5 The Heat Equation

We have been studying conservation laws which, for a conserved quantity or set of quantities u with corresponding fluxes f, adopt the general form

$$u_t + \nabla \cdot f = 0.$$

So far, we have closed these systems with relations between the fluxes and the conserved quantities of the form

$$f = f(u).$$

However, in many scenarios, f may depend not just on u but also on its gradient. The simplest case, where u is a scalar and $f = -\nabla u$, gives rise to the heat equation:

$$u_t = \Delta u \tag{108}$$

that we study in this section.

5.1 Physical origin

The heat equation appears in models in a multitude of ways. Fourier first introduced it to describe heat transfer. Here u represents the temperature, which is conserved as a manifestation of energy. Then Fourier argued the very intuitive principle that heat flows from warmer to cooler areas, so, in the simplest, linear model, it follows the direction of minus the temperature gradient, yielding (108).

The equation appears also in mechanical scenarios involving friction. For instance, we can describe the flow of water through a porous media in terms of the height h of the water column and its mean speed u. From conservation of mass (linearized), we have

$$h_t + \nabla \cdot u = 0. \tag{109}$$

To close this equation, we invoke Newton's law. The forces acting on the fluid are the pressure, which is proportional to h, and the drag, proportional to minus u. Then we can write

$$u_t + \nabla h = -u.$$

If the motion evolves slowly, as is often the case in porous media, we can neglect the inertial term, and obtain

$$\nabla h = -u.$$

Combining this with (109), we obtain again the heat equation

$$h_t = \Delta h.$$

The heat equation models *diffusive* processes, which rule for instance the evolution of the concentration of ink in water. To see this, think of the interface between two regions with different concentrations. The fluid's turbulence or, in a smaller scale, just its molecules moving back and forth, will make mass cross

the interface in both directions. For an incompressible fluid in a fixed domain, the net mass flux is zero. Yet the net amount of ink carried through the interface dos not vanish: mass coming from the area with the greater concentration will carry more ink than the same amount of mass going the other way. Since the net ink flow is proportional to the difference of concentrations, we obtain again the heat equation, once we take the limit of small traveled distances, so that concentration differences become derivatives.

This argument, refined and formalized, makes the heat equation appear prominently in stochastic calculus: random, *Brownian* motion, gives rise to diffusion.

The traffic model of our first few classes also yields the heat equation, if we have the drivers look ahead and react to the gradient rather than to the car density itself.

Finally, the heat equation also appears describing not natural phenomena but algorithms: *descent* algorithms in optimization often evolve a field by following its gradient. Also numerical approximations to partial differential equations often involve local averages of the solution –Godunov's method is a clear example of this. Such local averages, which act to reduce the gradients, obey variations of the heat equation.

5.2 The fundamental solution

We start by solving the initial value problem

$$u_t = \Delta u\,, \quad u(x,0) = u_0(x) \tag{110}$$

in all of n-dimensional space. Using the fact that the equation is linear, if we divide our initial value data into pieces, say

$$u_0(x) = u_1(x) + u_2(x)\,,$$

then $u(x,t)$ will be the sum of the solutions to the initial value problems with u_1 and u_2 as initial data. More generally, if

$$u_0(x) = \int a(y) g(x,y)\, dy$$

and $G(x,t,y)$ satisfies

$$G_t = \Delta_x G\,, \quad G(x,0,y) = g(x,y)\,,$$

then

$$u(x,t) = \int a(y) G(x,t,y)\, dy$$

solves (110). There are many ways of decomposing quite arbitrary functions $u_0(x)$ into elementary components; the Fourier Transform, for instance, will do the job. This we will do a bit later (we need to honor Fourier when solving his heat equation!), but presently we take a different approach, thinking of a

function $u_0(x)$ as the sum of infinitely many functions, each giving us its value at one point and zero elsewhere:

$$u_0(x) = \int u_0(y)\delta(x-y)\,dy\,,$$

where δ stands for the n-dimensional δ-function.

