Part IB — Methods

Based on lectures by D. B. Skinner
Notes taken by Dexter Chua

Michaelmas 2015

These notes are not endorsed by the lecturers, and I have modified them (often significantly) after lectures. They are nowhere near accurate representations of what was actually lectured, and in particular, all errors are almost surely mine.

Self-adjoint ODEs
Periodic functions. Fourier series: definition and simple properties; Parseval's theorem. Equations of second order. Self-adjoint differential operators. The Sturm-Liouville equation; eigenfunctions and eigenvalues; reality of eigenvalues and orthogonality of eigenfunctions; eigenfunction expansions (Fourier series as prototype), approximation in mean square, statement of completeness. [5]

PDEs on bounded domains: separation of variables
Physical basis of Laplace's equation, the wave equation and the diffusion equation. General method of separation of variables in Cartesian, cylindrical and spherical coordinates. Legendre's equation: derivation, solutions including explicit forms of P_0, P_1 and P_2, orthogonality. Bessel's equation of integer order as an example of a self-adjoint eigenvalue problem with non-trivial weight.

Examples including potentials on rectangular and circular domains and on a spherical domain (axisymmetric case only), waves on a finite string and heat flow down a semi-infinite rod. [5]

Inhomogeneous ODEs: Green's functions
Properties of the Dirac delta function. Initial value problems and forced problems with two fixed end points; solution using Green's functions. Eigenfunction expansions of the delta function and Green's functions. [4]

Fourier transforms
Fourier transforms: definition and simple properties; inversion and convolution theorems. The discrete Fourier transform. Examples of application to linear systems. Relationship of transfer function to Green's function for initial value problems. [4]

PDEs on unbounded domains
Classification of PDEs in two independent variables. Well posedness. Solution by the method of characteristics. Green's functions for PDEs in 1, 2 and 3 independent variables; fundamental solutions of the wave equation, Laplace's equation and the diffusion equation. The method of images. Application to the forced wave equation, Poisson's equation and forced diffusion equation. Transient solutions of diffusion problems: the error function. [6]

Contents

0   Introduction

In the previous courses, the (partial) differential equations we have seen are mostly linear. For example, we have Laplace's equation:

$$\frac{\partial^2 \phi}{\partial x^2} + \frac{\partial \phi}{\partial y^2} = 0,$$

and the heat equation:

$$\frac{\partial \phi}{\partial t} = \kappa \left(\frac{\partial^2 \phi}{\partial x^2} + \frac{\partial^2 \phi}{\partial y^2} \right).$$

The Schrödinger' equation in quantum mechanics is also linear:

$$i\hbar \frac{\partial \Phi}{\partial t} = -\frac{\hbar^2}{2m} \frac{\partial^2 \phi}{\partial x^2} + V(x)\Phi(x).$$

By being linear, these equations have the property that if ϕ_1, ϕ_2 are solutions, then so are $\lambda_1 \phi_1 + \lambda_2 \phi_2$ (for any constants λ_i).

Why are all these linear? In general, if we just randomly write down a differential equation, most likely it is not going to be linear. So where did the linearity of equations of physics come from?

The answer is that the real world is *not* linear in general. However, often we are not looking for a completely accurate and precise description of the universe. When we have low energy/speed/whatever, we can often quite accurately approximate reality by a linear equation. For example, the equation of general relativity is very complicated and nowhere near being linear, but for small masses and velocities, they reduce to Newton's law of gravitation, which is linear.

The only exception to this seems to be Schrödinger's equation. While there are many theories and equations that superseded the Schrödinger equation, these are all still linear in nature. It seems that linearity is the thing that underpins quantum mechanics.

Due to the prevalence of linear equations, it is rather important that we understand these equations well, and this is our primary objective of the course.

1 Vector spaces

When dealing with functions and differential equations, we will often think of the space of functions as a vector space. In many cases, we will try to find a "basis" for our space of functions, and expand our functions in terms of the basis. Under different situations, we would want to use a different basis for our space. For example, when dealing with periodic functions, we will want to pick basis elements that are themselves periodic. In other situations, these basis elements would not be that helpful.

A familiar example would be the Taylor series, where we try to approximate a function f by

$$f(x) = \sum_{n=0}^{\infty} \frac{f^{(n)}(0)}{n!} x^n.$$

Here we are thinking of $\{x^n : n \in \mathbb{N}\}$ as the basis of our space, and trying to approximate an arbitrary function as a sum of the basis elements. When writing the function f as a sum like this, it is of course important to consider whether the sum converges, and when it does, whether it actually converges back to f.

Another issue of concern is if we have a general set of basis functions $\{y_n\}$, how can we find the coefficients c_n such that $f(x) = \sum c_n y_n(x)$? This is the bit where linear algebra comes in. Finding these coefficients is something we understand well in linear algebra, and we will attempt to borrow the results and apply them to our space of functions.

Another concept we would want to borrow is eigenvalues and eigenfunctions, as well as self-adjoint ("Hermitian") operators. As we go along the course, we will see some close connections between functions and vector spaces, and we can often get inspirations from linear algebra.

Of course, there is no guarantee that the results from linear algebra would apply directly, since most of our linear algebra results was about finite basis and finite linear sums. However, it is often a good starting point, and usually works when dealing with sufficiently nice functions.

We start with some preliminary definitions, which should be familiar from IA Vectors and Matrices and/or IB Linear Algebra.

Definition (Vector space). A *vector space* over $\mathbb{C}$ (or $\mathbb{R}$) is a set V with an operation $+$ which obeys

- (i) $\mathbf{u} + \mathbf{v} = \mathbf{v} + \mathbf{u}$ (commutativity)

- (ii) $(\mathbf{u} + \mathbf{v}) + \mathbf{w} = \mathbf{u} + (\mathbf{v} + \mathbf{w})$ (associativity)

- (iii) There is some $\mathbf{0} \in V$ such that $\mathbf{0} + \mathbf{u} = \mathbf{u}$ for all $\mathbf{u}$ (identity)

We can also multiply vectors by a scalars $\lambda \in \mathbb{C}$, which satisfies

- (i) $\lambda(\mu \mathbf{v}) = (\lambda\mu)\mathbf{v}$ (associativity)

- (ii) $\lambda(\mathbf{u} + \mathbf{v}) = \lambda\mathbf{u} + \lambda\mathbf{v}$ (distributivity in V)

- (iii) $(\lambda + \mu)\mathbf{u} = \lambda\mathbf{u} + \lambda\mathbf{v}$ (distributivity in $\mathbb{C}$)

- (iv) $1\mathbf{v} = \mathbf{v}$ (identity)

Often, we wouldn't have *just* a vector space. We usually give them some additional structure, such as an inner product.

Definition (Inner product). An *inner product* on V is a map $(\,\cdot\,,\,\cdot\,) : V \times V \to \mathbb{C}$ that satisfies

(i) $(\mathbf{u}, \lambda\mathbf{v}) = \lambda(\mathbf{u}, \mathbf{v})$ (linearity in second argument)

(ii) $(\mathbf{u}, \mathbf{v} + \mathbf{w}) = (\mathbf{u}, \mathbf{v}) + (\mathbf{u}, \mathbf{w})$ (additivity)

(iii) $(\mathbf{u}, \mathbf{v}) = (\mathbf{v}, \mathbf{u})^*$ (conjugate symmetry)

(iv) $(\mathbf{u}, \mathbf{u}) \geq 0$, with equality iff $\mathbf{u} = \mathbf{0}$ (positivity)

Note that the positivity condition makes sense since conjugate symmetry entails that $(\mathbf{u}, \mathbf{u}) \in \mathbb{R}$.

The inner product in turn defines a norm $\|\mathbf{u}\| = \sqrt{(\mathbf{u}, \mathbf{u})}$ that provides the notion of length and distance.

It is important to note that we only have linearity in the *second argument*. For the first argument, we have $(\lambda\mathbf{u}, \mathbf{v}) = (\mathbf{v}, \lambda\mathbf{u})^* = \lambda^*(\mathbf{v}, \mathbf{u})^* = \lambda^*(\mathbf{u}, \mathbf{v})$.

Definition (Basis). A set of vectors $\{\mathbf{v}_1, \mathbf{v}_2, \cdots, \mathbf{v}_n\}$ form a *basis* of V iff any $\mathbf{u} \in V$ can be uniquely written as a linear combination

$$\mathbf{u} = \sum_{i=1}^{n} \lambda_i \mathbf{v}_i$$

for some scalars λ_i. The *dimension* of a vector space is the number of basis vectors in its basis.

A basis is *orthogonal* (with respect to the inner product) if $(\mathbf{v}_i, \mathbf{v}_j) = 0$ whenever $i \neq j$.

A basis is *orthonormal* (with respect to the inner product) if it is orthogonal and $(\mathbf{v}_i, \mathbf{v}_i) = 1$ for all i.

Orthonormal bases are the nice bases, and these are what we want to work with.

Given an orthonormal basis, we can use the inner product to find the expansion of any $\mathbf{u} \in V$ in terms of the basis, for if

$$\mathbf{u} = \sum_i \lambda_i \mathbf{v}_i,$$

taking the inner product with $\mathbf{v}_j$ gives

$$(\mathbf{v}_j, \mathbf{u}) = \left(\mathbf{v}_j, \sum_i \lambda_i \mathbf{v}_i\right) = \sum_i \lambda_i(\mathbf{v}_j, \mathbf{v}_i) = \lambda_j,$$

using additivity and linearity. Hence we get the general formula

$$\lambda_i = (\mathbf{v}_i, \mathbf{u}).$$

We have seen all these so far in IA Vectors and Matrices, where a vector is a list of finitely many numbers. However, *functions* can also be thought of as elements of an (infinite dimensional) vector space.

Suppose we have $f, g : \Omega \to \mathbb{C}$. Then we can define the sum $f + g$ by $(f + g)(x) = f(x) + g(x)$. Given scalar λ, we can also define $(\lambda f)(x) = \lambda f(x)$.

This also makes intuitive sense. We can simply view a functions as a list of numbers, where we list out the values of f at each point. The list could be infinite, but a list nonetheless.

Most of the time, we don't want to look at the set of *all* functions. That would be too huge and uninteresting. A natural class of functions to consider would be the set of solutions to some particular differential solution. However, this doesn't always work. For this class to actually be a vector space, the sum of two solutions and the scalar multiple of a solution must also be a solution. This is exactly the requirement that the differential equation is linear. Hence, the set of solutions to a linear differential equation would form a vector space. Linearity pops up again.

Now what about the inner product? A natural definition is

$$(f, g) = \int_\Sigma f(x)^* g(x) \, \mathrm{d}\mu,$$

where μ is some measure. For example, we could integrate $\mathrm{d}x$, or $\mathrm{d}x^2$. This measure specifies how much weighting we give to each point x.

Why does this definition make sense? Recall that the usual inner product on finite-dimensional vector spaces is $\sum v_i^* w_i$. Here we are just summing the different components of v and w. We just said we can think of the function f as a list of all its values, and this integral is just the sum of all components of f and g.

Example. Let $\Sigma = [a, b]$. Then we could take

$$(f, g) = \int_a^b f(x)^* g(x) \, \mathrm{d}x.$$

Alternatively, let $\Sigma = D^2 \subseteq \mathbb{R}^2$ be the unit disk. Then we could have

$$(f, g) = \int_0^1 \int_0^{2\pi} f(r, \theta)^* g(r, \theta) \, \mathrm{d}\theta \, r \, \mathrm{d}r$$

Note that we were careful and said that $\mathrm{d}\mu$ is "some measure". Here we are integrating against $\mathrm{d}\theta \, r \, \mathrm{d}r$. We will later see cases where this can be even more complicated.

If Σ has a boundary, we will often want to restrict our functions to take particular values on the boundary, known as boundary conditions. Often, we want the boundary conditions to preserve linearity. We call these nice boundary conditions *homogeneous* conditions.

Definition (Homogeneous boundary conditions)**.** A boundary condition is *homogeneous* if whenever f and g satisfy the boundary conditions, then so does $\lambda f + \mu g$ for any $\lambda, \mu \in \mathbb{C}$ (or $\mathbb{R}$).

Example. Let $\Sigma = [a, b]$. We could require that $f(a) + 7f'(b) = 0$, or maybe $f(a) + 3f''(a) = 0$. These are examples of homogeneous boundary conditions. On the other hand, the requirement $f(a) = 1$ is *not* homogeneous.
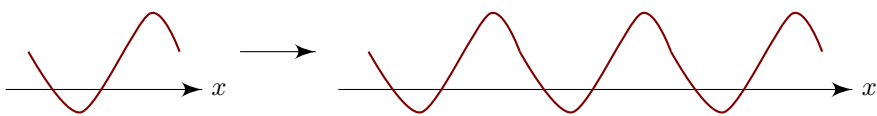
2 Fourier series

The first type of functions we will consider is periodic functions.

Definition (Periodic function)**.** A function f is *periodic* if there is some fixed R such that $f(x + R) = f(x)$ for all x.

However, it is often much more convenient to think of this as a function $f : S^1 \to \mathbb{C}$ from unit circle to $\mathbb{C}$, and parametrize our function by an angle θ ranging from 0 to 2π.

Now why do we care about periodic functions? Apart from the fact that genuine periodic functions exist, we can also use them to model functions on a compact domain. For example, if we have a function defined on $[0, 1]$, we can pretend it is a periodic function on $\mathbb{R}$ by making infinitely many copies of the function to the intervals $[1, 2]$, $[2, 3]$ etc.

2.1 Fourier series

So. As mentioned in the previous chapter, we want to find a set of "basis functions" for periodic functions. We could go with the simplest case of periodic functions we know of — the exponential function $e^{in\theta}$. These functions have a period of 2π, and are rather easy to work with. We all know how to integrate and differentiate the exponential function.

More importantly, this set of basis functions is orthogonal. We have

$$(e^{im\theta}, e^{in\theta}) = \int_{-\pi}^{\pi} e^{-im\theta} e^{in\theta} \, \mathrm{d}\theta = \int_{-\pi}^{\pi} e^{i(n-m)\theta} \, \mathrm{d}\theta = \begin{cases} 2\pi & n = m \\ 0 & n \neq m \end{cases} = 2\pi \delta_{nm}$$

We can normalize these to get a set of orthonormal functions $\left\{ \frac{1}{\sqrt{2\pi}} e^{in\theta} : n \in \mathbb{Z} \right\}$.

Fourier's idea was to use this as a basis for *any* periodic function. Fourier claimed that any $f : S^1 \to \mathbb{C}$ can be expanded in this basis:

$$f(\theta) = \sum_{n \in \mathbb{Z}} \hat{f}_n e^{in\theta},$$

where

$$\hat{f}_n = \frac{1}{2\pi}(e^{in\theta}, f) = \frac{1}{2\pi} \int_{-\pi}^{\pi} e^{-in\theta} f(\theta) \, \mathrm{d}\theta.$$

These really should be defined as $f(\theta) = \sum \hat{f}_n \frac{e^{in\theta}}{\sqrt{2\pi}}$ with $\hat{f}_n = \left(\frac{e^{in\theta}}{\sqrt{2\pi}}, f \right)$, but for convenience reasons, we move all the constant factors to the $\hat{f}_n$ coefficients.

We can consider the special case where $f : S^1 \to \mathbb{R}$ is a real function. We might want to make our expansion look a bit more "real". We get

$$(\hat{f}_n)^* = \left(\frac{1}{2\pi} \int_{-\pi}^{\pi} e^{-in\theta} f(\theta) \, \mathrm{d}\theta \right)^* = \frac{1}{2\pi} \int_{-\pi}^{\pi} e^{in\theta} f(\theta) \, \mathrm{d}\theta = \hat{f}_{-n}.$$

So we can replace our Fourier series by

$$f(\theta) = \hat{f}_0 + \sum_{n=1}^{\infty} \left(\hat{f}_n e^{in\theta} + \hat{f}_{-n} e^{-in\theta} \right) = \hat{f}_0 + \sum_{n=1}^{\infty} \left(\hat{f}_n e^{in\theta} + \hat{f}_n^* e^{-in\theta} \right).$$

Setting $\hat{f}_n = \dfrac{a_n + ib_n}{2}$, we can write this as

$$f(\theta) = \hat{f}_0 + \sum_{n=1}^{\infty} (a_n \cos n\theta + b_n \sin n\theta)$$

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

Here the coefficients are

$$a_n = \frac{1}{\pi} \int_{-\pi}^{\pi} \cos n\theta f(\theta) \, \mathrm{d}\theta, \quad b_n = \frac{1}{\pi} \int_{-\pi}^{\pi} \sin n\theta f(\theta) \, \mathrm{d}\theta.$$

This is an alternative formulation of the Fourier series in terms of sin and cos.

So when given a real function, which expansion should we use? It depends. If our function is odd (or even), it would be useful to pick the sine/cosine expansion, since the cosine (or sine) terms will simply disappear. On the other hand, if we want to stick our function into a differential equation, exponential functions are usually more helpful.

2.2   Convergence of Fourier series

When Fourier first proposed the idea of a Fourier series, people didn't really believe in him. How can we be sure that the infinite series actually converges? It turns out that in many cases, they don't.

To investigate the convergence of the series, we define the *partial Fourier sum* as

$$S_n f = \sum_{m=-n}^{n} \hat{f}_m e^{im\theta}.$$

The question we want to answer is whether $S_n f$ "converges" to f. Here we have to be careful with what we mean by convergence. As we (might) have seen in Analysis, there are many ways of defining convergence of functions. If we have a "norm" on the space of functions, we can define convergence to mean $\lim_{n\to\infty} \|S_n f - f\| = 0$. Our "norm" can be defined as

$$\|S_n f - f\| = \frac{1}{2\pi} \int_{-\pi}^{\pi} |S_n f(\theta) - f(\theta)|^2 \, \mathrm{d}\theta.$$

However, this doesn't really work if $S_n f$ and f can be arbitrary functions, as in this does not necessarily have to be a norm. Indeed, if $\lim_{n\to\infty} S_n f$ differs from f on finitely or countably many points, the integral will still be zero. In particular, $\lim S_n f$ and f can differ at all rational points, but this definition will say that $S_n f$ converges to f.

Hence another possible definition of convergence is to require

$$\lim_{n \to \infty} S_n f(\theta) - f(\theta) = 0$$

for all θ. This is known as *pointwise convergence.* However, this is often too weak a notion. We can ask for more, and require that the *rate* of convergent is independent of θ. This is known as uniform convergence, and is defined by

$$\lim_{n \to \infty} \sup_\theta |S_n f(\theta) - f(\theta)| = 0.$$

Of course, with different definitions of convergence, we can get different answers to whether it converges. Unfortunately, even if we manage to get our definition of convergence right, it is still difficult to come up with a criterion to decide if a function has a convergent Fourier Series.

It is not difficult to prove a general criterion for convergence in norm, but we will not do that, since that is analysis. If interested, one should take the IID Analysis of Functions course. In this course, instead of trying to come up with something general, let's look at an example instead.

Example. Consider the sawtooth function $f(\theta) = \theta$.
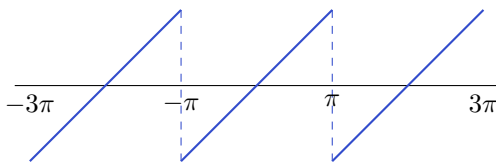
Note that this function is discontinuous at odd multiples of π.

The Fourier coefficients (for $n \neq 0$) are

$$\hat{f}_n = \frac{1}{2\pi} \int_{-\pi}^{\pi} e^{-in\theta} \theta \, \mathrm{d}\theta = \left[-\frac{1}{2\pi in} e^{-in\theta} \theta \right]_{-\pi}^{\pi} + \frac{1}{2\pi in} \int_{-\pi}^{\pi} e^{in\theta} \, \mathrm{d}\theta = \frac{(-1)^{n+1}}{in}.$$

We also have $\hat{f}_0 = 0$.

Hence we have

$$\theta = \sum_{n \neq 0} \frac{(-1)^{n+1}}{in} e^{in\theta}.$$

It turns out that this series converges to the sawtooth for all $\theta \neq (2m+1)\pi$, i.e. everywhere that the sawtooth is continuous.

Let's look explicitly at the case where $\theta = \pi$. Each term of the partial Fourier series is zero. So we can say that the Fourier series converges to 0. This is the *average* value of $\lim_{\varepsilon \to 0} f(\pi \pm \varepsilon)$.

This is typical. At an isolated discontinuity, the Fourier series is the average of the limiting values of the original function as we approach from either side.

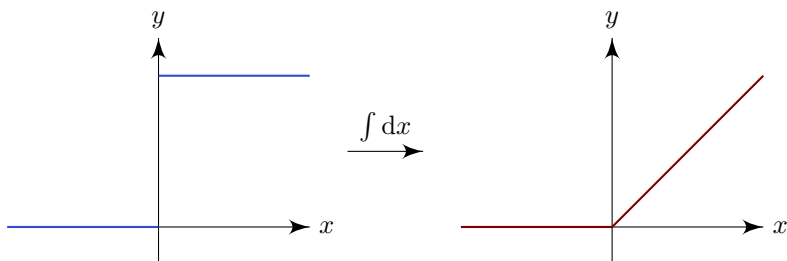## 2.3   Differentiability and Fourier series

Integration is a smoothening operator. If we take, say, the step function

$$\Theta(x) = \begin{cases} 1 & x > 0 \\ 0 & x < 0 \end{cases}.$$

Then the integral is given by

$$\int \Theta(x)\, dx = \begin{cases} x & x > 0 \\ 0 & x < 0 \end{cases}.$$

This is now a continuous function. If we integrate it again, we make the positive side look like a quadratic curve, and it is now differentiable.

On the other hand, when we differentiate a function, we generally make it worse. For example, when we differentiate the continuous function $\int \Theta(x)\, dx$, we obtain the discontinuous function $\Theta(x)$. If we attempt to differentiate this again, we get the δ-(non)-function.

Hence, it is often helpful to characterize the "smoothness" of a function by how many times we can differentiate it. It turns out this is rather relevant to the behaviour of the Fourier series.

Suppose that we have a function that is itself continuous and whose first $m - 1$ derivatives are also continuous, but $f^{(m)}$ has isolated discontinuities at $\{\theta_1, \theta_2, \theta_3, \cdots, \theta_r\}$.

We can look at the kth Fourier coefficient:

$$\hat{f}_k = \frac{1}{2\pi} \int_{-\pi}^{\pi} e^{-ik\theta} f(\theta)\, d\theta$$

$$= \left[-\frac{1}{2\pi ik} e^{-ik\theta} f(\theta) \right]_{-\pi}^{\pi} + \frac{1}{2\pi ik} \int_{-\pi}^{\pi} e^{-ik\theta} f'(\theta)\, d\theta$$

The first term vanishes since $f(\theta)$ is continuous and takes the same value at π and $-\pi$. So we are left with

$$= \frac{1}{2\pi ik} \int_{-\pi}^{\pi} e^{-ik\theta} f'(\theta)\, d\theta$$

$$= \cdots$$

$$= \frac{1}{(ik)^m} \frac{1}{2\pi} \int_{-\pi}^{\pi} e^{ik\theta} f^{(m)}(\theta)\, d\theta.$$

Now we have to be careful, since $f^{(m)}$ is no longer continuous. However provided that $f^{(m)}$ is everywhere finite, we can approximate this by removing small strips $(\theta_i - \varepsilon, \theta_i + \varepsilon)$ from our domain of integration, and take the limit $\varepsilon \to 0$. We can write this as

$$= \lim_{\varepsilon \to 0} \frac{1}{(ik)^m} \frac{1}{2\pi} \int_{-\pi}^{\theta_1 - \varepsilon} + \int_{\theta_1 + \varepsilon}^{\theta_2 - \varepsilon} + \cdots + \int_{\theta_r + \varepsilon}^{\pi} e^{ik\theta} f^{(m)}(\theta)\, d\theta$$

$$= \frac{1}{(ik)^{m+1}} \frac{1}{2\pi} \left[\sum_{s=1}^{r} (f^{(m)}(\theta_s^+) - f^{(m)}(\theta_s^-)) + \int_{(-\pi,\pi)\backslash\theta} e^{-ik\theta} f^{(m+1)}(\theta)\, d\theta \right].$$

We now have to stop. So $\hat{f}_k$ decays like $\left(\frac{1}{k}\right)^{m+1}$ if our function and its $(m-1)$ derivatives are continuous. This means that if a function is more differentiable, then the coefficients decay more quickly.

This makes sense. If we have a rather smooth function, then we would expect the first few Fourier terms (with low frequency) to account for most of the variation of f. Hence the coefficients decay really quickly.

However, if the function is jiggly and bumps around all the time, we would expect to need some higher frequency terms to account for the minute variation. Hence the terms would not decay away that quickly. So in general, if we can differentiate it more times, then the terms should decay quicker.

An important result, which is in some sense a Linear Algebra result, is *Parseval's theorem.*

Theorem (Parseval's theorem).

$$(f, f) = \int_{-\pi}^{\pi} |f(\theta)|^2 \, \mathrm{d}\theta = 2\pi \sum_{n \in \mathbb{Z}} |\hat{f}_n|^2$$

Proof.

$$
\begin{aligned}
(f, f) &= \int_{-\pi}^{\pi} |f(\theta)|^2 \, \mathrm{d}\theta \\
&= \int_{-\pi}^{\pi} \left(\sum_{m \in \mathbb{Z}} \hat{f}_m^* e^{-im\theta} \right) \left(\sum_{n \in \mathbb{Z}} \hat{f}_n e^{in\theta} \right) \, \mathrm{d}\theta \\
&= \sum_{m,n \in \mathbb{Z}} \hat{f}_m^* \hat{f}_n \int_{-\pi}^{\pi} e^{i(n-m)\theta} \, \mathrm{d}\theta \\
&= 2\pi \sum_{m,n \in \mathbb{Z}} \hat{f}_m^* \hat{f}_n \delta_{mn} \\
&= 2\pi \sum_{n \in \mathbb{Z}} |\hat{f}_n|^2 \qquad\qquad\qquad \square
\end{aligned}
$$

Note that this is not a fully rigorous proof, since we assumed not only that the Fourier series converges to the function, but also that we could commute the infinite sums. However, for the purposes of an applied course, this is sufficient.

Last time, we computed that the sawtooth function $f(\theta) = \theta$ has Fourier coefficients

$$\hat{f}_0 = 0, \quad \hat{f}_n = \frac{i(-1)^n}{n} \text{ for } n \neq 0.$$

But why do we care? It turns out this has some applications in number theory. You might have heard of the Riemann ζ-function, defined by

$$\zeta(s) = \sum_{n=1}^{\infty} \frac{1}{n^s}.$$

We will show that this obeys the property that for any m, $\zeta(2m) = \pi^{2m} q$ for some $q \in \mathbb{Q}$. This may not be obvious at first sight. So let's apply Parseval's theorem for the sawtooth defined by $f(\theta) = \theta$. By direct computation, we know that

$$(f, f) = \int_{-\pi}^{\pi} \theta^2 \, \mathrm{d}\theta = \frac{2\pi^3}{3}.$$

However, by Parseval's theorem, we know that

$$(f, f) = 2\pi \sum_{n \in \mathbb{Z}} |\hat{f}_n|^2 = 4\pi \sum_{n=1}^{\infty} \frac{1}{n^2}.$$

Putting these together, we learn that

$$\sum_{n=1}^{\infty} \frac{1}{n^2} = \zeta(2) = \frac{\pi^2}{6}.$$

We have just done it for the case where $m = 1$. But if we integrate the sawtooth function repeatedly, then we can get the general result for all m.

At this point, we might ask, why are we choosing these $e^{im\theta}$ as our basis? Surely there are a lot of different sets of basis we can use. For example, in finite dimensions, we can just arbitrary choose random vectors (that are not linearly dependent) to get a set of basis vectors. However, in practice, we don't pick them randomly. We often choose a basis that can reveal the symmetry of a system. For example, if we have a rotationally symmetric system, we would like to use polar coordinates. Similarly, if we have periodic functions, then $e^{im\theta}$ is often a good choice of basis.

3 Sturm-Liouville Theory

3.1 Sturm-Liouville operators

In finite dimensions, we often consider *linear maps* $M : V \to W$. If $\{\mathbf{v}_i\}$ is a basis for V and $\{\mathbf{w}_i\}$ is a basis for W, then we can represent the map by a matrix with entries

$$M_{ai} = (\mathbf{w}_a, M\mathbf{v}_i).$$

A map $M : V \to V$ is called *self-adjoint* if $M^\dagger = M$ as matrices. However, it is not obvious how we can extend this notion to arbitrary maps between arbitrary vector spaces (with an inner product) when they cannot be represented by a matrix.

Instead, we make the following definitions:

Definition (Adjoint and self-adjoint)**.** The *adjoint B* of a map $A : V \to V$ is a map such that

$$(B\mathbf{u}, \mathbf{v}) = (\mathbf{u}, A\mathbf{v})$$

for all vectors $\mathbf{u}, \mathbf{v} \in V$. A map is then *self-adjoint* if

$$(M\mathbf{u}, \mathbf{v}) = (\mathbf{u}, M\mathbf{v}).$$

Self-adjoint matrices come with a natural basis. Recall that the *eigenvalues* of a matrix are the roots of $\det(M - \lambda I) = 0$. The *eigenvector* corresponding to an eigenvalue λ is defined by $M\mathbf{v}_i = \lambda_i \mathbf{v}_i$.

