Math 18500 Week 7: Partial Differential Equations and Fourier Transforms

Partial Differential Equations. Up to this point we have only considered *ordinary* differential equations. By definition, an ordinary differential equation is a relationship between an unknown function $y = y(x)$ and its derivatives:

$$F(x, y, y', \dots, y^{(n-1)}, y^{(n)}) = 0$$

Note that in this context the unknown function is always a function of a *single* variable.

Sometimes it is necessary to solve for an unknown function of more than one variable,

$$y = y(x_1, x_2, \dots, x_n),$$

given a relationship between that function and its *partial* derivatives,

$$F\left(x_i, \frac{\partial y}{\partial x_i}, \frac{\partial^2 y}{\partial x_i \partial x_j}, \dots\right) = 0$$

An equation like this is called a *partial* differential equation.

If you only know about three partial differential equations, the ones you should know are:

(1) The *heat* or *diffusion* equation:

$$\frac{\partial u}{\partial t} - \alpha^2 \frac{\partial^2 u}{\partial x^2} = \rho(x, t)$$

(2) The *wave* equation:

$$\frac{\partial^2 u}{\partial t^2} - \frac{1}{c^2} \frac{\partial^2 u}{\partial x^2} = \rho(x, t)$$

(3) The *Poisson* equation:

$$\frac{\partial^2 \phi}{\partial x^2} + \frac{\partial^2 \phi}{\partial y^2} + \frac{\partial^2 \phi}{\partial z^2} = -\rho(x, y, z)$$

In the case of both the heat and the wave equation, we are trying to solve for an unknown function $u = u(x, t)$ whose variables are x and t. You can think of x as representing position and t as representing time. Each of these equations can be thought of as modelling a function of position, which is changing over time.

For example, the heat equation can be used to model the *temperature* of a metal rod. Imagine putting one end of the rod into a fire. The end which is in the fire will have a higher temperature than the end you are holding, and there will be a range of temperatures inbetween. That is, the temperature is a function of position within the rod.

Now suppose you remove the rod from the fire. In this case, the heat which is concentrated at one end of the rod will *diffuse* throughout the rod, and eventually the entire rod will have an even temperature. This shows that the temperature is also a function of time, so we need a partial differential equation to model it.

In the heat equation, the scalar α^2 is a constant, which represents the rate of heat diffusion. The function $\rho(x, t)$ represents an external source of heat (such as the fire) which is transferring heat to the rod.

The wave equation can be used to model the propagation of waves. You can think of the function $u(x, t)$ as representing the height of the wave. The constant c represents the speed of propagation - it is measured in units of distance over time. The function $\rho(x, t)$ represents a source which is generating the waves.

In the Poisson equation all of the variables should be interpreted spatially - the unknown function $u(x, y, z)$ is a function of position only. One possible interpretation is that $u(x, y, z)$ represents the *electric potential* due to a collection of charges - with this interpretation, the function $\rho(x, y, z)$ on the right hand side of the equation can be thought of as the *charge density* of a static charge distribution which generates the electric field.

The operator

$$\vec{\nabla}^2 = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} + \frac{\partial^2}{\partial z^2}$$

which appears on the left hand side of the Poisson equation is called the Laplace operator, and it appears in many other partial differential equations. For example, the heat and wave equations both have 3-dimensional versions, which involve a Laplace operator instead of a second derivative:

$$\frac{\partial u}{\partial t} - \alpha^2 \vec{\nabla}^2 u = \rho(x, y, z, t)$$

$$\frac{\partial^2 u}{\partial t^2} - c^2 \vec{\nabla}^2 u = \rho(x, y, z, t)$$

From this point of view, solutions of the Laplace equation represent *steady state solutions* of the heat or wave equation, i.e. solutions which satisfy $\frac{\partial u}{\partial t} = 0$ and therefore do not change over time.

Just as we need to impose initial conditions in order to uniquely specify a solution of an ordinary differential equations, we typically need to impose *boundary conditions* in order to uniquely specify a solution of a partial differential equation.

For example, when we use the heat equation to model the temperature distribution in an insulated metal rod, one possible choice of boundary conditions is

$$\frac{\partial u}{\partial x}(0, t) = 0 \ , \quad \frac{\partial u}{\partial x}(L, t) = 0 \ , \quad u(x, 0) = f(x)$$

Here L is the length of the rod - we have placed one end at $x = 0$ and the other end at $x = L$. To understand why these are called boundary conditions, you have to think about the two-dimensional region on which the function $u(x, t)$ is defined. We are only considering values of x such that $0 \leq x \leq L$, corresponding to locations in the rod, and values t such that $t \geq 0$, i.e. all times after some initial time. Overall, these inequalities define an infinite vertical strip in the xt plane:

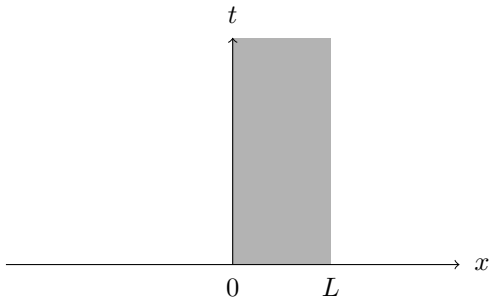

The boundary conditions give us various kinds of information about the function's behavior on the three parts of the boundary of this region (the two vertical lines, and the horizontal segment).