Then our problem for $G(x,t,y)$, the *Green*'s function or *fundamental solution* to the heat equation, is

$$G_t = \Delta_x G\,, \quad G(x,0,y) = \delta(x-y).$$

Since the heat equation is invariant under translation, we have that

$$G(x,t,y) = G(x-y,t,0)\,,$$

which, abusing language, we shall simply denote $G(x-y,t)$. Then $G(x,t)$ satisfies

$$G_t = \Delta G\,, \quad G(x,0) = \delta(x)\,, \tag{111}$$

and $u(x,t)$ can be written in the form

$$u(x,t) = \int u_0(y)G(x-y,t)\,dy.$$

Our next task, therefore, is to find G.

5.2.1 The one-dimensional case

In one dimension, (111) becomes

$$G_t = G_{xx}\,, \quad G(x,0) = \delta(x).$$

We will find G following an instructive path. First, notice that, again because of linearity, if $u(x,t)$ satisfies the initial value problem

$$u_t = u_{xx}\,, \quad u(x,0) = u_0(x)\,,$$

then its derivative

$$v = u_x$$

satisfies

$$v_t = v_{xx}\,, \quad v(x,0) = u_0'(x)\,.$$

Therefore, instead of solving the problem for G, we can solve the simpler one that its integral F satisfies:

$$F_t = F_{xx}\,, \quad F(x,0) = H(x),$$

where $H(x)$ is the Heaviside function, zero to the left of $x=0$ and one to its right. Once we find F, we can simply write $G(x,t) = F_x(x,t)$.

The problem for F should remind us of the Riemann problem of some classes ago, that also had piece-wise constant initial data; we will solve it using a similar trick. First, we notice that the initial data are invariant under the stretching

$$x \to \lambda x$$

for any positive λ. The heat equation, on the other hand, is invariant under the following combined stretching of x and t:

$$x \to \lambda x\,, \quad t \to \lambda^2 t.$$

Hence the solution F needs to satisfy

$$F(\lambda x, \lambda^2 t) = F(x,t)$$

for any positive value of λ. In particular, adopting $\lambda = \frac{1}{\sqrt{t}}$, we have that

$$F(x,t) = F\left(\frac{x}{\sqrt{t}}, 1\right) = \Phi(\xi),$$

where

$$\xi = \frac{x}{\sqrt{t}}.$$

Hence

$$F_t = -\frac{1}{2}\left(\frac{x}{t^{\frac{3}{2}}}\right)\Phi'(\xi) \quad \text{and} \quad F_{xx} = \left(\frac{1}{t}\right)\Phi''(\xi)\,,$$

and the heat equation becomes the ODE

$$\Phi'' + \frac{\xi}{2}\Phi' = 0\,,$$

with solution

$$\Phi(\xi) = \frac{1}{\sqrt{4\pi}}\int_{-\infty}^{\xi} e^{-\frac{s^2}{4}}\,ds\,,$$

the error function. The lower limit of integration and the constant factor outside have been picked so that F agrees with the Heaviside function when $t \to 0$, and so $\xi \to \pm\infty$. Finally, we obtain our Green's function

$$G(x,t) = F_x(x,t) = \frac{1}{\sqrt{4\pi t}} e^{-\frac{x^2}{4t}}\,. \tag{112}$$

5.2.2 The multidimensional fundamental solution

In dimensions $n > 1$ we need to change our argument, since we can no longer think of the delta function as a derivative of a Heaviside. Instead, we jump directly to the stretching argument. The problem for G is spherically symmetric, so the solution can depend only on r and t. As in 1D, the heat equation is invariant under the combined stretching

$$r \to \lambda r\,, \quad t \to \lambda^2 t.$$

The initial data, however, is not. The n-dimensional δ function satisfies

$$\delta(\lambda x) = \frac{1}{\lambda^n}\delta(x),$$

as can be easily verified by integrating both sides over a domain including the origin, and changing variables on the left to $y = \lambda x$.