In general, eigenvalues can be any complex number. However, self-adjoint maps have *real* eigenvalues. Suppose

$$M\mathbf{v}_i = \lambda_i \mathbf{v}_i.$$

Then we have

$$\lambda_i(\mathbf{v}_i, \mathbf{v}_i) = (\mathbf{v}_i, M\mathbf{v}_i) = (M\mathbf{v}_i, \mathbf{v}_i) = \lambda_i^*(\mathbf{v}_i, \mathbf{v}_i).$$

So $\lambda_i = \lambda_i^*$.

Furthermore, eigenvectors with distinct eigenvalues are orthogonal with respect to the inner product. Suppose that

$$M\mathbf{v}_i = \lambda_i \mathbf{v}_i, \quad M\mathbf{v}_j = \lambda_j \mathbf{v}_j.$$

Then

$$\lambda_i(\mathbf{v}_j, \mathbf{v}_i) = (\mathbf{v}_j, M\mathbf{v}_i) = (M\mathbf{v}_j, \mathbf{v}_i) = \lambda_j(\mathbf{v}_j, \mathbf{v}_i).$$

Since $\lambda_i \neq \lambda_j$, we must have $(\mathbf{v}_j, \mathbf{v}_i) = 0$.

Knowing eigenvalues and eigenvalues gives a neat way so solve linear equations of the form

$$M\mathbf{u} = \mathbf{f}.$$

Here we are given M and $\mathbf{f}$, and want to find $\mathbf{u}$. Of course, the answer is $\mathbf{u} = M^{-1}\mathbf{f}$. However, if we expand in terms of eigenvectors, we obtain

$$M\mathbf{u} = M \sum u_i \mathbf{v}_i = \sum u_i \lambda_i \mathbf{v}_i.$$

Hence we have

$$\sum u_i \lambda_i \mathbf{v}_i = \sum f_i \mathbf{v}_i.$$

Taking the inner product with $\mathbf{v}_j$, we know that

$$u_j = \frac{f_j}{\lambda_j}.$$

So far, these are all things from IA Vectors and Matrices. Sturm-Liouville theory is the infinite-dimensional analogue.

In our vector space of differentiable functions, our "matrices" would be linear differential operators $\mathcal{L}$. For example, we could have

$$\mathcal{L} = A_p(x)\frac{\mathrm{d}^p}{\mathrm{d}x^p} + A_{p-1}(x)\frac{\mathrm{d}^{p-1}}{\mathrm{d}x^{p-1}} + \cdots + A_1(x)\frac{\mathrm{d}}{\mathrm{d}x} + A_0(x).$$

It is an easy check that this is in fact linear.

We say $\mathcal{L}$ has order p if the highest derivative that appears is $\frac{\mathrm{d}^p}{\mathrm{d}x^p}$.

In most applications, we will be interested in the case $p = 2$. When will our $\mathcal{L}$ be self-adjoint?

In the $p = 2$ case, we have

$$\begin{aligned}
\mathcal{L}y &= P\frac{\mathrm{d}^2 y}{\mathrm{d}x^2} + R\frac{\mathrm{d}y}{\mathrm{d}x} - Qy \\
&= P\left[\frac{\mathrm{d}^2 y}{\mathrm{d}x^2} + \frac{R}{P}\frac{\mathrm{d}y}{\mathrm{d}x} - \frac{Q}{P}y\right] \\
&= P\left[e^{-\int \frac{R}{P}\,\mathrm{d}x}\frac{\mathrm{d}}{\mathrm{d}x}\left(e^{\int \frac{R}{P}\,\mathrm{d}x}\frac{\mathrm{d}y}{\mathrm{d}x}\right) - \frac{Q}{P}y\right]
\end{aligned}$$

Let $p = \exp\left(\int \frac{R}{P}\,\mathrm{d}x\right)$. Then we can write this as

$$= Pp^{-1}\left[\frac{\mathrm{d}}{\mathrm{d}x}\left(p\frac{\mathrm{d}y}{\mathrm{d}x}\right) - \frac{Q}{P}py\right].$$

We further define $q = \frac{Q}{P}p$. We also drop a factor of Pp^{-1}. Then we are left with

$$\mathcal{L} = \frac{\mathrm{d}}{\mathrm{d}x}\left(p(x)\frac{\mathrm{d}}{\mathrm{d}x}\right) - q(x).$$

This is the *Sturm-Liouville form* of the operator. Now let's compute $(f, \mathcal{L}g)$. We integrate by parts numerous times to obtain

$$\begin{aligned}
(f, \mathcal{L}g) &= \int_a^b f^*\left(\frac{\mathrm{d}}{\mathrm{d}x}\left(p\frac{\mathrm{d}g}{\mathrm{d}x}\right) - qg\right)\,\mathrm{d}x \\
&= [f^* p g']_a^b - \int_a^b \left(\frac{\mathrm{d}f^*}{\mathrm{d}x}p\frac{\mathrm{d}g}{\mathrm{d}x} + f^* qg\right)\,\mathrm{d}x \\
&= [f^* pg' - f'^* pg]_a^b + \int_a^b \left(\frac{\mathrm{d}}{\mathrm{d}x}\left(p\frac{\mathrm{d}f^*}{\mathrm{d}x}\right) - qf^*\right)g\,\mathrm{d}x \\
&= [(f^* g' - f'^* g)p]_a^b + (\mathcal{L}f, g),
\end{aligned}$$

assuming that p, q are real.

So 2nd order linear differential operators are self-adjoint with respect to this norm if p, q are real and the boundary terms vanish. When do the boundary terms vanish? One possibility is when p is periodic (with the right period), or if we constrain f and g to be periodic.

Example. We can consider a simple case, where

$$\mathcal{L} = \frac{\mathrm{d}^2}{\mathrm{d}x^2}.$$

Here we have $p = 1, q = 0$. If we ask for functions to be periodic on $[a, b]$, then

$$\int_a^b f^* \frac{\mathrm{d}^2 g}{\mathrm{d}x^2} \, \mathrm{d}x = \int_a^b \frac{\mathrm{d}^2 f^*}{\mathrm{d}x^2} g \, \mathrm{d}x.$$

Note that it is important that we have a *second-order* differential operator. If it is first-order, then we would have a negative sign, since we integrated by parts once.

Just as in finite dimensions, self-adjoint operators have eigenfunctions and eigenvalues with special properties. First, we define a more sophisticated inner product.

Definition (Inner product with weight)**.** An *inner product with weight w*, written $(\,\cdot\,, \,\cdot\,)_w$, is defined by

$$(f, g)_w = \int_a^b f^*(x) g(x) w(x) \, \mathrm{d}x,$$

where w is real, non-negative, and has only finitely many zeroes.

Why do we want a weight $w(x)$? In the future, we might want to work with the unit disk, instead of a square in $\mathbb{R}^2$. When we want to use polar coordinates, we will have to integrate with $r \, \mathrm{d}r \, \mathrm{d}\theta$, instead of just $\mathrm{d}r \, \mathrm{d}\theta$. Hence we need the weight of r. Also, we allow it to have finitely many zeroes, so that the radius can be 0 at the origin.

Why can't we have more zeroes? We want the inner product to keep the property that $(f, f)_w = 0$ iff $f = 0$ (for continuous f). If w is zero at too many places, then the inner product could be zero without f being zero.

We now define what it means to be an eigenfunction.

Definition (Eigenfunction with weight)**.** An *eigenfunction with weight w* of $\mathcal{L}$ is a function $y : [a, b] \to \mathbb{C}$ obeying the differential equation

$$\mathcal{L}y = \lambda w y,$$

where $\lambda \in \mathbb{C}$ is the eigenvalue.

This might be strange at first sight. It seems like we can take any nonsense y, apply $\mathcal{L}$, to get some nonsense $\mathcal{L}y$. But then it is fine, since we can write it as some nonsense w times our original y. So any function is an eigenfunction? No! There are many constraints w has to satisfy, like being positive, real and having finitely many zeroes. It turns out this severely restraints what values y can take, so not everything will be an eigenfunction. In fact we can develop this theory without the weight function w. However, weight functions are much more convenient when, say, dealing with the unit disk.

Proposition. The eigenvalues of a Sturm-Liouville operator are real.

Proof. Suppose $\mathcal{L}y_i = \lambda_i w y_i$. Then

$$\lambda_i(y_i, y_i)_w = \lambda_i(y_i, wy_i) = (y_i, \mathcal{L}y_i) = (\mathcal{L}y_i, y_i) = (\lambda_i w y_i, y_i) = \lambda_i^*(y_i, y_i)_w.$$

Since $(y_i, y_i)_w \neq 0$, we have $\lambda_i = \lambda_i^*$. $\qquad\square$

Note that the first and last terms use the weighted inner product, but the middle terms use the unweighted inner product.

Proposition. Eigenfunctions with different eigenvalues (but same weight) are orthogonal.

Proof. Let $\mathcal{L}y_i = \lambda_i w y_i$ and $\mathcal{L}y_j = \lambda_j w y_j$. Then

$$\lambda_i(y_j, y_i)_w = (y_j, \mathcal{L}y_i) = (\mathcal{L}y_j, y_i) = \lambda_j(y_j, y_i)_w.$$

Since $\lambda_i \neq \lambda_j$, we must have $(y_j, y_i)_w = 0$. $\qquad\square$

Those were pretty straightforward manipulations. However, the main results of Sturm–Liouville theory are significantly harder, and we will not prove them. We shall just state them and explore some examples.

Theorem. On a *compact* domain, the eigenvalues $\lambda_1, \lambda_2, \cdots$ form a countably infinite sequence and are discrete.

This will be a rather helpful result in quantum mechanics, since in quantum mechanics, the possible values of, say, the energy are the eigenvalues of the Hamiltonian operator. Then this result says that the possible values of the energy are discrete and form an infinite sequence.

Note here the word *compact*. In quantum mechanics, if we restrict a particle in a well $[0, 1]$, then it will have quantized energy level since the domain is compact. However, if the particle is free, then it can have any energy at all since we no longer have a compact domain. Similarly, angular momentum is quantized, since it describe rotations, which takes values in S^1, which is compact.

Theorem. The eigenfunctions are complete: any function $f : [a, b] \to \mathbb{C}$ (obeying appropriate boundary conditions) can be expanded as

$$f(x) = \sum_n \hat{f}_n y_n(x),$$

where

$$\hat{f}_n = \int y_n^*(x) f(x) w(x) \, \mathrm{d}x.$$

Example. Let $[a, b] = [-L, L]$, $\mathcal{L} = \frac{\mathrm{d}^2}{\mathrm{d}x^2}$, $w = 1$, restricting to periodic functions Then our eigenfunction obeys

$$\frac{\mathrm{d}^2 y_n}{\mathrm{d}x^2} = \lambda_n y_n(x),$$

Then our eigenfunctions are

$$y_n(x) = \exp\left(\frac{in\pi x}{L}\right)$$

with eigenvalues

$$\lambda_n = -\frac{n^2\pi^2}{L^2}$$

for $n \in \mathbb{Z}$. This is just the Fourier series!

Example (Hermite polynomials)**.** We are going to cheat a little bit and pick our domain to be $\mathbb{R}$. We want to study the differential equation

$$\frac{1}{2}H'' - xH' = \lambda H,$$

with $H : \mathbb{R} \to \mathbb{C}$. We want to put this in Sturm-Liouville form. We have

$$p(x) = \exp\left(-\int_0^x 2t\,\mathrm{d}t\right) = e^{-x^2},$$

ignoring constant factors. Then $q(x) = 0$. We can rewrite this as

$$\frac{\mathrm{d}}{\mathrm{d}x}\left(e^{-x^2}\frac{\mathrm{d}H}{\mathrm{d}x}\right) = -2\lambda e^{-x^2}H(x).$$

So we take our weight function to be $w(x) = e^{-x^2}$.

We now ask that $H(x)$ grows at most polynomially as $|x| \to \infty$. In particular, we want $e^{-x^2}H(x)^2 \to 0$. This ensures that the boundary terms from integration by parts vanish at the infinite boundary, so that our Sturm-Liouville operator is self-adjoint.

The eigenfunctions turn out to be

$$H_n(x) = (-1)^n e^{x^2}\frac{\mathrm{d}^n}{\mathrm{d}x^n}\left(e^{-x^2}\right).$$

These are known as the *Hermite polynomials*. Note that these are indeed polynomials. When we differentiate the e^{-x^2} term many times, we get a lot of things from the product rule, but they will always keep an e^{-x^2}, which will ultimately cancel with e^{x^2}.

Just as for matrices, we can use the eigenfunction expansion to solve forced differential equations. For example, if might want to solve

$$\mathcal{L}g = f(x),$$

where $f(x)$ is a forcing term. We can write this as

$$\mathcal{L}g = w(x)F(x).$$

We expand our g as

$$g(x) = \sum_{n\in\mathbb{Z}}\hat{g}_n y_n(x).$$

Then by linearity,

$$\mathcal{L}g = \sum_{n\in\mathbb{Z}}\hat{g}_n\mathcal{L}y_n(x) = \sum_{n\in\mathbb{Z}}\hat{g}_n\lambda_n w(x)y_n(x).$$

We can also expand our forcing function as

$$w(x)F(x) = w(x)\left(\sum_{n\in\mathbb{Z}} \hat{F}_n y_n(x)\right).$$

Taking the (regular) inner product with $y_m(x)$ (and noting orthogonality of eigenfunctions), we obtain

$$w(x)\hat{g}_m\lambda_m = w(x)\hat{F}_m$$

This tells us that

$$\hat{g}_m = \frac{\hat{F}_m}{\lambda_m}.$$

So we have

$$g(x) = \sum_{n\in\mathbb{Z}} \frac{\hat{F}_n}{\lambda_n} y_n(x),$$

provided all λ_n are non-zero.

This is a systematic way of solving forced differential equations. We used to solve these by "being smart". We just looked at the forcing term and tried to guess what would work. Unsurprisingly, this approach does not succeed all the time. Thus it is helpful to have a systematic way of solving the equations.

It is often helpful to rewrite this into another form, using the fact that $\hat{F}_n = (y_n, F)_w$. So we have

$$g(x) = \sum_{n\in\mathbb{Z}} \frac{1}{\lambda_n}(y_n, F)_w y_n(x) = \int_a^b \sum_{n\in\mathbb{Z}} \frac{1}{\lambda_n} y_n^*(t) y_n(x) w(t) F(t) \, \mathrm{d}t.$$

Note that we swapped the sum and the integral, which is in general a dangerous thing to do, but we don't really care because this is an applied course. We can further write the above as

$$g(x) = \int_a^b G(x,t)F(t)w(t) \, \mathrm{d}t,$$

where $G(x,t)$ is the infinite sum

$$G(x,t) = \sum_{n\in\mathbb{Z}} \frac{1}{\lambda_n} y_n^*(t) y_n(x).$$

We call this the Green's function. Note that this depends on λ_n and y_n only. It depends on the differential operator $\mathcal{L}$, but not the forcing term F. We can think of this as something like the "inverse matrix", which we can use to solve the forced differential equation for any forcing term.

Recall that for a matrix, the inverse exists if the determinant is non-zero, which is true if the eigenvalues are all non-zero. Similarly, here a necessary condition for the Green's function to exist is that all the eigenvalues are non-zero.

We now have a second version of Parseval's theorem.

Theorem (Parseval's theorem II)**.**

$$(f,f)_w = \sum_{n\in\mathbb{Z}} |\hat{f}_n|^2.$$

Proof. We have

$$
\begin{aligned}
(f, f)_w &= \int_\Omega f^*(x) f(x) w(x) \, \mathrm{d}x \\
&= \sum_{n,m \in \mathbb{Z}} \int_\Omega \hat{f}_n^* y_n^*(x) \hat{f}_m y_m(x) w(x) \, \mathrm{d}x \\
&= \sum_{n,m \in \mathbb{Z}} \hat{f}_n^* \hat{f}_m (y_n, y_m)_w \\
&= \sum_{n \in \mathbb{Z}} |\hat{f}_n|^2. \qquad \square
\end{aligned}
$$

3.2   Least squares approximation

So far, we have expanded functions in terms of infinite series. However, in real life, when we ask a computer to do this for us, it is incapable of storing and calculating an infinite series. So we need to truncate it at some point.

Suppose we approximate some function $f : \Omega \to \mathbb{C}$ by a *finite* set of eigenfunctions $y_n(x)$. Suppose we write the approximation g as

$$
g(x) = \sum_{k=1}^n c_k y_k(x).
$$

The objective here is to figure out what values of the coefficients c_k are "the best", i.e. can make g represent f as closely as possible.

Firstly, we have to make it explicit what we mean by "as closely as possible". Here we take this to mean that we want to minimize the w-norm $(f - g, f - g)_w$. By linearity of the norm, we know that

$$
(f - g, f - g)_w = (f, f)_w + (g, g)_w - (f, g)_w - (g, f)_w.
$$

To minimize this norm, we want

$$
0 = \frac{\partial}{\partial c_j}(f - g, f - g)_w = \frac{\partial}{\partial c_j}[(f, f)_w + (g, g)_w - (f, g)_w - (g, f)_w].
$$

We know that the $(f, f)_w$ term vanishes since it does not depend on c_k. Expanding our definitions of g, we can get

$$
0 = \frac{\partial}{\partial c_j} \left(\sum_{i=1}^\infty |c_k|^n - \sum_{k=1}^n \hat{f}_k^* c_k - \sum_{k=1}^n c_k^* \hat{f}_k \right) = c_j^* - \hat{f}_j^*.
$$

Note that here we are treating c_j^* and c_j as distinct quantities. So when we vary c_j, c_j^* is unchanged. To formally justify this treatment, we can vary the real and imaginary parts separately.

Hence, the extremum is achieved at $c_j^* = \hat{f}_j^*$. Similarly, varying with respect to c_j^*, we get that $c_j = \hat{f}_j$.

To check that this is indeed an minimum, we can look at the second-derivatives

$$
\frac{\partial^2}{\partial c_i \partial c_j}(f - g, f - g)_w = \frac{\partial^2}{\partial c_i^* c_j^*}(f - g, f - g)_w = 0,
$$

while

$$\frac{\partial^2}{\partial c_i^* \partial c_j}(f - g, f - g)_w = \delta_{ij} \geq 0.$$

Hence this is indeed a minimum.

Thus we know that $(f - g, f - g)_w$ is minimized over all $g(x)$ when

$$c_k = \hat{f}_k = (y_k, f)_w.$$

These are exactly the coefficients in our infinite expansion. Hence if we truncate our series at an arbitrary point, it is still the best approximation we can get.

4   Partial differential equations

So far, we have come up with a lot of general theory about functions and stuff. We are going to look into some applications of these. In this section, we will develop the general technique of *separation of variables*, and apply this to many many problems.

4.1   Laplace's equation

Definition (Laplace's equation)**.** *Laplace's equation* on $\mathbb{R}^n$ says that a (twice-differentiable) equation $\phi : \mathbb{R}^n \to \mathbb{C}$ obeys

$$\nabla^2 \phi = 0,$$

where

$$\nabla^2 = \sum_{i=1}^{n} \frac{\partial^2}{\partial x_i^2}.$$

This equation arises in many situations. For example, if we have a conservative force $\mathbf{F} = -\nabla\phi$, and we are in a region where the force obeys the source-free Gauss' law $\nabla \cdot \mathbf{F} = 0$ (e.g. when there is no electric charge in an electromagnetic field), then we get Laplace's equation.

It also arises as a special case of the heat equation $\frac{\partial \phi}{\partial t} = \kappa \nabla^2 \phi$, and wave equation $\frac{\partial^2 \phi}{\partial t^2} = c^2 \nabla^2 \phi$, when there is no time dependence.

Definition (Harmonic functions)**.** Functions that obey Laplace's equation are often called *harmonic functions*.

Proposition. Let Ω be a compact domain, and $\partial\Omega$ be its boundary. Let $f : \partial\Omega \to \mathbb{R}$. Then there is a unique solution to $\nabla^2 \phi = 0$ on Ω with $\phi|_{\partial\Omega} = f$.

Proof. Suppose ϕ_1 and ϕ_2 are both solutions such that $\phi_1|_{\partial\Omega} = \phi_2|_{\partial\Omega} = f$. Then let $\Phi = \phi_1 - \phi_2$. So $\Phi = 0$ on the boundary. So we have

$$0 = \int_\Omega \Phi \nabla^2 \Phi \, \mathrm{d}^n x = -\int_\Omega (\nabla\Phi) \cdot (\nabla\Phi) \, \mathrm{d}x + \int_{\partial\Omega} \Phi \nabla\Phi \cdot \mathbf{n} \, \mathrm{d}^{n-1} x.$$

We note that the second term vanishes since $\Phi = 0$ on the boundary. So we have

$$0 = -\int_\Omega (\nabla\Phi) \cdot (\nabla\Phi) \, \mathrm{d}x.$$

However, we know that $(\nabla\Phi) \cdot (\nabla\Phi) \geq 0$ with equality iff $\nabla\Phi = 0$. Hence Φ is constant throughout Ω. Since at the boundary, $\Phi = 0$, we have $\Phi = 0$ throughout, i.e. $\phi_1 = \phi_2$. $\qquad\square$

Asking that $\phi|_{\partial\Omega} = f(x)$ is a *Dirichlet* boundary condition. We can instead ask $\mathbf{n} \cdot \nabla\phi|_{\partial\Omega} = g$, a *Neumann* condition. In this case, we have a unique solution up to a constant factor.

These we've all seen in IA Vector Calculus, but this time in arbitrary dimensions.

4.2   Laplace's equation in the unit disk in $\mathbb{R}^2$

Let $\Omega = \{(x, y) \in \mathbb{R}^2 : x^2 + y^2 \leq 1\}$. Then Laplace's equation becomes

$$0 = \nabla^2 \phi = \frac{\partial^2 \phi}{\partial x^2} + \frac{\partial^2 \phi}{\partial y^2} = \frac{\partial^2 \phi}{\partial z \partial \bar{z}},$$

where we define $z = x + iy, \bar{z} = x - iy$. Then the general solution is

$$\phi(z, \bar{z}) = \psi(z) + \chi(\bar{z})$$

for some ψ, χ.

Suppose we want a solution obeying the Dirichlet boundary condition $\phi(z, \bar{z})|_{\partial \Omega} = f(\theta)$, where the boundary $\partial \Omega$ is the unit circle. Since the domain is the unit circle S^1, we can Fourier-expand it! We can write

$$f(\theta) = \sum_{n \in \mathbb{Z}} \hat{f}_n e^{in\theta} = \hat{f}_0 + \sum_{n=1}^{\infty} \hat{f}_n e^{in\theta} + \sum_{n=1}^{\infty} \hat{f}_{-n} e^{-in\theta}.$$

However, we know that $z = re^{i\theta}$ and $\bar{z} = re^{i\theta}$. So on the boundary, we know that $z = e^{i\theta}$ and $\bar{z} = e^{-i\theta}$. So we can write

$$f(\theta) = \hat{f}_0 + \sum_{n=1}^{\infty} (\hat{f}_n z^n + \hat{f}_{-n} \bar{z}^n)$$

This is defined for $|z| = 1$. Now we can define ϕ on the unit disk by

$$\phi(z, \bar{z}) = \hat{f}_0 + \sum_{n=1}^{\infty} \hat{f}_n z^n + \sum_{n=1}^{\infty} \hat{f}_{-n} \bar{z}^n.$$

It is clear that this obeys the boundary conditions by construction. Also, $\phi(z, \bar{z})$ certainly converges when $|z| < 1$ if the series for $f(\theta)$ converged on $\partial \Omega$. Also, $\phi(z, \bar{z})$ is of the form $\psi(z) + \chi(\bar{z})$. So it is a solution to Laplace's equation on the unit disk. Since the solution is unique, we're done!

4.3   Separation of variables

Unfortunately, the use of complex variables is very special to the case where $\Omega \subseteq \mathbb{R}^2$. In higher dimensions, we proceed differently.

We let $\Omega = \{(x, y, z) \in \mathbb{R}^3 : 0 \leq x \leq a, 0 \leq y \leq b, z \geq 0\}$. We want the following boundary conditions:

$$\psi(0, y, z) = \psi(a, y, z) = 0$$
$$\psi(x, 0, z) = \psi(x, b, z) = 0$$
$$\lim_{z \to \infty} \psi(x, y, z) = 0$$
$$\psi(x, y, 0) = f(x, y),$$

where f is a given function. In other words, we want our function to be f when $z = 0$ and vanish at the boundaries.

The first step is to look for a solution of $\nabla^2 \psi(x, y, z) = 0$ of the form

$$\psi(x, y, z) = X(x)Y(y)Z(z).$$

Then we have

$$0 = \nabla^2 \psi = YZX'' + XZY'' + XYZ'' = XYZ\left(\frac{X''}{X} + \frac{Y''}{Y} + \frac{Z''}{Z}\right).$$

As long as $\psi \neq 0$, we can divide through by ψ and obtain

$$\frac{X''}{X} + \frac{Y''}{Y} + \frac{Z''}{Z} = 0.$$

The key point is each term depends on only *one* of the variables (x, y, z). If we vary, say, x, then $\frac{Y''}{Y} + \frac{Z''}{Z}$ does not change. For the total sum to be 0, $\frac{X''}{X}$ must be constant. So each term has to be separately constant. We can thus write

$$X'' = -\lambda X, \quad Y'' = -\mu Y, \quad Z'' = (\lambda + \mu)Z.$$

The signs before λ and μ are there just to make our life easier later on.

We can solve these to obtain

$$X = a \sin\sqrt{\lambda}x + b \cos\sqrt{\lambda}x,$$
$$Y = c \sin\sqrt{\lambda}y + d \cos\sqrt{\lambda}x,$$
$$Z = g \exp(\sqrt{\lambda + \mu}z) + h \exp(-\sqrt{\lambda + \mu}z).$$

We now impose the homogeneous boundary conditions, i.e. the conditions that ψ vanishes at the walls and at infinity.

- At $x = 0$, we need $\psi(0, y, z) = 0$. So $b = 0$.

- At $x = a$, we need $\psi(a, y, z) = 0$. So $\lambda = \left(\frac{n\pi}{a}\right)^2$.

- At $y = 0$, we need $\psi(x, 0, z) = 0$. So $d = 0$.

- At $y = b$, we need $\psi(x, b, z) = 0$. So $\mu = \left(\frac{m\pi}{b}\right)^2$.

- As $z \to \infty$, $\psi(x, y, z) \to 0$. So $g = 0$.

Here we have $n, m \in \mathbb{N}$.

So we found a solution

$$\psi(x, y, z) = A_{n,m} \sin\left(\frac{n\pi x}{a}\right) \sin\left(\frac{m\pi y}{b}\right) \exp(-s_{n,m}z),$$

where $A_{n,m}$ is an arbitrary constant, and

$$s_{n,m}^2 = \left(\frac{n^2}{a^2} + \frac{m^2}{b^2}\right)\pi^2.$$

This obeys the homogeneous boundary conditions for any $n, m \in \mathbb{N}$ but not the inhomogeneous condition at $z = 0$.

By linearity, the general solution obeying the homogeneous boundary conditions is

$$\psi(x, y, z) = \sum_{n,m=1}^{\infty} A_{n,m} \sin\left(\frac{n\pi x}{a}\right) \sin\left(\frac{m\pi y}{b}\right) \exp(-s_{n,m}z)$$

The final step is to impose the inhomogeneous boundary condition $\psi(x, y, 0) = f(x, y)$. In other words, we need

$$\sum_{n,m=1}^{\infty} A_{n,m} \sin\left(\frac{n\pi x}{a}\right) \sin\left(\frac{m\pi y}{b}\right) = f(x, y). \tag{$\dagger$}$$

The objective is thus to find the $A_{n,m}$ coefficients.

We can use the orthogonality relation we've previously had:

$$\int_0^a \sin\left(\frac{k\pi x}{a}\right) \sin\left(\frac{n\pi x}{a}\right) \, \mathrm{d}x = \delta_{k,n} \frac{a}{2}.$$

So we multiply ($\dagger$) by $\sin\left(\frac{k\pi x}{a}\right)$ and integrate:

$$\sum_{n,m=1}^{\infty} A_{n,m} \int_0^a \sin\left(\frac{k\pi x}{a}\right) \sin\left(\frac{n\pi x}{a}\right) \, \mathrm{d}x \sin\left(\frac{m\pi y}{b}\right) = \int_0^a \sin\left(\frac{k\pi x}{a}\right) f(x, y) \, \mathrm{d}x.$$

Using the orthogonality relation, we have

$$\frac{a}{2} \sum_{m=1}^{\infty} A_{k,m} \sin\left(\frac{m\pi y}{b}\right) = \int_0^a \sin\left(\frac{k\pi x}{a}\right) f(x, y) \, \mathrm{d}x.$$

We can perform the same trick again, and obtain

$$\frac{ab}{4} A_{k,j} = \int_{[0,a] \times [0,b]} \sin\left(\frac{k\pi x}{a}\right) \sin\left(\frac{j\pi y}{b}\right) f(x, y) \, \mathrm{d}x \, \mathrm{d}y. \tag{$*$}$$

So we found the solution

$$\psi(x, y, z) = \sum_{n,m=1}^{\infty} A_{n,m} \sin\left(\frac{n\pi x}{a}\right) \sin\left(\frac{m\pi y}{b}\right) \exp(-s_{n,m} z)$$

where $A_{m,n}$ is given by ($*$). Since we have shown that there is a unique solution to Laplace's equation obeying Dirichlet boundary conditions, we're done.