The first two conditions relate to the vertical lines. Intuitively, they tell us that the ends of the rod are insulated and do not conduct any heat.[1] Boundary conditions of this type (which specify the behavior of derivatives of u at a point on the boundary of the region) are called *Neumann* conditions.

The third condition is more analogous to the usual concept of specifying the initial value of a solution to a first order ordinary differential equation - it means that we are specifying an initial temperature distribution within the rod at time $t = 0$. Boundary conditions of this type (which specify the value of the function at a point on the boundary of the region) are usually called *Dirichlet* conditions.

---

[1] More specifically, the total heat within the rod at time t can be defined by

$$Q = \int_0^L u(x, t) dx$$

The heat equation tells us that

$$\frac{dQ}{dt} = \int_0^L \frac{\partial u}{\partial t}(x, t) dx = \alpha^2 \int_0^L \frac{\partial^2 u}{\partial x^2}(x, t) dx = \alpha^2 \frac{\partial u}{\partial x}(L, t) - \alpha^2 \frac{\partial u}{\partial x}(0, t)$$

We can interpret the two terms on the right hand side as the rates at which heat is being lost from the two ends of the rod - the Neumann conditions say that heat is not being lost from either end.

Typically, when you solve partial differential equations, you need to apply some combination of these two types of boundary conditions. Exactly how this is done will depend not only on the type of equation you are trying to solve, but also the specific application you have in mind.

In the context of the heat equation (or any other equation which involves a time variable), conditions satisfied at time $t = 0$ are usually referred to as *initial* conditions rather than *boundary* conditions. So, we would usually say that the boundary value problem

$$\frac{\partial u}{\partial t} - \alpha^2 \frac{\partial^2 u}{\partial x^2} = 0$$

$$\frac{\partial u}{\partial x}(0,t) = 0 \ , \quad \frac{\partial u}{\partial x}(L,t) = 0 \ , \quad u(x,0) = f(x)$$

has Neumann boundary conditions and an initial condition.

The Heat Equation on an Interval. The reason why we are discussing partial differential equations now is that most of the important applications of Fourier series involve partial differential equations. Indeed, when Fourier invented Fourier series, his goal was to solve the heat equation!

To explain what Fourier was doing, consider the heat equation without a source term,

$$\frac{\partial u}{\partial t} - \alpha^2 \frac{\partial^2 u}{\partial x^2} = 0$$

and the Dirichlet boundary conditions

$$u(0,t) = 0 \ , \quad u(L,t) = 0$$
$$u(x,0) = f(x)$$

Solving the heat equation with these boundary conditions can be interpreted as finding the temperature distribution in an insulated metal rod as a function of time, where the ends of the rod are maintained at room temperature (the arbitrarily chosen 0 of our temperature scale).

The process which Fourier used to solve this equation involved three steps:

(1) Find all nonzero *separated* solutions.
(2) Among the separated solutions, determine which ones satisfy the boundary conditions,

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

but not necessarily the initial condition

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

(3) *Superimpose* the separated solutions to obtain a solution which *does* satisfy the initial condition.

To explain the final step, notice that the operator on the left hand side of the heat equation,

$$\frac{\partial}{\partial t} - \alpha^2 \frac{\partial^2}{\partial x^2}$$

is a *linear* operator, the overall equation

$$\frac{\partial u}{\partial t} - \alpha^2 \frac{\partial^2 u}{\partial x^2} = 0$$

is homogeneous, and the boundary conditions

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

are *also* homogeneous (the right hand side is 0). Therefore, given solutions $u_n(x,t)$ which satisfy the boundary conditions,

$$\frac{\partial u_n}{\partial t} - \alpha^2 \frac{\partial^2 u_n}{\partial x^2} = 0$$
$$u_n(0,t) = 0 \ , \quad u_n(L,t) = 0,$$

any linear combination

$$u(x,t) = \sum_n c_n u_n(x,t)$$

will also be a solution satisfying the boundary conditions:

$$\frac{\partial u}{\partial t} - \alpha^2 \frac{\partial^2 u}{\partial x^2} = \sum_n c_n \left(\frac{\partial u_n}{\partial t} - \alpha^2 \frac{\partial^2 u_n}{\partial x^2} \right) = \sum_n 0 = 0$$
$$u(0,t) = \sum_n c_n u_n(0,t) = \sum_n 0 = 0$$
$$u(L,t) = \sum_n c_n u_n(L,t) = \sum_n 0 = 0$$

