often called *generalized functions* in physics.) A distribution is nothing but a linear transformation

$$L^1(\mathbb{R}) \to \mathbb{R}.$$

(Linear transformation from a vector space to the scalar field are also called linear *functionals*, and their study constitutes the subject of *functional analysis*.) The Dirac delta function is thus the distribution defined by sending $f(x)$ to $f(0)$. But, in what sense is this a "generalized function" and why do use integral notation to denote its effect on $f(x)$, as we did above? The answer is that *actual* functions do indeed give rise to distributions! Indeed, to an honest function $g(x)$ (say L^1) we can associate the distribution $L^1(\mathbb{R}) \to \mathbb{R}$ defined by

$$f(x) \mapsto \int_{-\infty}^{\infty} f(x)g(x)\,dx.$$

The right side is simply a number in $\mathbb{R}$, and this operation is linear since integration is linear, so this does give a distribution. Because of this, we thus often use integration notation in general to denote the effect of applying a distribution to a given function $f(x)$, as in the definition of $\delta(x)$.

Back to fundamental solutions. So, using the notion of the Dirac delta function (a distribution, not a function), we can give meaning to various things we considered leading up to this point. For example, we say that

$$\lim_{T\to\infty} \frac{2\sin sT}{s} = 2\pi\delta(x),$$

not as an equality among functions, but as an equality among distributions instead; and we say that the Fourier transform of 1 is $2\pi\delta(x)$, or in other words

$$\int_{-\infty}^{\infty} e^{-isx}\,dx = 2\pi\delta(x),$$

again not as an equality of functions but as an equality of distributions. (So, for example, something like

$$\int_{-\infty}^{\infty} f(x)\left[\int_{-\infty}^{\infty} e^{-isp}\,dp\right]\,dx,$$

which equals $2\pi\int_{-\infty}^{\infty} f(x)\delta(x)\,dx$, is, by definition, $2\pi f(0)$.)

Moreover, and finally, the "value" of

$$h(z,t) = \frac{1}{\sqrt{4\pi\alpha^2 t}} e^{-z^2/4\alpha^2 t}$$

at $t = 0$ is $\delta(x)$, as a distribution. This is thus the initial condition we wanted, so that this $h(z,t)$—which gives rise to the heat kernel—is the solution of the heat initial value problem

$$u_t = \alpha^2 u_{xx},\ u(x,0) = \delta(x).$$

This initial condition models the setup where we have the entirety of our initial heat—infinite in value—concentrated all at $x = 0$, which then "diffuses" (quickly!) into the fundamental solution $h(z,t)$ above as t increases. (Good stuff!)

The integral of sinc. We will do more with the delta function next time, but we finish for now with saying something about the integral of the sinc function:

$$\int_{-\infty}^{\infty} \frac{\sin x}{x}\, dx = \pi.$$

This is an integral which cannot be easily computed using real-analytic methods alone, but is precisely the type of thing which complex analysis makes doable. (Note that this integral involves no complex numbers at all, so it may seem surprising that introducing complex numbers actually makes things simpler! This is a phenomenon which occurs repeatedly in complex analysis.)

The key is to view $\sin x$ as the imaginary part of e^{ix}, so that the integral we want is nothing but the imaginary part of

$$\int_{-\infty}^{\infty} \frac{e^{ix}}{x}\, dx.$$

To compute this, we then consider the *complex* function $\frac{e^{iz}}{z}$ and integrate it over the following curve $\Gamma_{R,\epsilon}$ in the complex plane:

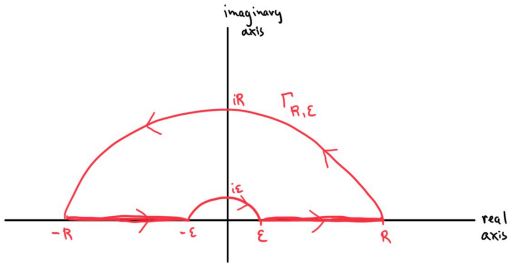

(We will not define what it means to integrate a complex function over a curve in the complex plane here, but in some sense, it is somehow analogous to integrating a vector field over a curve.) This curves follows the upper-half of a circle of radius R, then a segment along the real axis from $-R$ to $-\epsilon$, then the upper-half a of circle of radius ϵ, and finally a segment on the real axis from ϵ to R. (The small circle in the middle is there to make sure we avoid the singularity of $\frac{e^{iz}}{z}$ at $z = 0$.) If we take the limit of the integral

$$\int_{\Gamma_{R,\epsilon}} \frac{e^{iz}}{z}\, dz$$

as $R \to \infty$ and $\epsilon \to 0$, the bounds on the x-axis approach those on the integral

$$\int_{-\infty}^{\infty} \frac{e^{ix}}{x}\, dx$$

we want, so the idea is that we can extract this integral we want by indeed taking the limit of the complex integral above.

By some crucial theorem in complex analysis (what's called *Cauchy's theorem*, which is in some ways analogous to *Green's theorem* for vector fields), the integral of $\frac{e^{iz}}{z}$ over $\Gamma_{R,\epsilon}$ is zero, essentially because the singularity of this function does not occur within the region bounded by this curve. But the integral over the entire curves splits up into four pieces: the integral over the larger semi-circle, the integral over the smaller semi-circle, and integrals over each segment on the x-axis. Thus we get

$$0 = \int_{\Gamma_{R,\epsilon}} \frac{e^{iz}}{z}\,dz = \int_{\text{large circle}} \frac{e^{iz}}{z}\,dz + \int_{\text{small circle}} \frac{e^{iz}}{z}\,dz + \int_{\text{segment 1}} \frac{e^{iz}}{z}\,dz + \int_{\text{segment 2}} \frac{e^{iz}}{z}\,dz.$$

The integrals over the two segments occur only along the real axis, so after rearranging we get

$$\int_{-R}^{-\epsilon} \frac{e^{ix}}{x}\,dx + \int_{\epsilon}^{R} \frac{e^{ix}}{x}\,dx = -\int_{\text{large circle}} \frac{e^{iz}}{z}\,dz - \int_{\text{small circle}} \frac{e^{iz}}{z}\,dz.$$

In the limit as $R \to \infty$ and $\epsilon \to 0$, the left side becomes precisely the integral we want

$$\int_{-\infty}^{\infty} \frac{e^{ix}}{x}\,dx,$$

while the integral over the larger circle turns out to go to zero (this is not obvious!) and the integral over the smaller circle turns out to approach $-i\pi$ (also not obvious!). Thus we end up with

$$\int_{-\infty}^{\infty} \frac{e^{ix}}{x}\,dx = i\pi,$$

and taking imaginary parts gives the sinc integral we wanted:

$$\int_{-\infty}^{\infty} \frac{\sin x}{x}\,dx = \pi.$$

Hooray! (Again, you need a complex analysis course like MATH 382 to see what is really going on here, namely in terms of how the complex integral is actually defined and why the limits of the integrals over the two circles have the values we said they do. Using complex analysis to compute many of the types of integrals which show up in Fourier analysis is one of the main practical reasons why MATH 382 is required for ISP—that and understanding more properties of harmonic functions, i.e. solutions of the Laplace equation, in terms of their complex-analytic interpretation as the real and imaginary parts of complex-differentiable functions.)

Lecture 21: Fundamental Solutions

Warm-Up. We finally justify the Fourier inversion formula:

$$f(x) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \hat{f}(p)\, e^{ipx}\,dp,$$

assuming that f is C^1. (If f is only piecewise C^1, the left side should be $\frac{1}{2}(f(x^-) + f(x^+))$. Technically we should also assume f and $\hat{f}$ are in $L^1(\mathbb{R})$.) This is now simply an exercise in working with the delta function. First we can put in the definition of $\hat{f}$:

$$\frac{1}{2\pi} \int_{-\infty}^{\infty} \left[\int_{-\infty}^{\infty} f(s) e^{-isp}\,ds \right] e^{ipx}\,dp.$$

After switching the order of integration in this double integral and rearranging terms we get

$$\frac{1}{2\pi}\int_{-\infty}^{\infty} f(s)\left[\int_{-\infty}^{\infty} e^{-i(s-x)p}\,dp\right]\,ds.$$

Now, the inner integral

$$\int_{-\infty}^{\infty} e^{-i(s-x)p}\,dp$$

is the Fourier transform of 1 evaluated at $s-x$, and as we argued last time this can be given meaning as (2π times) the Dirac delta function at $s-x$:

$$\int_{-\infty}^{\infty} e^{-i(s-x)p}\,dp = 2\pi\delta(s-x).$$

Thus the double integral above becomes

$$\frac{1}{2\pi}\int_{-\infty}^{\infty} f(s)\left[\int_{-\infty}^{\infty} e^{-i(s-x)p}\,dp\right]\,ds = \frac{1}{2\pi}\int_{-\infty}^{\infty} f(s)[2\pi\delta(s-x)]\,ds.$$

The 2π's cancel, and by the way in which the Dirac delta function behaves or is to be interpreted, this final integral should pick out the value of $f(s)$ at the value of s where the input $s-x$ into the delta function is zero, meaning at $s=x$:

$$\int_{-\infty}^{\infty} f(s)\delta(s-x)\,ds = f(x).$$

Thus we get

$$\frac{1}{2\pi}\int_{-\infty}^{\infty} \hat{f}(p)\,e^{ipx}\,dp = f(x)$$

as the Fourier inversion formula claims.

More on delta. The fact we used above that

$$\int_{-\infty}^{\infty} f(s)\delta(s-x)\,ds = f(x)$$

is, as said, essentially just the *definition* of what $\delta(s-x)$ means: it is the "distribution" defining the linear transformation which sends $f(s)$ to $f(x)$. This type of identity, and its use above, might seem quite magical, and indeed in some ways it is. But, we emphasize again (as we described last time when "constructing" the delta function) that what this identity *really* means is that the following limit identity holds:

$$\lim_{T\to\infty}\int_{-\infty}^{\infty} f(s)\left[\frac{2\sin((s-x)T)}{s-x}\right]\,ds = f(x).$$

(Last time we only wrote this down in the case where $x=0$, but the same reasoning we outlined there applies here simply by "shitfting" the center.) It is *this* identity that underlies all computations we do with the delta function, and in particular provides the *formal* approach to the Fourier inversion formula. Working with the delta function as we did above is simply a convenient way to work with this limit without having to mention it directly.

Using the notion of convolution we gave on a previous discussion sheet, we can recast

$$\int_{-\infty}^{\infty} f(s)\delta(s-x)\,ds = f(x)$$

as the statement that $f * \delta = f$. It is also true that $\delta * f = f$, so this says that δ is something like an "identity" for the convolution operation. If we consider the sequence of functions defined by

$$\delta_n(s) = \frac{\sin(sn)}{\pi s},$$

then the statement that $f * \delta = f$ is really the statement that

$$\lim_{n\to\infty} f * \delta_n = f.$$

Such a sequence of functions is called an *approximate identity*, since it becomes an "identity" in the limit. The existence of such approximate identities is the key to making this all rigorous, and are crucial component of the study of functional analysis. (Incidentally, the sequence of Dirchlet kernels D_N also form an approximate identity, which is essentially the reason why the Fourier convergence theorem for Fourier series holds.)

Fundamental solutions of PDEs. Let us wrap up our discussion of PDEs on $\mathbb{R}$ and using the Fourier transform to solve them by elaborating on the notion of a fundamental solution. To be clear, a *fundamental solution* of a (linear) PDE is a solution satisfying an initial or boundary condition given by a delta function. The point is that from such a solution we can construct *every* solution via convolution. (A fundamental solution is also called a *Green's function* for the PDE.)

We have already seen the fundamental solution

$$G(x,t) = \frac{1}{\sqrt{4\pi\alpha^2 t}} e^{-x^2/4\alpha^2 t}$$

for the heat equation $u_t = \alpha^2 u_{xx}$ on $\mathbb{R}$, which satisfies the initial condition $u(x,0) = \delta(x)$. To be clear, the fact that this is indeed the appropriate initial condition for this solution is just the statement that

$$\int_{-\infty}^{\infty} \delta(s)G(x-s,t)\,ds = G(x,t).$$

(Here, $G(x-s,t)$ is the heat kernel, which via our derivation of it turns the integral on the left into the function satisfying the heat equation with particular initial function given by the function we are multiplying by the heat kernel in this integral, which is the delta function in this.)