Note that if we imposed a boundary condition at finite z, say $0 \leq z \leq c$, then both sets of exponential $\exp(\pm\sqrt{\lambda + \mu} z)$ would have contributed. Similarly, if ψ does not vanish at the other boundaries, then the cos terms would also contribute.

In general, to actually find our ψ, we have to do the horrible integral for $A_{m,n}$, and this is not always easy (since integrals are *hard*). We will just look at a simple example.

Example. Suppose $f(x, y) = 1$. Then

$$A_{m,n} = \frac{4}{ab} \int_0^a \sin\left(\frac{n\pi x}{a}\right) \, \mathrm{d}x \int_0^b \sin\left(\frac{m\pi y}{b}\right) \, \mathrm{d}y = \begin{cases} \frac{16}{\pi^2 mn} & n, m \text{ both odd} \\ 0 & \text{otherwise} \end{cases}$$

Hence we have

$$\psi(x, y, z) = \frac{16}{\pi^2} \sum_{n,m \text{ odd}} \frac{1}{nm} \sin\left(\frac{n\pi x}{a}\right) \sin\left(\frac{m\pi y}{b}\right) \exp(-s_{m,n} z).$$

4.4   Laplace's equation in spherical polar coordinates

4.4.1   Laplace's equation in spherical polar coordinates

We'll often also be interested in taking

$$\Omega = \{(s, y, z) \in \mathbb{R}^3 : \sqrt{x^2 + y^2 + z^2} \leq a\}.$$

Since our domain is spherical, it makes sense to use some coordinate system with spherical symmetry. We use spherical coordinates (r, θ, ϕ), where

$$x = r \sin\theta \cos\phi, \quad y = r \sin\theta \sin\phi, \quad z = r \cos\theta.$$

The Laplacian is

$$\nabla^2 = \frac{1}{r^2} \frac{\partial}{\partial r} \left(r^2 \frac{\partial}{\partial r} \right) + \frac{1}{r^2 \sin\theta} \frac{\partial}{\partial \theta} \left(\sin\theta \frac{\partial}{\partial \theta} \right) + \frac{1}{r^2 \sin^2\theta} \frac{\partial^2}{\partial \phi^2}.$$

Similarly, the volume element is

$$dV = dx \, dy \, dz = r^2 \sin\theta \, dr \, d\theta \, d\phi.$$

In this course, we'll only look at *axisymmetric solutions*, where $\psi(r, \theta, \phi) = \psi(r, \theta)$ does not depend on ϕ.

We perform separation of variables, where we look for any set of solution of $\nabla^2 \psi = 0$ inside Ω where $\psi(r, \theta) = R(r)\Theta(\theta)$.

Laplace's equation then becomes

$$\frac{\psi}{r^2} \left[\frac{1}{R} \frac{d}{dr}(r^2 R') + \frac{1}{\Theta \sin\theta} \frac{d}{d\theta} \left(\sin\theta \frac{d\Theta}{d\theta} \right) \right] = 0.$$

Similarly, since each term depends only on one variable, each must separately be constant. So we have

$$\frac{d}{dr} \left(r^2 \frac{dR}{dr} \right) = \lambda R, \quad \frac{d}{d\theta} \left(\sin\theta \frac{d\Theta}{d\theta} \right) = -\lambda \sin\theta \Theta.$$

Note that both these equations are eigenfunction equations for some Sturm-Liouville operator and the $\Theta(\theta)$ equation has a non-trivial weight function $w(\theta) = \sin\theta$.

4.4.2   Legendre Polynomials

The angular equation is known as the *Legendre's equation*. It's standard to set $x = \cos\theta$ (which is *not* the Cartesian coordinate, but just a new variable). Then we have

$$\frac{d}{d\theta} = -\sin\theta \frac{d}{dx}.$$

So the Legendre's equation becomes

$$-\sin\theta \frac{d}{dx} \left[\sin\theta(-\sin\theta) \frac{d\Theta}{dx} \right] + \lambda \sin\theta \Theta = 0.$$

Equivalently, in terms of x, we have

$$\frac{d}{dx} \left[(1 - x^2) \frac{d\Theta}{dx} \right] = -\lambda \Theta.$$

Note that since $0 \leq \theta \leq \pi$, the domain of x is $-1 \leq x \leq 1$.

This operator is a Sturm–Liouville operator with

$$p(x) = 1 - x^2, \quad q(x) = 0.$$

For the Sturm–Liouville operator to be self-adjoint, we had

$$(g, \mathcal{L}f) = (\mathcal{L}g, f) + [p(x)(g^* f' - g'^* f)]^1_{-1}.$$

We want the boundary term to vanish. Since $p(x) = 1 - x^2$ vanishes at our boundary $x = \pm 1$, the Sturm-Liouville operator is self-adjoint provided our function $\Theta(x)$ is regular at $x = \pm 1$. Hence we look for a set of Legendre's equations inside $(-1, 1)$ that remains regular including at $x = \pm 1$. We can try a power series

$$\Theta(x) = \sum_{n=0}^{\infty} a_n x^n.$$

The Legendre's equation becomes

$$(1 - x^2) \sum_{n=0}^{\infty} a_n n(n-1) x^{n-2} - 2 \sum_{n=0}^{\infty} a_n n x^n + \lambda \sum_{n=0}^{\infty} a_n x^n = 0.$$

For this to hold for all $x \in (-1, 1)$, the equation must hold for each coefficient of x separately. So

$$a_n(\lambda - n(n+1)) + a_{n-2}(n+2)(n+1) = 0.$$

This requires that

$$a_{n+2} = \frac{n(n+1) - \lambda}{(n+2)(n+1)} a_n.$$

This equation relates a_{n+2} to a_n. So we can choose a_0 and a_1 freely. So we get two linearly independents solutions $\Theta_0(x)$ and $\Theta_1(x)$, where they satisfy

$$\Theta_0(-x) = \Theta_0(x), \quad \Theta_1(-x) = -\Theta_1(x).$$

In particular, we can expand our recurrence formula to obtain

$$\Theta_0(x) = a_0 \left[1 - \frac{\lambda}{2} x^2 - \frac{(6 - \lambda)\lambda}{4!} x^4 + \cdots \right]$$
$$\Theta_1(x) = a_1 \left[x + \frac{(2 - \lambda)}{3!} x^2 + \frac{(12 - \lambda)(2 - \lambda)}{5!} x^5 + \cdots \right].$$

We now consider the boundary conditions. We know that $\Theta(x)$ must be regular (i.e. finite) at $x = \pm 1$. This seems innocent, However, we have

$$\lim_{n \to \infty} \frac{a_{n+2}}{a_n} = 1.$$

This is fine if we are inside $(-1, 1)$, since the power series will still converge. However, *at* ± 1, the ratio test is not conclusive and does not guarantee convergence. In fact, more sophisticated convergence tests show that the infinite series would diverge on the boundary! To avoid this problem, we need to choose λ such

that the series truncates. That is, if we set $\lambda = \ell(\ell + 1)$ for some $\ell \in \mathbb{N}$, then our power series will truncate. Then since we are left with a finite polynomial, it of course converges.

In this case, the finiteness boundary condition fixes our possible values of eigenvalues. This is how quantization occurs in quantum mechanics. In this case, this process is why angular momentum is quantized.

The resulting polynomial solutions $P_\ell(x)$ are called *Legendre polynomials*. For example, we have

$$P_0(x) = 1$$
$$P_1(x) = x$$
$$P_2(x) = \frac{1}{2}(3x^2 - 1)$$
$$P_3(x) = \frac{1}{2}(5x^3 - 3x),$$

where the overall normalization is chosen to fix $P_\ell(1) = 1$.

It turns out that

$$P_\ell(x) = \frac{1}{2^\ell \ell!} \frac{\mathrm{d}^\ell}{\mathrm{d}x^\ell}(x^2 - 1)^\ell.$$

The constants in front are just to ensure normalization. We now check that this indeed gives the desired normalization:

$$P_\ell(1) = \frac{1}{2^\ell} \frac{\mathrm{d}^\ell}{\mathrm{d}x^\ell}[(x - 1)^\ell (x + 1)^\ell]\Big|_{x=1}$$
$$= \frac{1}{2^\ell \ell!} \left[\ell!(x + 1)^\ell + (x - 1)(\text{stuff})\right]_{x=1}$$
$$= 1$$

We have previously shown that the eigenfunctions of Sturm-Liouville operators are orthogonal. Let's see that directly now for this operator. The weight function is $w(x) = 1$.

We first prove a short lemma: for $0 \le k \le \ell$, we have

$$\frac{\mathrm{d}^k}{\mathrm{d}x^k}(x^2 - 1)^\ell = Q_{\ell,k}(x^2 - 1)^{\ell-k}$$

for some degree k polynomial $Q_{\ell,k}(x)$.

We show this by induction. This is trivially true when $k = 0$. We have

$$\frac{\mathrm{d}^{k+1}}{\mathrm{d}x^{k+1}}(x^2 + 1)^\ell = \frac{\mathrm{d}}{\mathrm{d}x}\left[Q_{\ell,k}(x)(x^2 - 1)^{\ell-k}\right]$$
$$= Q'_{\ell,k}(x^2 - 1)^{\ell-k} + 2(\ell - k)Q_{\ell,k}x(x^2 - 1)^{\ell-k-1}$$
$$= (x^2 - 1)^{\ell-k-1}\left[Q'_{\ell,k}(x^2 - 1) - 2(\ell - k)Q_{\ell,k}x\right].$$

Since $Q_{\ell,k}$ has degree k, $Q'_{\ell,k}$ has degree $k - 1$. So the right bunch of stuff has degree $k + 1$.

Now we can show orthogonality. We have

$$
\begin{aligned}
(P_\ell, P_{\ell'}) &= \int_{-1}^{1} P_\ell(x) P_{\ell'}(x) \; \mathrm{d}x \\
&= \frac{1}{2^\ell \ell!} \int_{-1}^{1} \frac{\mathrm{d}^\ell}{\mathrm{d}x^\ell}(x^2 - 1)^\ell P_{\ell'}(x) \; \mathrm{d}x \\
&= \frac{1}{2^\ell \ell!} \left[\frac{\mathrm{d}^{\ell-1}}{\mathrm{d}x^{\ell-1}}(x^2 - 1)^2 P_{\ell'}(x) \right] - \frac{1}{2^\ell \ell!} \int_{-1}^{1} \frac{\mathrm{d}^{\ell-1}}{\mathrm{d}x^{\ell-1}}(x^2 - 1)^\ell \frac{\mathrm{d}P_{\ell'}}{\mathrm{d}x} \; \mathrm{d}x \\
&= -\frac{1}{2^\ell \ell!} \int_{-1}^{1} \frac{\mathrm{d}^{\ell-1}}{\mathrm{d}x^{\ell-1}}(x^2 - 1)^\ell \frac{\mathrm{d}P_{\ell'}}{\mathrm{d}x} \; \mathrm{d}x.
\end{aligned}
$$

Note that the first term disappears since the $(\ell - 1)^{\text{th}}$ derivative of $(x^2 - 1)^\ell$ still has a factor of $x^2 - 1$. So integration by parts allows us to transfer the derivative from $x^2 - 1$ to $P_{\ell'}$.

Now if $\ell \neq \ell'$, we can wlog assume $\ell' < \ell$. Then we can integrate by parts ℓ' times until we get the ℓ'^{th} derivative of $P_{\ell'}$, which is zero.

In fact, we can show that

$$
(P_\ell, P_{\ell'}) = \frac{2\delta_{\ell\ell'}}{2\ell + 1}.
$$

hence the $P_\ell(x)$ are orthogonal polynomials.

By the fundamental theorem of algebra, we know that $P_\ell(x)$ has ℓ roots. In fact, they are always real, and lie in $(-1, 1)$.

Suppose not. Suppose only $m < \ell$ roots lie in $(-1, 1)$. Then let $Q_m(x) = \prod_{r=1}^{m}(x - x_r)$, where $\{x_1, x_2, \cdots, x_m\}$ are these m roots. Consider the polynomial $P_\ell(x)Q_m(x)$. If we factorize this, we get $\prod_{r=m+1}^{\ell}(x - r)\prod_{r=1}^{m}(x - x_r)^2$. The first few terms have roots outside $(-1, 1)$, and hence do not change sign in $(-1, 1)$. The latter terms are always non-negative. Hence for some appropriate sign, we have

$$
\pm \int_{-1}^{1} P_\ell(x) Q_m(x) \; \mathrm{d}x > 0.
$$

However, we can expand

$$
Q_m(x) = \sum_{r=1}^{m} q_r P_r(x),
$$

but $(P_\ell, P_r) = 0$ for all $r < \ell$. This is a contradiction.

4.4.3 Solution to radial part

In our original equation $\nabla^2 \phi = 0$, we now set $\Theta(\theta) = P_\ell(\cos\theta)$. The radial equation becomes

$$
(r^2 R')' = \ell(\ell + 1)R.
$$

Trying a solution of the form $R(r) = r^\alpha$, we find that $\alpha(\alpha + 1) = \ell(\ell + 1)$. So the solution is $\alpha = \ell$ or $-(\ell + 1)$. Hence

$$
\phi(r, \theta) = \sum \left(a_\ell r^\ell + \frac{b_\ell}{r^{\ell+1}} \right) P_\ell(\cos\theta)
$$

is the general solution to the Laplace equation $\nabla^2\phi = 0$ on the domain. In we want regularity at $r = 0$, then we need $b_\ell = 0$ for all ℓ.

The a_ℓ are fixed if we require $\phi(a, \theta) = f(\theta)$ on $\partial\Omega$, i.e. when $r = 1$.

We expand

$$f(\theta) = \sum_{\ell=0}^{\infty} F_\ell P_\ell(\cos\theta),$$

where

$$F_\ell = (P_\ell, f) = \int_0^\pi P_\ell(\cos\theta) f(\theta) \, \mathrm{d}(\cos\theta).$$

So

$$\phi(r, \theta) = \sum_{\ell=0}^{\infty} F_\ell \left(\frac{r}{a}\right)^\alpha P_\ell(\cos\theta).$$

4.5 Multipole expansions for Laplace's equation

One can quickly check that

$$\phi(\mathbf{r}) = \frac{1}{|\mathbf{r} - \mathbf{r}'|}$$

solves Laplace's equation $\nabla^2\phi = 0$ for all $\mathbf{r} \in \mathbb{R}^3 \setminus \mathbf{r}'$. For example, if $\mathbf{r}' = \hat{\mathbf{k}}$, where $\hat{\mathbf{k}}$ is the unit vector in the z direction, then

$$\frac{1}{|\mathbf{r} - \hat{\mathbf{k}}|} = \frac{1}{\sqrt{r^2 + 1 - 2r\cos\theta}} = \sum_{\ell=0}^{\infty} c_\ell r^\ell P_\ell(\cos\theta).$$

To find these coefficients, we can employ a little trick. Since $P_\ell(1) = 0$, at $\theta = 0$, we have

$$\sum_{\ell=0}^{\infty} c_\ell r^\ell = \frac{1}{1-r} = \sum_{\ell=0}^{\infty} r^\ell.$$

So all the coefficients must be 1. This is valid for $r < 1$.

More generally, we have

$$\frac{1}{|\mathbf{r} - \mathbf{r}'|} = \frac{1}{r'} \sum_{\ell=0}^{\infty} \left(\frac{r}{r'}\right)^\ell P_\ell(\hat{\mathbf{r}} \cdot \hat{\mathbf{r}}').$$

This is called the *multiple expansion*, and is valid when $r < r'$. Thus

$$\frac{1}{|\mathbf{r} - \mathbf{r}'|} = \frac{1}{r'} + \frac{r}{r'^2}\hat{\mathbf{r}} \cdot \hat{\mathbf{r}}' + \cdots.$$

The first term $\frac{1}{r'}$ is known as the *monopole*, and the second term is known as the *dipole*, in analogy to charges in electromagnetism. The monopole term is what we get due to a single charge, and the second term is the result of the interaction of two charges.

4.6   Laplace's equation in cylindrical coordinates

We let
$$\Omega = \{(r, \theta, z) \in \mathbb{R}^3 : r \leq a, z \geq 0\}.$$

In cylindrical polars, we have

$$\nabla^2 \phi = \frac{1}{r} \frac{\partial}{\partial r} \left(r \frac{\partial \phi}{\partial r} \right) + \frac{1}{r^2} \frac{\partial^2 \phi}{\partial \theta^2} + \frac{\partial^2 \phi}{\partial z^2} = 0.$$

We look for a solution that is regular inside Ω and obeys the boundary conditions

$$\phi(a, \theta, z) = 0$$
$$\phi(r, \theta, 0) = f(r, \theta)$$
$$\lim_{z \to \infty} \phi(r, \theta, z) = 0,$$

where f is some fixed function.

Again, we start by separation of variables. We try

$$\phi(r, \theta, z) = R(r)\Theta(\theta)Z(z).$$

Then we get

$$\frac{1}{rR}(rR')' + \frac{1}{r^2} \frac{\Theta''}{\Theta} + \frac{Z''}{Z} = 0.$$

So we immediately know that

$$Z'' = \mu Z$$

We now replace Z''/Z by μ, and multiply by r^2 to obtain

$$\frac{r}{R}(rR')' + \frac{\Theta''}{\Theta} + \mu r^2 = 0.$$

So we know that

$$\Theta'' = -\lambda \Theta.$$

Then we are left with

$$r^2 R'' + r R' + (\mu r^2 - \lambda)R = 0.$$

Since we require periodicity in θ, we have

$$\Theta(\theta) = a_n \sin n\theta + b_n \cos n\theta, \quad \lambda = n^2, \quad n \in \mathbb{N}.$$

Since we want the solution to decay as $z \to \infty$, we have

$$Z(z) = c_\mu e^{-\sqrt{\mu} z}$$

The radial equation is of Sturm-Liouville type since it can be written as

$$\frac{\mathrm{d}}{\mathrm{d}r} \left(r \frac{\mathrm{d}R}{\mathrm{d}r} \right) - \frac{n^2}{r} R = -\mu r R.$$

Here we have

$$p(r) = r, q(r) = -\frac{n^2}{r}, w(r) = r.$$

Introducing $x = r\sqrt{\mu}$, we can rewrite this as

$$x^2 \frac{d^2 R}{dx^2} + x \frac{dR}{dx} + (x^2 - n^2)R = 0.$$

This is Bessel's equation. Note that this is actually a whole family of differential equations, one for each n. The n is not the eigenvalue we are trying to figure out in this equation. It is already fixed by the Θ equation. The actual eigenvalue we are finding out is μ, not n.

Since Bessel's equations are second-order, there are two independent solution for each n, namely $J_n(x)$ and $Y_n(x)$. These are called Bessel functions of order n of the first (J_n) or second (Y_n) kind. We will not study these functions' properties, but just state some useful properties of them.

The J_n's are all regular at the origin. In particular, as $x \to 0$ $J_n(x) \sim x^n$. These look like decaying sine waves, but the zeroes are not regularly spaced.

On the other hand, the Y_n's are similar but singular at the origin. As $x \to 0$, $Y_0(x) \sim \ln x$, while $Y_n(x) \sim x^{-n}$.

Now we can write our separable solution as

$$\phi(r, \theta, z) = (a_n \sin n\theta + b_n \cos n\theta)e^{-\sqrt{\mu}z}[c_{\mu,n}J_n(r\sqrt{\mu}) + d_{\mu,n}Y_n(r, \sqrt{\mu})].$$

Since we want regularity at $r = 0$, we don't want to have the Y_n terms. So $d_{\mu,n} = 0$.

We now impose our first boundary condition $\phi(a, \theta, z) = 0$. This demands $J_n(a\sqrt{\mu}) = 0$. So we must have

$$\sqrt{\mu} = \frac{k_{ni}}{a},$$

where k_{ni} is the ith root of $J_n(x)$ (since there is not much pattern to these roots, this is the best description we can give!).

So our general solution obeying the homogeneous conditions is

$$\phi(r, \theta, z) = \sum_{n=0}^{\infty} \sum_{i \in \text{roots}} (A_{ni} \sin n\theta + B_{ni} \cos n\theta) \exp\left(-\frac{k_{ni}}{a}z\right) J_n\left(\frac{k_{ni}r}{a}\right).$$

We finally impose the inhomogeneous boundary condition $\phi(r, \theta, 0) = f(r, \theta)$. To do this, we use the orthogonality condition

$$\int_0^a J_n\left(\frac{k_{nj}r}{a}\right) J_n\left(\frac{k_{ni}r}{a}\right) r \, dr = \frac{a^2}{2}\delta_{ij}[J_n'(k_{ni})]^2,$$

which is proved in the example sheets. Note that we have a weight function r here. Also, remember this is the orthogonality relation for different roots of Bessel's functions of the *same order*. It does not relate Bessel's functions of different orders.

We now integrate our general solution with respect to $\cos m\theta$ to obtain

$$\frac{1}{\pi} \int_{-\pi}^{\pi} f(r, \theta) \cos m\theta \, d\theta = \sum_{i \in \text{roots}} J_m\left(\frac{k_{mi}r}{a}\right).$$

Now we just have Bessel's function for a single order j. So we can use the orthogonality relation for the Bessel's functions to obtain

$$B_{mj} = \frac{1}{[J'_m(k_{mj})]} \frac{2}{\pi a^2} \int_0^a \int_{-\pi}^{\pi} \cos m\theta J_m\left(\frac{k_{mj}r}{a}\right) f(r,\theta) r \, dr \, d\theta.$$

How can we possibly perform this integral? We don't know how J_m looks like, what the roots k_{mj} are, and we multiply all these complicated things together and integrate! Often, we just don't. We ask our computers to do this numerically for us. There are indeed some rare cases where we are indeed able to get an explicit solution, but these are rare.

4.7   The heat equation

We choose $\Omega = \mathbb{R}^n \times [0, \infty)$. We treat $\mathbb{R}^n$ as the "space" variable, and $[0, \infty)$ as our time variable.

Definition (Heat equation)**.** The *heat equation* for a function $\phi : \Omega \to \mathbb{C}$ is

$$\frac{\partial \phi}{\partial t} = \kappa \nabla^2 \phi,$$

where $\kappa > 0$ is the *diffusion constant*.

We can think of this as a function ϕ representing the "temperature" at different points in space and time, and this equation says the rate of change of temperature is proportional to $\nabla^2 \phi$.

Our previous study of the Laplace's equation comes in handy. When our system is in equilibrium and does not change with time, then we are left with Laplace's equation.

Let's first look at some interesting properties of the heat equation.

The first most important property of the heat equation is that the total "heat" is conserved under evolution by the heat equation, i.e.

$$\frac{d}{dt} \int_{\mathbb{R}^n} \phi(\mathbf{x}, t) \, d^n x = 0,$$

since we can write this as

$$\frac{d}{dt} \int_{\mathbb{R}^n} \phi(\mathbf{x}, t) \, d^n x = \int_{\mathbb{R}^n} \frac{\partial \phi}{\partial t} \, d^n x = \kappa \int_{\mathbb{R}^n} \nabla^2 \phi \, d^n x = 0,$$

provided $\nabla \phi(\mathbf{x}, t) \to 0$ as $|\mathbf{x}| \to \infty$.

In particular, on a compact space, i.e. $\Omega = M \times [0, \infty)$ with M compact, there is no boundary term and we have

$$\frac{d}{dt} \int_M \phi \, d\mu = 0.$$

The second useful property is if $\phi(\mathbf{x}, t)$ solves the heat equation, so do $\phi_1(\mathbf{x}, t) = \phi(\mathbf{x} - \mathbf{x}_0, t - t_0)$, and $\phi_2(\mathbf{x}, t) = A\phi(\lambda \mathbf{x}, \lambda^2 t)$.

Let's try to choose A such that

$$\int_{\mathbb{R}^n} \phi_2(\mathbf{x}, t) \, d^n x = \int_{\mathbb{R}^n} \phi(\mathbf{x}, t) \, d^n x.$$

We have

$$\int_{\mathbb{R}^n} \phi_2(\mathbf{x}, t)\, \mathrm{d}^n x = A \int_{\mathbb{R}^n} \phi(\lambda \mathbf{x}, \lambda^2 t)\, \mathrm{d}^n x = A\lambda^{-n} \int_{\mathbb{R}^n} \phi(\mathbf{y}, \lambda^2 t)\, \mathrm{d}^n y,$$

where we substituted $\mathbf{y} = \lambda \mathbf{x}$.

So if we set $A = \lambda^n$, then the total heat in ϕ_2 at time $\lambda^2 t$ is the same as the total heat in ϕ at time t. But the total heat is constant in time. So the total heat is the same all the time.

Note again that if $\frac{\partial \phi}{\partial t} = \kappa \nabla^2 \phi$ on $\Omega \times [0, \infty)$, then so does $\lambda^n \phi(\lambda x, \lambda^2 t)$. This means that nothing is lost by letting $\lambda = \frac{1}{\sqrt{t\kappa}}$ and consider solutions of the form

$$\phi(\mathbf{x}, t) = \frac{1}{(\kappa t)^{n/2}} F\left(\frac{\mathbf{x}}{\sqrt{\kappa t}}\right) = \frac{1}{(\kappa t)^{n/2}} F(\boldsymbol{\eta}),$$

where

$$\boldsymbol{\eta} = \frac{\mathbf{x}}{\sqrt{\kappa t}}.$$

For example, in $1 + 1$ dimensions, we can look for a solution of the form $\frac{1}{\sqrt{\kappa t}} F(\eta)$. We have

$$\frac{\partial \phi}{\partial t} = \frac{\partial}{\partial t}\left(\frac{1}{\sqrt{\kappa t}} F(\eta)\right) = \frac{-1}{2\sqrt{\kappa t^3}} F(\eta) + \frac{1}{\sqrt{\kappa t}} \frac{\mathrm{d}\eta}{\mathrm{d}t} F'(\eta) = \frac{-1}{2\sqrt{\kappa t^3}} [F + \eta F']$$

On the other side of the equation, we have

$$\kappa \frac{\partial^2}{\partial x^2}\left(\frac{1}{\sqrt{\kappa t}} F(\eta)\right) = \frac{\kappa}{\sqrt{\kappa t}} \frac{\partial}{\partial x}\left(\frac{\partial \eta}{\partial x} F'\right) = \frac{1}{\sqrt{\kappa t^3}} F''.$$

So the equation becomes

$$0 = 2F'' + \eta F' + F = (2F' + \eta F)'.$$

So we have

$$2F' + \eta F = \text{const},$$

and the boundary conditions require that the constant is zero. So we get

$$F' = \frac{-\eta}{2} F.$$

We now see a solution

$$F(\eta) = a \exp\left(\frac{-\eta^2}{4}\right).$$

By convention, we now normalize our solution $\phi(x, t)$ by requiring

$$\int_{-\infty}^{\infty} \phi(x, t)\, \mathrm{d}x = 1.$$

This gives

$$a = \frac{1}{\sqrt{4\pi}},$$

and the full solution is

$$\phi(x,t) = \frac{1}{\sqrt{4\pi\kappa t}} \exp\left(-\frac{x^2}{4\kappa t}\right).$$

More generally, on $\mathbb{R}^n \times [0, \infty)$, we have the solution

$$\phi(\mathbf{x}, t) = \frac{1}{(4\pi(t - t_0))^{n/2}} \exp\left(\frac{-(\mathbf{x} - \mathbf{x}_0)^2}{4\kappa(t - t_0)}\right).$$

At any fixed time $t \neq t_0$, the solutions are Gaussians centered on x_0 with standard deviation $\sqrt{2\kappa(t - t_0)}$, and the height of the peak at $x = x_0$ is $\sqrt{\frac{1}{4\pi\kappa(t-t_0)}}$.

We see that the solution gets "flatter" as t increases. This is a general property of the evolution by the heat equation. Suppose we have an eigenfunction $y(\mathbf{x})$ of the Laplacian on Ω, i.e. $\nabla^2 y = -\lambda y$. We want to show that in certain cases, λ must be positive. Consider

$$-\lambda \int_\Omega |y|^2 \mathrm{d}^n x = \int_\Omega y^*(x)\nabla^2 y(x) \, \mathrm{d}^n x = \int_{\partial\Omega} y^*\mathbf{n} \cdot \nabla y \, \mathrm{d}^{n-1}x - \int_\Omega |\nabla y|^2 \, \mathrm{d}^n x.$$

If there is no boundary term, e.g. when Ω is closed and compact (e.g. a sphere), then we get

$$\lambda = \frac{\int |\nabla y|^2 \, \mathrm{d}^n x}{\int_\Omega |y|^2 \, \mathrm{d}^n x}.$$

Hence the eigenvalues are all positive. What does this mean for heat flow? Suppose $\phi(\mathbf{x}, t)$ obeys the heat equation for $t > 0$. At time $t = 0$, we have $\phi(\mathbf{x}, 0) = f(\mathbf{x})$. We can write the solution (somehow) as

$$\phi(\mathbf{x}, t) = e^{t\nabla^2}\phi(\mathbf{x}, 0),$$

where we (for convenience) let $\kappa = 1$. But we can expand our initial condition

$$\phi(\mathbf{x}, 0) = f(\mathbf{x}) = \sum_I c_I y_I(\mathbf{x})$$

in a complete set $\{y_I(\mathbf{x})\}$ of eigenfunctions for ∇^2.