So, we can hope to construct a solution satisfying the initial conditions by finding enough particular solutions u_n through guesswork. Specifically, we must find enough particular solutions that an arbitrary function $f(x)$ can be expressed as a linear combination of their initial values, i.e. there are values of c_n such that

$$\sum_n c_n u_n(x,0) = f(x)$$

Following Fourier's first step, we look for particular solutions which are *separated*. By definition, a separated solution is an ansatz of the special form

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

where T is a function of t and X is a function of x.[2] You can see why this kind of ansatz is a good idea if you substitute it into the heat equation:

$$\frac{\partial}{\partial t}\left[T(t)X(x)\right] = \alpha^2 \frac{\partial^2}{\partial x^2}\left[T(t)X(x)\right]$$
$$T'(t)X(x) = \alpha^2 T(t)X''(x)$$

Assuming that u is nonzero, we know that both T and X must be nonzero, so we can divide by them:

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

Notice a strange property of this equation: on the left side we have a function which is independent of x, and on the right side we have a function which is independent of t.

So, in order for the two sides of the equation to be equal, they must *both* be independent of *both* x and t! This implies that both sides are equal to a *constant*, which is conventionally denote by λ:

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

Therefore, to find the separated solutions we must solve the separate equations

$$X''(x) = \lambda X(x)$$
$$T'(t) = \alpha^2 \lambda T(t)$$

At this point we know the solutions of these equations by heart:

$$X = ae^{\sqrt{\lambda}x} + be^{-\sqrt{\lambda}x}$$
$$T = ce^{\alpha^2 \lambda t}$$

At least, these are the solutions if $\lambda \neq 0$. In the case $\lambda = 0$, the equations become

$$X''(x) = \lambda X(x) = 0$$
$$T'(t) = \alpha^2 \lambda T(t) = 0$$

and therefore

$$X = a + bx$$
$$T = c$$

Multiplying these together, we obtain

$$u(x,t) = T(t)X(x) = ac + bcx = A + Bx$$

Note that we have absorbed the constants a and c into a single constant A, and we have absorbed the constants b and c into a single constant B.

For the case $\lambda \neq 0$, the square root in the formula for X is a bit annoying, so we introduce the change of variables

$$\lambda = -\mu^2$$

With this notation, we have

$$X = ae^{i\mu x} + be^{-i\mu x}$$
$$T = ce^{-\alpha^2 \mu^2 t}$$

Multiplying these together, we obtain

$$u(x,t) = T(t)X(x) = e^{-\alpha^2 \mu^2 t}\left(Ae^{i\mu x} + Be^{-i\mu x}\right)$$

where we have again combined some arbitrary constants.

[2] Admittedly, it is a bit confusing to use the same letter for the function and the variable! Unfortunately, this notation is standard for applications of the separation of variables technique, so you have to get used to it.

Having completed the first step and found all of the nonzero separated solutions, we now impose our desired boundary conditions,

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

In the case $\mu = 0$ the only way these conditions can be satisfied is if

$$A + Bx|_{x=0} = A + B \cdot 0 = A = 0$$
$$A + Bx|_{x=1} = A + B \cdot 1 = A + B = 0$$

These equations imply that $A = B = 0$, hence $u = 0$. So, if $\mu = 0$ there is no non-zero separated solution which satisfies the boundary conditions.

In the case $\mu \neq 0$ (or equivalently $\lambda \neq 0$), the only way the boundary conditions can be satisfied is if

$$e^{-\alpha^2\mu^2 t}\left(Ae^{i\mu x} + Be^{-i\mu x}\right)\Big|_{x=0} = (A+B)e^{-\alpha^2\mu^2 t} = 0$$

and

$$e^{-\alpha^2\mu^2 t}\left(Ae^{i\mu x} + Be^{-i\mu x}\right)\Big|_{x=0} = (Ae^{i\mu L} + Be^{-i\mu L})e^{-\alpha^2\mu^2 t} = 0$$

for all values of t. Dividing by the exponential term (which is always nonzero), these equations simplify to

$$A + B = 0$$
$$Ae^{i\mu L} + Be^{-i\mu L} = 0.$$

Replacing B with $-A$ in the second equation, and then dividing by A, we obtain

$$e^{i\mu L} - e^{-i\mu L} = 0$$
$$e^{i\mu L} = e^{-i\mu L}$$
$$e^{2i\mu L} = 1$$

The only way the last identity can be satisfied is if

$$2\mu L = 2\pi n,$$

for some integer n. Rearranging, we have

$$\mu = \frac{n\pi}{L}$$

and therefore the most general separated solution which satisfies the boundary conditions is

$$u_n(x,t) = b_n e^{-\alpha^2\left(\frac{n\pi}{L}\right)^2 t}\left(\frac{e^{\frac{\pi i n x}{L}} - e^{-\frac{\pi i n x}{L}}}{2i}\right) = b_n e^{\frac{-\alpha^2\pi^2 n^2 t}{L^2}}\sin\left(\frac{\pi n x}{L}\right),$$

where $b_n = 2iA$ is an arbitrary constant.