Let us check that we do indeed get every other solution by convolving with this one. Given the initial condition $u(x,0) = f(x)$, the claim is that the convolution $f(x) * G(x,t)$ (with respect to the x-coordinate) should given the solution of the heat equation satisfying this particular initial condition. The convolution is

$$f(x) * G(x,t) = \int_{-\infty}^{\infty} f(s)G(x-s,t)\,ds.$$

Evaluating at $t = 0$ and using the initial condition for $G(x,t)$ gives

$$f(x) * G(x,0) = \int_{-\infty}^{\infty} f(s)G(x-s,0)\,ds$$

$$= \int_{-\infty}^{\infty} f(s)\delta(x-s)\,ds$$
$$= f(x),$$

so that this convolution really does satisfy $u(x,0) = f(x)$.

Fundamental wave solution. Now we consider the fundamental solution of the wave equation $u_{tt} = a^2 u_{xx}$ satisfying the initial conditions

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

(The initial condition $u(x,0) = 0, u_t(x,0) = g(x)$ gives rise to its own fundamental solution, and the fundamental solution for the general initial conditions $u(x,0) = f(x), u_t(x,0) = g(x)$ will be a sum of these two.) Actually, you worked out a recent discussion sheet, using Fourier transform methods, that the solution of the initial value problem above is

$$u(x,t) = \frac{1}{2}[f(x-at) + f(x+at)].$$

This is just a generalization to $\mathbb{R}$ of something we saw previously on $[0,L]$ using Fourier series methods: the wave solution with given initial position and zero velocity is formed by taking the average of a rightward moving wave with a leftward moving wave.

We claim that on $\mathbb{R}$ the fundamental solution can be taken to be

$$G(x,t) = \frac{1}{2}[\delta(x-at) + \delta(x+at)],$$

which does satisfy the appropriate initial condition: $G(x,0) = \delta(x)$. Let us convolve G with any f:

$$\begin{aligned} f(x) * G(x,t) &= \int_{-\infty}^{\infty} f(s)G(x-s,t)\,ds \\ &= \frac{1}{2}\int_{-\infty}^{\infty} f(s)\delta(x-s-at)\,ds + \frac{1}{2}\int_{-\infty}^{\infty} f(s)\delta(x-s+at)\,ds \\ &= \frac{1}{2}f(x-at) + \frac{1}{2}f(x+at) = \frac{1}{2}[f(x-at) + f(x+at)]. \end{aligned}$$

This thus agrees with the solution found in discussion.

Laplace fundamental solution. As a final example, let us describe a certain type of fundamental solution for the Laplace equation. We consider the "radially symmetric" Laplace equation, where we consider only functions that are independent of the polar variable θ. Our functions then can be thought of as depending on the radial variable r alone, so the fundamental solution we want is the one satisfying

$$\Delta u = \frac{\delta(r)}{2\pi}.$$

(We will not explain here why the delta function takes the form $\frac{1}{2\pi}\delta(r)$ in polar coordinates—it comes from a change of variables computation. Also, since $\delta(r)$ "equals" 0 for $r \neq 0$, this fundamental solution will be a harmonic function away from the origin.) In polar form this equation is

$$u_{rr} + \frac{1}{r}u_r + \frac{1}{r^2}u_{\theta\theta} = \frac{\delta(r)}{2\pi},$$

but since u is assumed to be independent of θ in this case, this simplifies to

$$u_{rr} + \frac{1}{r}u_r = \frac{\delta(r)}{2\pi}.$$

The left side can be written as $\frac{1}{r}(ru_r)'$ (where the prime denotes a derivative with respect to r), and

$$\frac{1}{r}(ru_r)' = \frac{\delta(r)}{2\pi}$$

and then be solved for u by integrating a couple of times. The answer you get is

$$u(r) = \frac{1}{2\pi}\ln r, \text{ or in other words } u(x,y) = \frac{1}{2\pi}\ln\sqrt{x^2+y^2}.$$

This is the fundamental solution of the radially-symmetric Laplace equation.

Then given the PDE $\Delta u = f$ (this is called *Poisson's equation* when $f \neq 0$), the solution is then obtained by convolving (in two variables) with the fundamental solution above:

$$u(x,y) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} f(s,t)\left[\frac{1}{2\pi}\ln\sqrt{(x-s)^2+(y-t)^2}\right] ds\, dt.$$

Verifying that this does satisfy Poisson's equation requires working with *derivatives* of delta functions, which is indeed a thing (!), but beyond the scope of this course, so we leave it at that.

The heat kernel via Fourier series. Finally, we describe how the heat kernel on $\mathbb{R}$ *can* actually be obtained via Fourier series methods, making the connection between Fourier series and Fourier transforms clearer. Consider the heat equation $u_t = u_{xx}$ (let us just take $\alpha^2 = 1$ to simplify notation) on the interval $[-L, L]$ with boundary conditions

$$u(-L,t) = u(L,t) \quad \text{and} \quad u_x(-L,t) = u_x(L,t).$$

(These just say that u, u_x should be periodic.) Using separation of variables, the general solution in complex form can be found to be

$$\sum_{n=-\infty}^{\infty} c_n e^{-\frac{n^2\pi^2 t}{L^2}} e^{\frac{in\pi x}{L}}.$$

(The ODE $X'' - \lambda X = 0$ for X in this case has no solutions for $\lambda > 0$, nonzero constant solutions for $\lambda = 0$, and then nonzero solutions for $\lambda = -\frac{n^2\pi^2}{L^2} < 0$, all which are included above by the complex exponentials $e^{\frac{in\pi x}{L}}$ for varying integers n, including $n = 0$.) Given an initial condition $u(x,0) = f(x)$, we can put in the formula for the Fourier coefficients c_n to get

$$\sum_{n=-\infty}^{\infty}\left(\frac{1}{2L}\int_{-L}^{L} f(s)e^{-\frac{in\pi s}{L}}\, ds\right) e^{-\frac{n^2\pi^2 t}{L^2}} e^{\frac{in\pi x}{L}}.$$

Now, in order to move from the interval $[-L, L]$ to all of $\mathbb{R}$, we take the limit as $L \to \infty$. Note that the integral from $-L$ to L above will become an integral from $-\infty$ to ∞, precisely the type of thing showing up in solutions to the heat equation on $\mathbb{R}$. To make this limit precise, let us set $y_n = \frac{n\pi}{L}$. So far these only take on a discrete set of values, but as $L \to \infty$ these discrete values

start to fill up more and more of $\mathbb{R}$, so that in the limit we expect these to cover all of $\mathbb{R}$. The change Δy_n in y_n, i.e. the distance from one y_n to the next is

$$\Delta y_n = y_{n+1} - y_n = \frac{(n+1)\pi}{L} - \frac{n\pi}{L} = \frac{\pi}{L}.$$

With this we can write the $\frac{1}{2L}$ term in the series above as

$$\frac{1}{2L} = \frac{\Delta y_n}{2\pi},$$

and so in this notation the entire series expression above becomes

$$\sum_{n=-\infty}^{\infty} \left(\frac{1}{2\pi} \int_{-L}^{L} f(s) e^{-iy_n s}\, ds \right) e^{-y_n^2 t} e^{iy_n x} \Delta y_n.$$

But, this exact series is nothing but a *Riemann* sum for a certain integral! As we take $L \to \infty$, the y_n's fill out all values y in $\mathbb{R}$, Δy_n becomes dy, the bounds $-\infty$ to ∞ on n becomes bounds on y, and the bounds on the inner integral become infinite, so that the limit gives

$$\frac{1}{2\pi} \int_{-\infty}^{\infty} \left(\int_{-\infty}^{\infty} f(s) e^{-iys}\, ds \right) e^{-y^2 t} e^{iyx} dy.$$

If you go back and look at our derivation of the heat kernel, you will see that this is precisely the integral we obtained there when recovering $u(x,t)$ from its Fourier transform $\hat{u}(y,t)$. The computation there then gives the integral involving the heat kernel, and that is the point: the heat kernel integral arises by taking a limit of the solutino of the heat equation over a bounded interval as that interval grows in size! The upshot is that Fourier series and Fourier transforms really are two sides of the same coin, where the "continuous" transform arises from the "discrete" Fourier coefficients simply by taking an appropriate limit.

Lecture 22: Eigenvalue Problems

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

Eigenvalue problems. We now head into our final topic of the quarter, where we will see, even more clearly than before, how much of the study of linear differential equations really does amount to infinite-dimensional linear algebra. Let us return to the heat equation on $[0, L]$. When solving for the separated solution X we end up solving

$$X'' - \lambda X = 0.$$

We can write this equation as $X'' = \lambda X$, and so if we denote by $L = \frac{d^2}{dx^2}$ the linear operator that sends a function to its second derivative, we see that the ODE for X takes the form

$$L[X] = \lambda X.$$

But, in linear-algebraic terms, this equation says precisely that X is an *eigenvector* of L with *eigenvalue* λ! If we impose boundary conditions, say, $X(0) = 0 = X(L)$, then we are looking for eigenvectors of L in the vector space

$$V = \{\text{functions on } [0, L] \text{ such that } X(0) = 0 = X(L)\}.$$

In this setting the eigenvectors are called *eigenfunctions*, and so the upshot is that solving the boundary value problem

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

comes down to finding eigenfunctions of the linear operator $L = \frac{d^2}{dx^2}$ in V. We call such a boundary value problem, where the goal is to determine the λ for which a nonzero solution exists and to determine those nonzero solutions, an *eigenvalue problem.*

Examples. For the specific eigenvalue problem

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

the process we went through before in finding solutions shows that the eigenvalues are $\lambda_n = -\frac{n^2\pi^2}{L^2}$ and corresponding eigenfunctions are $X_n = \sin(\frac{n\pi x}{L})$. Any eigenfunction with eigenvalue λ_n is a scalar multiple of this specific X_n, so this specific X_n *spans* the entire eigenspace corresponding to λ_n and hence forms a basis for this eigenspace.

We previously also considered the eigenvalue problem

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

when looking at the heat equation with one insulated end. The analysis we went through before showed the the eigenvalues where

$$\lambda_n = -\frac{(2n-1)^2\pi^2}{4L^2}$$

and corresponding eigenfunctions were

$$X_n = \sin\frac{(2n-1)\pi x}{2L}.$$

Each of these again spans the entire eigenspace for that eigenvalue.

Another example. Consider the eigenvalue problem

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

These are not boundary conditions we have carefully considered before, so we run through our usual analysis. First, for $\lambda > 0$ the general solution of the given ODE is

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

The boundary conditions give

$$\sqrt{\lambda}(c_1 - c_2) = 0 \quad \text{and} \quad \underbrace{(c_1 e^{\sqrt{\lambda}} + c_2 e^{-\sqrt{\lambda}})}_{X(1)} + \underbrace{\sqrt{\lambda}(c_1 e^{\sqrt{\lambda}} - c_2 e^{-\sqrt{\lambda}})}_{X'(1)} = 0.$$

Thus $c_1 = c_2$, and then the second condition reduces to

$$c_1(1 + \sqrt{\lambda})e^{\sqrt{\lambda}} = c_1(\sqrt{\lambda} - 1)e^{-\sqrt{\lambda}}.$$

Solving this further yields $c_1 = 0$, so $c_2 = c_1 = 0$ as well and there are no nonzero solutions. Hence the given problem has no positive eigenvalues. For $\lambda = 0$ we get $X = c_1 + c_2 x$, and the boundary conditions force $X = 0$, so that 0 is also not an eigenvalue.