By linearity, we have

$$\begin{aligned}
\phi(\mathbf{x}, t) &= e^{t\nabla^2}\left(\sum_I c_I y_I(\mathbf{x})\right) \\
&= \sum_I c_I e^{t\nabla^2} y_I(\mathbf{x}) \\
&= \sum_I c_I e^{-\lambda_2 t} y_I(\mathbf{x}) \\
&= c_I(t) y_I(\mathbf{x}),
\end{aligned}$$

where $c_I(t) = c_I e^{-\lambda_2 t}$. Since λ_I are all positive, we know that $c_I(t)$ decays exponentially with time. In particular, the coefficients corresponding to the largest eigenvalues $|\lambda_I|$ decays most rapidly. We also know that eigenfunctions

with larger eigenvalues are in general more "spiky". So these smooth out rather quickly.

When people first came up with the heat equation to describe heat loss, they were slightly skeptical — this flattening out caused by the heat equation is not time-reversible, while Newton's laws of motions are all time reversible. How could this smoothening out arise from reversible equations?

Einstein came up with an example where the heat equation can come out of apparently naturally from seemingly reversible laws, namely Brownian motion. The idea is that a particle will randomly jump around in space, where the movement is independent of time and position.

Let the probability that a dust particle moves through a step y over a time Δt be $p(y, \Delta t)$. For any fixed Δt, we must have

$$\int_{-\infty}^{\infty} p(y, \Delta t) \, \mathrm{d}y = 1.$$

We also assume that $p(y, \Delta t)$ is independent of time, and of the location of the dust particle. We also assume $p(y, \Delta t)$ is strongly peaked around $y = 0$ and $p(y, \Delta t) = p(-y, \Delta t)$. Now let $P(x, t)$ be the probability that the dust particle is located at x at time t. Then at time $t + \delta t$, we have

$$P(x, t + \Delta t) = \int_{-\infty}^{\infty} P(x - y, t) p(y, \Delta t) \, \mathrm{d}y.$$

For $P(x - y, t)$ sufficiently smooth, we can write

$$P(x - y, t) \approx P(x, t) - y \frac{\partial P}{\partial x}(x, t) + \frac{y^2}{2} \frac{\partial^2 P}{\partial x^2}(x, t).$$

So we get

$$P(x, t + \Delta t) \approx P(x, t) - \frac{\partial P}{\partial x}(x, t) \langle y \rangle + \frac{1}{2} \langle y^2 \rangle \frac{\partial^2 P}{\partial x^2} P(x, t) + \cdots,$$

where

$$\langle y^r \rangle = \int_{-\infty}^{\infty} y^r p(y, \Delta t) \, \mathrm{d}y.$$

Since p is even, we expect $\langle y^r \rangle$ to vanish when r is odd. Also, since y is strongly peaked at 0, we expect the higher order terms to be small. So we can write

$$P(x, t + \Delta t) - P(x, t) = \frac{1}{2} \langle y^2 \rangle \frac{\partial^2 P}{\partial x^2}.$$

Suppose that as we take the limit $\Delta t \to 0$, we get $\frac{\langle y^2 \rangle}{2\Delta t} \to \kappa$ for some κ. Then this becomes the heat equation

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

Proposition. Suppose $\phi : \Omega \times [0, \infty) \to \mathbb{R}$ satisfies the heat equation $\frac{\partial \phi}{\partial t} = \kappa \nabla^2 \phi$, and obeys

- Initial conditions $\phi(\mathbf{x}, 0) = f(x)$ for all $x \in \Omega$

– Boundary condition $\phi(\mathbf{x}, t)|_{\partial\Omega} = g(\mathbf{x}, t)$ for all $t \in [0, \infty)$.

Then $\phi(\mathbf{x}, t)$ is unique.

Proof. Suppose ϕ_1 and ϕ_2 are both solutions. Then define $\Phi = \phi_1 - \phi_2$ and

$$E(t) = \frac{1}{2} \int_\Omega \Phi^2 \, \mathrm{d}V.$$

Then we know that $E(t) \geq 0$. Since ϕ_1, ϕ_2 both obey the heat equation, so does Φ. Also, on the boundary and at $t = 0$, we know that $\Phi = 0$. We have

$$\begin{aligned}
\frac{\mathrm{d}E}{\mathrm{d}t} &= \int_\Omega \Phi \frac{\mathrm{d}\Phi}{\mathrm{d}t} \, \mathrm{d}V \\
&= \kappa \int_\Omega \Phi \nabla^2 \Phi \, \mathrm{d}V \\
&= \kappa \int_{\partial\Omega} \Phi \nabla\Phi \cdot \mathrm{d}\mathbf{S} - \kappa \int_\Omega (\nabla\Phi)^2 \, \mathrm{d}V \\
&= -\kappa \int_\Omega (\nabla\Phi)^2 \, \mathrm{d}V \\
&\leq 0.
\end{aligned}$$

So we know that E decreases with time but is always non-negative. We also know that at time $t = 0$, $E = \Phi = 0$. So $E = 0$ always. So $\Phi = 0$ always. So $\phi_1 = \phi_2$. $\qquad\square$

Example (Heat conduction in uniform medium)**.** Suppose we are on Earth, and the Sun heats up the surface of the Earth through sun light. So the sun will maintain the soil at some fixed temperature. However, this temperature varies with time as we move through the day-night cycle and seasons of the year.

We let $\phi(x, t)$ be the temperature of the soil as a function of the depth x, defined on $\mathbb{R}^{\geq 0} \times [0, \infty)$. Then it obeys the heat equation with conditions

(i) $\phi(0, t) = \phi_0 + A \cos\left(\frac{2\pi t}{t_D}\right) + B \cos\left(\frac{2\pi t}{t_Y}\right)$.

(ii) $\phi(x, t) \to$ const as $x \to \infty$.

We know that

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

We try the separation of variables

$$\phi(x, t) = T(t)X(x).$$

Then we get the equations

$$T' = \lambda T, \quad X'' = \frac{\lambda}{K}X.$$

From the boundary solution, we know that our things will be oscillatory. So we let λ be imaginary, and set $\lambda = i\omega$. So we have

$$\phi(x, t) = e^{i\omega t}\left(a_\omega e^{-\sqrt{i\omega/K}\,x} + b_\omega e^{\sqrt{i\omega/K}\,x}\right).$$

Note that we have

$$\sqrt{\frac{i\omega}{K}} = \begin{cases} (1+i)\sqrt{\frac{|\omega|}{2K}} & \omega > 0 \\ (i-1)\sqrt{\frac{|\omega|}{2K}} & \omega < 0 \end{cases}$$

Since $\phi(x,t) \to$ constant as $x \to \infty$, we don't want our ϕ to blow up. So if $\omega < 0$, we need $a_\omega = 0$. Otherwise, we need $b_\omega = 0$.

To match up at $x = 0$, we just want terms with

$$|\omega| = \omega_D = \frac{2\pi}{t_D}, \quad |\omega| = \omega_Y = \frac{2\pi}{t_Y}.$$

So we can write out solution as

$$\phi(x,t) = \phi_0 + A \exp\left(-\sqrt{\frac{\omega_D}{2K}}\right) \cos\left(\omega_D t - \sqrt{\frac{\omega_D}{2K}}x\right)$$
$$+ B \exp\left(-\sqrt{\frac{\omega_Y}{2K}}\right) \cos\left(\omega_Y t - \sqrt{\frac{\omega_Y}{2K}}x\right)$$

We can notice a few things. Firstly, as we go further down, the effect of the sun decays, and the temperature is more stable. Also, the effect of the day-night cycle decays more quickly than the annual cycle, which makes sense. We also see that while the temperature does fluctuate with the same frequency as the day-night/annual cycle, as we go down, there is a phase shift. This is helpful since we can store things underground and make them cool in summer, warm in winter.

Example (Heat conduction in a finite rod)**.** We now have a finite rod of length $2L$.

$$x = -L \qquad x = 0 \qquad x = L$$

We have the initial conditions

$$\phi(x,0) = \Theta(x) = \begin{cases} 1 & 0 < x < L \\ 0 & -L < x < 0 \end{cases}$$

and the boundary conditions

$$\phi(-L,t) = 0, \quad \phi(L,t) = 1$$

So we start with a step temperature, and then maintain the two ends at fixed temperatures 0 and 1.

We are going to do separation of variables, but we note that all our boundary and initial conditions are inhomogeneous. This is not helpful. So we use a little trick. We first look for *any* solution satisfying the boundary conditions $\phi_S(-L,t) = 0$, $\phi_S(L,t) = 1$. For example, we can look for time-independent solutions $\phi_S(x,t) = \phi_S(x)$. Then we need $\frac{\mathrm{d}^2\phi_S}{\mathrm{d}x^2} = 0$. So we get

$$\phi_S(x) = \frac{x+L}{2L}.$$

By linearity, $\psi(x,t) = \phi(x,t) - \phi_s(x)$ obeys the heat equation with the conditions

$$\psi(-L,t) = \psi(L,t) = 0,$$

which is homogeneous! Our initial condition now becomes

$$\psi(x,0) = \Theta(x) - \frac{x+L}{2L}.$$

We now perform separation of variables for

$$\psi(x,t) = X(x)T(t).$$

Then we obtain the equations

$$T' = -\kappa\lambda T, \quad X' = -\lambda X.$$

Then we have

$$\psi(x,t) = \left[a\sin(\sqrt{\lambda}x) + b\cos(\sqrt{\lambda}x)\right]e^{-\kappa\lambda t}.$$

Since initial condition is odd, we can eliminate all cos terms. Our boundary conditions also requires

$$\lambda = \frac{n^2\pi^2}{L^2}, \quad n = 1, 2, \cdots.$$

So we have

$$\phi(x,t) = \phi_s(x) + \sum_{n=1}^{\infty} a_n \sin\left(\frac{n\pi x}{L}\right) \exp\left(-\frac{\kappa n^2 \pi^2}{L^2}t\right),$$

where a_n are the Fourier coefficients

$$a_n = \frac{1}{L}\int_{-L}^{L} \left[\Theta(x) - \frac{x+L}{2L}\right]\sin\left(\frac{n\pi x}{L}\right)\,\mathrm{d}x = \frac{1}{n\pi}.$$

Example (Cooling of a uniform sphere). Once upon a time, Charles Darwin went around the Earth, looked at species, and decided that evolution happened. When he came up with his theory, it worked quite well, except that there was one worry. Was there enough time on Earth for evolution to occur?

This calculation was done by Kelvin. He knew well that the Earth started as a ball of molten rock, and obviously life couldn't have evolved when the world was still molten rock. So he would want to know how long it took for the Earth to cool down to its current temperature, and if that was sufficient for evolution to occur.

We can, unsurprisingly, use the heat equation. We assume that at time $t = 0$, the temperature is ϕ_0, the melting temperature of rock. We also assume that the space is cold, and we let the temperature on the boundary of Earth as 0. We then solve the heat equation on a sphere (or ball), and see how much time it takes for Earth to cool down to its present temperature.

We let $\Omega = \{(x, y, z) \in \mathbb{R}^3, r \leq R\}$, and we want a solution $\phi(r, t)$ of the heat equation that is spherically symmetric and obeys the conditions

– $\phi(R, t) = 0$

– $\phi(r, 0) = \phi_0,$

We can write the heat equation as

$$\frac{\partial \phi}{\partial t} = \kappa \nabla^2 \phi = \frac{\kappa}{r^2} \frac{\partial}{\partial r} \left(r^2 \frac{\partial \phi}{\partial r} \right).$$

Again, we do the separation of variables.

$$\phi(r, t) = R(r)T(t).$$

So we get

$$T' = -\lambda^2 \kappa T, \quad \frac{\mathrm{d}}{\mathrm{d}r} \left(r^2 \frac{\mathrm{d}R}{\mathrm{d}r} \right) = -\lambda^2 r^2 R.$$

We can simplify this a bit by letting $R(r) = \frac{S(r)}{r}$, then our radial equation just becomes

$$S'' = -\lambda^2 S.$$

We can solve this to get

$$R(r) = A_\lambda \frac{\sin \lambda r}{r} + B_\lambda \frac{\cos \lambda r}{r}.$$

We want a regular solution at $r = 0$. So we need $B_\lambda = 0$. Also, the boundary condition $\phi(R, t) = 0$ gives

$$\lambda = \frac{n\pi}{R}, \quad n = 1, 2, \cdots$$

So we get

$$\phi(r, t) = \sum_{n \in \mathbb{Z}} \frac{A_n}{r} \sin \left(\frac{n\pi r}{R} \right) \exp \left(\frac{-\kappa n^2 \pi^2 t}{R^2} \right),$$

where our coefficients are

$$A_n = (-1)^{n+1} \frac{\phi_0 R}{n\pi}.$$

We know that the Earth isn't just a cold piece of rock. There are still volcanoes. So we know many terms still contribute to the sum nowadays. This is rather difficult to work with. So we instead look at the temperature gradient

$$\frac{\partial \phi}{\partial r} = \frac{\phi_0}{r} \sum_{n \in \mathbb{Z}} (-1)^{n+1} \cos \left(\frac{n\pi r}{R} \right) \exp \left(\frac{-\kappa n^2 \pi^2 t}{R^2} \right) + \sin \text{stuff}.$$

We evaluate this at the surface of the Earth, $R = r$. So we get the gradient

$$\frac{\partial \phi}{\partial r} \bigg|_R = -\frac{\phi_0}{R} \sum_{n \in \mathbb{Z}} \exp \left(-\frac{\kappa n^2 \pi^2 t}{R^2} \right) \approx \frac{\phi_0}{R} \int_{-\infty}^{\infty} \exp \left(-\frac{\kappa \pi y^2 t}{R^2} \right) \, \mathrm{d}y = \phi_0 \sqrt{\frac{1}{\pi \kappa t}}.$$

So the age of the Earth is approximately

$$t \approx \left(\frac{\phi_0}{V} \right)^2 \frac{1}{4\pi K},$$

where

$$V = \frac{\partial \phi}{\partial r} \bigg|_R.$$

We can plug the numbers in, and get that the Earth is 100 million years. This is not enough time for evolution.

Later on, people discovered that fission of metals inside the core of Earth produce heat, and provides an alternative source of heat. So problem solved! The current estimate of the age of the universe is around 4 billion years, and evolution did have enough time to occur.

4.8   The wave equation

Consider a string $x \in [0, L]$ undergoing small oscillations described by $\phi(x, t)$.
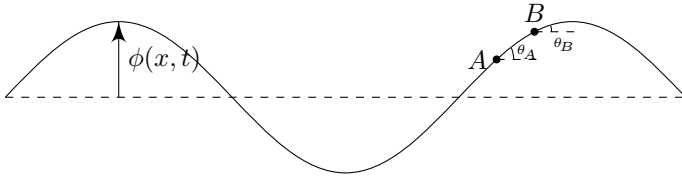
Consider two points A, B separated by a small distance δx. Let T_A (T_B) be the outward pointing tension tangent to the string at A (B). Since there is no sideways (x) motion, there is no net horizontal force. So

$$T_A \cos \theta_A = T_B \cos \theta_B = T. \tag{$*$}$$

If the string has mass per unit length μ, then in the vertical direction, Newton's second law gives

$$\mu \delta x \frac{\partial^2 \phi}{\partial t^2} = T_B \sin \theta_B - T_A \sin \theta_A.$$

We now divide everything by T, noting the relation $(*)$, and get

$$\begin{aligned}
\mu \frac{\delta x}{T} \frac{\partial^2 \phi}{\partial t^2} &= \frac{T_B \sin \theta_B}{T_B \cos \theta_B} - \frac{T_A \sin \theta_A}{T_A \cos \theta_A} \\
&= \tan \theta_B - \tan \theta_A \\
&= \left. \frac{\partial \phi}{\partial x} \right|_B - \left. \frac{\partial \phi}{\partial x} \right|_A \\
&\approx \delta x \frac{\partial^2 \phi}{\partial x^2}.
\end{aligned}$$

Taking the limit $\delta x \to 0$ and setting $c^2 = T/\mu$, we have that $\phi(x, t)$ obeys the wave equation

$$\frac{1}{c^2} \frac{\partial^2 \phi}{\partial t^2} = \frac{\partial^2 \phi}{\partial x^2}.$$

From IA Differential Equations, we've learnt that the solution can be written as $f(x - ct) + g(x + ct)$. However, this does not extend to higher dimensions, but separation of variables does. So let's do that.

Assume that the string is fixed at both ends. Then $\phi(0, t) = \phi(L, t) = 0$ for all t. The we can perform separation of variables, and the general solution can then be written

$$\phi(x, t) = \sum_{n=1}^{\infty} \sin \left(\frac{n\pi x}{L} \right) \left[A_n \cos \left(\frac{n\pi ct}{L} \right) + B_n \sin \left(\frac{n\pi ct}{L} \right) \right].$$

The coefficients A_n are fixed by the initial profile $\phi(x, 0)$ of the string, while the coefficients B_n are fixed by the initial string velocity $\frac{\partial \phi}{\partial t}(x, 0)$. Note that we need *two* sets of initial conditions, since the wave equation is second-order in time.

Energetics and uniqueness

An oscillating string contains has some sort of *energy*. The *kinetic energy* of a small element δx of the string is

$$\frac{1}{2}\mu \delta x \left(\frac{\partial \phi}{\partial t}\right)^2.$$

The *total kinetic energy* of the string is hence the integral

$$K(t) = \frac{\mu}{2} \int_0^L \left(\frac{\partial \phi}{\partial t}\right)^2 \, dx.$$

The string also has potential energy due to tension. The *potential energy* of a small element δx of the string is

$$\begin{aligned}
T \times \text{extension} &= T(\delta s - \delta x) \\
&= T(\sqrt{\delta x^2 + \delta \phi^2} - \delta x) \\
&\approx T\delta x \left(1 + \frac{1}{2}\left(\frac{\delta \phi}{\delta x}\right)^2 + \cdots\right) - T\delta x \\
&= \frac{T}{2}\delta x \left(\frac{\delta \phi}{\delta x}\right)^2.
\end{aligned}$$

Hence the total potential energy of the string is

$$V(t) = \frac{\mu}{2} \int_0^L c^2 \left(\frac{\partial \phi}{\partial x}\right)^2 \, dx,$$

using the definition of c^2.

Using our series expansion for ϕ, we get

$$K(t) = \frac{\mu \pi^2 c^2}{4L} \sum_{n=1}^{\infty} n^2 \left[A_n \sin\left(\frac{n\pi ct}{L}\right) - B_n \cos\left(\frac{n\pi ct}{L}\right)\right]^2$$

$$V(t) = \frac{\mu \pi^2 c^2}{4L} \sum_{n=1}^{\infty} n^2 \left[A_n \cos\left(\frac{n\pi ct}{L}\right) + B_n \sin\left(\frac{n\pi ct}{L}\right)\right]^2$$

The *total energy* is

$$E(t) = \frac{n\pi^2 c^2}{4L} \sum_{n=1}^{\infty} n^2 (A_n^2 + B_n^2).$$

What can we see here? Our solution is essentially an (infinite) sum of independent harmonic oscillators, one for each n. The period of the fundamental mode ($n = 1$) is $\frac{2\pi}{\omega} = 2\pi \cdot \frac{L}{\pi c} = \frac{2L}{c}$. Thus, averaging over a period, the average kinetic energy is

$$\bar{K} = \frac{c}{2L} \int_0^{2L/c} K(t) \, dt = \bar{V} = \frac{c}{2L} \int_0^{2L/c} V(t) \, dt = \frac{E}{2}.$$

Hence we have an equipartition of the energy between the kinetic and potential energy.

The energy also allows us to prove a uniqueness theorem. First we show that energy is conserved in general.

Proposition. Suppose $\phi : \Omega \times [0, \infty) \to \mathbb{R}$ obeys the wave equation $\frac{\partial^2 \phi}{\partial t^2} = c^2 \nabla^2 \phi$ inside $\Omega \times (0, \infty)$, and is fixed at the boundary. Then E is constant.

Proof. By definition, we have

$$\frac{\mathrm{d}E}{\mathrm{d}t} = \int_\Omega \frac{\partial^2 \psi}{\partial t^2} \frac{\partial \psi}{\partial t} + c^2 \nabla \left(\frac{\partial \phi}{\partial t} \right) \cdot \nabla \phi \, \mathrm{d}V.$$

We integrate by parts in the second term to obtain

$$\frac{\mathrm{d}E}{\mathrm{d}t} = \int_\Omega \frac{\mathrm{d}\phi}{\mathrm{d}t} \left(\frac{\partial^2 \phi}{\partial t^2} - c^2 \nabla^2 \phi \right) \, \mathrm{d}V + c^2 \int_{\partial \Omega} \frac{\partial \phi}{\partial t} \nabla \phi \cdot \mathrm{d}S.$$

Since $\frac{\partial^2 \phi}{\partial t^2} = c^2 \nabla^2 \phi$ by the wave equation, and ϕ is constant on $\partial \Omega$, we know that

$$\frac{\mathrm{d}E_\phi}{\mathrm{d}t} = 0. \qquad \square$$

Proposition. Suppose $\phi : \Omega \times [0, \infty) \to \mathbb{R}$ obeys the wave equation $\frac{\partial^2 \phi}{\partial t^2} = c^2 \nabla^2 \phi$ inside $\Omega \times (0, \infty)$, and obeys, for some f, g, h,

  (i) $\phi(x, 0) = f(x)$;

  (ii) $\frac{\partial \phi}{\partial t}(x, 0) = g(x)$; and

  (iii) $\phi|_{\partial \Omega \times [0, \infty)} = h(x)$.

Then ϕ is unique.

Proof. Suppose ϕ_1 and ϕ_2 are two such solutions. Then $\psi = \phi_1 - \phi_2$ obeys the wave equation

$$\frac{\partial^2 \psi}{\partial t^2} = c^2 \nabla^2 \psi,$$

and

$$\psi|_{\partial \Omega \times [0, \infty)} = \psi|_{\Omega \times \{0\}} = \left. \frac{\partial \psi}{\partial t} \right|_{\Omega \times \{0\}} = 0.$$

We let

$$E_\psi(t) = \frac{1}{2} \int_\Omega \left[\left(\frac{\partial \psi}{\partial t} \right)^2 + c^2 \nabla \psi \cdot \nabla \psi \right] \, \mathrm{d}V.$$

Then since ψ obeys the wave equation with fixed boundary conditions, we know E_ψ is constant.

Initially, at $t = 0$, we know that $\psi = \frac{\partial \psi}{\partial t} = 0$. So $E_\psi(0) = 0$. At time t, we have

$$E_\psi = \frac{1}{2} \int_\Omega \left(\frac{\partial \psi}{\partial t} \right)^2 + c^2 (\nabla \psi) \cdot (\nabla \psi) \, \mathrm{d}V = 0.$$

Hence we must have $\frac{\partial \psi}{\partial t} = 0$. So ψ is constant. Since it is 0 at the beginning, it is always 0. $\qquad \square$

Example. Consider $\Omega = \{(x, y) \in \mathbb{R}^2, x^2 + y^2 \leq 1\}$, and let $\phi(r, \theta, t)$ solve

$$\frac{1}{c^2}\frac{\partial^2 \phi}{\partial t^2} = \nabla^2 \phi = \frac{1}{r}\left(r\frac{\partial \phi}{\partial r}\right) + \frac{1}{r^2}\frac{\partial^2 \phi}{\partial \theta^2},$$

with the boundary condition $\phi|_{\partial\Omega} = 0$. We can imagine this as a drum, where the membrane can freely oscillate with the boundary fixed.

Separating variables with $\phi(r, \theta, t) = T(t)R(r)\Theta(\theta)$, we get

$$T'' = -c^2\lambda T, \quad \Theta'' = -\mu\Theta, \quad r(R')' + (r^2\lambda - \mu)R = 0.$$

Then as before, T and Θ are both sine and cosine waves. Since we are in polars coordinates, we need $\phi(t, r, \theta + 2\pi) = \phi(t, r, \theta)$. So we must have $\mu = m^2$ for some $m \in \mathbb{N}$. Then the radial equation becomes

$$r(rR')' + (r^2\lambda - m^2)R = 0,$$

which is Bessel's equation of order m. So we have

$$R(r) = a_m J_m(\sqrt{\lambda}r) + b_m Y_m(\sqrt{\lambda}r).$$

Since we want regularity at $r = 0$, we need $b_m = 0$ for all m. To satisfy the boundary condition $\phi|_{\partial\Omega} = 0$, we must choose $\sqrt{\lambda} = k_{mi}$, where k_{mi} is the ith root of J_m.

Hence the general solution is

$$\phi(t, r, \theta) = \sum_{i=0}^{\infty}[A_{0i}\sin(k_{0i}ct) + B_{0i}\cos(k_{0i}ct)]J_0(k_{0i}r)$$

$$+ \sum_{m=1}^{\infty}\sum_{i=0}^{\infty}[A_{mi}\cos(m\theta) + B_{mi}\sin(m\theta)]\sin k_{mi}ct J_m(k_{mi}r)$$

$$+ \sum_{m=1}^{\infty}\sum_{i=0}^{\infty}[C_{mi}\cos(m\theta) + D_{mi}\sin(m\theta)]\cos k_{mi}ct J_m(k_{mi}r)$$

For example, suppose we have the initial conditions $\phi(0, r, \theta) = 0$, $\partial_t\phi(0, r, \theta) = g(r)$. So we start with a flat surface and suddenly hit it with some force. By symmetry, we must have $A_{mi}, B_{mi}, C_{mi}, D_{mi} = 0$ for $m \neq 0$. If this does not seem obvious, we can perform some integrals with the orthogonality relations to prove this.

At $t = 0$, we need $\phi = 0$. So we must have $B_{0i} = 0$. So we are left with

$$\phi = \sum_{i=0}^{\infty} A_{0i}\sin(k_{0i}ct)J_0(k_{0j}r).$$

We can differentiate this, multiply with $J_0(k_{0j}r)r$ to obtain

$$\int_0^1 \sum_{i=0}^{\infty} k_{0i}cA_{0i}J_0(k_{0i}r)J_0(k_{0j}r)r \,\mathrm{d}r = \int_0^1 g(r)J_0(k_{0j}r)r \,\mathrm{d}r.$$

Using the orthogonality relations for J_0 from the example sheet, we get

$$A_{0i} = \frac{2}{ck_{0i}}\frac{1}{[J_0'(k_{0i})]^2}\int_0^1 g(r)J_0(k_{0j}r)r \,\mathrm{d}r.$$

Note that the frequencies come from the roots of the Bessel's function, and are not evenly spaced. This is different from, say, string instruments, where the frequencies are evenly spaced. So drums sound differently from strings.

So far, we have used separation of variables to solve our differential equations. This is quite good, as it worked in our examples, but there are a few issues with it. Of course, we have the problem of whether it converges or not, but there is a more fundamental problem.

To perform separation of variables, we need to pick a good coordinate system, such that the boundary conditions come in a nice form. We were able to perform separation of variables because we can find some coordinates that fit our boundary conditions well. However, in real life, most things are not nice. Our domain might have a weird shape, and we cannot easily find good coordinates for it.

Hence, instead of solving things directly, we want to ask some more general questions about them. In particular, Kac asked the question "can you hear the shape of a drum?" — suppose we know all the frequencies of the modes of oscillation on some domain Ω. Can we know what Ω is like?

The answer is no, and we can explicitly construct two distinct drums that sound the same. Fortunately, we get an affirmative answer if we restrict ourselves a bit. If we require Ω to be convex, and has a real analytic boundary, then yes! In fact, we have the following result: let $N(\lambda_0)$ be the number of eigenvalues less than λ_0. Then we can show that

$$4\pi^2 \lim_{\lambda_0 \to \infty} \frac{N(\lambda_0)}{\lambda_0} = \text{Area}(\Omega).$$
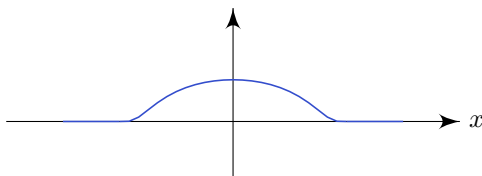
5   Distributions

5.1   Distributions

When performing separation of variables, we also have the problem of convergence as well. To perform separation of variables, we first find some particular solutions of the form, say, $X(x)Y(y)Z(z)$. We know that these solve, say, the wave equation. However, what we do next is take an *infinite* sum of these functions. First of all, how can we be sure that this converges at all? Even if it did, how do we know that the sum satisfies the differential equation? As we have seen in Fourier series, an infinite sum of continuous functions can be discontinuous. If it is not even continuous, how can we say it is a solution of a differential equation?