Then the symmetry that leaves both the equation and the initial data invariant is

$$r \to \lambda r\,, \quad t \to \lambda^2 t \quad G \to \lambda^n G.$$

It follows that

$$\lambda^n G(\lambda r, \lambda^2 t) = G(r,t),$$

so picking $\lambda = \frac{1}{\sqrt{t}}$, we have

$$G(r,t) = \frac{1}{t^{\frac{n}{2}}} G\left(\frac{r}{\sqrt{t}}, 1\right) = \frac{1}{t^{\frac{n}{2}}}\Gamma(\xi),$$

with

$$\xi = \frac{r}{\sqrt{t}}.$$

It is most convenient to write the heat equation in polar coordinates in the form

$$\left(r^{n-1}G\right)_t = \left(r^{n-1}G_r\right)_r.$$

which in terms of Γ becomes the ODE

$$\left(\xi^{n-1}\Gamma' + \frac{\xi^n}{2}\Gamma\right)' = 0.$$

Integrating this once yields

$$\Gamma' + \frac{\xi}{2}\Gamma = 0$$

–the constant of integration needs to vanish if we want sensible behavior as $\xi \to 0$–, so

$$\Gamma = \frac{1}{(4\pi)^{\frac{n}{2}}} e^{-\frac{\xi^2}{4}}$$

–this constant picked so that Γ integrates to one–, yielding the fundamental solution

$$G(x,t) = \frac{1}{(4\pi t)^{\frac{n}{2}}} e^{-\frac{|x|^2}{4t}}. \tag{113}$$

Finally, the solution to the initial value problem (110) for $u(x,t)$ takes the form

$$u(x,t) = \frac{1}{(4\pi t)^{\frac{n}{2}}} \int e^{-\frac{|x-y|^2}{4t}}\, u_0(y)\, dy : \tag{114}$$

each point in the initial data irradiates a spherically symmetric Gaussian whose radius expands as the square root of time and whose amplitude decays accordingly so as to preserve mass.

5.3 Solution through Fourier analysis

In the analysis of the prior subsection, we used the linearity of the heat equation to solve it as a superposition of infinitely many initial value problems, each of which represented the real initial data at one point. The advantage of this approach is that each of the auxiliary initial value problems has data that is *localized* in physical space. We saw how each of these localized data evolves, as a spherically symmetric Gaussian whose variance grows linearly with time.

Yet there is another natural choice for the representation of the initial data: one that is local not in real but in Fourier space. We have touched briefly on this in the section on dispersive waves, where we saw that sinusoidal solutions are typical in many systems. In general, we should expect to find solutions of the form

$$a_k(t)e^{ik\cdot x} \tag{115}$$

in any linear system of partial differential equations that has constant coefficients –more precisely, coefficients that do not depend on the spatial coordinates x. The reason is that the *Fourier modes* (115) are eigenfunctions of all spatial derivatives. In particular, plugging the ansatz (115) into the heat equation (108) yields

$$a_k'(t) + \|k\|^2 a(t) = 0\,,$$

with solution

$$a_k(t) = a_k(0)e^{-\|k\|^2 t}$$

that decays exponentially, faster for the more rapidly oscillating modes.

Superposing many if these solutions –one for each possible value of the wave-vector k– and fitting the $a_k(0)$ to the initial values for $u_0(x)$, we obtain

$$u(x,t) = \int \hat{u}_0(k)e^{ik\cdot x - \|k\|^2 t}dk, \tag{116}$$

where

$$\hat{u}_0(k) = \frac{1}{(2\pi)^n}\int u_0(y)e^{-ik\cdot y}dy \tag{117}$$

is the Fourier transform of the initial data $u_0(x)$.