Finally, we consider $\lambda = -k^2 < 0$, with k positive. Then

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

The boundary condition $X'(0) = 0$ forces $c_2 = 0$. Then $X = c_1 \cos(kx)$, and the second boundary condition gives

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

Hence k must satisfy $k \sin k = \cos k$, or in other words $\cot k = k$. A sketch of $y = \cot x$ and $y = x$

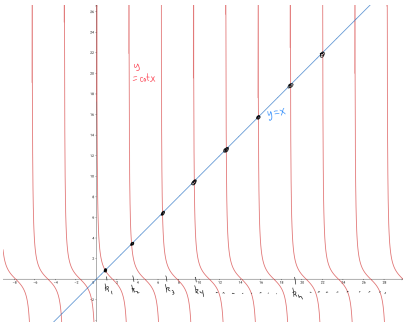

shows that there are infinitely many such positive k, so if we denote these in increasing order by

$$k_1 < k_2 < k_3 < \dots,$$

then the eigenvalues for this problem are $\lambda_n = -k_n^2$, which occur in decreasing order

$$\lambda_1 > \lambda_2 > \lambda_3 > \cdots.$$

(Note that if had written our eigenvalue problem as $X'' + \lambda X = 0$ instead of $X'' - \lambda X = 0$, the eigenvalues would change sign, so that for $X'' + \lambda X = 0$ we would have eigenvalues $\lambda_i = k_i^2$ in *increasing* order. This is a minor point but will be important in clarifying an observation we'll make later.) The corresponding eigenfunctions are

$$X_n = \sin(k_n x) = \sin(\sqrt{-\lambda_n} x),$$

where we use $-\lambda_n$ because our eigenvalues are negative. (If we use $X'' + \lambda X = 0$ instead with $\lambda_n = k_n^2$, the eigenfunctions are $\sin(\sqrt{\lambda_n} x$.)

Self-adjointness. Now, was the existence of eigenvalues and eigenvectors in each of the boundary value problems above a coincidence, or is there something deeper at play? The answer is, of course, yes there is something deeper going on. If we think back to the setting of finite-dimensional linear algebra and matrices, you might recall that a key way to guarantee that a matrix has "enough" real eigenvalues and eigenvectors is to assume that is is *symmetric*, meaning that it equals its own transpose. The *spectral theorem* says that any real symmetric $n \times n$ matrix is *orthogonally diagonalizable*, which means that there exists a basis for $\mathbb{R}^n$ consisting of orthogonal eigenvectors of that matrix. We claim that it is this type of result which accounts for the "eigen-observations" in the boundary value problems above.

Namely, we claim that the linear differential operator $L = \frac{d^2}{dx^2}$ is "symmetric" in an appropriate sense. To say what we really mean by this, we recall the following characterization of symmetric matrices: A is symmetric if and only if

$$A\mathbf{x} \cdot \mathbf{y} = \mathbf{x} \cdot A\mathbf{y}$$

for all vectors $\mathbf{x}, \mathbf{y}$, where $\cdot$ denotes the standard dot product. The point is that in general it is the transpose A^t of A that satisfies

$$A\mathbf{x} \cdot \mathbf{y} = \mathbf{x} \cdot A^t\mathbf{y},$$

so A being symmetric means that $A = A^t$ itself satisfies this dot product equation. If you have ever seen a proof of the spectral theorem for matrices, or of the ingredients which go into the proof (we'll see some of these later in the infinite-dimensional setting), you will notice that it is precisely the property $A\mathbf{x} \cdot \mathbf{y} = \mathbf{x} \cdot A\mathbf{y}$ that makes everything work out.

But we know how to generalize the dot product to infinite dimensions, by considering the inner product of *functions* defined on, say, $[0, L]$:

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

Thus, to say that $L = \frac{d^2}{dx^2}$ is "symmetric" means that L should satisfy the same type of inner product equation as symmetric matrices do, namely:

$$(L[f], g) = (f, L[g]).$$

To be more precise, we only consider functions satisfying whatever boundary conditions we want, say for example $f(0) = 0 = f(L)$. The claim is that $L = \frac{d^2}{dx^2}$, viewed as a linear operator on the vector space of functions satisfying these boundary conditions, does indeed satisfy the inner product equation above. For a general inner product on a vector space, we say that linear operations satisfying this equation are *self-adjoint*, which is the general version of "symmetric". (Thus, symmetric matrices are ones which are self-adjoint with respect to the standard dot product on $\mathbb{R}^n$. The term "adjoint" refers to the linear operator L^* satisfying

$$(L[f], g) = (f, L^*[g])$$

in general, which is called the *adjoint* of L. Self-adjoint operators are thus ones which equal their own adjoints. Adjoints are a generalization of matrix transposes.)

To see that $L = \frac{d^2}{dx^2}$ is self-adjoint, we first determine the adjoint of $\frac{d}{dx}$ using integration by parts. We have:

$$\begin{aligned}
(f', g) &= \int_0^L f'(x)g(x)\,dx \\
&= f(x)g(x)\Big|_0^L - \int_0^L f(x)g'(x)\,dx \\
&= 0 + \int_0^L f(x)[-g'(x)]\,dx \\
&= (f, -g'),
\end{aligned}$$

where the boundary term $f(x)g(x)\Big|_0^L$ vanishes due to the boundary conditions $f(0) = 0 = f(L)$ we are imposing. This says that the adjoint of $\frac{d}{dx}$ is $-\frac{d}{dx}$, since we can replace the application of $\frac{d}{dx}$ in the first argument of (f', g) by an application of $-\frac{d}{dx}$ in the second argument instead. But, now we can apply this operator twice to get L, and each time we do we pick up an extra negative, so that the negatives cancel out in the end:

$$(L[f], g) = (\tfrac{d^2f}{dx^2}, g) = (\tfrac{d}{dx}\left(\tfrac{df}{dx}\right), g) = (\tfrac{df}{dx}, -\tfrac{dg}{dx}) = (f, -\tfrac{d}{dx}(-\tfrac{dg}{dx})) = (f, \tfrac{d^2g}{dx^2}) = (f, L[g]).$$

Thus, L is self-adjoint with respect to this inner product as claimed. We'll see that this the key reason why we can expect to find eigenvalues and eigenfunctions for ODEs involving L, and other more general operators.

Lecture 23: Sturm-Liouville

Warm-Up. We solve the following eigenvalue problem:

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

(In other words, we are finding the eigenvalues and eigenfunctions of $L = \frac{d^2}{dx^2} + 1$ on the vector space of functions satisfying the boundary conditions above, where $\frac{d^2}{dx^2} + 1$ is linear operator that computes a second derivative and adds on 1 times the input: $L[y] = \frac{d^2}{dx^2}y + 1y = y'' + y$.) Let us rewrite our equation as

$$y'' + (1 - \lambda)y = 0.$$

For $1 - \lambda = -k^2 < 0$, the general solution is

$$y = c_1 e^{kx} + c_2 e^{-kx},$$

and a direct (a bit tedious, so we skip it) computation shows that this satisfies

$$y(0) + y'(0) = 0 = y(\pi) + y'(\pi)$$

only when $c_1 = c_2 = 0$, so there are no eigenvalues here. For $1 - \lambda = 0$, we have

$$y = c_1 + c_2 x,$$

and again the boundary conditions force $c_1 = c_2 = 0$.

So, we assume $1 - \lambda = k^2 > 0$, so that

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

The boundary condition $y(0) + y'(0) = 0$ gives

$$c_1 + kc_2 = 0,$$

and the boundary condition $y(\pi) + y'(\pi) = 0$ 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 n, and thus the eigenvalues are $\lambda_n = 1 - n^2$. The solution of our ODE in this case is

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

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

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

(We will see that, just as functions can be written as, say, Fourier sine series expansions, they can also be written as "expansions" in terms of *these* eigenfunctions above! This will give more general types of "Fourier-like" series, suitable to study the specific boundary value problem at hand.)

Sturm-Liouville theory. The proper setting in which to discuss all of the eigen/self-adjoint observations we made last time is that of *Sturm-Liouville theory*. Sturm-Liouville theory deals with second-order linear differential equations of the form

$$(p(x)y')' + q(x)y = \lambda r(x)y.$$

Here y is the unknown function (of x), and such an ODE is said to be in *Sturm-Liouville form*. We will consider boundary conditions of the form

$$\alpha_1 y(0) + \alpha_2 y'(0) = 0 = \beta_1 y(L) + \beta_2 y'(L),$$

which are called *separated* boundary conditions since there is no interaction between the terms at 0 and the terms at L, so that the boundary points are "separated" from one another.

Here are two remarks. First, note that our sign conventions here are a bit different than the book's. The book writes the Sturm-Liouville ODE as

$$(p(x)y')' - q(x)y + \lambda r(x)y = 0, \text{ or equivalently } (p(x)y')' - q(x)y = -\lambda r(x)y.$$

So, the signs of $q(x)$ and λ in my version differ from that in the book. The different signs in $q(x)$ doesn't matter at all (just replace what I'm calling $q(x)$ by the function $-q(x)$ instead), and the change in λ just means that my eigenvalues will be the negative of those in the book (recall the difference in eigenvalues of $X'' - \lambda X = 0$ vs $X'' + \lambda X = 0$ we mentioned last time), but importantly the *eigenfunctions* will be the same either way. The convention we're using here will just make the connection to linear algebra simply to see, but in the end it has no real impact on the theory, apart from changing the signs of the eigenvalues.

The second remark is that, although the Sturm-Liouville ODE form above might seem to be quite restrictive at first, it in fact covers pretty much *all* second-order linear homogeneous ODEs. Indeed, you will show on a discussion problem that essentially any second-order linear homogeneous ODE

$$P(x)y'' + Q(x)y' + R(x)y = 0$$

can be put in Sturm-Liouville form, at least when the coefficient functions are "nice". Certainly, for most ODEs you will ever see in any real practical context this will be true. Second-order ODEs arise all the times in physics and chemistry via Newton's laws and other considerations, so Sturm-Liouville theory is indeed quite broadly applicable.

Operators on Hilbert spaces. Introduce the *Sturm-Liouville* operator:

$$L[y] = (p(x)y')' + q(x)y,$$

which is a linear differential operator acting on some vector space of functions. (Note again the difference in signs here vs the book.) With this notation, the Sturm-Liouville ODE becomes

$$L[y] = \lambda r(x)y.$$

This looks close to an eigenvector equation, except for the nonconstant $r(x)$ term on the right. If we assume that $r(x) > 0$, for example, then we can rewrite the Sturm-Liouville equation as

$$\frac{L[y]}{r(x)} = \lambda y,$$

which now *is* an eigenvector equation. The conclusion is that solutions of the Sturm-Liouville equation

$$(p(x)y')' + q(x)y = \lambda r(x)y$$

are precisely eigenfunctions of the modified Sturm-Liouville operator

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

at least when $r(x) > 0$. The case where $r(x) > 0$ is called the *regular* case of Sturm-Liouville theory, or more precisely the regular case is when we assume that p, p', q, r are all C^1 and that p, r are always positive. The positivity assumption on $r(x)$ is the most important for us, and we will clarify the role that $r(x) > 0$ plays in the theory later.

So, if $L[y] = (p(x)y')' + q(x)y$ is the linear operator we care about, on which vector space is it meant to act? Let us recall the notation $L^2(X)$ as denoting the space of all real-valued *square-integrable* functions on X, which are functions $f : X \to \mathbb{R}$ such that the integral of the square of f over X exists and is finite:

$$\int_X f(x)^2\,dx < \infty.$$

In the case at hand, we are working with $L^2([0, L])$ (the second "L" is different than the first!), so functions for which $\int_0^L f(x)^2\,dx$ is finite. $L^2(X)$ is a vector space since adding and scalar multiplying L^2 functions results in an L^2 function, and is the prototypical example of what's called a *Hilbert space*. What makes this a "Hilbert space" is the fact that $L^2(X)$ comes equipped with an inner product

$$(f, g) = \int_X f(x)g(x)\,dx,$$

and that it satisfies some type of convergence property, which we will say a bit about later. So, we have our Sturm-Liouville operator L acting on some Hilbert space of functions $L^2([0, L])$ (all three "L"'s mean different things here!), and it in this setting that Sturm-Liouville theory takes place. (Actually, we should restrict our domain to consider only those square-integral functions which are, say, C^1, in order to guarantee that their derivatives exist, as required in the very definition of the operator $L[y] = (p(x)y')' + q(x)y$. This will be a standing assumption going forward, so we won't clarify it every single time we need it.)

Lagrange's identity. We now want to show that the Sturm-Liouville operator L is self-adjoint, at least on an appropriate subspace of our Hilbert space, since it is this property which will give

rise to all the eigen-goodness we expect. We begin by determining the adjoint of L with respect to our integral inner product. The computation again comes down to integration by parts. Recall that our operator is defined by $L[y] = (p(x)y')' + q(x)y$.