Hence, at first people were rather skeptical of this method. They thought that while these methods worked, they only work on a small, restricted domain of problems. However, later people realized that this method is indeed rather general, as long as we allow some more "generalized" functions that are not functions in the usual sense. These are known as distributions.

To define a distribution, we first pick a class of "nice" test functions, where "nice" means we can do whatever we want to them (e.g. differentiate, integrate etc.) A main example is the set of infinitely smooth functions with compact support on some set $K \subseteq \Omega$, written $C_{\text{cpt}}^\infty(\Omega)$ or $D(\Omega)$. For example, we can have the bump function defined by

$$\phi(x) = \begin{cases} e^{-1/(1-x^2)} & |x| < 1 \\ 0 & \text{otherwise.} \end{cases}$$

We now define a distribution T to be a linear map $T : D(\Omega) \to \mathbb{R}$. For those who are doing IB Linear Algebra, this is the dual space of the space of ("nice") real-valued functions. For $\phi \in D(\Omega)$, we often write its image under T as $T[\phi]$.

Example. The simplest example is just an ordinary function that is integrable over any compact region. Then we define the distribution T_f as

$$T_f[\phi] = \int_\Omega f(x)\phi(x) \, \mathrm{d}x.$$

Note that this is a linear map since integration is linear (and multiplication is commutative and distributes over addition). Hence every function "is" a distribution.

Of course, this would not be interesting if we only had ordinary functions. The most important example of a distribution (that is not an ordinary function) is the Dirac delta "function".

Definition (Dirac-delta)**.** The *Dirac-delta* is a distribution defined by

$$\delta[\phi] = \phi(0).$$

By analogy with the first case, we often abuse notation and write

$$\delta[\phi] = \int_\Omega \delta(x)\phi(x) \, \mathrm{d}x,$$

and pretend $\delta(x)$ is an actual function. Of course, it is not really a function, i.e. it is not a map $\Omega \to \mathbb{R}$. If it were, then we must have $\delta(x) = 0$ whenever $x \neq 0$, since $\delta[\phi] = \phi(0)$ only depends on what happens at 0. But then this integral will just give 0 if $\delta(0) \in \mathbb{R}$. Some people like to think of this as a function that is zero anywhere and "infinitely large" everywhere else. Formally, though, we have to think of it as a distribution.

Although distributions can be arbitrarily singular and insane, we can nonetheless define *all* their derivatives, defined by

$$T'[\phi] = -T[\phi'].$$

This is motivated by the case of regular functions, where we get, after integrating by parts,

$$\int_\Omega f'(x)\phi(x) \, \mathrm{d}x = -\int_\Omega f(x)\phi'(x) \, \mathrm{d}x,$$

with no boundary terms since we have a compact support. Since ϕ is infinitely differentiable, we can take arbitrary derivatives of distributions.

So even though distributions can be crazy and singular, everything can be differentiated. This is good.

Generalized functions can occur as limits of sequences of normal functions. For example, the family of functions

$$G_n(x) = \frac{n}{\sqrt{\pi}} \exp(-n^2 x^2)$$

are smooth for any finite n, and $G_n[\phi] \to \delta[\phi]$ for any ϕ.
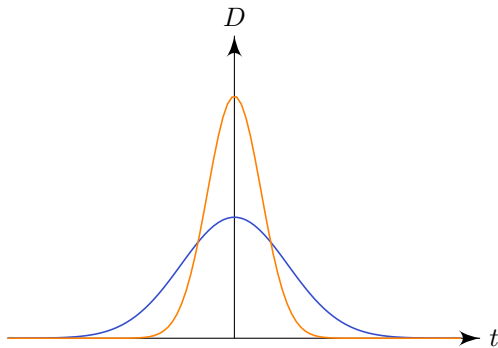
It thus makes sense to define

$$\delta'(\phi) = -\delta[\phi'] = -\phi'(0),$$

as this is the limit of the sequence

$$\lim_{n \to \infty} \int_\Omega G_n'(x)\phi(x) \, \mathrm{d}x.$$

It is often convenient to think of $\delta(x)$ as $\lim\limits_{n\to\infty} G_n(x)$, and $\delta'(x) = \lim\limits_{n\to\infty} G'_n(x)$ etc, despite the fact that these limits do not exist as functions.

We can look at some properties of $\delta(x)$:

– Translation:
$$\int_{-\infty}^{\infty} \delta(x-a)\phi(x)\,\mathrm{d}x = \int_{-\infty}^{\infty} \delta(y)\phi(y+a)\,\mathrm{d}x.$$

– Scaling:
$$\int_{-\infty}^{\infty} \delta(cx)\phi(x)\,\mathrm{d}x = \int_{-\infty}^{\infty} \delta(y)\phi\left(\frac{y}{c}\right)\frac{\mathrm{d}y}{|c|} = \frac{1}{|c|}\phi(0).$$

– These are both special cases of the following: suppose $f(x)$ is a continuously differentiable function with isolated simple zeros at x_i. Then near any of its zeros x_i, we have $f(x) \approx (x-x_i)\left.\frac{\partial f}{\partial x}\right|_{x_i}$. Then

$$\int_{-\infty}^{\infty} \delta(f(x))\phi(x)\,\mathrm{d}x = \sum_{i=1}^{n} \int_{-\infty}^{\infty} \delta\left((x-x_i)\left.\frac{\partial f}{\partial x}\right|_{x_i}\right)\phi(x)\,\mathrm{d}x$$
$$= \sum_{i=1}^{n} \frac{1}{|f'(x_i)|}\phi(x_i).$$

We can also expand the δ-function in a basis of eigenfunctions. Suppose we live in the interval $[-L, L]$, and write a Fourier expansion

$$\delta(x) = \sum_{n\in\mathbb{Z}} \hat{\delta}_n e^{in\pi x/L}$$

with

$$\hat{\delta}_n = \frac{1}{2L}\int_{-L}^{L} e^{in\pi x/L}\delta(x)\,\mathrm{d}x = \frac{1}{2L}$$

So we have

$$\delta(x) = \frac{1}{2L}\sum_{n\in\mathbb{Z}} e^{in\pi x/L}.$$

This *does* make sense as a distribution. We let

$$S_N\delta(x) = \frac{1}{2L}\sum_{n=-N}^{N} e^{in\pi x/L}.$$

Then

$$\lim_{N\to\infty} \int_{-L}^{L} S_N \delta(x)\phi(x)\,\mathrm{d}x = \lim_{N\to\infty} \frac{1}{2L} \int_{L}^{L} \sum_{n=-N}^{N} e^{in\pi x/L}\phi(x)\,\mathrm{d}x$$

$$= \lim_{N\to\infty} \sum_{n=-N}^{N} \left[\frac{1}{2L} \int_{-L}^{L} e^{in\pi x/L}\phi(x)\,\mathrm{d}x \right]$$

$$= \lim_{N\to\infty} \sum_{n=-N}^{N} \hat{\phi}_{-n}$$

$$= \lim_{N\to\infty} \sum_{n=-N}^{N} \hat{\phi}_n e^{in\pi 0/L}$$

$$= \phi(0),$$

since the Fourier series of the smooth function $\phi(x)$ *does* converge for all $x \in [-L, L]$.

We can equally well expand $\delta(x)$ in terms of any other set of orthonormal eigenfunctions. Let $\{y_n(x)\}$ be a complete set of eigenfunctions on $[a, b]$ that are orthogonal with respect to a weight function $w(x)$. Then we can write

$$\delta(x - \xi) = \sum_n c_n y_n(x),$$

with

$$c_n = \int_a^b y_n^*(x)\delta(x - \xi)w(x)\,\mathrm{d}x = y_n^*(\xi)w(\xi).$$

So

$$\delta(x - \xi) = w(\xi) \sum_n y_n^*(\xi)y_n(x) = w(x) \sum_n y_n^*(\xi)y_n(x),$$

using the fact that

$$\delta(x - \xi) = \frac{w(\xi)}{w(x)}\delta(x - \xi).$$

5.2   Green's functions

One of the main uses of the δ function is the Green's function. Suppose we wish to solve the 2nd order ordinary differential equation $\mathcal{L}y = f$, where $f(x)$ is a bounded forcing term, and $\mathcal{L}$ is a bounded differential operator, say

$$\mathcal{L} = \alpha(x)\frac{\partial^2}{\partial x^2} + \beta(x)\frac{\partial}{\partial x} + \gamma(x).$$

As a first step, we try to find the Green's function for $\mathcal{L}$, which we shall write as $G(x, \xi)$, which obeys

$$\mathcal{L}G(x, \xi) = \delta(x - \xi).$$

Given $G(x, \xi)$, we can then define

$$y(x) = \int_a^b G(x, \xi)f(\xi)\,\mathrm{d}\xi.$$

Then we have

$$\mathcal{L}y = \int_a^b \mathcal{L}G(x,\xi)f(\xi)\,\mathrm{d}\xi = \int_a^b \delta(x-\xi)f(\xi)\,\mathrm{d}\xi = f(x).$$

We can formally say $y = \mathcal{L}^{-1}f$, and the Green's function shows that $\mathcal{L}^{-1}$ means an integral.

Hence if we can solve for Green's function, then we have solved the differential equation, at least in terms of an integral.

To find the Green's function $G(x,\xi)$ obeying the boundary conditions

$$G(a,\xi) = G(b,\xi) = 0,$$

first note that $\mathcal{L}G = 0$ for all $x \in [a,\xi) \cup (\xi,b]$, i.e. everywhere except ξ itself. Thus we must be able to expand it in a basis of solutions in these two regions.

Suppose $\{y_1(x), y_2(x)\}$ are a basis of solutions to $\mathcal{L}G = 0$ everywhere on $[a,b]$, with boundary conditions $y_1(a) = 0, y_2(b) = 0$. Then we must have

$$G(x,\xi) = \begin{cases} A(\xi)y_1(x) & a \le x < \xi \\ B(\xi)y_2(x) & \xi < x \le b \end{cases}$$

So we have a whole family of solutions. To fix the coefficients, we must decide how to join these solutions together over $x = \xi$.

If $G(x,\xi)$ were discontinuous at $x = \xi$, then $\partial_x G|_{x=\xi}$ would involve a δ function, while $\partial_x^2 G|_{x=\xi}$ would involve the derivative of the δ function. This is not good, since nothing in $\mathcal{L}G = \delta(x-\xi)$ would balance a δ'. So $G(x,\xi)$ must be everywhere continuous. Hence we require

$$A(\xi)y_1(\xi) = B(\xi)y_2(\xi). \tag{$*$}$$

Now integrate over a small region $(\xi - \varepsilon, \xi + \varepsilon)$ surrounding ξ. Then we have

$$\int_{\xi-\varepsilon}^{\xi+\varepsilon} \left[\alpha(x)\frac{\mathrm{d}^2 G}{\mathrm{d}x^2} + \beta(x)\frac{\mathrm{d}G}{\mathrm{d}x} + \gamma(x)G \right]\,\mathrm{d}x = \int_{\xi-\varepsilon}^{\xi+\varepsilon} \delta(x-\xi)\,\mathrm{d}x = 1.$$

By continuity of G, we know that the γG term does not contribute. While G' is discontinuous, it is still finite. So the $\beta G'$ term also does not contribute. So we have

$$\lim_{\varepsilon \to 0} \int_{\xi-\varepsilon}^{\xi+\varepsilon} \alpha G''\,\mathrm{d}x = 1.$$

We now integrate by parts to obtain

$$\lim_{\varepsilon \to 0} [\alpha G']_{\xi-\varepsilon}^{\xi+\varepsilon} + \int_{\xi-\varepsilon}^{\xi+\varepsilon} \alpha' G'\,\mathrm{d}x = 1.$$

Again, by finiteness of G', the integral does not contribute. So we know that

$$\alpha(\xi)\left(\left.\frac{\partial G}{\partial x}\right|_{\xi+} - \left.\frac{\partial G}{\partial x}\right|_{\xi-} \right) = 1$$

Hence we obtain

$$B(\xi)y_2'(\xi) - A(\xi)y_1'(\xi) = \frac{1}{\alpha(\xi)}.$$

Together with $(*)$, we know that

$$A(\xi) = \frac{y_2(\xi)}{\alpha(\xi)W(\xi)}, \quad B(\xi) = \frac{y_1(\xi)}{\alpha(\xi)W(\xi)},$$

where W is the *Wronskian*

$$W = y_1 y_2' - y_2 y_1'.$$

Hence, we know that

$$G(x, \xi) = \frac{1}{\alpha(\xi)W(\xi)} \begin{cases} y_2(\xi)y_1(x) & a \le x \le \xi \\ y_1(\xi)y_2(x) & \xi < x \le b. \end{cases}$$

Using the step function Θ, we can write this as

$$G(x, \xi) = \frac{1}{\alpha(\xi)W(\xi)}[\Theta(\xi - x)y_2(\xi)y_1(x) + \Theta(x - \xi)y_1(\xi)y_2(x)].$$

So our general solution is

$$\begin{aligned} y(x) &= \int_a^b G(x, \xi)f(\xi) \, d\xi \\ &= \int_x^b \frac{f(\xi)}{\alpha(\xi)W(\xi)}y_2(\xi)y_1(x) \, d\xi + \int_a^x \frac{f(\xi)}{\alpha(\xi)W(\xi)}y_1(\xi)y_2(x) \, d\xi. \end{aligned}$$

Example. Consider

$$\mathcal{L}y = -y'' - y = f$$

for $x \in (0, 1)$ with $y(0) = y(1) = 0$. We choose our basis solution to satisfy $y'' = -y$ as $\{\sin x, \sin(1 - x)\}$. Then we can compute the Wronskian

$$W(x) = -\sin x \cos(1 - x) - \sin(1 - x)\cos x = -\sin 1.$$

Our Green's function is

$$G(x, \xi) = \frac{1}{\sin 1}[\Theta(\xi - x)\sin(1 - \xi)\sin x + \Theta(x - \xi)\sin \xi \sin(1 - x)].$$

Hence we get

$$y(x) = \sin x \int_x^1 \frac{\sin(1 - \xi)}{\sin 1}f(\xi) \, d\xi + \sin(1 - x)\int_0^x \frac{\sin \xi}{\sin 1}f(\xi) \, d\xi.$$

Example. Suppose we have a string with mass per unit length $\mu(x)$ and held under a tension T. In the presence of gravity, Newton's second law gives

$$\mu \frac{\partial^2 y}{\partial t^2} = T\frac{\partial^2 y}{\partial x^2} + \mu g.$$
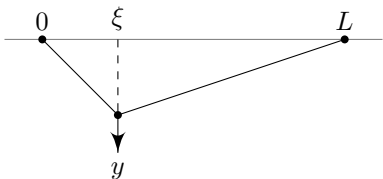
We look for the steady state solution $\dot{y} = 0$ shape of the string, assuming $y(0, t) = y(L, t) = 0$. So we get

$$\frac{\partial^2 y}{\partial x^2} = -\frac{\mu(x)}{T}g.$$

We look for a Green's function obeying

$$\frac{\partial^2 y}{\partial x^2} = -\delta(x - \xi).$$

This can be interpreted as the contribution of a pointlike mass located at $x = \xi$.

The homogeneous equation $y'' = 0$ gives

$$y = Ax + B(x - L).$$

So we get

$$G(x, \xi) = \begin{cases} A(\xi x) & 0 \le x < \xi \\ B(\xi)(x - L) & \xi < x \le L. \end{cases}$$

Continuity at $x = \xi$ gives

$$A(\xi)\xi = B(\xi)(\xi - L).$$

The jump condition on the derivative gives

$$A(\xi) - B(\xi) = 1.$$

We can solve these to get

$$A(\xi) = \frac{\xi - L}{L}, \quad B(\xi) = \frac{\xi}{L}.$$

Hence the Green's function is

$$G(x, \xi) = \frac{\xi - L}{L}\Theta(\xi - x) + \frac{\xi}{L}(x - L)\Theta(x - \xi).$$

Notice that ξ is always less that L. So the first term has a negative slope; while ξ is always positive, so the second term has positive slope.

For the general case, we can just substitute this into the integral. We can give this integral a physical interpretation. We can think that we have many pointlike particles m_i at small separations Δx along the string. So we get

$$g(x) = \sum_{i=1}^{n} G(x, x_i)\frac{m_i g}{T},$$

and in the limit $\Delta x \to 0$ and $n \to \infty$, we get

$$g(x) \to \frac{g}{T} \int_0^L G(x, \xi)\mu(\xi) \, \mathrm{d}\xi.$$

5.3 Green's functions for initial value problems

Consider $\mathcal{L}y = f(t)$ subject to $y(t = t_0) = 0$ and $y'(t = t_0) = 0$. So instead of having two boundaries, we have one boundary and restrict both the value and the derivative at this boundary.

As before, let $y_1(t), y_2(t)$ be *any* basis of solutions to $\mathcal{L}y = 0$. The Green's function obeys

$$\mathcal{L}(G) = \delta(t - \tau).$$

We can write our Green's function as

$$G(t, \tau) = \begin{cases} A(\tau)y_1(t) + B(\tau)y_2(t) & t_0 \leq t < \tau \\ C(\tau)y_1(t) + D(\tau)y_2(t) & t > \tau. \end{cases}$$

Our initial conditions require that

$$\begin{pmatrix} y_1(t_0) & y_2(t_0) \\ y_1'(t_0) & y_2'(t_0) \end{pmatrix} \begin{pmatrix} A \\ B \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$$

However, we know that the matrix is non-singular (by definition of y_1, y_2 being a basis). So we must have $A, B = 0$ (which makes sense, since $G = 0$ is obviously a solution for $t_0 \leq t < \tau$).

Our continuity and jump conditions now require

$$0 = C(\tau)t_1(t) + D(\tau)y_2(t)$$

$$\frac{1}{\alpha(\tau)} = C(\tau)y_1'(\tau) + D(\tau)y_2'(\tau).$$

We can solve this to get $C(\tau)$ and $D(\tau)$. Then we get

$$y(t) = \int_{t_0}^{\infty} G(t, \tau) f(\tau) \, \mathrm{d}\tau = \int_{t_0}^{t} G(t, \tau) f(\tau) \, \mathrm{d}\tau,$$

since we know that when $\tau > t$, the Green's function $G(t, \tau) = 0$. Thus the solution $y(t)$ depends on the forcing term term f only for times $< t$. This expresses causality!

Example. For example, suppose we have $\ddot{y} + y = f(t)$ with $f(0) = \dot{y}(0) = 0$. Then we have

$$G(t, \tau) = \Theta(t - \tau)[C(\tau)\cos(t - \tau) + D(\tau)\sin(t - \tau)]$$

for some $C(\tau), D(\tau)$. The continuity and jump conditions gives $D(\tau) = 1, C(\tau) = 0$. So we get

$$G(t, \tau) = \Theta(t - \tau)\sin t(\tau).$$

So we get

$$y(t) = \int_0^t \sin(t - \tau) f(\tau) \, \mathrm{d}\tau.$$

6   Fourier transforms

6.1   The Fourier transform

So far, we have considered writing a function $f : S^1 \to \mathbb{C}$ (or a periodic function) as a Fourier sum

$$f(\theta) = \sum_{n \in \mathbb{Z}} \hat{f}_n e^{in\theta}, \quad \hat{f}_n = \frac{1}{2\pi} \int_{-\pi}^{\pi} e^{-in\theta} f(\theta) \, \mathrm{d}\theta.$$

What if we don't have a periodic function, but just an arbitrary function $f : \mathbb{R} \to \mathbb{C}$? Can we still do Fourier analysis? Yes!

To start with, instead of $\hat{f}_n = \frac{1}{2\pi} \int_{-\pi}^{\pi} e^{-in\theta} f(\theta) \, \mathrm{d}\theta$, we define the Fourier transform as follows:

Definition (Fourier transform)**.** The *Fourier transform* of an (absolutely integrable) function $f : \mathbb{R} \to \mathbb{C}$ is defined as

$$\tilde{f}(k) = \int_{-\infty}^{\infty} e^{-ikx} f(x) \, \mathrm{d}x$$

for *all* $k \in \mathbb{R}$. We will also write $\tilde{f}(k) = \mathcal{F}[f(x)]$.

Note that for any k, we have

$$|\tilde{f}(k)| = \left| \int_{-\infty}^{\infty} e^{-ikx} f(x) \, \mathrm{d}x \right| \leq \int_{-\infty}^{\infty} |e^{-ikx} f(x)| \, \mathrm{d}x = \int_{-\infty}^{\infty} |f(x)| \, \mathrm{d}x.$$

Since we have assumed that our function is absolutely integrable, this is finite, and the definition makes sense.

We can look at a few basic properties of the Fourier transform.

(i) Linearity: if $f, g : \mathbb{R} \to \mathbb{C}$ are absolutely integrable and c_1, c_2 are constants, then
$$\mathcal{F}[c_1 f(x) + c_2 g(x)] = c_1 \mathcal{F}[f(x)] + c_2 \mathcal{F}[g(x)].$$

This follows directly from the linearity of the integral.

(ii) Translation:

$$\mathcal{F}[f(x - a)] = \int_{\mathbb{R}} e^{-ikx} f(x - a) \, \mathrm{d}x$$

Let $y = x - a$. Then we have

$$= \int_{\mathbb{R}} e^{-ik(y+a)} f(y) \, \mathrm{d}y$$
$$= e^{-ika} \int_{\mathbb{R}} e^{-iky} f(y) \, \mathrm{d}y$$
$$= e^{-ika} \mathcal{F}[f(x)],$$

So a translation in x-space becomes a re-phasing in k-space.

(iii) Re-phasing:

$$\mathcal{F}[e^{-i\ell x}f(x)] = \int_{-\infty}^{\infty} e^{-ikx}e^{-i\ell x}f(x)\ \mathrm{d}x$$

$$= \int_{-\infty}^{\infty} e^{-i(k+\ell)x}f(x)\ \mathrm{d}x$$

$$= \tilde{f}(k+\ell),$$

where $\ell \in \mathbb{R}$ and $\tilde{f}(x) = \mathcal{F}[f(x)]$.

(iv) Scaling:

$$\mathcal{F}[f(cx)] = \int_{-\infty}^{\infty} e^{-ikx}f(cx)\ \mathrm{d}x$$

$$= \int_{-\infty}^{\infty} e^{-iky/c}f(y)\frac{\mathrm{d}y}{|c|}$$

$$= \frac{1}{|c|}\tilde{f}\left(\frac{k}{c}\right).$$

Note that we have to take the absolute value of c, since if we replace x by y/c and c is negative, then we will flip the bounds of the integral. The extra minus sign then turns c into $-c = |c|$.

(v) Convolutions: for functions $f, g : \mathbb{R} \to \mathbb{C}$, we define the *convolution* as

$$f * g(x) = \int_{-\infty}^{\infty} f(x-y)g(y)\ \mathrm{d}y.$$

We then have

$$\mathcal{F}[f * g(x)] = \int_{-\infty}^{\infty} e^{-ikx}\left[\int_{-\infty}^{\infty} f(x-y)g(y)\ \mathrm{d}y\right]\ \mathrm{d}x$$

$$= \int_{\mathbb{R}^2} e^{ik(x-y)}f(x-y)e^{-iky}g(y)\ \mathrm{d}y\ \mathrm{d}x$$

$$= \int_{\mathbb{R}} e^{-iku}f(u)\ \mathrm{d}u \int_{\mathbb{R}} e^{-iky}g(y)\ \mathrm{d}y$$

$$= \mathcal{F}[f]\mathcal{F}[g],$$

where $u = x - y$. So the Fourier transform of a convolution is the product of individual Fourier transforms.

(vi) The most useful property of the Fourier transform is that it "turns differentiation into multiplication". Integrating by parts, we have

$$\mathcal{F}[f'(x)] = \int_{-\infty}^{\infty} e^{-ikx}\frac{\mathrm{d}f}{\mathrm{d}x}\ \mathrm{d}x$$

$$= \int_{-\infty}^{\infty} -\frac{\mathrm{d}}{\mathrm{d}x}(e^{-ikx})f(x)\ \mathrm{d}x$$

Note that we don't have any boundary terms since for the function to be absolutely integrable, it has to decay to zero as we go to infinity. So we have

$$= ik \int_{-\infty}^{\infty} e^{-ikx} f(x) \, dx$$

$$= ik\mathcal{F}[f(x)]$$

Conversely,

$$\mathcal{F}[xf(x)] = \int_{-\infty}^{\infty} e^{-ikx} xf(x) \, dx = i\frac{d}{dk} \int_{-\infty}^{\infty} e^{-ikx} f(x) \, dx = i\tilde{f}'(k).$$

This are useful results, since if we have a complicated function to find the Fourier transform of, we can use these to break it apart into smaller pieces and hopefully make our lives easier.

Suppose we have a differential equation

$$\mathcal{L}(\partial)y = f,$$

where

$$\mathcal{L}(\partial) = \sum_{r=0}^{p} c_r \frac{d^r}{dx^r}$$

is a differential operator of pth order with constant coefficients.

Taking the Fourier transform of both sides of the equation, we find

$$\mathcal{F}[\mathcal{L}(\partial)y] = \mathcal{F}[f(x)] = \tilde{f}(k).$$

The interesting part is the left hand side, since the Fourier transform turns differentiation into multiplication. The left hand side becomes

$$c_0\tilde{y}(k) + c_1 ik\tilde{y}(k) + c_2(ik)^2\tilde{y}(k) + \cdots + c_p(ik)^p\tilde{y}(k) = \mathcal{L}(ik)\tilde{y}(k).$$

Here $\mathcal{L}(ik)$ is a polynomial in ik. Thus taking the Fourier transform has changed our ordinary differential equation into the algebraic equation

$$\mathcal{L}(ik)\tilde{y}(k) = \tilde{f}(k).$$

Since $\mathcal{L}(ik)$ is just multiplication by a polynomial, we can immediately get

$$\tilde{y}(k) = \frac{\tilde{f}(k)}{\mathcal{L}(ik)}.$$

Example. For $\phi : \mathbb{R}^n \to \mathbb{C}$, suppose we have the equation

$$\nabla^2\phi - m^2\phi = \rho(x),$$

where

$$\nabla^2 = \sum_{i=1}^{n} \frac{d^2}{dx_i^2}$$

is the Laplacian on $\mathbb{R}^n$.

We define the n dimensional Fourier transform by

$$\tilde{\phi}(\mathbf{k}) = \int_{\mathbb{R}^n} e^{-i\mathbf{k}\cdot\mathbf{x}} f(\mathbf{x}) \, \mathrm{d}\mathbf{x} = \mathcal{F}[\phi(x)]$$

where now $\mathbf{k} \in \mathbb{R}^n$ is an n-dimensional vector. So we get

$$\mathcal{F}[\nabla^2\phi - m^2\phi] = \mathcal{F}[\rho] = \tilde{\rho}(\mathbf{k}).$$

The first term is

$$\int_{\mathbb{R}^n} e^{-i\mathbf{k}\cdot\mathbf{x}} \nabla \cdot \nabla\phi \, \mathrm{d}^n x = -\int_{\mathbb{R}^n} \nabla(e^{-i\mathbf{k}\cdot\mathbf{x}}) \cdot \nabla\phi \, \mathrm{d}^n x$$

$$= \int_{\mathbb{R}^n} \nabla^2(e^{-i\mathbf{k}\cdot\mathbf{x}})\phi \, \mathrm{d}^n x$$

$$= -\mathbf{k}\cdot\mathbf{k} \int_{\mathbb{R}^n} e^{-i\mathbf{k}\cdot\mathbf{x}} \phi(\mathbf{x}) \, \mathrm{d}^n x$$

$$= -\mathbf{k}\cdot\mathbf{k}\tilde{\phi}(\mathbf{k}).$$

So our equation becomes

$$-\mathbf{k}\cdot\mathbf{k}\tilde{\phi}(\mathbf{k}) - m^2\tilde{\phi}(\mathbf{k}) = \tilde{\rho}(\mathbf{k}).$$

So we get

$$\tilde{\phi}(\mathbf{k}) = -\frac{\tilde{\rho}(\mathbf{k})}{|\mathbf{k}|^2 + m^2}.$$

So differential equations are trivial in k space. The problem is, of course, that we want our solution in x space, not k space. So we need to find a way to restore our function back to x space.

6.2   The Fourier inversion theorem

We need to be able to express our original function $f(x)$ in terms of its Fourier transform $\tilde{f}(k)$. Recall that in the periodic case where $f(x) = f(x + L)$, we have

$$f(x) = \sum_{n\in\mathbb{Z}} \hat{f}_n e^{2inx\pi/L}.$$

We can try to obtain a similar expression for a non-periodic function $f : \mathbb{R} \to \mathbb{C}$ by taking the limit $L \to \infty$. Recall that our coefficients are defined by

$$\hat{f}_n = \frac{1}{L} \int_{-L/2}^{L/2} e^{-2in\pi x/L} f(x) \, \mathrm{d}x.$$

Now define

$$\Delta k = \frac{2\pi}{L}.$$

So we have

$$f(x) = \sum_{n\in\mathbb{Z}} e^{inx\Delta k} \frac{\Delta k}{2\pi} \int_{-L/2}^{L/2} e^{-inx\Delta k} f(u) \, \mathrm{d}u$$

As we take the limit $L \to \infty$, we can approximate

$$\int_{-L/2}^{L/2} e^{-ixn\Delta k} f(x) \, \mathrm{d}x = \int_{-\infty}^{\infty} e^{-ix(n\Delta k)} f(x) \, \mathrm{d}x = \tilde{f}(n\Delta k).$$

Then for a non-periodic function, we have

$$f(x) = \lim_{\Delta k \to \infty} \sum_{n \in \mathbb{Z}} \frac{\Delta k}{2\pi} e^{in\Delta kx} \tilde{f}(n\Delta k) = \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{ikx} \tilde{f}(k) \, \mathrm{d}k.$$

So

$$f(x) = \mathcal{F}^{-1}[f(k)] = \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{ikx} \tilde{f}(k) \, \mathrm{d}k.$$

This is a really dodgy argument. We first took the limit as $L \to \infty$ to turn the integral from $\int_{-L}^{L}$ to $\int_{-\infty}^{\infty}$. *Then* we take the limit as $\Delta k \to 0$. However, we can't really do this, since Δk is defined to be $2\pi/L$, and both limits have to be taken at the same time. So we should really just take this as a heuristic argument for why this works, instead of a formal proof.