Superimposing the separated solutions, we obtain

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

Notice that we have not included $u_0(x,t)$, because it is equal to 0, and would not contribute to the sum. Also, we have not included terms $u_n(x,t)$ with $n < 0$, because $u_n(x,t)$ and $u_{-n}(x,t)$ are scalar multiples of each other, and can be combined into a single term.

Substituting $t = 0$, imposing the initial condition $u(x,0) = f(x)$, we find that

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

As we have seen, the values of b_n are uniquely determined by this equation! In particular, we know that they can be found by evaluating the following integrals:

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

Evaluating these integrals and substituting their values back into the formula for $u(x,t)$ results in a solution that satisfies the initial condition $u(x,0) = f(x)$. But notice that it is only logically valid to draw this conclusion *if we make the assumption that it is possible to write $f(x)$ as a sine series in the first place!*

At this point, Fourier went ahead and boldly asserted that *every* solution of the heat equation could be written as a sum of separated solutions. This assumption *implied* that any function could be written as a sine series, because presumably nature knows how to solve the heat equation with arbitrary initial conditions!

To be clear, there was a certain amount of hubris in Fourier's assertion - just because you happen to have guessed a whole bunch of solutions of an equation, it doesn't mean you've found *all possible* solutions of the equation! Because of this logical leap, Fourier's assertion was rejected by many mathematicians of his time.

Later, mathematicians resolved the controversy in Fourier's favor, and gave a rigorous proof that every periodic function has a Fourier series. But it is interesting to think that this remarkable fact might never have been discovered had there not been a concrete physical problem whose solution required it!

Fourier Transforms. We have seen that any 2π-periodic function $f(t)$ can be written as a Fourier series,

$$f(t) = \sum_{n=-\infty}^{\infty} c_n e^{int}$$

where the coefficients are given by

$$c_n = \frac{1}{2\pi} \int_{-\pi}^{\pi} f(t) e^{-int} dx$$

This allows us to think of any periodic function in the "frequency domain", i.e. as a superposition of sinusoids with various amplitudes and frequencies.

Unfortunately, Fourier series are not sufficient for many applications, for two reasons:

(1) It is often necessary to consider functions which are *not* periodic.
(2) It is often necessary to consider superpositions of sinusoids with *arbitrary* frequencies (i.e. a set of frequencies which are not integer multiples of a single base frequency).

These reasons are related - if we superimpose two sinusoids whose frequencies are *not* integer multiples of a common base frequency, then the result is *never* a periodic function. For example, try graphing the function

$$f(t) = \sin(t) + \sin(\pi t)$$

on your computer - you will see that it never repeats itself!

Fortunately there is a generalization of Fourier series which is valid even for non-periodic functions. Given a function $f(t)$, we define its **Fourier transform** to be the function

$$\hat{f}(k) = \int_{-\infty}^{\infty} f(t) e^{-ikt} dt.$$

Notice that $\hat{f}$ is a function of a new variable k, which takes arbitrary real values (not just integers).

Also notice that the Fourier transform is *linear*. This means that if $f(t)$ is a linear combination of functions,

$$f(t) = c_1 f_1(t) + c_2 f_2(t)$$

then its Fourier transform is the corresponding linear combination of the Fourier transforms,

$$\hat{f}(t) = c_1 \hat{f}_1(t) + c_2 \hat{f}_2(t).$$

When we want to emphasize this property of the Fourier transform, we will use the following notation:

$$\mathcal{F}[f(t)] = \hat{f}(k)$$

With this notation, the linearity of the Fourier transform can be expressed as follows:

$$\mathcal{F}[c_1 f_1(t) + c_2 f_2(t)] = c_1 \mathcal{F}[f_1(t)] + c_2 \mathcal{F}[f_2(t)]$$

Fourier transforms are useful because of the **Fourier inversion formula**,

$$f(t) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \hat{f}(k) e^{ikt} dk,$$

which is valid as long as $f(t)$ satisfies certain technical conditions (to be described later).