For two functions u and v, we have:

$$\begin{aligned}(L[u], v) &= \int_0^L L[u]v\,dx \\ &= \int_0^L [(p(x)u')' + q(x)u]v\,dx \\ &= \int_0^L (p(x)u')'v\,dx + \int_0^L q(x)uv\,dx.\end{aligned}$$

We will leave the second term as is, and integrate by parts in the first to "undo" the outer derivative in $(p(x)u')'$:

$$\int_0^L (p(x)u')'v\,dx = p(x)u'v\Big|_0^L - \int_0^L p(x)u'v'\,dx.$$

Now we group the $p(x)$ and v' terms in the second piece together and integrate by parts again:

$$\begin{aligned}\int_0^L p(x)u'v'\,dx &= \int_0^L u'(p(x)v')\,dx \\ &= up(x)v'\Big|_0^L - \int_0^L u(p(x)v')'\,dx.\end{aligned}$$

After putting this into the expression we obtained after the first integration by parts, and bringing back in the $\int_0^L q(x)uv\,dx$ term, we altogether get:

$$\begin{aligned}(L[u], v) &= \int_0^L (p(x)u')'v\,dx + \int_0^L q(x)uv\,dx \\ &= p(x)u'v\Big|_0^L - \int_0^L p(x)u'v'\,dx + \int_0^L q(x)uv\,dx \\ &= p(x)u'v\Big|_0^L - up(x)v'\Big|_0^L + \int_0^L u(p(x)v')'\,dx + \int_0^L q(x)uv\,dx \\ &= p(x)(u'v - uv')\Big|_0^L + \int_0^L u[(p(x)v')' + q(x)v]\,dx \\ &= p(x)(u'v - uv')\Big|_0^L + \int_0^L u[v]\,dx = p(x)(u'v - uv')\Big|_0^L + (u, L[v]).\end{aligned}$$

This final identity

$$(L[u], v) = p(x)(u'v - uv')\Big|_0^L + (u, L[v])$$

is known as *Lagrange's identity*. As we will see next time, the point is that the separated boundary conditions we consider will force the "boundary term"

$$p(x)(u'v - uv')\Big|_0^L$$

to be zero, so that Lagrange's identity says that L will indeed be self-adjoint on the space of functions satisfying these boundary conditions.

Lecture 24: More on Sturm-Liouville

Warm-Up. We show that if L is a self-adjoint operator with respect to some inner product, meaning

$$(L[u], v) = (u, L[v]),$$

then all eigenvalues of L are real and eigenvectors that correspond to different eigenvalues are orthogonal. (You might recall that these same properties are true for symmetric matrices.) Of course, we will apply this to the case of a Sturm-Liouville operator where the inner product is given by some integral, but this is a general fact that works for other inner products as well. The important properties needed are how inner products behave with respect to scalar multiplication, and the fact that the inner product (u, u) of a vector with itself is always bigger than or equal to zero (in particular, it is always real), and equals 0 only for $u = 0$. We should clarify here that we are really considering a *complex* inner product, where we have

$$(\lambda u, v) = \lambda(u, v) \text{ for a complex scalar } \lambda, \text{ but } (u, \lambda v) = \overline{\lambda}(u, v).$$

That is, when pulling a complex scalar out of the second argument, it becomes the *conjugate* complex scalar. We can see that this is true for the complex integral inner product we considered when discussing complex Fourier series:

$$(f, g) = \int_0^L f(x)\overline{g(x)}\, dx,$$

where $\overline{\lambda g(x)} = \overline{\lambda}\,\overline{g(x)}$ results in "pulling out" the conjugate scalar as claimed. (We are then really working in the Hilbert space of *complex-valued* functions on $[0, L]$, of which the Hilbert space of real-valued functions is a subspace.)

Thus, suppose λ is a complex eigenvalue of L, with (complex) eigenvector u. Since L is self-adjoint, we have

$$(L[u], u) = (u, L[u]).$$

The left side is $(\lambda u, u) = \lambda(u, u)$, while the right side is $(u, \lambda u) = \overline{\lambda}(u, u)$, so

$$\lambda(u, u) = \overline{\lambda}(u, u).$$

Since eigenvectors are always nonzero, $(u, u) > 0$, so this forces $\lambda = \overline{\lambda}$, and hence λ is real.

Now suppose u and v are eigenvectors corresponding to different eigenvalues λ and μ, respectively. Then $(L[u], v) = (u, L[v])$, so

$$(\lambda u, v) = (u, \mu v).$$

The left side is $\lambda(u, v)$ and the right side is $\overline{\mu}(u, v)$, but we already know that μ has to be real, so

$$\lambda(u, v) = \mu(u, v).$$

Since $\lambda \neq \mu$, this forces $(u, v) = 0$, so u and v are orthogonal as claimed.

Self-adjoint Sturm-Liouville. Let us return to Lagrange's identity:

$$(L[u], v) = p(x)(u'v - uv')\Big|_0^L + (u, L[v]).$$

Now we impose our separated boundary conditions:

$$\alpha_1 y(0) + \alpha_2 y'(0) = 0 = \beta_1 y(L) + \beta_2 y'(L).$$

Note that we expect at least one of α_1, α_2 to be nonzero, since otherwise the first boundary condition is no condition at all, and similarly for β_1, β_2.

For functions satisfying these boundary conditions, we have

$$y'(0) = -\frac{\alpha_1}{\alpha_2} y(0).$$

(At least, assuming $\alpha_2 \neq 0$; if α_2 is zero then we solve for $y(0)$ in terms of $y'(0)$ instead). Thus the boundary term

$$p(x)(u'v - uv')$$

in Lagrange's identity upon evaluating at 0 becomes

$$p(0)(u'(0)v(0) - u(0)v'(0)) = p(0)\left[-\frac{\alpha_1}{\alpha_2}u(0)v(0) - u(0)\left(-\frac{\alpha_1}{\alpha_2}\right)v(0)\right] = 0,$$

and the same happens when evaluating at L. (Again, perhaps we have to divide by β_1 instead of β_2, or by α_1 instead of α_2 here depending on which scalars are nonzero, but the boundary term vanishes either way.) Thus, among the functions in the subspace of $L^2([0, L])$ carved out by our separated boundary conditions, we have

$$(L[u], v) = (u, L[v]),$$

so L is self-adjoint on this subspace. Immediately we thus know that all eigenvalues of L on this subspace are real and that eigenfunctions for different eigenvalue are orthogonal.

Changing the weight. Recall that using the operator $L[y] = (p(x)y')' + q(x)y$, our Sturm-Liouville ODE was

$$L[y] = \lambda r(x) y.$$

As we mentioned before, this is not quite an eigenvector equation due to the nonconstant $r(x)$, but (assuming regularity) to fix this we instead consider the modified Sturm-Liouville operator obtained by dividing by $r(x)$:

$$\frac{L[y]}{r(x)} = \lambda y.$$

Solutions of the Sturm-Liouville ODE are then indeed eigenfunctions of $\frac{L}{r(x)}$.

If $r(x) = 1$ (as is true in all examples we've seen so far, as we'll soon clarify), then this operator is just L, which is self-adjoint with respect to the usual integral inner product. If $r(x) \neq 1$, however, we lose self-adjointness with respect to this usual inner product, but can recover it using a "weighted" inner product instead. Indeed, we have

$$\int_0^L L[u]v\, dx = \int_0^L uL[v]\, dx,$$

so if we divide and multiply both sides by $r(x) > 0$, we can rewrite this as

$$\int_0^L r(x)\left(\frac{L[u]}{r(x)}\right)v\, dx = \int_0^L r(x)u\left(\frac{L[v]}{r(x)}\right)\, dx.$$

Thus if we introduce the *weighted* inner product

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

we have

$$\left\langle \frac{L[u]}{r(x)}, v \right\rangle = \left\langle u, \frac{L[v]}{r(x)} \right\rangle,$$

so that $\frac{L}{r(x)}$ becomes self-adjoint. The fact that $r(x) > 0$ (i.e. regularity) is needed here to guarantee that this weighted inner product is still an inner product, meaning in particular that it still satisfies $\langle u, u \rangle \geq 0$ and $\langle u, u \rangle = 0$ only for $u = 0$.

(Weighted inner products are used in finite-dimensions as well, although possibly not in your previous linear algebra. For example, modifying the usual dot product on $\mathbb{R}^2$ by putting in some weight terms like $(\mathbf{x}, \mathbf{y}) = 2x_1y_1 + 3x_2y_2$ also gives an inner product on $\mathbb{R}^2$. The idea here is that, for whatever reason, perhaps we want to give more weight to the second coordinates of our vectors instead of the first, and using weighted inner products gives a formal way to do so.)

Sturm-Liouville Spectral Theorem. Thus, finally, we have that solutions of a regular Sturm-Liouville ODE

$$\underbrace{(p(x)y')' + q(x)y}_{L[y]} = \lambda r(x) y$$

with separated boundary conditions are the same as eigenfunctions of the "weighted" Sturm-Liouville operator $\frac{L}{r(x)}$ on the space of functions satisfying those boundary conditions, which is self-adjoint with respect to an $r(x)$-weighted inner product. For matrices, it is the spectral theorem that then gives the eigen-properties of symmetric matrices, so now we give our infinite-dimemsional analog, which we call the *Sturm-Liouville Spectral Theorem*.

The claim is that, with the setup above, the following are true:

- $\frac{L}{r(x)}$ has infinitely many real eigenvalues, which are *discrete* in the sense that they can be arranged in (decreasing) order:

  $$\lambda_1 > \lambda_2 > \lambda_3 > \dots.$$

  (Here is one place where the difference in sign convention between my Sturm-Liouville equation and the book's pops up: with the book's convention of using $-\lambda$ instead of λ, the eigenvalues are instead *increasing*. As we said before, this does not lead to any real material difference, so we sweep it under the rug.)

- All eigenspaces of $\frac{L}{r(x)}$ are one-dimensional, meaning that all eigenfunctions for a specific eigenvalue are scalar multiples of a single nonzero one.

- Finally, and most importantly, eigenfunctions ϕ_n corresponding to the eigenvalues λ_n form an orthogonal *basis* for the Hilbert space $L^2([0, L])$, or at least the subspace of C^1 functions in this Hilbert space, meaning that any such f can be written as

  $$f(x) = \sum_{n=1}^{\infty} c_n \phi_n(x)$$

  for an appropriate choices of scalars c_n. We call such a series expression the *eigenfunction expansion* of f relative to this basis.

So, just as with symmetric matrices, self-adjoint Sturm-Liouville operators do give rise to a basis of orthogonal eigenvectors of our vector space. In the first part, we already know that the eigenvalues are real from the Warm-Up, so the important part is there are infinitely many of them and that they

are discrete. The second part can be proven using the notion of the *Wronskian* of two functions, but we will skips this since this part is not so crucial for our purposes.

The fact that the eigenfunctions in the third part are orthogonal is just what we showed in the Warm-Up, so the real claim here is that they form a basis for our space. There are two remarks we should make here. First, it is true that these give a basis for all of L^2, not just the subspace of C^1 functions, but the meaning behind what the equality

$$f(x) = \sum_{n=1}^{\infty} c_n \phi_n(x)$$

actually means changes for L^2 as opposed to the C^1 case. We won't go into the details since we only really care about the C^1 case, but this comes down to the type of "convergence" we care about. (This is not about the distinction between pointwise vs uniform convergence—there is yet another type of convergence, namely what's called L^2-convergence, we comes into play. For our purposes, all that matters is that uniform convergence implies L^2-convergence.) Note well that this applies even if f does *not* satisfy our separated boundary conditions: *any* appropriately nice function can be "eigenfunction expanded". (In the case where f is only piecewise C^1, the eigenfunction series will converge to $\frac{1}{2}(f(x^-) + f(x^+))$ at points where f is not continuous—this possibly includes the endpoints of $[0, L]$—just as we saw for Fourier series.)