Nevertheless, note that the inverse Fourier transform looks very similar to the Fourier transform itself. The only differences are a factor of $\frac{1}{2\pi}$, and the sign of the exponent. We can get rid of the first difference by evenly distributing the factor among the Fourier transform and inverse Fourier transform. For example, we can instead define

$$\mathcal{F}[f(x)] = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{-ikx} f(x) \, \mathrm{d}x.$$

Then $\mathcal{F}^{-1}$ looks just like $\mathcal{F}$ apart from having the wrong sign.

However, we will stick to our original definition so that we don't have to deal with messy square roots. Hence we have the duality

$$\mathcal{F}[f(x)] = \mathcal{F}^{-1}[f(-x)](2\pi).$$

This is useful because it means we can use our knowledge of Fourier transforms to compute inverse Fourier transform.

Note that this does not occur in the case of the Fourier *series*. In the Fourier series, we obtain the coefficients by evaluating an integral, and restore the original function by taking a discrete sum. These operations are not symmetric.

Example. Recall that we defined the convolution as

$$f * g(x) = \int_{-\infty}^{\infty} f(x - y) g(y) \, \mathrm{d}y.$$

We then computed

$$\mathcal{F}[f * g(x)] = \tilde{f}(k) \tilde{f}(g).$$

It follows by definition of the inverse that

$$\mathcal{F}^{-1}[\tilde{f}(k) \tilde{g}(k)] = f * g(x).$$

Therefore, we know that

$$\mathcal{F}[f(x)g(x)] = 2\pi \int_{-\infty}^{\infty} \tilde{f}(k - \ell) \tilde{g}(\ell) \, \mathrm{d}\ell = 2\pi \tilde{f} * \tilde{g}(k).$$

Example. Last time we saw that if $\nabla^2\phi - m^2\phi = \rho(x)$, then

$$\tilde{\phi}(\mathbf{k}) = -\frac{\tilde{\rho}(\mathbf{k})}{|\mathbf{k}|^2 + m^2}.$$

Hence we can retrieve our ϕ as

$$\phi(\mathbf{x}) = \mathcal{F}^{-1}[\tilde{\phi}(k)] = -\frac{1}{(2\pi)^n} \int_{\mathbb{R}^n} e^{i\mathbf{k}\cdot\mathbf{x}} \frac{\tilde{\rho}(\mathbf{k})}{|\mathbf{k}|^2 + m^2} \, \mathrm{d}^n x.$$

Note that we have $(2\pi)^n$ instead of 2π since we get a factor of 2π for each dimension (and the negative sign was just brought down from the original expression).

Since we are taking the inverse Fourier transform of a product, we know that the result is just a convolution of $\mathcal{F}^{-1}[\tilde{\rho}(k)] = \rho(x)$ and $\mathcal{F}^{-1}\left(\frac{1}{|\mathbf{k}|^2+m^2}\right)$.

Recall that when we worked with Green's function, if we have a forcing term, then the final solution is the convolution of our Green's function with the forcing term. Here we get a similar expression in higher dimensions.

6.3   Parseval's theorem for Fourier transforms

Theorem (Parseval's theorem (again))**.** Suppose $f, g : \mathbb{R} \to \mathbb{C}$ are sufficiently well-behaved that $\tilde{f}$ and $\tilde{g}$ exist and we indeed have $\mathcal{F}^{-1}[\tilde{f}] = f, \mathcal{F}^{-1}[\tilde{g}] = g$. Then

$$(f, g) = \int_{\mathbb{R}} f^*(x)g(x) \, \mathrm{d}x = \frac{1}{2\pi}(\tilde{f}, \tilde{g}).$$

In particular, if $f = g$, then

$$\|f\|^2 = \frac{1}{2\pi}\|\tilde{f}\|^2.$$

So taking the Fourier transform preserves the L^2 norm (up to a constant factor of $\frac{1}{2\pi}$).
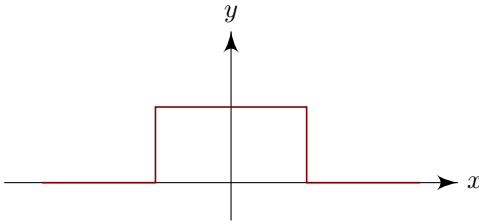
Proof.

$$\begin{aligned}
(f, g) &= \int_{\mathbb{R}} f^*(x)g(x) \, \mathrm{d}x \\
&= \int_{-\infty}^{\infty} f^*(x) \left[\frac{1}{2\pi} \int_{-\infty}^{\infty} e^{ikx}\tilde{g}(x) \, \mathrm{d}k\right] \, \mathrm{d}x \\
&= \frac{1}{2\pi} \int_{-\infty}^{\infty} \left[\int_{-\infty}^{\infty} f^*(x)e^{ikx} \, \mathrm{d}x\right] \tilde{g}(k) \, \mathrm{d}k \\
&= \frac{1}{2\pi} \int_{-\infty}^{\infty} \left[\int_{-\infty}^{\infty} f(x)e^{-ikx} \, \mathrm{d}x\right]^* \tilde{g}(k) \, \mathrm{d}k \\
&= \frac{1}{2\pi} \int_{-\infty}^{\infty} \tilde{f}^*(k)\tilde{g}(k) \, \mathrm{d}k \\
&= \frac{1}{2\pi}(\tilde{f}, \tilde{g}). \hspace{3cm} \square
\end{aligned}$$

6.4   A word of warning

Suppose $f(x)$ is defined by

$$f(x) = \begin{cases} 1 & |x| < 1 \\ 0 & |x| \geq 1. \end{cases}$$

This function looks rather innocent. Sure it has discontinuities at ± 1, but these are not too absurd, and we are just integrating. This is certainly absolutely integrable. We can easily compute the Fourier transform as

$$\tilde{f}(k) = \int_{-\infty}^{\infty} e^{-ikx} f(x) \, \mathrm{d}x = \int_{-1}^{1} e^{-ikx} \, \mathrm{d}x = \frac{2 \sin k}{k}.$$

Is our original function equal to the integral $\mathcal{F}^{-1}[\tilde{f}(l)]$? We have

$$\mathcal{F}^{-1} \left[\frac{2 \sin k}{k} \right] = \frac{1}{\pi} \int_{-\infty}^{\infty} e^{ikx} \frac{\sin k}{k} \, \mathrm{d}k.$$

This is hard to do. So let's first see if this function is absolutely integrable. We have

$$\int_{-\infty}^{\infty} \left| e^{ikx} \frac{\sin k}{k} \right| \, \mathrm{d}x = \int_{-\infty}^{\infty} \left| \frac{\sin k}{k} \right| \, \mathrm{d}k$$

$$= 2 \int_{0}^{\infty} \left| \frac{\sin k}{k} \right| \, \mathrm{d}k$$

$$\geq 2 \sum_{n=0}^{N} \int_{(n+1/4)\pi}^{(n+3/4)\pi} \left| \frac{\sin k}{k} \right| \, \mathrm{d}k.$$

The idea here is instead of looking at the integral over the whole real line, we just pick out segments of it. In these small segments, we know that $|\sin k| \geq \frac{1}{\sqrt{2}}$. So we can bound this by

$$\int_{-\infty}^{\infty} \left| e^{ikx} \frac{\sin k}{k} \right| \, \mathrm{d}x \geq 2 \sum_{n=0}^{N} \frac{1}{\sqrt{2}} \int_{(n+1/4)\pi}^{(n+3/4)\pi} \frac{\mathrm{d}k}{k}$$

$$\geq \sqrt{2} \sum_{n=0}^{N} \int_{(n+1/4)\pi}^{(n+3/4)\pi} \frac{\mathrm{d}k}{(n+1)\pi}$$

$$= \frac{\sqrt{2}}{2} \sum_{n=0}^{N} \frac{1}{n+1},$$

and this diverges as $N \to \infty$. So we don't have absolute integrability. So we have to be careful when we are doing these things. Well, in theory. We actually won't. We'll just ignore this issue.

6.5   Fourier transformation of distributions

We have

$$\mathcal{F}[\delta(x)] = \int_{-\infty}^{\infty} e^{ikx}\delta(x)\,\mathrm{d}x = 1.$$

Hence we have

$$\mathcal{F}^{-1}[1] = \frac{1}{2\pi}\int_{-\infty}^{\infty} e^{ikx}\,\mathrm{d}k = \delta(x).$$

Of course, it is extremely hard to make sense of this integral. It quite obviously diverges as a normal integral, but we can just have faith and believe this makes sense as long as we are talking about distributions.

Similarly, from our rules of translations, we get

$$\mathcal{F}[\delta(x - a)] = e^{-ika}$$

and

$$\mathcal{F}[e^{-i\ell x}] = 2\pi\delta(k - \ell),$$

Hence we get

$$\mathcal{F}[\cos(\ell x)] = \mathcal{F}\left[\frac{1}{2}(e^{i\ell x} + e^{-i\ell x})\right] = \frac{1}{2}\mathcal{F}[e^{i\ell x}] + \frac{1}{2}\mathcal{F}[e^{-i\ell x}] = \pi[\delta(k-\ell)+\delta(k+\ell)].$$

We see that highly localized functions in x-space have very spread-out behaviour in k-space, and vice versa.

6.6   Linear systems and response functions

Suppose we have an amplifier that modifies an input signal $I(t)$ to produce an output $O(t)$. Typically, amplifiers work by modifying the amplitudes and phases of specific frequencies in the output. By Fourier's inversion theorem, we know

$$I(t) = \frac{1}{2\pi}\int e^{i\omega t}\tilde{I}(\omega)\,\mathrm{d}\omega.$$

This $\tilde{I}(\omega)$ is the *resolution* of $I(t)$.

We specify what the amplifier does by the *transfer function* $\tilde{R}(\omega)$ such that the output is given by

$$O(t) = \frac{1}{2\pi}\int_{-\infty}^{\infty} e^{i\omega t}\tilde{R}(\omega)\tilde{I}(\omega)\,\mathrm{d}\omega.$$

Since this $\tilde{R}(\omega)\tilde{I}(\omega)$ is a product, on computing $O(t) = \mathcal{F}^{-1}[\tilde{R}(\omega)\tilde{I}(\omega)]$, we obtain a *convolution*

$$O(t) = \int_{-\infty}^{\infty} R(t - u)I(u)\,\mathrm{d}u,$$

where

$$R(t) = \frac{1}{2\pi}\int_{-\infty}^{\infty} e^{i\omega t}\tilde{R}(\omega)\,\mathrm{d}\omega$$

is the *response function*. By plugging it directly into the equation above, we see that $R(t)$ is the response to an input signal $I(t) = \delta(t)$.

Now note that causality implies that the amplifier cannot "respond" before any input has been given. So we must have $R(t) = 0$ for all $t < 0$.

Assume that we only start providing input at $t = 0$. Then

$$O(t) = \int_{-\infty}^{\infty} R(t-u)I(u) \ \mathrm{d}u = \int_0^t R(t-u)I(u) \ \mathrm{d}u.$$

This is exactly same form of solution as we found for initial value problems with the response function $R(t)$ playing the role of the Green's function.

6.7 General form of transfer function

To model the situation, suppose the amplifier's operation is described by the ordinary differential equation

$$I(t) = \mathcal{L}_m O(t),$$

where

$$\mathcal{L}_m = \sum_{i=0}^m a_i \frac{\mathrm{d}^i}{\mathrm{d}t^i}$$

with $a_i \in \mathbb{C}$. In other words, we have an mth order ordinary differential equation with constant coefficients, and the input is the forcing function. Notice that we are not saying that $O(t)$ can be expressed as a function of derivatives of $I(t)$. Instead, $I(t)$ influences $O(t)$ by being the forcing term the differential equation has to satisfy. This makes sense because when we send an input into the amplifier, this input wave "pushes" the amplifier, and the amplifier is forced to react in some way to match the input.

Taking the Fourier transform and using $\mathcal{F}\left[\frac{\mathrm{d}O}{\mathrm{d}t}\right] = i\omega \tilde{O}(\omega)$, we have

$$\tilde{I}(\omega) = \sum_{j=0}^m a_j (i\omega)^j \tilde{O}(\omega).$$

So we get

$$\tilde{O}(\omega) = \frac{\tilde{I}(\omega)}{a_0 + i\omega a_1 + \cdots + (i\omega)^m a_m}.$$

Hence, the transfer function is just

$$\tilde{R}(\omega) = \frac{1}{a_0 + i\omega a_1 + \cdots + (i\omega)^m a_m}.$$

The denominator is an nth order polynomial. By the fundamental theorem of algebra, it has m roots, say $c_j \in \mathbb{C}$ for $j = 1, \cdots, J$, where c_j has multiplicity k_j. Then we can write our transfer function as

$$\tilde{R}(\omega) = \frac{1}{a_m} \prod_{j=1}^J \frac{1}{(i\omega - c_j)^{k_j}}.$$

By repeated use of partial fractions, we can write this as

$$\tilde{R}(\omega) = \sum_{j=1}^J \sum_{r=1}^{k_j} \frac{\Gamma_{rj}}{(i\omega - c_j)^r}$$

for some constants $\Gamma_{rj} \in \mathbb{C}$.

By linearity of the (inverse) Fourier transform, we can find $O(t)$ if we know the inverse Fourier transform of all functions of the form

$$\frac{1}{(i\omega - \alpha)^p}.$$

To compute this, there are many possible approaches. We can try to evaluate the integral directly, learn some fancy tricks from complex analysis, or cheat, and let the lecturer tell you the answer. Consider the function

$$h_0(t) = \begin{cases} e^{\alpha t} & t > 0 \\ 0 & \text{otherwise} \end{cases}$$

The Fourier transform is then

$$\tilde{h}_o(\omega) = \int_{-\infty}^{\infty} e^{-i\omega t} h_0(t) \, \mathrm{d}t = \int_0^{\infty} e^{(\alpha - i\omega)t} \, \mathrm{d}t = \frac{1}{i\omega - \alpha}$$

provided $\mathrm{Re}(\alpha) < 0$ (so that $e^{(\alpha - i\omega)t} \to 0$ as $t \to \infty$). Similarly, we can let

$$h_1(t) = \begin{cases} te^{\alpha t} & t > 0 \\ 0 & \text{otherwise} \end{cases}$$

Then we get

$$\tilde{h}_1(\omega) = \mathcal{F}[th_0(t)] = i\frac{\mathrm{d}}{\mathrm{d}\omega}\mathcal{F}[h_0(t)] = \frac{1}{(i\omega - \alpha)^2}.$$

Proceeding inductively, we can define

$$h_p(t) = \begin{cases} \frac{t^p}{p!}e^{\alpha t} & t > 0 \\ 0 & \text{otherwise} \end{cases},$$

and this has Fourier transform

$$\tilde{h}_p(\omega) = \frac{1}{(i\omega - \alpha)^{p+1}},$$

again provided $\mathrm{Re}(\alpha) < 0$. So the response function is a linear combination of these functions $h_p(t)$ (if any of the roots c_j have non-negative real part, then it turns out the system is unstable, and is better analysed by the Laplace transform). We see that the response function does indeed vanish for all $t < 0$. In fact, each term (except h_0) increases from zero at $t = 0$ to rise to some maximum before eventually decaying exponentially.

Fourier transforms can also be used to solve ordinary differential equations on the whole real line $\mathbb{R}$, provided the functions are sufficiently well-behaved. For example, suppose $y : \mathbb{R} \to \mathbb{R}$ solves

$$y'' - A^2 y = -f(x)$$

with y and $y' \to 0$ as $|x| \to \infty$. Taking the Fourier transform, we get

$$\tilde{y}(k) = \frac{\tilde{f}(k)}{k^2 + A^2}.$$

Since this is a product, the inverse Fourier transform will be a convolution of $f(x)$ and an object whose Fourier transform is $\frac{1}{k^2+A^2}$. Again, we'll cheat. Consider

$$h(x) = \frac{1}{2\mu}e^{-\mu|x|}.$$

with $\mu > 0$. Then

$$\tilde{f}(k) = \frac{1}{2\mu} \int_{-\infty}^{\infty} e^{-ikx}e^{-\mu|x|} \; \mathrm{d}x$$

$$= \frac{1}{\mu} \operatorname{Re} \int_{0}^{\infty} e^{-(\mu+ik)x} \; \mathrm{d}x$$

$$= \frac{1}{\mu} \operatorname{Re} \frac{1}{ik + \mu}$$

$$= \frac{1}{\mu^2 + k^2}.$$

Therefore we get

$$y(x) = \frac{1}{2A} \int_{-\infty}^{\infty} e^{-A|x-u|}f(u) \; \mathrm{d}u.$$

So far, we have done Fourier analysis over some abelian groups. For example, we've done it over S^1, which is an abelian group under multiplication, and $\mathbb{R}$, which is an abelian group under addition. We will next look at Fourier analysis over another abelian group, $\mathbb{Z}_m$, known as the discrete Fourier transform.

6.8   The discrete Fourier transform

Recall that the Fourier transform is defined as

$$\tilde{f}(k) = \int_{\mathbb{R}} e^{-ikx} f(x) \; \mathrm{d}x.$$

To find the Fourier transform, we have to know how to perform the integral. If we cannot do the integral, then we cannot do the Fourier transform. However, it is usually difficult to perform integrals for even slightly more complicated functions.

A more serious problem is that we usually don't even have a closed form of the function f for us to perform the integral. In real life, f might be some radio signal, and we are just given some data of what value f takes at different points.. There is no hope that we can do the integral exactly.

Hence, we first give ourselves a simplifying assumption. We suppose that f is mostly concentrated in some finite region $[-R, S]$. More precisely, $|f(x)| \ll 1$ for $x \notin [-R, S]$ for some $R, S > 0$. Then we can approximate our integral as

$$\tilde{f}(k) = \int_{-R}^{S} e^{-ikx} f(x) \; \mathrm{d}x.$$

Afterwards, we perform the integral numerically. Suppose we "sample" $f(x)$ at

$$x = x_j = -R + j\frac{R + S}{N}$$

for N a large integer and $j = 0, 1, \cdots, N - 1$. Then

$$\tilde{f}(k) \approx \frac{R + S}{N} \sum_{j=0}^{N-1} f(x_j) e^{-ikx_j}.$$

This is just the Riemann sum.

Similarly, our computer can only store the result $\tilde{f}(k)$ for some finite list of k. Let's choose these to be at

$$k = k_m = \frac{2\pi m}{R + S}.$$

Then after some cancellation,

$$\tilde{f}(k_m) \approx \frac{R + S}{N} e^{\frac{2\pi i m R}{R + S}} \sum_{j=0}^{N-1} f(x_j) e^{\frac{2\pi i}{N} jm}$$

$$= (R + S) e^{\frac{2\pi i m R}{R + S}} \left[\frac{1}{N} \sum_{j=0}^{N-1} f(x_j) \omega^{-jm} \right],$$

where

$$\omega = e^{\frac{2\pi i}{N}}$$

is an Nth root of unity. We call the thing in the brackets

$$F(m) = \frac{1}{N} \sum_{j=0}^{N-1} f(x_j) \omega^{-jm}.$$

Of course, we've thrown away lots of information about our original function $f(x)$, since we made approximations all over the place. For example, we have already lost all knowledge of structures varying more rapidly than our sampling interval $\frac{R+S}{N}$. Also, $F(m + N) = F(m)$, since $\omega^N = 1$. So we have "forced" some periodicity into our function F, while $\tilde{f}(k)$ itself was not necessarily periodic.

For the usual Fourier transform, we were able to re-construct our original function from the $\tilde{f}(k)$, but here we clearly cannot. However, if we know the $F(m)$ for all $m = 0, 1, \cdots, N - 1$, then we *can* reconstruct the exact values of $f(x_j)$ for all j by just solving linear equations. To make this more precise, we want to put what we're doing into the linear algebra framework we've been using. Let $G = \{1, \omega, \omega^2, \cdots, \omega^{N-1}\}$. For our purposes below, we can just treat this as a discrete set, and ω^i are just meaningless symbols that happen to visually look like the powers of our Nth roots of unity.

Consider a function $g : G \to \mathbb{C}$ defined by

$$g(\omega^j) = f(x_j).$$

This is actually nothing but just a new notation for $f(x_j)$. Then using this new notation, we have

$$F(m) = \frac{1}{N} \sum_{j=0}^{N-1} \omega^{-jm} g(\omega^j).$$

The space of *all* functions $g : G \to \mathbb{C}$ is a finite-dimensional vector space, isomorphic to $\mathbb{C}^N$. This has an inner product

$$(f, g) = \frac{1}{N} \sum_{j=0}^{N-1} f^*(\omega^j) g(\omega^j).$$

This inner product obeys

$$(g, f) = (f, g)^*$$
$$(f, c_1 g_1 + c_2 g_2) = c_1(f, g_1) + c_2(f, g_2)$$
$$(f, f) \geq 0 \text{ with equality iff } f(\omega_j) = 0$$

Now let $e_n : G \to \mathbb{C}$ be the function

$$e_m(\omega^j) = \omega^{jm}.$$

Then the set of functions $\{e_m\}$ for $m = 0, \cdots, N - 1$ is an orthonormal basis with respect to our inner product. To show this, we can compute

$$(e_m, e_m) = \frac{1}{N} \sum_{j=0}^{N-1} e_m^*(\omega^j) e_m(\omega^j)$$
$$= \frac{1}{N} \sum_{j=0}^{N-1} \omega^{-jm} \omega^{jm}$$
$$= \frac{1}{N} \sum_{j=0}^{N-1} 1$$
$$= 1.$$

For $n \neq m$, we have

$$(e_n, e_m) = \frac{1}{N} \sum_{j=0}^{N-1} e_n^*(\omega^j) e^m(\omega^j)$$
$$= \frac{1}{N} \sum_{j=0}^{N-1} \omega^{j(m-n)}$$
$$= \frac{1}{N} \frac{1 - \omega^{(m-n)N}}{1 - \omega^{m-n}}.$$

Since $m - n$ is an integer and ω is an Nth root of unity, we know that $\omega^{(m-n)N} = 1$. So the numerator is 0. However, since $n \neq m$, $m - n \neq 0$. So the denominator is non-zero. So we get 0. Hence

$$(e_n, e_m) = 0.$$

We can now rewrite our $F(m)$ as

$$F(m) = \frac{1}{N} \sum_{j=0}^{N-1} \omega^{-jm} f(x_j) = \frac{1}{N} \sum_{j=0}^{N-1} e_m^*(\omega^j) g(\omega^j) = (e_m, g).$$

Hence we can expand our g as

$$g = \sum_{m=0}^{N-1} (e_m, g)e_m = \sum_{m=0}^{N-1} F(m)e_m.$$

Writing f instead of g, we recover the formula

$$f(x_j) = g(\omega^j) = \sum_{m=0}^{N-1} F(m)e_m(\omega^j).$$

If we forget about our fs and just look at the g, what we have effectively done is take the Fourier transform of functions taking values on $G = \{1, \omega, \cdots, \omega^{N-1}\} \cong \mathbb{Z}_N$.

This is exactly analogous to what we did for the Fourier transform, except that everything is now discrete, and we don't have to worry about convergence since these are finite sums.

6.9   The fast Fourier transform*

What we've said so far is we've defined

$$F(m) = \frac{1}{N} \sum_{j=0}^{N-1} \omega^{-mj} g(\omega^j).$$

To compute this directly, even given the values of ω^{-mj} for all j, m, this takes $N - 1$ additions and $N + 1$ multiplications. This is $2N$ operations for a *single* value of m. To know $F(m)$ for all m, this takes $2N^2$ operations.

This is a problem. Historically, during the cold war, especially after the Cuban missile crisis, people were in fear that the world will one day go into a huge nuclear war, and something had to be done to prevent this. Thus the countries decided to come up with a treaty to ensure people don't perform nuclear testings anymore.

However, the obvious problem is that we can't easily check if other countries are doing nuclear tests. Nuclear tests are easy to spot if the test is done above ground, since there will be a huge explosion. However, if it is done underground, it is much harder to test.

They then came up with the idea to put some seismometers all around the world, and record vibrations in the ground. To distinguish normal crust movements from nuclear tests, they wanted to take the Fourier transform and analyze the frequencies of the vibrations. However, they had a *large* amount of data, and the value of N is on the order of magnitude of a few million. So $2N^2$ will be a huge number that the computers at that time were not able to handle. So they need to find a trick to perform Fourier transforms quickly. This is known as the *fast Fourier transform*. Note that this is nothing new mathematically. This is entirely a computational trick.

Now suppose $N = 2M$. We can write

$$F(m) = \frac{1}{2M} \sum_{j=0}^{2M-1} \omega^{-jm} g(\omega^j)$$

$$= \frac{1}{2} \left[\frac{1}{M} \sum_{k=0}^{M-1} \omega^{-2km} g(\omega^{2k}) + \omega^{-(2k+1)m} g(\omega^{2k+1}) \right].$$

We now let η be an Mth root of unity, and define

$$G(\eta^k) = g(\omega^{2k}), \quad H(\eta^k) = g(\omega^{2k+1}).$$

We then have

$$F(m) = \frac{1}{2} \left[\frac{1}{M} \sum_{k=0}^{M-1} \eta^{-km} G(\eta^k) + \frac{\omega^{-m}}{M} \sum_{k=0}^{M-1} \eta^{-km} H(\eta^k) \right]$$

$$= \frac{1}{2} [\tilde{G}(m) + \omega^{-m} \tilde{H}(m)].$$

Suppose we are given the values of $\tilde{G}(m)$ and $\tilde{H}(m)$ and ω^{-m} for all $m = \{0, \cdots, N-1\}$. Then we can compute $F(m)$ using 3 operations per value of m. So this takes $3 \times N = 6M$ operations for all M.

We can compute ω^{-m} for all m using at most $2N$ operations, and suppose it takes P_M operations to compute $\tilde{G}(m)$ for all m. Then the number of operations needed to compute $F(m)$ is

$$P_{2M} = 2P_M + 6M + 2M.$$

Now let $N = 2^n$. Then by iterating this, we can find

$$P_N \leq 4N \log_2 N \ll N^2.$$

So by breaking the Fourier transform apart, we are able to greatly reduce the number of operations needed to compute the Fourier transform. This is used nowadays everywhere for everything.

7   More partial differential equations

7.1   Well-posedness

Recall that to find a unique solution of Laplace's equation $\nabla^2 \phi = 0$ on a bounded domain $\Omega \subseteq \mathbb{R}^n$, we imposed the Dirichlet boundary condition $\phi|_{\partial\Omega} = f$ or the Neumann boundary condition $\mathbf{n} \cdot \nabla\phi|_{\partial\Omega} = g$.

For the heat equation $\frac{\partial\phi}{\partial t} = \kappa \nabla^2 \phi$ on $\Omega \times [0, \infty)$, we asked for $\partial|_{\partial\Omega \times [0,\infty)} = f$ and also $\phi|_{\Omega \times \{0\}} = g$.

For the wave equation $\frac{\partial^2 \phi}{\partial t^2} = c^2 \nabla^2 \phi$ on $\Omega \times [0, \infty)$, we imposed $\phi|_{\partial\Omega \times [0,\infty)} = f$, $\phi|_{\Omega \times \{0\}} = g$ and also $\partial_t\phi|_{\Omega \times \{0\}} = h$.

All the boundary and initial conditions restrict the value of ϕ on some co-dimension 1 surface, i.e. a surface whose dimension is one less than the domain. This is known as *Cauchy data* for our partial differential equation.

Definition (Well-posed problem)**.** A partial differential equation problem is said to be well-posed if its Cauchy data means

(i) A solution exists;

(ii) The solution is unique;

(iii) A "small change" in the Cauchy data leads to a "small change" in the solution.

To understand the last condition, we need to make it clear what we mean by "small change". To do this properly, we need to impose some topology on our space of functions, which is some technicalities we will not go into. Instead, we can look at a simple example.

Suppose we have the heat equation $\partial_t\phi = \kappa \nabla^2 \phi$. We know that whatever starting condition we have, the solution quickly smooths out with time. Any spikiness of the initial conditions get exponentially suppressed. Hence this is a well-posed problem — changing the initial condition slightly will result in a similar solution.