To express this result in operator notation, we can define the *inverse Fourier transform*,

$$\mathcal{F}^{-1}[g(k)] = \frac{1}{2\pi} \int_{-\infty}^{\infty} g(k) e^{ikt} dk$$

Then the Fourier inversion formula can be stated as follows:

$$f(t) = \mathcal{F}^{-1}[\mathcal{F}[f(t)]]$$

Similarly, we also have

$$g(k) = \mathcal{F}\left[\mathcal{F}^{-1}[g(k)]\right]$$

Applying the Fourier inversion formula is analogous to writing a periodic function as a Fourier series

$$f(t) = \sum_{n=-\infty}^{\infty} c_n e^{int}$$

but instead of *summing* over all *integer* frequencies, we are *integrating* over all *real* frequencies.

To better understand the inversion formula, suppose that $f(t)$ is zero except within an interval of finite size:

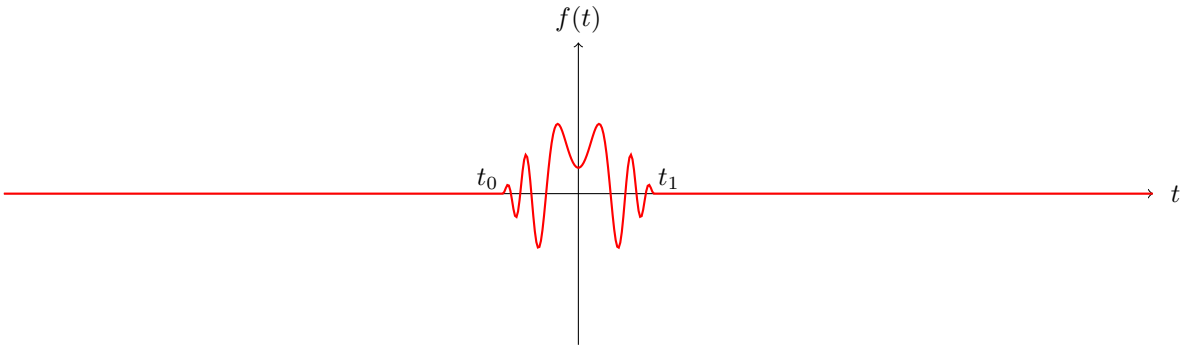

For example, $f(t)$ could represent a radio signal starting at time t_0 and ending at time t_1.

In a situation like this, we can always pick an arbitrarily large number T such that

$$-\frac{T}{2} < t_0 < t_1 < \frac{T}{2}$$

and then consider the T-periodic function $\tilde{f}(t)$ which is equal to $f(t)$ for $-\frac{T}{2} < t < \frac{T}{2}$:

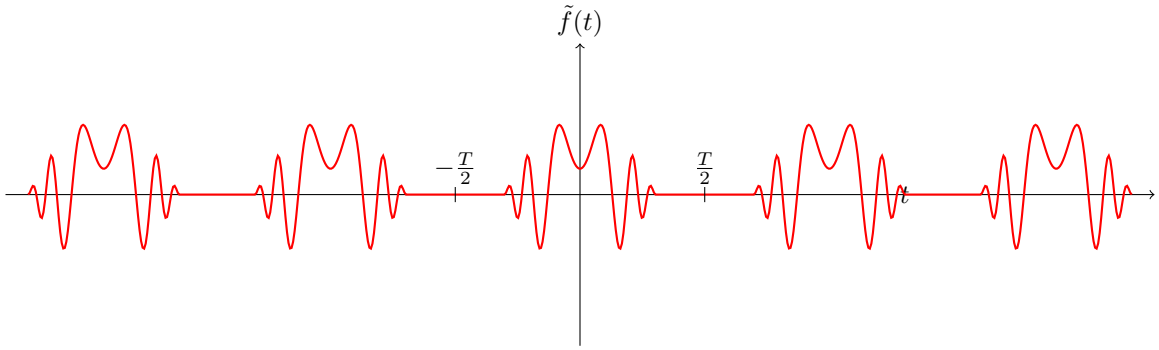

Then the Fourier series of $\tilde{f}(t)$,

$$\tilde{f}(t) = \sum_{n=-\infty}^{\infty} c_n e^{\frac{2\pi i n}{T}},$$

is equal to $f(t)$ on the interval $[-\frac{T}{2}, \frac{T}{2}]$, and its coefficients are given by

$$c_n = \frac{1}{T} \int_{-\frac{T}{2}}^{\frac{T}{2}} f(t) e^{-\frac{2\pi i n t}{L}} dt.$$

As long as the original function $f(t)$ is equal to zero outside of the interval $[-T, T]$, we can rewrite this as

$$c_n = \frac{1}{T} \int_{-\infty}^{\infty} f(t) e^{-\frac{2\pi i n t}{L}} dt.$$

At this point, we can define

$$k_n = \frac{2\pi n}{T}$$

to be the frequency of the n^{th} term in the sum. Then the spacing between consecutive frequencies,

$$\Delta k = k_{n+1} - k_n = \frac{2\pi(n+1)}{T} - \frac{2\pi n}{T} = \frac{2\pi}{T},$$

gets extremely small as $T \to \infty$, and we can rewrite the formula for c_n as

$$c_n = \frac{\Delta k}{2\pi} \int_{-\infty}^{\infty} f(t) e^{-ik_n t} dt = \frac{1}{2\pi} \Delta k \hat{f}(k_n)$$