The second remark is that we have to be a bit careful about what we mean by "basis" here. In the standard linear-algebraic usage, when saying that a given vector is a linear combination of basis vectors we always mean a linear combination of *finitely many* basis vectors, since "infinite linear combinations" do not make sense in arbitrary vector spaces. But, in the third part of this spectral theorem, we are indeed allowing infinite linear combinations expressed as an infinite series. Technically, what we have here then is what's called a *Schauder basis* for our space, which is the name for bases where infinite linear combinations are used. The previous notion of basis you would have seen, which only allows finite linear combinations, is called a *Hamel basis*. So, the third part is really about the existence of a Schauder basis. In this setting, we also say that the set of eigenfunctions forms a *complete* set of orthogonal eigenvectors for our Hilbert space.

Examples. Let us finish by putting previous examples we've seen now into the context of this spectral theorem. The eigenvalue problem

$$y'' - \lambda y = 0,\ y(0) = 0 = y(1)$$

is in regular Sturm-Liouville form by taking $p(x) = r(x) = 1$ and $q(x) = 0$:

$$(1y')' + 0y = \lambda 1 y.$$

The eigenvalues $\lambda_n = -n^2\pi^2$ are discrete and decreasing

$$-\pi^2 > -4\pi^2 > -9\pi^2 > \dots$$

(they would be increasing if we used the book's convention with $y'' + \lambda y = 0$ where $\lambda_n = n^2\pi^2$ instead), and the eigenspace corresponding to λ_n is spanned by $\phi_n = \sin(n\pi x)$. The fact these eigenfunctions are orthogonal is just the orthogonality relations for Fourier series we saw back in the first week (the inner product is not weighted in this case since $r(x) = 1$), and the corresponding eigenfunction expansion of a function

$$f(x) = \sum_{n=1}^{\infty} c_n \sin(n\pi x)$$

is just a Fourier sine series.

For $y'' - \lambda y = 0$ with separated boundary conditions $y(0) = 0 = y'(1)$, the eigenvalues were

$$-\pi^2 > -9\pi^2 > -25\pi^2 > \dots,$$

or $\lambda_n = -(2n-1)^2\pi^2$ in general, and the basis eigenfunctions were $\phi_n = \sin(\frac{(2n-1)\pi x}{2})$. The corresponding eigenfunctions expansions are the specific sine series of the form

$$f(x) = \sum_{n=1}^{\infty} c_n \sin\left(\frac{(2n-1)\pi x}{2}\right),$$

which we saw come up in the heat equation, for example, with one insulated end. The moral of the story is that given any nice enough ODE, we get a nice type of expansion uniquely suited to that ODE to work with. We'll soon put such expansions to use when solving general types of non-homogeneous ODEs.

Lecture 25: Eigenfunction Expansions

Warm-Up. We find the normalized eigenvectors for each of the following eigenvalue problems:

$$y'' + \lambda y = 0,\ y'(0) = 0 = y(1) + y'(1)$$

and

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

Note that in both cases the weight function is $r(x) = 1$. Now, by *normalized* eigenvectors we mean eigenvectors of norm 1, where the norm computed using the inner product: $\|f\| = \sqrt{(f,f)}$. To obtain a normalized eigenvector we simply take any eigenvector and divide it by its norm, just as you would have done in a previous linear algebra course.

We saw in an earlier example that the eigenvalues of the first problem are the positive solutions to $\cot\sqrt{\lambda} = \sqrt{\lambda}$. For the n-th eigenvalue λ_n, a corresponding eigenfunction is

$$\phi_n = \cos(\sqrt{\lambda_n}x).$$

We have

$$(\phi_n, \phi_n) = \int_0^1 \phi_n(x)^2\,dx = \int_0^1 \cos^2(\sqrt{\lambda_n}x)\,dx = \frac{1}{2}\left(1 + \frac{\sin(\sqrt{\lambda_n})\cos(\sqrt{\lambda_n})}{\sqrt{\lambda_n}}\right).$$

This can be cleaned up a bit using the fact that the eigenvalue satisfies $\cos\sqrt{\lambda} = \lambda$:

$$\frac{\cos(\sqrt{\lambda_n})}{\sin(\sqrt{\lambda_n})} = \sqrt{\lambda_n} \implies \cos(\sqrt{\lambda_n}) = \sqrt{\lambda_n}\sin(\sqrt{\lambda_n}),$$

so

$$(\phi_n, \phi_n) = \frac{1}{2}(1 + \sin^2(\sqrt{\lambda_n})).$$

Thus the normalized eigenfunctions are

$$\frac{\phi_n}{\|\phi_n\|} = \frac{\sqrt{2}\cos(\sqrt{\lambda_n}x)}{\sqrt{1 + \sin^2(\sqrt{\lambda_n})}}.$$

We saw previously that the eigenvalues of the second problem are $\lambda_n = 1 - n^2$, where n is a positive integer, with eigenfunctions

$$\phi_n = \sin(nx) - n\cos(nx).$$

Let us first emphasize that these eigenfunctions are orthogonal:

$$(\phi_n, \phi_m) = \int_0^\pi [\sin(nx) - n\cos(nx)][\sin(mx) - m\cos(mx)]\,dx = 0$$

for $m \neq n$. Of course this can be verified through a computation of these integrals, but the point is that we know it is true without computation purely as a consequence of the fact that these eigenfunctions arise from a self-adjoint operator. We also have

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

Thus the norm of ϕ_n is $\sqrt{\frac{\pi}{2}(n^2+1)}$, so the normalized eigenfunctions are

$$\frac{\phi_n}{\|\phi_n\|} = \frac{\sqrt{2}[\sin(nx) - n\cos(nx)]}{\sqrt{\pi(n^2+1)}}.$$

Existence of eigenvalues. We will not be able to prove the Sturm-Liouville spectral theorem in full in this course, but let us give a sense of what is needed in order to do so. To begin, we recall how the proof of the spectral theorem for matrices works. If A is an $n \times n$ symmetric matrix, the key facts we need are that A always has at least one real eigenvalue, and that A preserves orthogonal complements of invariant subspaces. The first fact follows from the fact that eigenvalues are roots of the characteristic polynomial: this polynomial for sure has a complex root, so A has a complex eigenvalue, and then the fact that A is symmetric guarantees that this complex root is actually real. By the second fact above we mean that if A sends every vector of a subspace into that same subspace, then it also sends every vector in orthogonal complement of that subspace into that orthogonal complement.

Here is how this works in the 3×3 case. First, A has at least one real eigenvalue, and so an eigenvector $\mathbf{v}_1$. The orthogonal complement property stated above then says that A will send anything in the plane *orthogonal* to $\mathbf{v}_1$ into that same plane:

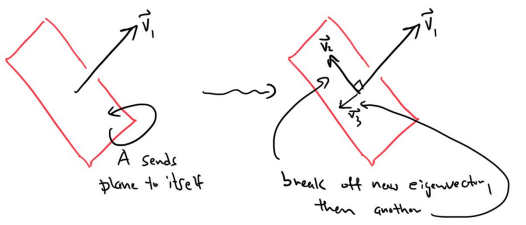

But now A gives a symmetric linear transformation from this plane to itself, so the same reasoning applies again: A has at least one eigenvector in this plane, call $\mathbf{v}_2$, and then A sends the space orthogonal to the span of $\mathbf{v}_1, \mathbf{v}_2$ into itself. At this final step this remaining orthogonal complement is a line, and a vector on this line gives a third eigenvector $\mathbf{v}_3$. These eigenvectors are all orthogonal by construction (each came from some orthogonal complement to the previous ones), and so they form an orthogonal eigenbasis of $\mathbb{R}^3$. The idea is that at each step we "break" off an orthogonal

complement, which then contains a new eigenvector we can add to the list of eigenvectors we are constructing, until we end up with three in total. This idea then works for larger matrices as well.

The details are not crucial to understand, but the upshot is that it is this same type of argument which essentially works in infinite dimensions too. The orthogonal complement preservation property above only depends on $A\mathbf{x} \cdot \mathbf{y} = \mathbf{x} \cdot A\mathbf{y}$, and so works for more general self-adjoint operators, at least once we deal with some subtle issues that arise infinite dimensions. The real hard part in infinite dimensions is in guaranteeing that any self-adjoint operator has at least one real eigenvalue. The issue is that there is no analog of the characteristic polynomial in infinite dimensions, since we cannot take the determinant of an "infinite" matrix. So, if want to guarantee the existence of an eigenvalue we need a different approach.

The solution comes from another topic you saw in a previous course: the method of *Lagrange multipliers*. The point is that there is another way to characterize eigenvalues (at least the largest and smallest ones) of a symmetric matrix: they are precisely the maximal and minimal values of the function

$$f(\mathbf{x}) = \mathbf{x} \cdot A\mathbf{x}$$

among all unit vectors! This can be shown by optimizing $f(\mathbf{x})$ subject to the constraint $\|\mathbf{x}\| = 1$ using Lagrange multipliers. Thus, even without using characteristic polynomials, we see that we have another way to guarantee that A has an eigenvalue via optimization instead. It turns out that this idea *does* work (essentially) in infinite-dimensions, and can be used to show that self-adjoint operators do have eigenvalues. There is still much work needed in making the jump to infinite dimensions here precise, but once we've done so, the proof of the spectral theorem works in a similar way as in the finite-dimensional case.

Back to Hilbert spaces. Let us take a brief aside to finish off the definition of the term "Hilbert space". We previously said that a Hilbert space was a vector space equipped with an inner product which satisfies some type of convergence property, with the key example being the space $L^2(X)$ of square-integrable functions on some space X. The convergence property we need can be stated in different ways, but here is one which likely the simplest to state: in a Hilbert space it should be true that a series $\sum \phi_n$ converges whenever the series of *norms* $\sum \|\phi_n\|$ converges. This latter series is just a series of numbers and convergence of this means convergence in the usual sense of calculus. The former series is a series in our Hilbert space, and convergence here really means "L^2-convergence", whatever that means. But, uniform convergence is a special case, and that's really what matters for us.

We will not need the precise definition in the rest of this course, but it's worth stating at least once. Hilbert spaces are of fundamental importance in many areas of mathematics (and in quantum mechanics!), essentially because they provide examples of infinite-dimensional spaces which behave, in many many ways, in some sense like $\mathbb{R}^n$. Take a course in "functional analysis" to learn more.

Weighted example. Let us now look at an example of a Sturm-Liouville problem with a weight. We consider the following:

$$(xy')' + \frac{1}{x}y = \lambda \frac{1}{x}y, \quad y(1) = 0 = y(e).$$

Note that this is in regular Sturm-Liouville form, with $p(x) = x, q(x) = \frac{1}{x}$, and with weight function $r(x) = \frac{1}{x} > 0$. (Note that these are all C^1 on the domain $[1, e]$ in question.) With this weight then we can rewrite our ODE as

$$\frac{(xy')' + \frac{1}{x}y}{\frac{1}{x}} = \lambda y,$$

so that we are considering eigenfunctions of the operator on the left.

This eigenvalue problem turns out to have eigenvalues $\lambda_n = 1 - n^2\pi^2$ with eigenfunctions

$$\phi_n = \sin(n\pi \ln x).$$

That these λ_n are the correct eigenvalues can be verified simply by plugging in $y = \phi_n$ and computing

$$\frac{(x\phi_n')' + \frac{1}{x}\phi_n}{\frac{1}{x}},$$

which (do it!) works out to be $(1-n^2\pi^2)\phi_n$. These specific eigenfunctions can be found by actually solving the ODE

$$(xy')' + \frac{1}{x}y = \lambda\frac{1}{x}y$$

with the given boundary conditions. Doing so is not hard, it just requires methods for solving what are called *Euler equations*, or more precisely Euler equations with repeated roots. We won't go through the process of actually solving this to find ϕ_n here.

These eigenfunctions are then orthogonal with respect to the weighted inner product

$$(f, g) = \int_1^e \frac{1}{x} f(x)g(x)\, dx.$$

This means that for $m \neq n$ we have

$$(\phi_n, \phi_m) = \int_1^e \frac{1}{x} \sin(n\pi \ln x) \sin(m\pi \ln x)\, dx = 0.$$

(This can be verified by direct computation using a clever substitution and trig identity, but of course we know it will be true because of general Sturm-Liouville results.) For $n = m$, we can compute that

$$(\phi_n, \phi_n) = \int_1^e \sin^2(n\pi \ln x)\, dx = \frac{1}{2},$$

so that each ϕ_n has norm $\frac{1}{\sqrt{2}}$. Thus, normalized eigenfunctions are given by $\widetilde{\phi_n} = \sqrt{2}\sin(n\pi \ln x)$.