However, if we take the heat equation but run it backwards in time, we get a non-well-posed problem. If we provide a tiny, small change in the "ending condition", as we go back in time, this perturbation grows *exponentially*, and the result could vary wildly.

Another example is as follows: consider the Laplace's equation $\nabla^2 \phi$ on the upper half plane $(x, y) \in \mathbb{R} \times \mathbb{R}^{\geq 0}$ subject to the boundary conditions

$$\phi(x, 0) = 0, \quad \partial_y\phi(x, 0) = g(x).$$

If we take $g(x) = 0$, then $\phi(x, y) = 0$ is the unique solution, obviously.

However, if we take $g(x) = \frac{\sin Ax}{A}$, then we get the unique solution

$$\phi(x, y) = \frac{\sin(Ax)\sinh(Ay)}{A^2}.$$

So far so good. However, now consider the limit as $A \to \infty$. Then

$$g(x) = \frac{\sin(Ax)}{A} \to 0.$$

for all $x \in \mathbb{R}$. However, at the special point $\phi\left(\frac{\pi}{2A}, y\right)$, we get

$$\phi\left(\frac{\pi}{2A}, y\right) = \frac{\sinh(Ay)}{A^2} \to A^{-2}e^{Ay},$$

which is unbounded. So as we take the limit as our boundary conditions $g(x) \to 0$, we get an unbounded solution.

How can we make sure this doesn't happen? We're first going to look at first order equations.

7.2   Method of characteristics

Curves in the plane

Suppose we have a curve on the plane $\mathbf{x} : \mathbb{R} \to \mathbb{R}^2$ given by $s \mapsto (x(s), y(s))$.

Definition (Tangent vector). The tangent vector to a smooth curve C given by $\mathbf{x} : \mathbb{R} \to \mathbb{R}^2$ with $\mathbf{x}(s) = (x(s), y(s))$ is

$$\left(\frac{\mathrm{d}x}{\mathrm{d}s}, \frac{\mathrm{d}y}{\mathrm{d}s}\right).$$

If we have some quantity $\phi : \mathbb{R}^2 \to \mathbb{R}$, then its value along the curve C is just $\phi(x(s), y(s))$.

A vector field $\mathbf{V}(x, y) : \mathbb{R}^2 \to \mathbb{R}^2$ defines a family of curves. To imagine this, suppose we are living in a river. The vector field tells us the how the water flows at each point. We can then obtain a curve by starting a point and flow along with the water. More precisely,

Definition (Integral curve). Let $\mathbf{V}(x, y) : \mathbb{R}^2 \to \mathbb{R}^2$ be a vector field. The *integral curves* associated to $\mathbf{V}$ are curves whose tangent $\left(\frac{\mathrm{d}x}{\mathrm{d}s}, \frac{\mathrm{d}y}{\mathrm{d}s}\right)$ is just $\mathbf{V}(x, y)$.

For sufficiently regular vector fields, we can fill the space with different curves. We will parametrize which curve we are on by the parameter t. More precisely, we have a curve $B = (x(t), y(t))$ that is *transverse* (i.e. nowhere parallel) to our family of curves, and we can label the members of our family by the value of t at which they intersect B.
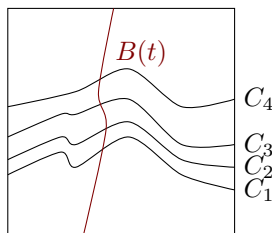
We can thus label our family of curves by $(x(s, t), y(s, t))$. If the Jacobian

$$J = \frac{\partial x}{\partial s}\frac{\partial y}{\partial t} - \frac{\partial x}{\partial t}\frac{\partial y}{\partial s} \neq 0,$$

then we can invert this to find (s, t) as a function of (x, y), i.e. at any point, we know which curve we are on, and how far along the curve we are.

This means we now have a new coordinate system (s, t) for our points in $\mathbb{R}^2$. It turns out by picking the right vector field $\mathbf{V}$, we can make differential equations much easier to solve in this new coordinate system.

The method of characteristics

Suppose $\phi : \mathbb{R}^2 \to \mathbb{R}$ obeys

$$a(x, y)\frac{\partial \phi}{\partial x} + b(x, y)\frac{\partial \phi}{\partial y} = 0.$$

We can define our vector field as

$$\mathbf{V}(x, y) = \begin{pmatrix} a(x, y) \\ b(x, y) \end{pmatrix}.$$

Then we can write the differential equation as

$$\mathbf{V} \cdot \nabla \phi = 0.$$

Along any particular integral curve of $\mathbf{V}$, we have

$$\frac{\partial \phi}{\partial s} = \frac{\mathrm{d}x(s)}{\mathrm{d}s}\frac{\partial \phi}{\partial x} + \frac{\mathrm{d}y(s)}{\mathrm{d}s}\frac{\partial \phi}{\partial y} = \mathbf{V} \cdot \nabla \phi,$$

where the integral curves of $\mathbf{V}$ are determined by

$$\left.\frac{\partial x}{\partial s}\right|_t = a(x, y), \quad \left.\frac{\partial y}{\partial s}\right|_t = b(x, y).$$

This is known as the characteristic equation.

Hence our partial differential equation just becomes the equation

$$\left.\frac{\partial \phi}{\partial s}\right|_t = 0.$$

To get ourselves a well-posed problem, we have to specify our boundary data along a transverse curve B. We pick our transverse curve as $s = 0$, and we suppose we are given the initial data

$$\phi(0, t) = h(t).$$

Since ϕ does not vary with s, our solution automatically is

$$\phi(s, t) = h(t).$$

Then $\phi(x, y)$ is given by inverting the characteristic equations to find $t(x, y)$.

We do a few examples to make this clear.

Example (Trivial)**.** Suppose ϕ obeys

$$\left.\frac{\partial \phi}{\partial x}\right|_y = 0.$$

We are given the boundary condition

$$\phi(0, y) = f(y).$$

The solution is obvious, since the differential equation tells us the function is just a function of y, and then the solution is obviously $\phi(x, y) = h(y)$. However,

we can try to use the method of characteristics to practice our skills. Our vector field is given by

$$\mathbf{V} = \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} \frac{\mathrm{d}x}{\mathrm{d}s} \\ \frac{\mathrm{d}y}{\mathrm{d}s} \end{pmatrix}.$$

Hence, we get

$$\frac{\mathrm{d}x}{\mathrm{d}s} = 1, \quad \frac{\mathrm{d}y}{\mathrm{d}s} = 0.$$

So we have

$$x = s + c, \quad y = d.$$

Our initial curve is given by $y = t, x = 0$. Since we want this to be the curve $s = 0$, we find our integral curves to be

$$x = s, \quad y = t.$$

Now for each fixed t, we have to solve

$$\left. \frac{\partial \phi}{\partial s} \right|_t = 0, \quad \phi(s,t) = h(t) = f(t).$$

So we know that

$$\phi(x,y) = f(y).$$

Example. Consider the equation

$$e^x \frac{\partial \phi}{\partial x} + \frac{\partial \phi}{\partial y} = 0.$$

with the boundary condition

$$\phi(x,0) = \cosh x.$$

We find that the vector field is

$$\mathbf{V} = \begin{pmatrix} e^x \\ 1 \end{pmatrix}$$

This gives

$$\frac{\mathrm{d}x}{\mathrm{d}s} = e^x, \quad \frac{\mathrm{d}y}{\mathrm{d}s} = 1.$$

Thus we have

$$e^{-x} = -s + c, \quad y = s + d.$$

We pick our curve as $x = t, y = 0$. So our relation becomes

$$e^{-x} = -s + e^{-t}, \quad y = s.$$

We thus have

$$\left. \frac{\partial \phi}{\partial s} \right|_t = 0, \quad \phi(s,t) = \phi(0,t) = \cosh(t) = \cosh[-\ln(y + e^{-x})]$$

So done.

Example. Let $\phi : \mathbb{R}^2 \to \mathbb{R}$ solve the inhomogeneous partial differential equation

$$\partial_x \phi + 2\partial_y \phi = ye^x$$

with $\phi(x, x) = \sin x$.

We can still use the method of characteristics. We have

$$\mathbf{u} = \begin{pmatrix} 1 \\ 2 \end{pmatrix}.$$

So the characteristic curves obey

$$\frac{\mathrm{d}x}{\mathrm{d}s} = 1, \quad \frac{\mathrm{d}y}{\mathrm{d}s} = 2.$$

This gives

$$x = s + t, \quad y = 2s + t$$

so that $x = y = t$ at $s = 0$. We can invert this to obtain the relations

$$s = y - x, \quad t = 2x - y.$$

The partial differential equation now becomes

$$\left.\frac{\mathrm{d}\phi}{\mathrm{d}s}\right|_t = \mathbf{u} \cdot \nabla \phi = ye^x = (2s + t)e^{s+t}$$

Note that this is just an ordinary differential equation in s because t is held constant. We have the boundary conditions

$$\phi(s = 0, t) = \sin t.$$

So we get

$$\phi(x(s, t), y(s, t)) = (2 - t)e^t(1 - e^s) + \sin t + 2se^{s+t}.$$

Putting it in terms of x and y, we have

$$\phi(x, y) = (2 - 2x + y)e^{2x-y} + \sin(2x - y) + (y - 2)e^x.$$

We see that our Cauchy data $\phi(x, x) = \sin x$ should be specified on a curve B that intersects each characteristic curve exactly once.

If we tried to use a characteristic curve B that intersects the characteristics multiple times, if we want a solution *at all*, we cannot specify data freely along B. its values at the points of intersection have to compatible with the partial differential equation.

For example, if we have a homogeneous equation, we saw the solution will be constant along the same characteristic. Hence if our Cauchy data requires the solution to take different values on the same characteristic, we will have no solution.

Moreover, even if it is compatible, if B does not intersect all characteristics, we will not get a unique solution. So the solution is not fixed along such characteristics.

On the other hand, if B intersects each characteristic curve transversely, the problem is well-posed. So there exists a unique solution, at least in the neighbourhood of B. Notice that data is never transferred between characteristic curves.

7.3   Characteristics for 2nd order partial differential equations

Whether the method of characteristics works for a 2nd order partial differential equation, or even higher order ones, depends on the "type" of the differential equation.

To classify these differential equations, we need to define

Definition (Symbol and principal part)**.** Let $\mathcal{L}$ be the general 2nd order differential operator on $\mathbb{R}^n$. We can write it as

$$\mathcal{L} = \sum_{i,j=1}^{n} a^{ij}(x) \frac{\partial^2}{\partial x^i \partial x^j} + \sum_{i=1}^{n} b^i(x) \frac{\partial}{\partial x^i} + c(X),$$

where $a^{ij}(x), b^i(x), c(x) \in \mathbb{R}$ and $a^{ij} = a^{ji}$ (wlog).

We define the *symbol* $\sigma(\mathbf{k}, x)$ of $\mathcal{L}$ to be

$$\sigma(\mathbf{k}, x) = \sum_{i,j=1}^{n} a^{ij}(x) k_i k_j + \sum_{i=1}^{n} b^i(x) k_i + c(x).$$

So we just replace the derivatives by the variable k.

The *principal part* of the symbol is the leading term

$$\sigma^p(\mathbf{k}, x) = \sum_{i,j=1}^{n} a^{ij}(x) k_i k_j.$$

Example. If $\mathcal{L} = \nabla^2$, then

$$\sigma(\mathbf{k}, x) = \sigma^p(\mathbf{k}, x) = \sum_{i=1}^{n} (k_i)^2.$$

If $\mathcal{L}$ is the heat operator (where we think of the last coordinate x^n as the time, and others as space), then the operator is given by

$$\mathcal{L} = \frac{\partial}{\partial x^n} - \sum_{i=1}^{n-1} \frac{\partial^2}{\partial x^{i2}}.$$

The symbol is then

$$\sigma(\mathbf{k}, x) = k_n - \sum_{i=1}^{n-1} (k_i)^2,$$

and the principal part is

$$\sigma^p(\mathbf{k}, x) = - \sum_{i=1}^{n-1} (k_i)^2.$$

Note that the symbol is closely related to the Fourier transform of the differential operator, since both turn differentiation into multiplication. Indeed, they are equal if the coefficients are constant. However, we define this symbol for arbitrary differential operators with non-constant coefficients.

In general, for each x, we can write

$$\sigma^p(\mathbf{k}, x) = \mathbf{k}^T A(x) \mathbf{k},$$

where $A(x)$ has elements $a^{ij}(x)$. Recall that a real symmetric matrix (such as A) has all real eigenvalues. So we define the following:

Definition (Elliptic, hyperbolic, ultra-hyperbolic and parabolic differential operators)**.** Let $\mathcal{L}$ be a differential operator. We say $\mathcal{L}$ is

- *elliptic at x* if all eigenvalues of $A(x)$ have the same sign. Equivalently, if $\sigma^p(\,\cdot\,, x)$ is a definite quadratic form;

- *hyperbolic at x* if all but one eigenvalues of $A(x)$ have the same sign;

- *ultra-hyperbolic at x* if $A(x)$ has more than one eigenvalues of each sign;

- *parabolic at x* if $A(x)$ has a zero eigenvalue, i.e. $\sigma^p(\,\cdot\,, x)$ is degenerate.

We say $\mathcal{L}$ is elliptic if $\mathcal{L}$ is elliptic at all x, and similarly for the other terms.

These names are supposed to remind us of conic sections, and indeed if we think of $\mathbf{k}^T A \mathbf{k}$ as an equation in $\mathbf{k}$, then we get a conic section.

Example. Let

$$\mathcal{L} = a(x,y)\frac{\partial^2}{\partial x^2} + 2b(x,y)\frac{\partial^2}{\partial x \partial y} + c(x,y)\frac{\partial^2}{\partial y^2} + d(x,y)\frac{\partial}{\partial x} + e(x,y)\frac{\partial}{\partial y} + f(x,y).$$

Then the principal part of the symbol is

$$\sigma^p(\mathbf{k}, x) = \begin{pmatrix} k_x & k_y \end{pmatrix} \begin{pmatrix} a(x,y) & b(x,y) \\ b(x,y) & c(x,y) \end{pmatrix} \begin{pmatrix} k_x \\ k_y \end{pmatrix}$$

Then $\mathcal{L}$ is elliptic at x if $b^2 - ac < 0$; hyperbolic if $b^2 - ac > 0$; and parabolic if $b^2 - ac = 0$.

Note that since we only have two dimensions and hence only two eigenvalues, we cannot possibly have an ultra-hyperbolic equation.

Characteristic surfaces

Definition (Characteristic surface)**.** Given a differential operator $\mathcal{L}$, let

$$f(x^1, x^2, \cdots, x^n) = 0$$

define a surface $C \subseteq \mathbb{R}^n$. We say C is *characteristic* if

$$\sum_{i,j=1}^n a^{ij}(x)\frac{\partial f}{\partial x^i}\frac{\partial f}{\partial x^j} = (\nabla f)^T A (\nabla f) = \sigma^p(\nabla f, x) = 0.$$

In the case where we only have two dimensions, a characteristic surface is just a curve.

We see that the characteristic equation restricts what values ∇f can take. Recall that ∇f is the normal to the surface C. So in general, at any point, we can find what the normal of the surface should be, and stitch these together to form the full characteristic surfaces.

For an elliptic operator, all the eigenvalues of A have the same sign. So there are no non-trivial real solutions to this equation $(\nabla f)^T A (\nabla f) = 0$. Consequently, elliptic operators have no real characteristics. So the method of characteristics would not be of any use when studying, say, Laplace's equation, at least if we want to stay in the realm of real numbers.

If $\mathcal{L}$ is parabolic, we for simplicity assume that A has exactly one zero eigenvector, say $\mathbf{n}$, and the other eigenvalues have the same sign. This is the case when, say, there are just two dimensions. So we have $A\mathbf{n} = \mathbf{n}^T A = \mathbf{0}$. We normalize $\mathbf{n}$ such that $\mathbf{n} \cdot \mathbf{n} = 1$.

For any ∇f, we can always decompose it as

$$\nabla f = \mathbf{n}(\mathbf{n} \cdot \nabla f) + [\nabla f - \mathbf{n} \cdot (\mathbf{n} \cdot \nabla f)].$$

This is a relation that is trivially true, since we just add and subtract the same thing. Note, however, that the first term points along $\mathbf{n}$, while the latter term is orthogonal to $\mathbf{n}$. To save some writing, we adopt the notation

$$\nabla_\perp f = \nabla f - \mathbf{n}(\mathbf{n} \cdot \nabla \mathbf{f}).$$

So we have

$$\nabla f = \mathbf{n}(\mathbf{n} \cdot \nabla f) + \nabla f_\perp.$$

Then we can compute

$$\begin{aligned}(\nabla f)^T A (\nabla f) &= [\mathbf{n}(\mathbf{n} \cdot \nabla f) + \nabla_\perp f]^T A [\mathbf{n}(\mathbf{n} \cdot \nabla f) + \nabla_\perp f] \\ &= (\nabla_\perp f)^T A (\nabla_\perp f).\end{aligned}$$

Then by assumption, $(\nabla_\perp f)^T A (\nabla_\perp f)$ is definite. So just as in the elliptic case, there are no non-trivial solutions. Hence, if f defines a characteristic surface, then $\nabla_\perp f = 0$. In other words, f is parallel to $\mathbf{n}$. So at any point, the normal to a characteristic surface must be $\mathbf{n}$, and there is only one possible characteristic.

If $\mathcal{L}$ is hyperbolic, we assume all but one eigenvalues are positive, and let $-\lambda$ be the unique negative eigenvalue. We let $\mathbf{n}$ be the corresponding unit eigenvector, where we normalize it such that $\mathbf{n} \cdot \mathbf{n} = 1$. We say f is characteristic if

$$\begin{aligned}0 &= (\nabla f)^T A (\nabla f) \\ &= [\mathbf{n}(\mathbf{n} \cdot \nabla f) + \nabla_\perp f]^T A [\mathbf{n}(\mathbf{n} \cdot \nabla f) + \nabla_\perp f] \\ &= -\lambda(\mathbf{n} \cdot \nabla f)^2 + (\nabla_\perp f)^T A (\nabla_\perp f).\end{aligned}$$

Consequently, for this to be a characteristic, we need

$$\mathbf{n} \cdot \nabla f = \pm \sqrt{\frac{(\nabla_\perp f)^T A (\nabla_\perp f)}{\lambda}}.$$

So there are two choices for $\mathbf{n} \cdot \nabla f$, given any $\nabla_\perp f$. So hyperbolic equations have two characteristic surfaces through any point.

This is not too helpful in general. However, in the case where we have two dimensions, we can find the characteristic curves explicitly. Suppose our curve is given by $f(x, y) = 0$. We can write $y = y(x)$. Then since f is constant along each characteristic, by the chain rule, we know

$$0 = \frac{\partial f}{\partial x} + \frac{\partial f}{\partial y}\frac{\mathrm{d}y}{\mathrm{d}x}.$$

Hence, we can compute

$$\frac{\mathrm{d}y}{\mathrm{d}x} = -\frac{-b \pm \sqrt{b^2 - ac}}{a}.$$

We now see explicitly how the type of the differential equation influences the number of characteristics — if $b^2 - ac > 0$, then we obtain two distinct differential equations and obtain two solutions; if $b^2 - ac = 0$, then we only have one equation; if $b^2 - ac < 0$, then there are no real characteristics.

Example. Consider

$$\partial_y^2 \phi - xy\partial_x^2 \phi = 0$$

on $\mathbb{R}^2$. Then $a = -xy, b = 0, c = 1$. So $b^2 - ac = xy$. So the type is elliptic if $xy < 0$, hyperbolic if $xy > 0$, and parabolic if $xy = 0$.

In the regions where it is hyperbolic, we find

$$\frac{-b \pm \sqrt{b^2 - ac}}{a} = \pm\frac{1}{\sqrt{xy}}.$$

Hence the two characteristics are given by

$$\frac{\mathrm{d}y}{\mathrm{d}x} = \pm\frac{1}{\sqrt{xy}}.$$

This has a solution

$$\frac{1}{3}y^{3/2} \pm x^{1/2} = c.$$

We now let

$$u = \frac{1}{3}y^{3/2} + x^{1/2},$$
$$v = \frac{1}{3}y^{3/2} - x^{1/2}.$$

Then the equation becomes

$$\frac{\partial^2 \phi}{\partial u \partial v} + \text{lower order terms} = 0.$$

Example. Consider the wave equation

$$\frac{\partial^2 \phi}{\partial t^2} - c^2\frac{\partial^2 \phi}{\partial x^2} = 0$$

on $\mathbb{R}^{1,1}$. Then the equation is hyperbolic everywhere, and the characteristic curves are $x \pm ct = \text{const}$. Let's look for a solution to the wave equation that obeys

$$\phi(x, 0) = f(x), \quad \partial_t\phi(x, 0) = g(x).$$

Now put $u = x - ct, v = x + ct$. Then the wave equation becomes

$$\frac{\partial^2}{\partial u \partial v} = 0.$$

So the general solution to this is

$$\phi(x, t) = G(u) + H(v) = G(x - ct) + H(x + ct).$$

The initial conditions now fix these functions

$$f(x) = G(x) + H(x), \quad g(x) = -cG'(x) + cH'(x).$$

Solving these, we find

$$\phi(x, t) = \frac{1}{2}[f(x - ct) + f(x + ct)] + \frac{1}{2c}\int_{x-ct}^{x+ct} g(y)\ \mathrm{d}y.$$

This is d'Alembert's solution to the $1 + 1$ dimensional wave equation.

Note that the value of ϕ at any point (x, t) is *completely* determined by f, g in the interval $[x - ct, x + ct]$. This is known as the *(past) domain of dependence* of the solution at (x, t), written $D^-(x, t)$. Similarly, at any time t, the initial data at $(x_0, 0)$ only affects the solution within the region $x_0 - ct \leq x \leq x_0 + ct$. This is the *range of influence* of data at x_0, written $D^+(x_0)$.
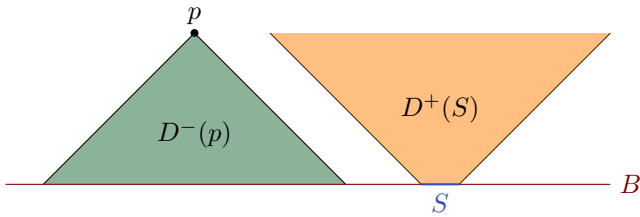
We see that disturbances in the wave equation propagate with speed c.

7.4 Green's functions for PDEs on $\mathbb{R}^n$

Green's functions for the heat equation

Suppose $\phi : \mathbb{R}^n \times [0, \infty) \to \mathbb{R}$ solves the heat equation

$$\partial_t \phi = D\nabla^2 \phi,$$

where D is the diffusion constant, subject to

$$\phi|_{\mathbb{R}^n \times \{0\}} = f.$$

To solve this, we take the Fourier transform of the heat equation in the spatial variables. For simplicity of notation, we take $n = 1$. Writing $\mathcal{F}[\phi(x, t)] = \tilde{\phi}(k, t)$, we get

$$\partial_t \tilde{\phi}(k, t) = -Dk^2 \tilde{\phi}(k, t)$$

with the boundary conditions

$$\tilde{\phi}(k, 0) = \tilde{f}(k).$$

Note that this is just an ordinary differential equation in t, and k is a fixed parameter for each k. So we have

$$\tilde{\phi}(k,t) = \tilde{f}(k)e^{-Dk^2t}.$$

We now take the inverse Fourier transform and get

$$\phi(x,t) = \frac{1}{2\pi} \int_{\mathbb{R}} e^{ikx} \left[\tilde{f}(k)e^{-Dk^2t} \right] \, dk$$

This is the inverse Fourier transform of a product. This thus gives the convolution

$$\phi(x,t) = f * \mathcal{F}^{-1}[e^{-Dk^2t}].$$

So we are done if we can find the inverse Fourier transform of the Gaussian. This is an easy exercise on example sheet 3, where we find

$$\mathcal{F}[e^{-a^2x^2}] = \frac{\sqrt{\pi}}{a} e^{-k^2/4a^2}.$$

So setting

$$a^2 = \frac{1}{4Dt},$$

So we get

$$\mathcal{F}^{-1}[e^{-Dk^2t}] = \frac{1}{\sqrt{4\pi Dt}} \exp\left(-\frac{x^2}{4Dt}\right)$$

We shall call this $S_1(x,t)$, where the subscript 1 tells us we are in $1+1$ dimensions. This is known as the *fundamental solution* of the heat equation. We then get

$$\phi(x,t) = \int_{-\infty}^{\infty} f(y)S_1(x-y,t) \, dy$$

Example. Suppose our initial data is

$$f(x) = \phi_0\delta(x).$$

So we start with a really cold room, with a huge spike in temperature at the middle of the room. Then we get

$$\phi(x,t) = \frac{\phi_0}{\sqrt{4\pi Dt}} \exp\left(-\frac{x^2}{4Dt}\right).$$

What this shows, as we've seen before, is that if we start with a delta function, then as time evolves, we get a Gaussian that gets shorter and broader.

Now note that if we start with a delta function, at $t = 0$, everywhere outside the origin is non-zero. However, after any infinitely small time t, ϕ becomes non-zero everywhere, instantly. Unlike the wave equation, information travels instantly to all of space, i.e. heat propagates arbitrarily fast according to this equation (of course in reality, it doesn't). This fundamentally, is because the heat equation is parabolic, and only has one characteristic.

Now suppose instead ϕ satisfies the inhomogeneous, forced, heat equation

$$\partial_t \phi - D\nabla^2 \phi = F(x, t),$$

but with homogeneous initial conditions $\phi|_{t=0} = 0$. Physically, this represents having an external heat source somewhere, starting at zero.

Note that if we can solve this, then we have completely solved the heat equation. If we have an inhomogeneous equation *and* an inhomogeneous initial condition, then we can solve this forced problem with homogeneous boundary conditions to get ϕ_F; and solve the unforced equation with homogeneous equation to get ϕ_H. Then the sum $\phi = \phi_F + \phi_H$ solves the full equation.

As before, we take the Fourier transform of the forced equation with respect to the spacial variables. As before, we will just do it in the cases with one spacial dimension. We find

$$\partial_t \tilde{\phi}(k, t) + Dk^2 \tilde{\phi}(k, t) = \tilde{F}(k, t),$$

with the initial condition

$$\tilde{\phi}(k, 0) = 0.$$

As before, we have reduced this to a first-order ordinary differential equation in t. Using an integrating factor, we can rewrite this as

$$\frac{\partial}{\partial t}[e^{Dk^2 t} \tilde{\phi}(k, t)] = e^{Dk^2 t} \tilde{F}(k, t).$$

The solution is them

$$\tilde{\phi}(k, t) = e^{-Dk^2 t} \int_0^t e^{Dk^2 u} \tilde{F}(k, u) \, \mathrm{d}u.$$

We define the Green's function $G(x, t; y, \tau))$ to be the solution to

$$[\partial_t - D\nabla_x^2] G(x, t; y, \tau) = \delta(x - y)\delta(t - \tau).$$

So the Fourier transform with respect to x gives

$$\tilde{G}(k, t, y, \tau) = e^{-Dk^2 t} \int_0^t e^{Dk^2 u} e^{iky} \delta(t - \tau) \, \mathrm{d}u,$$

where e^{iky} is just the Fourier transform of $\delta(x - y)$. This is equal to

$$\tilde{G}(k, t; y, \tau) = \begin{cases} 0 & t < \tau \\ e^{-iky} e^{-Dk^2(t-\tau)} & t > \tau \end{cases} = \Theta(t - \tau) e^{-iky} e^{-Dk^2(t-\tau)}.$$

Reverting the Fourier transform, we get

$$G(x, t; y, \tau) = \frac{\Theta(t - \tau)}{2\pi} \int_{\mathbb{R}} e^{ik(x-y)} e^{-Dk^2(t-\tau)} \, \mathrm{d}x$$

This integral is just the inverse Fourier transform of the Gaussian with a phase shift. So we end up with

$$G(x, t; y, \tau) = \frac{\Theta(t - \tau)}{\sqrt{4\pi D(t - \tau)}} \exp\left(-\frac{(x - y)^2}{4D(t - \tau)}\right) = \Theta(t - \tau) S_1(x - y; t - \tau).$$

The solution we seek is then

$$\phi(x,t) = \int_0^t \int_{\mathbb{R}} F(y,\tau)G(x,t;y,\tau) \, \mathrm{d}y \, \mathrm{d}\tau.$$

It is interesting that the solution to the forced equation involves the same function $S_1(x,t)$ as the homogeneous equation with inhomogeneous boundary conditions.