With this notation, we have

$$f(t) = \frac{1}{2\pi} \sum_{n=-\infty}^{\infty} \hat{f}(k_n) e^{ik_n t} \Delta k,$$

and in the limit as $T \to \infty$ the sum converges to an *integral*,

$$f(t) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \hat{f}(k) e^{-ikt} dk.$$

This proves that the Fourier inversion formula is valid for functions which are zero outside a finite interval.

Some care is required to extend this idea to functions which are nonzero for an *infinite* range of values of t, but it is still correct in this case, as long as

$$f(t) \to 0$$

sufficiently rapidly as $t \to \pm\infty$. More precisely, it is valid to use the Fourier inversion formula as long as

$$\int_{-\infty}^{\infty} |f(t)|^2 dt < \infty.$$

For functions like

$$f(t) = \frac{1}{\sqrt{1+|t|}}$$

which do not satisfy this condition, it is invalid to apply the Fourier inversion formula.

As an application of Fourier inversion, consider the function

$$f(t) = e^{-|t|} = \begin{cases} e^{-t} & t > 0 \\ e^{t} & t < 0 \end{cases}$$

Its Fourier transform is given by

$$\begin{aligned} \hat{f}(k) &= \int_{-\infty}^{0} e^{t} e^{-ikt} dt + \frac{1}{2\pi} \int_{0}^{\infty} e^{-t} e^{-ikt} dt \\ &= \left(\frac{e^{(1-ik)t}}{1-ik} \Big|_{-\infty}^{0} + \frac{e^{(-1-ik)t}}{-1-ik} \Big|_{0}^{\infty} \right) \\ &= \left(\frac{1}{1-ik} - \frac{1}{-1-ik} \right) \\ &= \left(\frac{2}{1+k^2} \right) \end{aligned}$$

Therefore, the Fourier inversion formula tells us that

$$f(t) = e^{-|t|} = \frac{1}{\pi} \int_{-\infty}^{\infty} \frac{e^{ikt}}{1+k^2} dk$$

Notice that the imaginary part of the integral on the right cancels out, because

$$\sin(kt)$$

is an odd function of k. Therefore, we could also write the above identity as

$$e^{-|t|} = \frac{1}{\pi} \int_{-\infty}^{\infty} \frac{\cos(kt)}{1+k^2} dk$$

The integral on the right is extremely difficult to evaluate directly, except in the special case $t = 0$:

$$\frac{1}{\pi} \int_{-\infty}^{\infty} \frac{1}{1+k^2} dk = \frac{1}{\pi} \tan^{-1}(k) \Big|_{-\infty}^{\infty} = \frac{\pi}{2\pi} - \left(-\frac{\pi}{2\pi} \right) = 1$$

In fact, it is usually rather difficult to calculate Fourier transforms and inverse Fourier transforms directly. Fortunately, the Fourier transform has a number of useful properties which can be applied to make this a bit easier. Here is one such property (arguably the most important one):

$(*)$ If $g(t) = \frac{d}{dt}[f(t)]$, then $\hat{g}(k) = ik\hat{f}(k)$.

This property can be derived using integration by parts:

$$\hat{g}(k) = \int_{-\infty}^{\infty} f'(t)e^{-ikt}dt = f(t)e^{-ikt}\big|_{-\infty}^{\infty} - \int_{-\infty}^{\infty} f(t)\left(-ike^{-ikt}\right)dt = ik\hat{f}(k)$$

Notice that this computation is only valid if we make the assumption

$$\lim_{t\to\pm\infty} f(t) = 0.$$

Again, this follows from the condition

$$\int_{-\infty}^{\infty} |f(t)|^2 dt < \infty$$

which we must always assume, in order to apply the Fourier transform.

As an example of applying this property, suppose we want to find a solution of the second order ODE

$$y'' + y' + y = e^{-|t|}.$$

To solve the equation, we could take the Fourier transform of both sides:

$$(ik)^2\hat{y}(k) + (ik)\hat{y}(k) + \hat{y}(k) = \frac{2}{(1+k^2)}$$

This gives the Fourier transform of the solution,

$$\hat{y}(k) = \frac{2}{(1+k^2)((1-k^2)+ik)}$$

and we can recover the solution via Fourier inversion:

$$y(t) = \frac{1}{2\pi}\int_{-\infty}^{\infty} \hat{y}(k)e^{ikt}dk = \frac{1}{2\pi}\int_{-\infty}^{\infty} \frac{2e^{ikt}}{(1+k^2)((1-k^2)+ik)}dt$$

This is a fine method, but it has some issues:

(1) It only gives solutions for which Fourier inversion is valid, i.e. solutions $y(t)$ which satisfy

$$\int_{-\infty}^{\infty} |y(t)|^2 dt < \infty.$$

We do not obtain *general* solutions of the equation, with arbitrary initial values.