Eigenfunction expansions. Because of Sturm-Liouville theory, we know that eigenfunctions ϕ_n above, or their normalized versions $\widetilde{\phi_n}$, form an orthogonal (orthonormal in the normalized case) basis for $L^2([1, e])$. In particular, we know that any C^1 function as an eigenfunction expansion of the form

$$f(x) = \sum_{n=1}^{\infty} c_n\phi_n(x).$$

Of course, the coefficients c_n can be derived just as we did for Fourier coefficients way back when. (In fact, we phrased the problem of deriving the Fourier coefficients in terms of inner products back in that first week precisely so that the same derivation would apply to *any* orthogonal basis.) The point was to take the inner product (weighted in this case) of both sides of this equation with some ϕ_m, and then use orthogonality to isolate c_m and solve. The answer is

$$c_n = \frac{(f, \phi_n)}{(\phi_n, \phi_n)},$$

which as before is some *projection* coefficient. If we use the orthonormal eigenfunctions $\widetilde{\phi_n}$ instead, the denominator becomes 1, so that

$$f(x) = \sum_{n=1}^{\infty} d_n \widetilde{\phi_n}(x) \text{ with } d_n = (f, \widetilde{\phi_n}).$$

(The benefit of using normalized eigenfunctions is thus to get a simpler expression for the expansion coefficient. However, note that we don't really save any work in doing so, since either way we have to compute (ϕ_n, ϕ_n): either it shows up in the formula for c_n in the non-normalized case, or it shows up in the expression for $\widetilde{\phi_n}$ in the normalized case. So, either way we get to the same point, and so we won't usually bother with requiring normalized eigenfunctions.) The overall point is that these eigenfunctions expansions are just exactly analogs of Fourier series, only using other types of "basis" functions. With the eigenfunctions in the weighted example above, we thus get the "Fourier-like" series

$$f(x) = \sum_{n=1}^{\infty} c_n \sin(n\pi \ln x)$$

where

$$c_n = \frac{(f, \phi_n)}{(\phi_n, \phi_n)} = \frac{\int_1^e \frac{1}{x} f(x) \sin(n\pi \ln x)\, dx}{\frac{1}{2}} = 2\int_1^e \frac{1}{x} f(x) \sin(n\pi \ln x)\, dx.$$

Back to weighted example. Let us see the eigenfunction expansion of $f(x) = 1$ in terms of these weighted eigenfunctions. We have

$$c_n = 2\int_1^e \frac{1}{x} \sin(n\pi \ln x)\, dx = \frac{2(1 - \cos n\pi)}{\pi n}.$$

Thus our expansion is

$$1 = \sum_{n=1}^{\infty} \frac{2(1 - \cos n\pi)}{\pi n} \sin(n\pi \ln x).$$

(If we had used the normalized eigenfunctions, the 2 in front would be a $\sqrt{2}$, but then another $\sqrt{2}$ would show up as the coefficient of $\sin(n\pi \ln x)$, so that overall the same factor of 2 appears.) Note that this equality really is only guaranteed to hold for $1 < x < e$, since strange things can happen at the boundary points 1 and e, where the convergence is actually to some average of right/left limiting values.

This might seem like a crazy identity, but looking at a plot shows that it is in fact not so crazy:

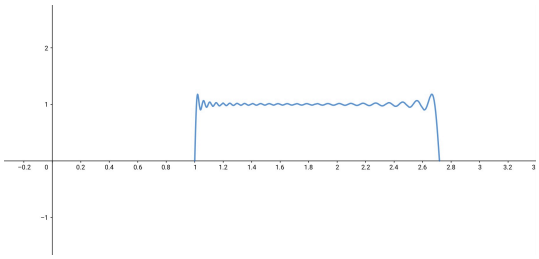

(This is the 50th-order partial sum.) Sure enough, this very strongly suggests that this series does in fact converge to the constant function 1, away from the endpoints. Such is the power of Sturm-Liouville!

Lecture 26: More on Expansions

Warm-Up. We determine the eigenfunction expansion of $f(x) = x$ in terms of the orthogonal eigenfunctions of

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

We have seen that these eigenfunctions are

$$\phi_n(x) = \sin(nx) - n\cos(nx),$$

with eigenvalues $\lambda_n = 1 - n^2$. The eigenfunction expansion of x is

$$x = \sum_{n=1}^{\infty} c_n \phi_n$$

where

$$c_n = \frac{(x, \phi_n)}{(\phi_n, \phi_n)}.$$

The weight for this Sturm-Liouville problem is $r(x) = 1$, so we are using the standard inner product and thus

$$c_n = (x, \phi_n) = \int_0^{\pi} x[\sin(nx) - n\cos(nx)]\, dx = \frac{1 - (1+\pi)\cos n\pi}{n}.$$

We previously saw that $\|\phi_n\|^2 = (\phi_n, \phi_n) = \frac{\pi}{2} = (n^2 + 1)$, so the desired eigenfunction expansion is

$$x = \sum_{n=1}^{\infty} \frac{2[(1 - (1+\pi)\cos n\pi]}{\pi n(n^2+1)}(\sin nx - n\cos nx),\ \text{for } 0 < x < \pi.$$

(Careful at the endpoints!)

Solving driven equations. Let us now put eigenfunction expansions to good use, and solve the nonhomogeneous boundary value problem given by

$$y'' - 3y = x,\ y(0) = 0 = y(1).$$

We call x here the "driving term", since it reflects some external factor that helps to "drive" the behavior of the system. The strategy is to expand everything in sight, meaning the driving term x and the unknown solution y, in terms of our eigenfunctions and see what we get. The eigenfunctions to consider are those of the associated homogeneous eigenvalue problem

$$y'' - \lambda y = 0,\ \text{or } y'' = \lambda y,\ \text{with } y(0) = 0 = y(1).$$

We have seen that this has eigenvalues $\lambda = -n^2\pi^2$ and eigenfunctions $\phi_n = \sin(n\pi x)$.

Now, expanding y gives

$$y = \sum_{n=1}^{\infty} c_n \phi_n$$

for some unknown coefficients, and expanding x gives

$$x = \sum_{n=1}^{\infty} -\frac{2\cos \pi n}{\pi n}\phi_n,$$

where the coefficients here come from the inner product $(x, \phi_n) = \int_0^1 x \sin(n\pi x)\, dx$. The goal is then to use the fact that these should satisfy our driven ODE in order to *determine* the unknown coefficients c_n. Once we know what c_n is explicitly, we will know our solution $y = \sum c_n \phi_n$. The benefit of expanding y in terms of eigenfunctions is that we can easily say what the result of plugging y into our differential equation is: this specific equation comes from the Sturm-Liouville operator $L = \frac{d^2}{dx^2}$, and the fact that $L[\phi_n] = \lambda_n \phi_n$ gives

$$y'' = L[y] = L\left[\sum_{n=1}^{\infty} c_n \phi_n\right] = \sum_{n=1}^{\infty} c_n L[\phi_n] = \sum_{n=1}^{\infty} c_n \lambda_n \phi_n = \sum_{n=1}^{\infty} c_n (-n^2\pi^2)\phi_n.$$

Note that the third equality here is just the fact that L is a linear transformation:

$$L[c_1\phi_1 + c_2\phi_2 + \cdots] = L[c_1\phi_1] + L[c_2\phi_2] + \cdots = c_1 L[\phi_1] + c_2 L[\phi_2] + \cdots .$$

(As usual, there are some convergence/continuity issues lurking in the background needed to guarantee that this linearity property not only applies to finite sums but also infinite sums as well, but we take this for granted.) Thus, in terms of eigenfunction expansions, our driven ODE $y'' - 3y = x$ thus becomes

$$\underbrace{\sum_{n=1}^{\infty} c_n(-n^2\pi^2)\phi_n}_{L[y]=y''} - \underbrace{\sum_{n=1}^{\infty} 3c_n\phi_n}_{3y} = \underbrace{\sum_{n=1}^{\infty} -\frac{2\cos\pi n}{\pi n}\phi_n}_{x}.$$

The beauty of having this identity in terms of the orthogonal eigen-expansion is that now we can directly compare coefficients on both sides: since eigenfunction expansions are unique, the coefficient of ϕ_n on the left side must agree with the coefficient of ϕ_n on the right, so

$$c_n(-n^2\pi^2) - 3c_n = c_n(-n^2\pi^2 - 3) = -\frac{2\cos\pi n}{\pi n}.$$

Thus we find that $c_n = \frac{2\cos\pi n}{(n^2\pi^2+3)\pi n}$, so the solution we desire is

$$y = \sum_{n=1}^{\infty} \frac{2\cos\pi n}{(n^2\pi^2+3)\pi n}\sin(n\pi x).$$

Essentially, phrasing everything in terms of an eigenfunction expansions turned all the calculus, i.e. computing y'', into algebra, i.e. finding c_n!

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. In our eigenfunction expansion of y:

$$y = \sum_{n=1}^{\infty} c_n \phi_n,$$

the point is that the c_n should be viewed as the "coordinates" of y with respect to the basis ϕ_n. We can encode these coordinates in an "infinite-dimensional" vector:

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

The fact that the ϕ_n are eigenvectors, so that

$$L[y] = \sum_{n=1}^{\infty} c_n \lambda_n \phi_n,$$

says that the coordinates of $L[y]$ are $\lambda_n c_n$. Thus, L has the following effect on "coordinate" vectors, where each coordinate is simply scaled 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 $L[y] - 3y = x$ amounts to solving for the c_n in the infinite-dimensional matrix equation

$$\begin{bmatrix} \lambda_1 - 3 & & & \\ & \lambda_2 - 3 & & \\ & & \lambda_3 - 3 & \\ & & & \ddots \end{bmatrix} \begin{bmatrix} c_1 \\ c_2 \\ c_3 \\ \vdots \end{bmatrix} = \begin{bmatrix} d_1 \\ d_2 \\ d_3 \\ \vdots \end{bmatrix}$$

where the d_n are the "coordinates" of x relative to the basis of eigenfunctions. But this matrix is invertible since none of the eigenvalues equal 3, and its inverse is still diagonal with diagonal entries being the reciprocals of the $\lambda_n - 3$, so we the solution we want is

$$\begin{bmatrix} c_1 \\ c_2 \\ c_3 \\ \vdots \end{bmatrix} = \begin{bmatrix} \lambda_1 - 3 & & & \\ & \lambda_2 - 3 & & \\ & & \lambda_3 - 3 & \\ & & & \ddots \end{bmatrix}^{-1} \begin{bmatrix} d_1 \\ d_2 \\ d_3 \\ \vdots \end{bmatrix} = \begin{bmatrix} \frac{1}{\lambda_1 - 3} & & & \\ & \frac{1}{\lambda_2 - 3} & & \\ & & \frac{1}{\lambda_3 - 3} & \\ & & & \ddots \end{bmatrix} \begin{bmatrix} d_1 \\ d_2 \\ d_3 \\ \vdots \end{bmatrix}.$$

This explains the values $c_n = \frac{d_n}{\lambda_n - 3}$ we derived before, so being able to "diagonalize" L has now allowed us to easily "invert" it in order to solve $L[y] - 3y = x \rightsquigarrow y = (L - 3I)^{-1}[x]$.

Orthogonality constraints. Now consider the driven equation

$$y'' - \mu y = x,\ y(0) = 0 = y(1)$$

for an arbitrary constant μ. We use the same eigenfunction expansions as before:

$$y = \sum_{n=1}^{\infty} c_n \sin(n\pi x),\ x = \sum_{n=1}^{\infty} \sin(n\pi x),\ \text{where } d_n = -\frac{2\cos \pi n}{\pi n}.$$

Plugging in our ODE gives

$$\sum_{n=1}^{\infty} -n^2\pi^2 c_n \sin(n\pi x) - \sum_{n=1}^{\infty} \mu c_n \sin(n\pi x) = \sum_{n=1}^{\infty} d_n \sin(n\pi x),$$

and comparing coefficients gives

$$-c_n(n^2\pi^2 + \mu) = d_n.$$

As before, we would like to divide by $-(n^2\pi^2 + \mu)$ in order to find the unknown c_n. The point is that this is valid, as long as μ is *not* an eigenvalue of the corresponding Sturm-Liouville problem! Indeed, if μ is not one of these eigenvalues $-n^2\pi^2$, then $n^2\pi^2 + \mu$ is never zero, and so

$$c_n = -\frac{d_n}{n^2\pi^2 + \mu} = \frac{2\cos \pi n}{\pi n(n^2\pi^2 + \mu)}$$

just as in the $\mu = 3$ example. Then $y = \sum_{n=1}^{\infty} c_n \sin(n\pi x)$ for these c_n is the solution we want.