In general, $S_n(\mathbf{x},t)$ solves

$$\frac{\partial S_n}{\partial t} - D\nabla^2 S_n = 0$$

with boundary conditions $S_n(\mathbf{x},0) = \delta^{(n)}(\mathbf{x} - \mathbf{y})$, and we can find

$$S_n(\mathbf{x},t) = \frac{1}{(4\pi Dt)^{n/2}} \exp\left(-\frac{|\mathbf{x} - \mathbf{y}|^2}{4Dt}\right).$$

Then in general, given an initial condition $\phi|_{t=0} = f(\mathbf{x})$, the solution is

$$\phi(\mathbf{x},t) = \int f(\mathbf{y})S(\mathbf{x} - \mathbf{y},t) \, \mathrm{d}^n y.$$

Similarly, $G_n(\mathbf{x},t;\mathbf{y},t)$ solves

$$\frac{\partial G_n}{\partial t} - D\nabla^2 G_n = \delta(t - \tau)\delta^{(n)}(\mathbf{x} - \mathbf{y}),$$

with the boundary condition $G_n(\mathbf{x},0;\mathbf{y},\tau) = 0$. The solution is

$$G(\mathbf{x},t;\mathbf{y},\tau) = \Theta(t - \tau)S_n(\mathbf{x} - \mathbf{y},t - \tau).$$

Given our Green's function, the general solution to the forced heat equation

$$\frac{\partial \phi}{\partial t} - D\nabla^2 \phi = F(\mathbf{x},t), \quad \phi(\mathbf{x},0) = 0$$

is just

$$\phi(\mathbf{x},t) = \int_0^\infty \int_{\mathbb{R}^n} F(\mathbf{y},\tau)G(\mathbf{x},t;\mathbf{y},\tau) \, \mathrm{d}^n y \, \mathrm{d}\tau$$

$$= \int_0^t \int_{\mathbb{R}^n} F(\mathbf{y},\tau)G(\mathbf{x},t;\mathbf{y},\tau) \, \mathrm{d}^n y \, \mathrm{d}\tau.$$

Duhamel noticed that we can write this as

$$\phi(\mathbf{x},t) = \int_0^t \phi_F(\mathbf{x},t;\tau) \, \mathrm{d}\tau,$$

where

$$\phi_F = \int_{\mathbb{R}} F(\mathbf{y},t)S_n(\mathbf{x} - \mathbf{y},t - \tau) \, \mathrm{d}^n y$$

solves the homogeneous heat equation with $\phi_F|_{t=\tau} = F(\mathbf{x},\tau)$.

Hence in general, we can think of the forcing term as providing a whole sequence of "initial conditions" for all $t > 0$. We then integrate over the times at which these conditions were imposed to find the full solution to the forced problem. This interpretation is called *Duhamel's principle*.

Green's functions for the wave equation

Suppose $\phi : \mathbb{R}^n \times [0, \infty) \to \mathbb{C}$ solves the inhomogeneous wave equation

$$\frac{\partial^2 \phi}{\partial t^2} - c^2 \nabla^2 \phi = F(\mathbf{x}, t)$$

with

$$\phi(\mathbf{x}, 0) = \frac{\partial}{\partial t}\phi(\mathbf{x}, 0) = 0.$$

We look for a Green's function $G_n(\mathbf{x}, t; \mathbf{y}, \tau)$ that solves

$$\frac{\partial G_n}{\partial t} - c^2 \nabla^2 G_n = \delta(t - \tau)\delta^{(n)}(\mathbf{x} - \mathbf{y}) \qquad (*)$$

with the same initial conditions

$$G_n(\mathbf{x}, 0, \mathbf{y}, \tau) = \frac{\partial}{\partial t}G_n(\mathbf{x}, 0, \mathbf{y}, \tau) = 0.$$

Just as before, we take the Fourier transform of this equation with respect the spacial variables $\mathbf{x}$. We get

$$\frac{\partial}{\partial t}\tilde{G}_n + c^2 |\mathbf{k}|^2 \tilde{G}_n = \delta(t - \tau)e^{-i\mathbf{k}\cdot\mathbf{y}}.$$

where $\tilde{G}_n = \tilde{G}_n(\mathbf{k}, t, \mathbf{y}, \tau)$.

This is just an ordinary differential equation from the point of view of t, and is of the same type of initial value problem that we studied earlier, and the solution is

$$\tilde{G}_n(\mathbf{k}, t, \mathbf{y}, \tau) = \Theta(t - \tau)e^{-i\mathbf{k}\cdot\mathbf{y}}\frac{\sin |\mathbf{k}|c(t - \tau)}{|\mathbf{k}|c}.$$

To recover the Green's function itself, we have to compute the inverse Fourier transform, and find

$$G_n(\mathbf{x}, t; \mathbf{y}, \tau) = \frac{1}{(2\pi)^n}\int_{\mathbb{R}^n} e^{i\mathbf{k}\cdot\mathbf{x}}\Theta(t - \tau)e^{-i\mathbf{k}\cdot\mathbf{y}}\frac{\sin |\mathbf{k}|c(t - \tau)}{|\mathbf{k}|c} \, \mathrm{d}^n k.$$

Unlike the case of the heat equation, the form of the answer we get here does depend on the number of spatial dimensions n. For definiteness, we look at the case where $n = 3$, since our world (probably) has three dimensions. Then our Green's function is

$$G(\mathbf{x}, t; \mathbf{y}, \tau) = \frac{\Theta(t - \tau)}{(2\pi)^3 c}\int_{\mathbb{R}^3} e^{i\mathbf{k}\cdot(\mathbf{x} - \mathbf{y})}\frac{\sin |\mathbf{k}|c(t - \tau)}{|\mathbf{k}|} \, \mathrm{d}^3 k.$$

We use spherical polar coordinates with the z-axis in k-space aligned along the direction of $\mathbf{x} - \mathbf{y}$. Hence $\mathbf{k} \cdot (\mathbf{x} - \mathbf{y}) = kr\cos\theta$, where $r = |\mathbf{x} - \mathbf{y}|$ and $k = |\mathbf{k}|$.

Note that nothing in our integral depends on φ, so we can pull out a factor of 2π, and get

$$G(\mathbf{x}, t; \mathbf{y}, \tau) = \frac{\Theta(t - \tau)}{(2\pi)^2 c}\int_0^\infty \int_0^\pi e^{ikr\cos\theta}\frac{\sin kc(t - \tau)}{k}k^2 \sin\theta \, \mathrm{d}\theta \, \mathrm{d}k$$

The next integral to do is the θ integral, which is straightforward since it is an exact differential. Setting $\alpha = c(t - \tau)$, we get

$$= \frac{\Theta(t-\tau)}{(2\pi)^2 c} \int_0^\infty \left[\frac{e^{ikr} - e^{-ikr}}{ikr} \right] \frac{\sin kc(t-\tau)}{k} k^2 \, \mathrm{d}k$$

$$= \frac{\Theta(t-\tau)}{(2\pi)^2 icr} \left[\int_0^\infty e^{ikr} \sin k\alpha \, \mathrm{d}k - \int_0^\infty e^{-ikr} \sin k\alpha \, \mathrm{d}k \right]$$

$$= \frac{\Theta(t-\tau)}{2\pi icr} \left[\frac{1}{2\pi} \int_{-\infty}^\infty e^{ikr} \sin k\alpha \, \mathrm{d}k \right]$$

$$= \frac{\Theta(t-\tau)}{2\pi icr} \mathcal{F}^{-1}[\sin k\alpha],$$

Now recall $\mathcal{F}[\delta(x - \alpha)] = e^{-ik\alpha}$. So

$$\mathcal{F}^{-1}[\sin k\alpha] = \mathcal{F}^{-1}\left[\frac{e^{ik\alpha} - e^{-ik\alpha}}{2i} \right] = \frac{1}{2i}[\delta(x + \alpha) - \delta(x - \alpha)].$$

Hence our Green's function is

$$G(\mathbf{x}, t; \mathbf{y}, \tau) = -\frac{\Theta(t-\tau)}{4\pi c|\mathbf{x} - \mathbf{y}|}\left[\delta\big(|\mathbf{x} - \mathbf{y}| + c(t - \tau)\big) - \delta\big(|\mathbf{x} - \mathbf{y}| - c(t - \tau)\big) \right].$$

Now we look at our delta functions. The step function is non-zero only if $t > \tau$. Hence $|\mathbf{x} - \mathbf{y}| + c(t - \tau)$ is always positive. So $\delta(|\mathbf{x} - \mathbf{y}| + c(t - \tau))$ does not contribute. On the other hand, $\delta(|\mathbf{x} - \mathbf{y}| - c(t - \tau))$ is non-zero only if $t > \tau$. So $\Theta(t - \tau)$ is always positive in this region. So we can write our Green's function as

$$G(\mathbf{x}, t; \mathbf{y}, \tau) = \frac{1}{4\pi c}\frac{1}{|\mathbf{x} - \mathbf{y}|}\delta(|\mathbf{x} - \mathbf{y}| - c(t - \tau)).$$

As always, given our Green's function, the general solution to the forced equation

$$\frac{\partial^2 \phi}{\partial t^2} - c^2 \nabla^2 \phi = F(\mathbf{x}, t)$$

is

$$\phi(\mathbf{x}, t) = \int_0^\infty \int_{\mathbb{R}^3} \frac{F(\mathbf{y}, \tau)}{4\pi c|\mathbf{x} - \mathbf{y}|}\delta(|\mathbf{x} - \mathbf{y}| - c(t - \tau)) \, \mathrm{d}^3y \, \mathrm{d}\tau.$$

We can use the delta function to do one of the integrals. It is up to us which integral we do, but we pick the time integral to do. Then we get

$$\phi(\mathbf{x}, t) = \frac{1}{4\pi c^2} \int_{\mathbb{R}^3} \frac{F(\mathbf{y}, t_{\mathrm{ret}})}{|\mathbf{x} - \mathbf{y}|} \, \mathrm{d}^3y,$$

where

$$t_{\mathrm{ret}} = t - \frac{|\mathbf{x} - \mathbf{y}|}{c}.$$

This shows that the effect of the forcing term at some point $\mathbf{y} \in \mathbb{R}^3$ affects the solution ϕ at some other point $\mathbf{x}$ not instantaneously, but only after time $|\mathbf{x} - \mathbf{y}|/c$ has elapsed. This is just as we saw for characteristics. This, again, tells us that information travels at speed c.

Also, we see the effect of the forcing term gets weaker and weaker as we move further away from the source. This dispersion depends on the number of

dimensions of the space. As we spread out in a three-dimensional space, the "energy" from the forcing term has to be distributed over a larger sphere, and the effect diminishes. On the contrary, in one-dimensional space, there is no spreading out to do, and we don't have this reduction. In fact, in one dimensions, we get

$$\phi(x, t) = \int_0^t \int_{\mathbb{R}} F(y, \tau) \frac{\Theta(c(t - \tau) - |x - y|)}{2c} \, \mathrm{d}y \, \mathrm{d}\tau.$$

We see there is now no suppression factor at the bottom, as expected.

7.5   Poisson's equation

Let $\phi : \mathbb{R}^3 \to \mathbb{R}$ satisfy the Poisson's equation

$$\nabla^2 \phi = -F,$$

where $F(\mathbf{x})$ is a forcing term.

The fundamental solution to this equation is defined to be $G_3(\mathbf{x}, \mathbf{y})$, where

$$\nabla^2 G_3(\mathbf{x}, \mathbf{y}) = \delta^{(3)}(\mathbf{x} - \mathbf{y}).$$

By rotational symmetry, $G_3(\mathbf{x}, \mathbf{y}) = G_3(|\mathbf{x} - \mathbf{y}|)$. Integrating over a ball

$$B_r = \{|\mathbf{x} - \mathbf{y}| \le r, \mathbf{x} \in \mathbb{R}^3\},$$

we have

$$
\begin{aligned}
1 &= \int_{B_r} \nabla^2 G_3 \, \mathrm{d}V \\
&= \int_{\partial B_r} \mathbf{n} \cdot \nabla G_3 \, \mathrm{d}S \\
&= \int_{S^2} \frac{\mathrm{d}G_3}{\mathrm{d}r} r^2 \sin \theta \, \mathrm{d}\theta \, \mathrm{d}\phi \\
&= 4\pi r^2 \frac{\mathrm{d}G_3}{\mathrm{d}r}.
\end{aligned}
$$

So we know

$$\frac{\mathrm{d}G_3}{\mathrm{d}r} = \frac{1}{4\pi r^2},$$

and hence

$$G_3(\mathbf{x}, \mathbf{y}) = -\frac{1}{4\pi |\mathbf{x} - \mathbf{y}|} + c.$$

We often set $c = 0$ such that

$$\lim_{|\mathbf{x}| \to \infty} G_3 = 0.$$

Green's identities

To make use of this fundamental solution in solving Poisson's equation, we first obtain some useful identities.

Suppose $\phi, \psi : \mathbb{R}^3 \to \mathbb{R}$ are both smooth everywhere in some region $\Omega \subseteq \mathbb{R}^3$ with boundary $\partial\Omega$. Then

$$\int_{\partial\Omega} \phi \mathbf{n} \cdot \nabla\psi \, \mathrm{d}S = \int_{\Omega} \nabla \cdot (\phi\nabla\psi) \, \mathrm{d}V = \int_{\Omega} \phi\nabla^2\psi + (\nabla\phi) \cdot (\nabla\psi) \, \mathrm{d}V.$$

So we get

Proposition (Green's first identity)**.**

$$\int_{\partial\Omega} \phi \mathbf{n} \cdot \nabla\psi \, \mathrm{d}S = \int_{\Omega} \phi\nabla^2\psi + (\nabla\phi) \cdot (\nabla\psi) \, \mathrm{d}V.$$

Of course, this is just an easy consequence of the divergence theorem, but when Green first came up with this, divergence theorem hasn't existed yet.

Similarly, we obtain

$$\int_{\partial\Omega} \psi \mathbf{n} \cdot \nabla\phi \, \mathrm{d}S = \int_{\Omega} \psi\nabla^2\phi + (\nabla\phi) \cdot (\nabla\psi) \, \mathrm{d}V.$$

Subtracting these two equations gives

Proposition (Green's second identity)**.**

$$\int_{\Omega} \phi\nabla^2\psi - \psi\nabla^2\phi \, \mathrm{d}V = \int_{\partial\Omega} \phi \mathbf{n} \cdot \nabla\psi - \psi \mathbf{n} \cdot \nabla\phi \, \mathrm{d}S.$$

Why is this useful? On the left, we have things like $\nabla^2\psi$ and $\nabla^2\phi$. These are things we are given by Poisson's or Laplace's equation. On the right, we have things on the boundary, and these are often the boundary conditions we are given. So this can be rather helpful when we have to solve these equations.

Using the Green's function

We wish to apply this result to the case $\psi = G_3(|\mathbf{x} - \mathbf{y}|)$. However, recall that when deriving Green's identity, we assumed ϕ and ψ are smooth everywhere in our domain. However, our Green's function is singular at $\mathbf{x} = \mathbf{y}$, but we want to integrate over this region as well. So we need to do this carefully. Because of the singularity, we want to take

$$\Omega = B_r - B_\varepsilon = \{\mathbf{x} \in \mathbb{R}^3 : \varepsilon \leq |\mathbf{x} - \mathbf{y}| \leq R\}.$$

In other words, we remove a small region of radius ε centered on $\mathbf{y}$ from the domain.

In this choice of Ω, it is completely safe to use Green's identity, since our Green's function is certainly regular everywhere in this Ω. First note that since $\nabla^2 G_3 = 0$ everywhere except at $\mathbf{x} = \mathbf{y}$, we get

$$\int_{\Omega} \phi\nabla^2 G_3 - G_3\nabla^2\phi \, \mathrm{d}V = -\int_{\Omega} G_3\nabla^2\phi \, \mathrm{d}V$$

Then Green's second identity gives

$$-\int_{\Omega} G_3\nabla^2\phi \, \mathrm{d}V = \int_{S_r^2} \phi(\mathbf{n} \cdot \nabla G_3) - G_3(\mathbf{n} \cdot \phi) \, \mathrm{d}S$$

$$+ \int_{S_\varepsilon^2} \phi(\mathbf{n} \cdot \nabla G_3) - G_3(\mathbf{n} \cdot \phi) \, \mathrm{d}S$$

Note that on the inner boundary, we have $\mathbf{n} = -\hat{\mathbf{r}}$. Also, at S_ε^2, we have

$$G_3|_{S_\varepsilon^2} = -\frac{1}{4\pi\varepsilon}, \quad \left.\frac{\mathrm{d}G_3}{\mathrm{d}r}\right|_{S_\varepsilon^2} = \frac{1}{4\pi\varepsilon^2}.$$

So the inner boundary terms are

$$\int_{S_\varepsilon^2} \Big(\phi(\mathbf{n} \cdot \nabla G_3) - G_3(\mathbf{n} \cdot \phi) \Big) \varepsilon^2 \sin\theta \, \mathrm{d}\theta \, \mathrm{d}\phi$$
$$= -\frac{\varepsilon^2}{4\pi\varepsilon^2} \int_{S_\varepsilon^2} \phi \sin\theta \, \mathrm{d}\theta \, \mathrm{d}\phi + \frac{\varepsilon^2}{4\pi\varepsilon} \int_{S_\varepsilon^2} (\mathbf{n} \cdot \nabla\phi) \sin\theta \, \mathrm{d}\theta \, \mathrm{d}\phi$$

Now the final integral is bounded by the assumption that ϕ is everywhere smooth. So as we take the limit $\varepsilon \to 0$, the final term vanishes. In the first term, the ε's cancel. So we are left with

$$= -\frac{1}{4\pi} \int_{S_\varepsilon^2} \phi \, \mathrm{d}\Omega$$
$$= -\bar{\phi}$$
$$\to -\phi(\mathbf{y})$$

where $\bar{\phi}$ is the average value of ϕ on the sphere.

Now suppose $\nabla^2\phi = -F$. Then this gives

Proposition (Green's third identity)**.**

$$\phi(\mathbf{y}) = \int_{\partial\Omega} \phi(\mathbf{n} \cdot \nabla G_3) - G_3(\mathbf{n} \cdot \nabla\phi) \, \mathrm{d}S - \int_\Omega G_3(\mathbf{x}, \mathbf{y})F(\mathbf{x}) \, \mathrm{d}^3x.$$

This expresses ϕ at any point $\mathbf{y}$ in Ω in terms of the fundamental solution G_3, the forcing term F and boundary data. In particular, if the boundary values of ϕ and $\mathbf{n} \cdot \nabla\phi$ vanish as we take $r \to \infty$, then we have

$$\phi(\mathbf{y}) = -\int_{\mathbb{R}^3} G_3(\mathbf{x}, \mathbf{y})F(\mathbf{x}) \, \mathrm{d}^3x.$$

So the fundamental solution *is* the Green's function for Poisson's equation on $\mathbb{R}^3$.

However, there is a puzzle. Suppose $F = 0$. So $\nabla^2\phi = 0$. Then Green's identity says

$$\phi(\mathbf{y}) = \int_{S_r^2} \left(\phi\frac{\mathrm{d}G_3}{\mathrm{d}r} - G_3\frac{\mathrm{d}\phi}{\mathrm{d}r} \right) \, \mathrm{d}S.$$

But we know there is a unique solution to Laplace's equation on every bounded domain once we specify the boundary value $\phi|_{\partial\Omega}$, or a unique-up-to-constant solution if we specify the boundary value of $\mathbf{n} \cdot \nabla\phi|_{\partial\Omega}$.

However, to get $\phi(\mathbf{y})$ using Green's identity, we need to know ϕ and *and* $\mathbf{n} \cdot \nabla\phi$ on the boundary. This is too much.

Green's third identity *is* a valid relation obeyed by solutions to Poisson's equation, but it is not constructive. We cannot specify ϕ *and* $\mathbf{n} \cdot \nabla\phi$ freely. What we would like is a formula of ϕ given, say, just the value of ϕ on the boundary.

Dirichlet Green's function

To overcome this (in the Dirichlet case), we seek to modify G_3 via

$$G_3 \to G = G_3 + H(\mathbf{x}, \mathbf{y}),$$

where $\nabla^2 H = 0$ everywhere in Ω, H is regular throughout Ω, and $G|_{\partial\Omega} = 0$. In other words, we find some H that does not affect the relations on G when acted on by ∇^2, but now our G will have boundary value 0. We will find this H later, but given such an H, we replace G_3 with $G - H$ in Green's third identity, and see that all the H terms fall out, i.e. G also satisfies Green's third identity. So

$$\phi(\mathbf{y}) = \int_{\partial\Omega} [\phi\mathbf{n} \cdot \nabla G - G\mathbf{n} \cdot \nabla\phi] \; \mathrm{d}S - \int FG \; \mathrm{d}V$$

$$= \int_{\partial\Omega} \phi\mathbf{n} \cdot \nabla G \; \mathrm{d}S - \int FG \; \mathrm{d}V.$$

So as long as we find H, we can express the value of $\phi(\mathbf{y})$ in terms of the values of ϕ on the boundary. Similarly, if we are given a Neumann condition, i.e. the value of $\mathbf{n} \cdot \nabla\phi$ on the boundary, we have to find an H that kills off $\mathbf{n} \cdot \nabla G$ on the boundary, and get a similar result.

In general, finding a harmonic $\nabla^2 H = 0$ with

$$H|_{\partial\Omega} = \left. \frac{1}{4\pi|\mathbf{x} - \mathbf{x}_0|} \right|_{\partial\Omega}$$

is a difficult problem. However, the *method of images* allows us to solve this in some special cases with lots of symmetry.

Example. Suppose
$$\Omega = \{(x, y, z) \in \mathbb{R}^3 : z \geq 0\}.$$
We wish to find a solution to $\nabla^2 = -F$ in Ω with $\phi \to 0$ rapidly as $|\mathbf{x}| \to \infty$ with boundary condition $\phi(x, y, 0) = g(x, y)$.

The fundamental solution

$$G_3(\mathbf{x}, \mathbf{x}_0) = -\frac{1}{4\pi} \frac{1}{|\mathbf{x} - \mathbf{x}_0|}$$

obeys all the conditions we need except

$$G_3|_{z=0} = -\frac{1}{4\pi} \frac{1}{[(x - x_0)^2 + (y - y_0)^2 + z_0^2]^{1/2}} \neq 0.$$

However, let $\mathbf{x}_0^R$ be the point $(x_0, y_0, -z_0)$. This is the reflection of x_0 in the boundary plane $z = 0$. Since the point $\mathbf{x}_0^R$ is outside our domain, $G_3(\mathbf{x}, \mathbf{x}_0^R)$ obeys

$$\nabla^2 G_3(\mathbf{x}, \mathbf{x}_0^R) = 0$$

for all $\mathbf{x} \in \Omega$, and also

$$G_3(\mathbf{x}, \mathbf{x}_0^R)|_{z=0} = G_3(\mathbf{x}, \mathbf{x}_0)|_{z=0}.$$

Hence we take

$$G(\mathbf{x}, \mathbf{x}_0) = G_3(\mathbf{x}, \mathbf{x}_0) - G_3(\mathbf{x}, \mathbf{x}_0^R).$$

The outward pointing normal to Ω at $z = 0$ is $\mathbf{n} = -\hat{\mathbf{z}}$. Hence we have

$$\mathbf{n} \cdot \nabla G|_{z=0} = \frac{1}{4\pi} \left[\frac{-(z - z_0)}{|\mathbf{x} - \mathbf{x}_0|^3} - \frac{-(z + z_0)}{|\mathbf{x} - \mathbf{x}_0^R|^3} \right]_{z=0}$$

$$= \frac{1}{2\pi} \frac{z_0}{[(x - x_0)^2 + (y - y_0)^2 + z_0^2]^{3/2}}.$$

Therefore our solution is

$$\phi(\mathbf{x}_0) = \frac{1}{4\pi} \int_\Omega \left[\frac{1}{|\mathbf{x} - \mathbf{x}_0|} - \frac{1}{|\mathbf{x} - \mathbf{x}_0^R|} \right] F(\mathbf{x}) \, \mathrm{d}^3 x$$

$$+ \frac{z_0}{2\pi} \int_{\mathbb{R}^2} \frac{g(x, y)}{[(x - x_0)^2 + (y - y_0)^2 + z_0^2]^{3/2}} \, \mathrm{d}x \, \mathrm{d}y.$$

What have we actually done here? The Green's function $G_3(\mathbf{x}, \mathbf{x}_0)$ in some sense represents a "charge" at $\mathbf{x}_0$. We can imagine that the term $G_3(\mathbf{x}, \mathbf{x}_0^R)$ represents the contribution to our solution from a point charge of opposite sign located at $\mathbf{x}_0^R$. Then by symmetry, G is zero at $z = 0$.

$$\overset{\displaystyle \bullet}{\mathbf{x}_0}$$

$$\underline{\hspace{5cm}} \; z = 0$$

$$\overset{\displaystyle \bullet}{}\\\mathbf{x}_0^R$$
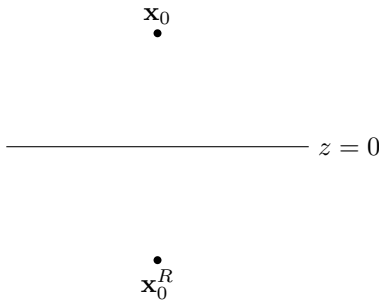
Our solution for $\phi(\mathbf{x}_0)$ is valid if we extend it to $\mathbf{x}_0 \in \mathbb{R}^3$ where $\phi(\mathbf{x}_0)$ is solved by choosing two mirror charge distributions. But this is irrelevant for our present purposes. We are only concerned with finding a solution in the region $z \geq 0$. This is just an artificial device we have in order to solve the problem.

Example. Suppose a chimney produces smoke such that the density $\phi(\mathbf{x}, t)$ of smoke obeys

$$\partial_t \phi - D\nabla^2 \phi = F(\mathbf{x}, t).$$

The left side is just the heat equation, modelling the diffusion of smoke, while the right forcing term describes the production of smoke by the chimney.

If this were a problem for $\mathbf{x} \in \mathbb{R}^3$, then the solution is

$$\phi(\mathbf{x}, t) = \int_0^t \int_{\mathbb{R}^3} F(\mathbf{y}, \tau) S_3(\mathbf{x} - \mathbf{y}, t - \tau) \, \mathrm{d}^3 \mathbf{y} \, \mathrm{d}\tau,$$

where

$$S_3(\mathbf{x} - \mathbf{y}, t - \tau) = \frac{1}{[4\pi D(t - \tau)]^{3/2}} \exp \left(-\frac{|\mathbf{x} - \mathbf{y}|^2}{4D(t - \tau)} \right).$$

This is true only if the smoke can diffuse in all of $\mathbb{R}^3$. However, this is not true for our current circumstances, since smoke does not diffuse into the ground.

To account for this, we should find a Green's function that obeys

$$\mathbf{n} \cdot \nabla G|_{z=0} = 0.$$

This says that there is no smoke diffusing in to the ground.

This is achieved by picking

$$G(\mathbf{x}, t; \mathbf{y}, \tau) = \Theta(t - \tau)[S_3(\mathbf{x} - \mathbf{y}, t - \tau) + S_3(\mathbf{x} - \mathbf{y}^R, t - \tau)].$$

We can directly check that this obeys

$$\partial_t D^2 \nabla^2 G = \delta(t - \tau)\delta^3(\mathbf{x} - \mathbf{y})$$

when $\mathbf{x} \in \Omega$, and also

$$\mathbf{n} \cdot \nabla G|_{z_0} = 0.$$

Hence the smoke density is given by

$$\phi(\mathbf{x}, t) = \int_0^t \int_\Omega F(\mathbf{y}, \tau)[S_3(\mathbf{x} - \mathbf{y}, t - \tau) + S_3(\mathbf{x} - \mathbf{y}^R, t - \tau)] \, d^3y.$$

We can think of the second term as the contribution from a "mirror chimney". Without a mirror chimney, we will have smoke flowing into the ground. With a mirror chimney, we will have equal amounts of mirror smoke flowing up from the ground, so there is no net flow. Of course, there are no mirror chimneys in reality. These are just artifacts we use to find the solution we want.
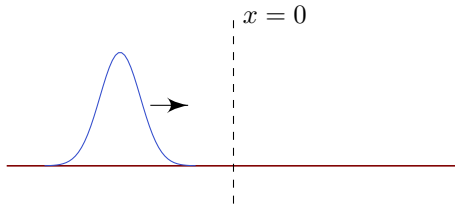
Example. Suppose we want to solve the wave equation in the region (x, t) such that $x > 0$ with boundary conditions

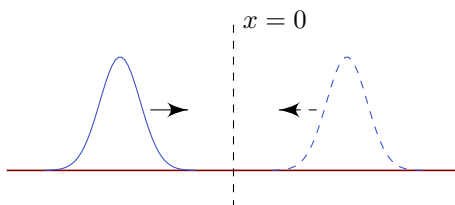$$\phi(x, 0) = b(x), \quad \partial_t(x, 0) = 0, \quad \partial_x \phi(0, t) = 0.$$

On $\mathbb{R}^{1,1}$ d'Alembert's solution gives

$$\phi(x, t) = \frac{1}{2}[b(x - ct) + b(x + ct)]$$

This is not what we want, since eventually we will have a wave moving past the $x = 0$ line.

To compensate for this, we introduce a mirror wave moving in the opposite direction, such that when as they pass through each other at $x = 0$, there is no net flow across the boundary.

More precisely, we include a *mirror initial condition* $\phi(x, 0) = b(x) + b(-x)$, where we set $b(x) = 0$ when $x < 0$. In the region $x > 0$ we are interested in, only the $b(x)$ term will contribute. In the $x < 0$ region, only $x > 0$ will contribute. Then the general solution is

$$\phi(x, t) = \frac{1}{2}[b(x + ct) + b(x + c) + b(-x - ct) + b(-x + ct)].$$