(2) In order to obtain the solution, we need to compute a nasty integral!

Next week we will introduce the Laplace transform, which is a modification of the Fourier transform that is specifically adapted to solving initial value problems. This will resolve difficulty (1) but not difficulty (2).

To resolve difficulty (2), we will derive a number of properties of Laplace transforms and use these properties to construct a table of Laplace transforms. This will allow us to reconstruct the solution from its Laplace transform - we just look it up in the table.

The Heat Equation on the Line. This section is optional but recommended - in it we will use Fourier transforms to solve the heat equation

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

on the entire real line, without any boundary conditions. In doing so we will encounter the concept of a *convolution*, which will play an important role in the context of Laplace transforms.

The idea we will use to solve the heat equation is to take the Fourier transform of both sides with respect to the x variable. The result will be an equation involving the Fourier transform

$$\hat{u}(k,t) = \int_{-\infty}^{\infty} u(x,t) e^{-ikx} dx$$

which we will be able to solve easily.

Before we try out this strategy, notice that the initial value of $u(x,t)$ determines the initial value of $\hat{u}(k,t)$:

$$\hat{u}(k,0) = \int_{-\infty}^{\infty} u(x,0) e^{-ikx} dx = \int_{-\infty}^{\infty} f(x) e^{-ikx} dx = \hat{f}(k)$$

We will use this fact momentarily.

Now let's try to implement the strategy. Taking the Fourier transform of both sides of the equation,

$$\mathcal{F}\left[\frac{\partial u}{\partial t} - \alpha^2 \frac{\partial^2 u}{\partial x^2}\right] = \mathcal{F}[0]$$

and applying the linearity property of the Fourier transform, we get

$$\mathcal{F}\left[\frac{\partial u}{\partial t}\right] - \alpha^2 \mathcal{F}\left[\frac{\partial^2 u}{\partial x^2}\right] = 0$$

To deal with the first term on the left, notice that

$$\mathcal{F}\left[\frac{\partial u}{\partial t}\right] = \int_{-\infty}^{\infty} \frac{\partial u}{\partial t}(x,t) e^{-ikx} dx = \frac{\partial}{\partial t}\left[\int_{-\infty}^{\infty} u(x,t) e^{-ikx} dx\right] = \frac{\partial}{\partial t}\left[\hat{u}(k,t)\right]$$

where in the second step we brought the derivative out from under the integral. To make the notation less terrible we can write this as follows:

$$\mathcal{F}\left[\frac{\partial u}{\partial t}\right] = \frac{\partial \hat{u}}{\partial t}.$$

To deal with the second term on the left, apply property $(*)$ of the Fourier transform, which was stated in the previous section. This gives:

$$\alpha^2 \mathcal{F}\left[\frac{\partial^2 u}{\partial x^2}\right] = \alpha^2 (ik)^2 \mathcal{F}[u] = -(\alpha k)^2 \hat{u}$$

Substituting, we find that

$$\frac{\partial \hat{u}}{\partial t} - \alpha^2 k^2 \hat{u} = 0$$

This is an ordinary differential equation, whose solution we know! The solution depends on the initial value $\hat{u}(k,0)$, which we have seen is equal to $\hat{f}(k)$:

$$\hat{u}(k,t) = \hat{u}(k,0) e^{-\alpha^2 k^2 t} = \hat{f}(k) e^{-\alpha^2 k^2 t}.$$

To obtain the solution $u(x,t)$, we can just apply the inverse Fourier transform:

$$u(x,t) = \mathcal{F}^{-1}\left[\hat{f}(k) e^{-\alpha^2 k^2 t}\right] = \frac{1}{2\pi} \int_{-\infty}^{\infty} \hat{f}(k) e^{-\alpha^2 k^2 t} e^{ikx} dk$$

Now, you might say this is the end of the story - we obtained a somewhat messy but explicit formula for the solution. But what we really want is a formula in terms of $f(x)$, not in terms of $\hat{f}(k)$!

This desire leads us to consider a more general problem. Suppose we have two functions $f(x)$ and $g(x)$, and we take their Fourier transforms:

$$\hat{f}(k) = \int_{-\infty}^{\infty} f(v)e^{-ikv}du$$

$$\hat{g}(k) = \int_{-\infty}^{\infty} g(u)e^{-iku}dv$$

In this situation, we can multiply the Fourier transforms,

$$\hat{f}(k)\hat{g}(k)$$

and ask, what is its inverse Fourier transform, in terms of $f(x)$ and $g(x)$?