However, if μ *is* an eigenvalue, say $\mu = -m^2\pi^2$, we have to be more careful. In this case the defining relation for c_m is

$$-c_m(n^2\pi^2 + \mu) = d_m,$$

where the left side is zero since $\mu = -m^2\pi^2$. Thus, in order for this relation to hold, we need d_m to be zero as well—if $d_m \neq 0$, for sure there is no solution to our driven ODE. In the specific case at hand, we have

$$d_m = -\frac{2\cos m\pi}{m\pi},$$

which is nonzero, so $y'' + m^2\pi^2 y = x$ has no solution satisfying $y(0) = 0 = y(1)$.

In the same way, for a general driving term f, the boundary value problem

$$y'' + m^2\pi^2 y = f(x),\ y(0) = 0 = y(1)$$

has a solution if and only if the inner product $(f, \sin(m\pi x))$—which is the coefficient of $\phi_m = \sin(m\pi x)$ in the eigenfunction expansion of f—is zero. To say that $(f, \phi_m) = 0$ means that f should be *orthogonal* to the eigenfunction ϕ_m, which says also that f is orthogonal to the entire eigenspace corresponding to the eigenvalue $\lambda_m = -m^2\pi^2$, since this entire eigenspace is spanned by ϕ_m alone. Thus, in general, the boundary value problem

$$y'' + \mu y = f(x),\ y(0) = 0 = y(1)$$

has a solution for sure if μ is not an eigenvalue of the associated Sturm-Liouville problem, and when μ is an eigenvalue, there is a solution if and only if f is orthogonal to the eigenspace corresponding to that eigenvalue. Thus, existence of solutions to these general nonhomogeneous ODEs comes down to certainly orthogonality relating the driving term to the eigenspaces.

If we phrase this in terms of trying to "invert" $(L-\mu I)[y] = f(x)$ as in the infinite-dimensional matrix approach mentioned above, the claim is that when μ is an eigenvalue of L, $L-\mu I$ is not invertible everywhere, but it *becomes* invertible when restricted to the *orthogonal complement* of the eigenspace corresponding to μ. We will say more about this idea next time, and see how it simply a reflection of some linear algebra you might have seen previously. In the case where the correct orthogonality constraint is satisfied, the driven equation $(L-\mu I)[y] = f(x)$ in fact has infinitely many solutions, as we will see next time in an example. (The point is that the relation you get for the coefficients in the eigenfunction expansions in general is

$$c_m(\lambda_m - \mu) = d_m,$$

so that when f is orthogonal to the eigenspace corresponding to $\mu = \lambda_m$, both sides here are automatically zero regardless of what c_m is. Thus any of the infinitely many choices for the specific coefficient c_m gives a valid solution.)

Lecture 27: Nonhomogeneous ODEs

Warm-Up. We determine the condition under which

$$y'' + y = f(x),\ y(0) + y'(0) = 0 = y(\pi) + y'(\pi)$$

has a solution, and find the solution when it does. We saw previously that the Sturm-Liouville operator $L[y] = y'' + y$ with these boundary conditions has eigenvalues $\lambda_n = 1 - n^2$ with eigenfunctions $\phi_n(x) = \sin(nx) - n\cos(nx)$. We expanding y and f in terms of these eigenfunctions:

$$y = \sum_{n=1}^{\infty} c_n\phi_n, \quad f(x) = \sum_{n=1}^{\infty} d_n\phi_n.$$

Then

$$L[y] = \sum_{n=1}^{\infty} c_n L[\phi_n] = \sum_{n=1}^{\infty} c_n(1-n^2)\phi_n,$$

so $y'' + y = f(x)$ becomes

$$\sum_{n=1}^{\infty} c_n(1-n^2)\phi_n = \sum_{n=1}^{\infty} d_n\phi_n.$$

Thus we require that $c_n(1-n^2) = d_n$ for each n. However, when $n = 1$ the left side is zero, so in order for there to be a solution it must be the case that $d_1 = \frac{(f,\phi_1)}{(\phi_1,\phi_1)}$ is zero, and thus we see that our given problem has a solution if and only if

$$(f, \phi_1) = \int_0^{\pi} f(x)(\sin x - \cos x)\, dx = 0.$$

Using the setup we described last time, the point is that if we write our ODE as

$$L[y] = 0y + f(x),$$

the fact that 0 is an eigenvalue of L ($\lambda_n = 1 - n^2$ is zero for $n = 1$) places restrictions on what f can be, so that f should be orthogonal to the eigenspace corresponding to 0. This is thus where the requirement that $(f, \phi_1) = 0$ comes from.

So, assuming that f is indeed orthogonal to ϕ_1, we can then describe our desired solution as follows. The requirement that

$$c_n(1-n^2) = d_n$$

is trivial when $n = 1$ since both sides are zero. For $n \geq 2$, we then get

$$c_n = \frac{d_n}{1-n^2}.$$

Thus, it is for these coefficients that $y = \sum_{n=1}^{\infty} c_n\phi_n$ gives a solution of our problem. (The d_n's are determined by $f(x)$, and so are known quantities given f.) Now, if we write the series expansion of y as follows

$$y = c_1\phi_1 + \sum_{n=2}^{\infty} c_n\phi_n,$$

we have that

$$y = c_1(\sin x - \cos x) + \sum_{n=2}^{\infty} \frac{d_n}{1-n^2}(\sin nx - n\cos nx)$$

is the desired solution. Here c_1 is an arbitrary constant, and so we actually get infinitely many solutions to $y'' + y = f(x),\ y(0) + y'(0) = 0 = y(\pi) + y'(\pi)$ when $\int_0^{\pi} f(x)(\sin x - \cos x)\,dx = 0$.

The form of the solution we get above makes sense from what you might recall about nonhomogeneous linear equations (either ODEs or systems of equations in the sense of linear algebra). The point is that the first term $c_1(\sin x - \cos x)$ is actually the general solution of the *homogeneous* equation $y'' + y = 0$ with these boundary conditions, and the second term (the series starting at $n = 2$) is a *particular* solution of our nonhomogeneous equation. Thus, the solution we have above is simply a reflection of the fact that solutions of nonhomogeneous linear equations in general are of the form

$$\text{(general homogeneous solution)} + \text{(particular nonhomogeneous solution)}.$$

For example, for a matrix equation $A\mathbf{x} = \mathbf{b}$, the general solution is formed by taking the general solution of $A\mathbf{x} = \mathbf{0}$ and adding to it *a* particular solution of $A\mathbf{x} = \mathbf{b}$.

If we had the equation $y'' + y = 2y + f(x)$ with these boundary conditions instead, then a solution always exist for any f since 2 is not an eigenvalue of our Sturm-Liouville operator, and the (unique) solution is

$$y = \sum_{n=1}^{\infty} \frac{d_n}{\lambda_n - 2}\phi_n$$

where the d_n are the expansion coefficients of $f(x)$.

Fredholm alternative. The fact that $L[y] = \mu y + f(x)$ has a solution in the case where μ is an eigenvalue of L if and only if f is orthogonal to the eigenspace corresponding to μ is an instance of something you might have seen in a previously linear algebra course. (Or maybe not! It might depend on which linear algebra course you had. Certainly if you saw the method of least squares in your prior linear algebra course, you likely saw what we're about to describe here.) The basic fact for matrices is that saying a linear equation $B\mathbf{x} = \mathbf{y}$ has a solution is the same as saying that $\mathbf{y}$ is in the image of B, and the image of B in fact is the same as the orthogonal complement of the kernel of B^T:

$$\operatorname{im} B = (\ker(B^T))^{\perp}.$$

Thus, $B\mathbf{x} = \mathbf{y}$ has a solution if and only if $\mathbf{y}$ is orthogogonal to $\ker(B^T)$. The same is true in our setting, so we have that

$$(L - \mu I)[y] = f(x) \text{ has a solution if and only if } f \text{ is orthogonal to } \ker(L - \mu I)^*,$$

where * denotes the adjoint. But L is self-adjoint in our case, so $(L - \mu I)^* = L - \mu I$, and thus our equation has a solution if and only if f is orthogonal to $\ker(L - \mu I)$, which is precisely the eigenspace of L corresponding to μ!

The fact that we can characterize whether $(L - \mu I)[y] = f(x)$ (or even just $B\mathbf{x} = \mathbf{y}$) has a solution or not in terms of an orthogonality condition on f (or $\mathbf{y}$) is an instance of what's called the *Fredholm alternative theorem*. The general form of this theorem is more broad than what we are describing here, but idea is the same: characterize existence of solutions in terms of orthogonality. Thus, once again, we see that what we are really doing in this course is just infinite-dimensional linear algebra!

Driven example with weights. For a final example before moving on to PDEs, we solve

$$(xy')' + \frac{1}{x}y = \frac{1}{x},\ y(1) = 0 = y(e).$$

If we rewrite the ODE as

$$\frac{(xy')' + \frac{1}{x}y}{\frac{1}{x}} = 1,$$

we see that this is in *weighted* Sturm-Liouville form, with weighted operator

$$L[y] = \frac{(xy')' + \frac{1}{x}y}{\frac{1}{x}}.$$

We saw before that the eigenvalues of L, with the given boundary conditions, are $\lambda_n = 1 - n^2\pi^2$, with eigenfunctions $\phi_n = \sin(n\pi \ln x)$. These are orthogonal with respect to the weighted inner product

$$(f, g) = \int_1^e \frac{1}{x} f(x)g(x)\, dx.$$

Now, if we think of our ODE as

$$\frac{(xy')' + \frac{1}{x}y}{\frac{1}{x}} = 0y + 1,$$

since 0 is not an eigenvalue of L there will be a unique solution. The expansion of the driving term 1 in terms of the ϕ_n is

$$1 = \sum_{n=1}^{\infty} \frac{2(1 - \cos n\pi)}{\pi n} \sin(n\pi \ln x),$$

as we worked out previously. Thus, a solution $y = \sum_{n=1}^{\infty} c_n\phi_n$ should satisfy

$$\underbrace{\sum_{n=1}^{\infty} c_n(1 - n^2\pi^2)\phi_n}_{L[y]} = \underbrace{\frac{2(1 - \cos n\pi)}{\pi n}\phi_n}_{1},$$

and so we need $c_n(1-n^2\pi^2) = \frac{2(1-\cos n\pi)}{\pi n}$. Hence the solution is

$$y = \sum_{n=1}^{\infty} \frac{2(1-\cos n\pi)}{(1-n^2\pi^2)\pi n} \sin(n\pi \ln x).$$

Nonhomogeneous heat equation. With all of this Sturm-Liouville theory built up, we now come to applying it to solve nonhomogeneous versions of PDEs we saw previously, in particular the heat and wave equations. As a first example, we take the following driven heat equation, which actually showed up on the last set of discussion section problems:

$$u_t = u_{xx} + xt,\ u(0,t) = 0 = u(\pi,t),\ u(x,0) = 0.$$

The driving term $f(x,t) = xt$ here describes some external heat source. The appropriate Sturm-Liouville problem for this equation is

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

which characterizes the separated solution X. This has eigenvalues $\lambda_n = -n^2$ and eigenfunctions $\phi_n = \sin(nx)$, as we have seen countless times before.

The strategy is to expand everything in sight in terms of these specific eigenfunctions. So, we expand the solution $u(x,t)$ as

$$u(x,t) = \sum_{n=1}^{\infty} U_n(t) \sin nx$$

for some functions $U_n(t)$. (In other words, for each t expand the function $u(x,t)$ of x in in terms of our eigenfunctions.) The expansion of the driving term xt is

$$xt = \sum_{n=1}^{\infty} c_n(t) \sin nx$$

where

$$c_n(t) = \frac{2}{\pi} \int_0^{\pi} xt \sin(nx)\, dx = -\frac{2t\cos n\pi}{n}.$$

(Again, think of t as fixed and expand xt to get coefficients $c_n(t)$ that depend on t.)

In order to satisfy $u_t = u_{xx} + xt$, we must have

$$\sum_{n=1}^{\infty} U_n'(t) \sin nx = \sum_{n=1}^{\infty} -n^2 U_n(t) \sin nx + \sum_{n=1}^{\infty} -\frac{2t\cos n\pi}{n} \sin nx.$$

(The first term on the right incorporates the eigenvalues $-n^2$ since this is the term obtained by applying the Sturm-Liouville operator $L[X] = X''$.) We can rewrite this as

$$\sum_{n=1}^{\infty} [U_n'(t) + n^2 U_n(t)] \sin nx = \sum_{n=1}^{\infty} -\frac{2t\cos n\pi}{n} \sin nx.$$

Thus, by comparing coefficients we see that the unknown $U_n(t)$ should satisfy

$$U_n'(t) + n^2 U_n(t) = -\frac{2t\cos n\pi}{n}.$$

Now, at $t = 0$ we have

$$u(x,0) = \sum_{n=1}^{\infty} U_n(0) \sin nx,$$

so in order to satisfy $u(x,0) = 0$ we must have that this series is the Fourier sine series of zero, so that $U_n(0) = 0$ for all n. Thus we find that $U_n(t)$ must satisfy the initial value problem

$$U_n'(t) + n^2 U_n(t) = -\frac{2t\cos n\pi}{n},\ U_n(0) = 0.$$

This can be solved using the method of integrating factors for first-order linear ODEs, and doing so will produce the explicit $U_n(t)$ we need so that

$$u(x,t) = \sum_{n=1}^{\infty} U_n(t) \sin(nx)$$

will be a solution of our "driven" heat equation. We'll wrap this up next time and look at more examples of the same process in action.

Lecture 28: Driven Heat and Waves

Warm-Up. We derive the ODE satisfied by the coefficients in an eigenfunction expansion solution of the driven heat conduction problem

$$u_t = \alpha^2 u_{xx} + 3e^{-2t} + x,\ u(0,t) = 0 = u(1,t),\ u(x,0) = 0.$$

Actually, let us go ahead and describe the approach for a general driving term $f(x,t)$ and initial condition $g(x)$, so

$$u_t = \alpha^2 u_{xx} + f(x,t),\ u(0,t) = 0 = u(1,t),\ u(x,0) = g(x).$$

(Different boundary conditions would just change the eigenfunctions we use, but otherwise the method is the same.) The corresponding eigenvalue problem for the Sturm-Liouville operator $L[X] = X''$ (which describes the homogeneous term on the right side of our PDE) has eigenfunctions

$$\phi_n = \sin(n\pi x) \text{ with eigenvalue } \lambda_n = -n^2\pi^2.$$

We expand our desired solution and driving term in terms of these eigenfunctions:

$$u(x,t) = \sum_{n=1}^{\infty} U_n(t)\phi_n(x)$$

and

$$f(x,t) = \sum_{n=1}^{\infty} b_n(t)\phi_n(x)$$

where

$$b_n(t) = 2\int_0^1 f(x,t)\phi_n(x)\,dx.$$

(Again, in each of these view t as fixed, so that the expansion is taken "with respect to x".)

Under these expansions our PDE becomes

$$\underbrace{\sum_{n=1}^{\infty} U_n'(t)\phi_n(x)}_{u_t} = \underbrace{\sum_{n=1}^{\infty} \alpha^2(-n^2\pi^2)U_n(t)\phi_n(x)}_{\alpha^2 u_{xx}} + \underbrace{\sum_{n=1}^{\infty} b_n(t)\phi_n(t)}_{f(x,t)}.$$

(The u_{xx} expression comes simply from scaling each ϕ_n by the appropriate eigenvalue.) After comparing coefficients we thus find that $U_n(t)$ must satisfy the non-homogeneous ODE

$$U_n'(t) = -\alpha^2 n^2\pi^2 U_n(t) + b_n(t).$$

The final piece of data comes from the initial condition, which says that

$$g(x) = u(x,0) = \sum_{n=1}^{\infty} U_n(0)\phi_n(x),$$

so that the $U_n(0)$ are the coefficients of $g(x)$ in its eigenfunction expansion:

$$U_n(0) = 2\int_0^1 g(x)\phi_n(x)\,dx.$$

Thus, $U_n(t)$ satisfies the initial value problem

$$U_n'(t) = -\alpha^2 n^2\pi^2 U_n(t) + 2\int_0^1 f(x,t)\phi_n(x)\,dx,\ U_n(0) = 2\int_0^1 g(x)\phi_n(x)\,dx.$$

Now, in the specific problem

$$u_t = \alpha^2 u_{xx} + 3e^{-2t} + x,\ u(0,t) = 0 = u(1,t),\ u(x,0) = 0,$$

we have

$$3e^{-2t} + x = \sum_{n=1}^{\infty} b_n(t)\phi_n(x)$$

with

$$b_n(t) = 2\int_0^1 (3e^{-2t} + x)\sin(n\pi x)\,dx = \frac{6e^{-2t}(1-\cos n\pi)}{n\pi} - \frac{2\cos n\pi}{n\pi}.$$

The coefficients in the eigenfunction expansion of $u(x,0) = 0$ are all zero, so we get the ODE

$$U_n'(t) = -\alpha^2 n^2\pi^2 U_n(t) + \frac{6e^{-2t}(1-\cos n\pi)}{n\pi} - \frac{2\cos n\pi}{n\pi},\ U_n(0) = 0.$$

This can be solved using the method of integrating factors for first-order linear ODEs from a previous course (or using a computer), and the solution turns out to be

$$U_n(t) = \left(\frac{2\cos n\pi}{\alpha^2 n^3\pi^3} - \frac{6(1-\cos n\pi)}{n\pi(-2+\alpha^2 n^2\pi^2)}\right)e^{-\alpha^2 n^2\pi^2 t} + \frac{6e^{-2t}(1-\cos n\pi)}{n\pi(-2+\alpha^2 n^2\pi^2)} - \frac{2\cos n\pi}{\alpha^2 n^3\pi^3}.$$

Thus the solution to our heat conduction problem is

$$u(x,t) = \sum_{n=1}^{\infty}\left[\left(\frac{2\cos n\pi}{Kn^3\pi^3} - \frac{6(1-\cos n\pi)}{n\pi(-2+Kn^2\pi^2)}\right)e^{-Kn^2\pi^2 t} + \frac{6e^{-2t}(1-\cos n\pi)}{n\pi(-2+Kn^2\pi^2)} - \frac{2\cos n\pi}{Kn^3\pi^3}\right]\sin n\pi x.$$

Example from last time. Let us return to the final example from last time, and finish it off. Recall that we were considering

$$u_t = u_{xx} + xt,\ u(0,t) = 0 = u(\pi,t),\ u(x,0) = 0,$$

and found that the eigenfunction-expanded solution was

$$u(x,t) = \sum_{n=1}^{\infty} U_n(t)\sin(nx)$$

where the $U_n(t)$ satisfy

$$U_n'(t) = -n^2 U_n(t) - \frac{2t\cos n\pi}{n},\ U_n(0) = 0.$$

Here, the $-\frac{2t\cos n\pi}{n}$ are the expansion coefficients of the driving term xt, $-n^2$ is the eigenvalue, and $U_n(0) = 0$ are the expansion coefficients of $u(x,0) = 0$. The solution of this first-order linear initial value problem is

$$U_n(t) = -\frac{2\cos n\pi}{n^5}e^{-n^2 t} + \frac{2\cos n\pi}{n^5} - \frac{2\cos n\pi}{n^3}t,$$

so that the solution of our heat conduction problem is

$$u(x,t) = \sum_{n=1}^{\infty}\left(-\frac{2\cos n\pi}{n^5}e^{-n^2 t} + \frac{2\cos n\pi}{n^5} - \frac{2\cos n\pi}{n^3}t\right)\sin nx.$$

Imagine if instead we had the problem

$$u_t = u_{xx} + u + xt,\ u(0,t) + u_x(0,t) = 0 = u(\pi,t) + u_x(\pi,t),$$

with some initial condition. The only difference here is that in this case we use the Sturm-Liouville problem

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

due to the homogeneous terms $u_{xx} + u$ on the right side of our PDE. This gives the eigenfunctions

$$\phi_n = \sin(nx) - n\cos(nx) \text{ with eigenvalues } \lambda_n = 1 - n^2,$$

as we worked our previously, but after this the process of solving is exactly the same. (We won't work out the details here, but do it for practice!)

Driven wave example. The same approach as in the heat conduction examples above work just as well for driven, nonhomogeneous wave equations. For example consider

$$u_{tt} = u_{xx} + 6x,\ u(0,t) = 0 = u(1,t),\ u(x,0) = 0,\ u_t(x,0) = 0.$$

We can view the $6x$ term as a contribution from an external motor, say, which affects the motion of the string, so that even if the initial position and velocity are both zero, wave-like motion will still occur. As before, we expand everything in terms of the eigenfunctions

$$\phi_n(x) = \sin(n\pi x) \quad (\text{eigenvalues } \lambda_n = -n^2\pi^2)$$

of the corresponding Sturm-Liouville problem $X'' = \lambda X, X(0) = 0 = X(1)$. The driving term has expansion

$$6x = \sum_{n=1}^{\infty} -\frac{12\cos n\pi}{n\pi}\sin n\pi x,$$

where the coefficients come from

$$2\int_0^1 6x\phi_n(x)\,dx = -\frac{12\cos n\pi}{n\pi}.$$

Thus for an expanded solution

$$u(x,t) = \sum_{n=1}^{\infty} U_n(t)\phi_n(x),$$

our PDE becomes

$$\sum_{n=1}^{\infty} U_n''(t)\phi_n(x) = \sum_{n=1}^{\infty} -n^2\pi^2 U_n(t)\phi_n(x) + \sum_{n=1}^{\infty} -\frac{12\cos n\pi}{n\pi}\sin n\pi x,$$

so we get the requirement

$$U_n''(t) + n^2\pi^2 U_n(t) = -\frac{12\cos n\pi}{n\pi}.$$

The initial conditions $u(x,0) = 0$ and $u_t(x,0) = 0$ give $U_n(0) = 0 = U_n'(0)$, we are left with the second-order initial value problem

$$U_n''(t) + n^2\pi^2 U_n(t) = -\frac{12\cos n\pi}{n\pi},\ U_n(0) = 0 = U_n'(0).$$

This can be solved using standard methods for solving second-order ODEs with constant coefficients, and the solution is

$$U_n(t) = \frac{12\cos n\pi}{n^3\pi^3}\cos n\pi t - \frac{12\cos n\pi}{n^3\pi^3}.$$

Hence the solution of our driven wave equation is

$$u(x,t) = \sum_{n=1}^{\infty}\left(\frac{12\cos n\pi}{n^3\pi^3}\cos n\pi t - \frac{12\cos n\pi}{n^3\pi^3}\right)\sin n\pi x.$$

Dampled example. We can also now easily consider other variations of the PDEs we've been seeing. For example,

$$u_{tt} = u_t + u_{xx} + 6x$$

is a *damped* driven wave equation, where the "damping" term u_t on the right describes, for example, the effect of "air drag" or some other medium through which the string moves. The process of solving this is, of course, the same. With the same boundary and initial conditions as before, we get

$$U_n''(t) = U_n'(t) - n^2\pi^2 U_n(t) - \frac{2t\cos n\pi}{n\pi}, U_n(0) = 0 = U_n'(0)$$

as the second-order ODE satisfied by the coefficients in an expand solution. Solving for $U_n(t)$ using standard techniques (computer) will yield our desired solution.

Lecture 29: Spectral Theory

The final day was devoted to putting essentially everything we've done the entire quarter into the context of *spectral theory*, which is the framework that puts Fourier series expansions and Fourier transforms on equal ground. The difference comes only from the difference between the *discrete* vs *continuous* spectrum of a linear operator, where the *spectrum* is some generalization of the set of eigenvalues. I'll include this all at some point later but will omit it for now since it is not required material for the final. Thanks for reading!