To find out, we can combine the product into a double integral:

$$\hat{f}(k)\hat{g}(k) = \left(\int_{-\infty}^{\infty} f(v)e^{-ikv}dv\right)\left(\int_{-\infty}^{\infty} g(u)e^{-iku}du\right) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} f(v)g(u)e^{-ik(v+u)}dvdu.$$

To simplify the double integral, we make the substitution $(u, v) = (x - v, v)$:

$$\hat{f}(k)\hat{g}(k) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} f(v)g(x-v)e^{-ikx}dvdx$$

and observe that the right hand side is a Fourier transform:

$$\hat{f}(k)\hat{g}(k) = \int_{-\infty}^{\infty}\left[\int_{-\infty}^{\infty} f(v)g(x-v)dv\right]e^{-ikx}dx = \mathcal{F}\left[\int_{-\infty}^{\infty} f(v)g(x-v)dv\right]$$

Taking the inverse Fourier transform of both sides, we conclude that

$$\mathcal{F}^{-1}\left[\hat{f}(k)\hat{g}(k)\right] = \int_{-\infty}^{\infty} f(v)g(x-v)dv$$

The expression on the right hand side is called the *convolution* of the functions $f(x)$ and $g(x)$. It is a bit frustrating that multiplying two Fourier transforms leads to a convoluted inverse transform, but that's life!

Anyway, applying the result above to our solution of the heat equation, we have

$$u(x,t) = \mathcal{F}^{-1}\left[\hat{f}(k)e^{-\alpha^2k^2t}\right] = \int_{-\infty}^{\infty} f(v)\phi(x-v,t)dv$$

where $\phi(x,t)$ is the function whose Fourier transform (with respect to x) is

$$\hat{\phi}(k,t) = e^{-\alpha^2k^2t}$$

To find this function explicitly, we need to apply the inverse Fourier transform. The result turns out to be

$$\phi(x,t) = \mathcal{F}^{-1}\left[e^{-\alpha^2k^2t}\right] = \frac{1}{\sqrt{4\pi\alpha^2t}}e^{-\frac{x^2}{4\alpha^2t}}$$

You can check that the function $\phi(x,t)$ is itself a solution of the heat equation - it is called the *fundamental solution*. If you plot it for a range of values of t you can see that it describes a heat distribution that is initially concentrated in an infinitely small interval centered at $x = 0$, and spreads out over time. Here is an animation:

https://www.geogebra.org/calculator/j52gabqp

To understand the meaning of the formula

$$u(x,t) = \int_{-\infty}^{\infty} f(v)\phi(x-v,t)dv,$$

you may find it more intuitive to think in terms of diffusion of a gas in a narrow tube (instead of diffusion of heat in a narrow rod). With this interpretation of the equation, $u(x,t)$ represents the concentration of the gas (number of molecules per unit length of the tube), and $f(x) = u(x,0)$ represents the initial concentration at time $t = 0$. With this interpretation, the quantity

$$f(v)dv$$

represents the number of gas molecules which are initially located in a small interval $[v, v + dv]$ on the x axis. As time passes, these molecules drift away and spread out, and after a time t they are distributed throughout space - the concentration at a given location x can be expressed as

$$f(v)\phi(x - v, t)dv.$$

But this represents just one contribution to the overall concentration, since we have only considered gas molecules which started at a particular initial location. To obtain the overall concentration, we must integrate over all possible initial locations. This is one interpretation of the formula

$$u(x, t) = \int_{-\infty}^{\infty} f(v)\phi(x - v, t)dv.$$

To read about an application of this formula in image processing, see

https://en.wikipedia.org/wiki/Gaussian_blur

The idea is to let the intensity of each pixel in a digital image diffuse outward to the neighboring pixels. The result is a blurry image - the greater the value of t, the blurrier the result.

When you do this you don't actually solve the diffusion equation, you just apply the convolution formula. But of course, you use the 2D generalization of the formula, not the 1D version given here - there is a formula in the linked article.

Now, it might seem like blurring an image would always be a bad idea in scientific applications, since blurring appears to destroy potentially important information. However, you might view this as a good thing in cases where the original image contains *too much* information.

To make this more precise, we can go back to our original formula for the Fourier transform of the solution:

$$\hat{u}(k, t) = \hat{f}(k)e^{-\alpha^2 k^2 t}.$$

Notice that the factor

$$e^{-\alpha^2 k^2 t}$$

is close to 1 for small values of k, and close to 0 for large values of k. So, blurring has the effect of reducing the intensity of the high frequency Fourier components (which are often regarded as unnecessary "noise"). From this point of view, blurring reduces noise and makes it easier to correctly analyze the image.