As an exercise, let us check that the general solution (116) through Fourier modes to the initial value problem agrees with (114), the one computed through the superposition of fundamental solutions: Plugging (117) into (116), we obtain

$$u(x,t) = \frac{1}{(2\pi)^n}\int\int u_0(y)\; e^{ik\cdot(x-y)-\|k\|^2 t}\; dy\; dk.$$

Exchanging the order of integration and completing squares yields

$$u(x,t) = \frac{1}{(2\pi)^n}\int\left(\int e^{-\left\|\sqrt{(t)}k - i\frac{x-y}{2\sqrt{t}}\right\|^2}\; dk\right) u_0(y)\; e^{-\frac{\|x-y\|^2}{4t}}\; dy.$$

The integral in k is that of a Gaussian displaced in the complex n-dimensional space; it can be computed exactly:

$$\int e^{-\left\|\sqrt{(t)}k - i\frac{x-y}{2\sqrt{t}}\right\|^2}\, dk = \left(\frac{\pi}{t}\right)^{\frac{n}{2}},$$

and we recover (114).

What do we learn from the two forms of the general solution to the initial value problem? The solution (114) in physical space tells us how point sources diffuse, expanding with a square-root dependence on time. Notice also that, although for any fixed time t the solution decays exponentially as $|x| \to \infty$, it is nowhere zero for any positive time, no matter how small. So, in a sense, the speed of diffusion is slow in the bulk, converging to zero as $t \to \infty$, while in another sense it is infinite, carrying information from the initial disturbance at $x = 0$ all the way to infinity at time $t = 0^+$. This can be contrasted with the hyperbolic scenarios we studied before, that had finite characteristic speeds at which information propagated, so that no information reached beyond the *area of influence* of the initial data but, once the information did reach a point, the local solution was affected by a sizable amount.

The solution (116) in Fourier space teaches something altogether different: small-scale features (with correspondingly large wavenumbers $\|k\|$) decay faster than those with longer scales. This is the effect of diffusion acting to suppress high gradients: it first eliminates the data's micro-structure, essentially replacing it with its local average, which evolves over a longer time-scale. Again, this can be compared with the dispersive scenario, where different wavenumbers do not decay, but instead travel at different speeds. In fact, it is instructive to compare the heat equation (108) with the Schroedinger equation for a free particle,

$$i\Psi_t + \Delta\Psi = 0. \tag{118}$$

Notice that the two are formally identical if one is allowed to introduce an imaginary time, $t \to it$, or equivalently an imaginary diffusion coefficient. Then the decay rate $\|k\|^2$ for the heat equation becomes the dispersion relation $\Omega = \|k\|^2$ for Schroedinger's. In fact, this formal analogy is at the heart of the FeynmanKac formula, linking the path integrals of quantum mechanics –Feinman's contribution– to those of Wiener processes –Kac's. Again, something we may come back to in due time ...

5.4 Arrows of time

Consider the integral $\Phi(t)$ over all of space of a function $\phi(u)$, where $u(x,t)$ satisfies the heat equation (108). Then

$$\Phi'(t) = \int \phi'(u)\, u_t\, dx = \int \phi'(u)\, \Delta u\, dx = \int \left[\nabla \cdot (\phi'(u)\nabla u) - \phi''(u)\, \|\nabla u\|^2\right] dx$$

so

$$\Phi'(t) = -\int \phi''(u)\, \|\nabla u\|^2\, dx \tag{119}$$

(Notice that this statement remains true for integrals over finite domains, provided that the domain has *insulated* boundaries, i.e. boundaries across which the normal flux of ϕ,

$$\frac{\partial \phi(u)}{\partial n}$$

vanishes.) If $\phi''(u)$ has a definite sign –i.e., if $\phi(u)$ is either strictly convex or strictly concave–, then (119) implies that $\Phi(t)$ is monotonically increasing or decreasing, providing an "arrow of time": a distinction between the past and the future. This arrow only stops once $u(x,t)$ is completely uniform throughout the domain.

To give more flavor to this principle, let us consider the case in which $u(x,t)$ represents a density, so it is everywhere positive (we'll see shortly that this property, if true for the initial data, is preserved by the heat equation.) Then the *entropy*

$$s = -u \log(u) \tag{120}$$

is a convex function of u, so it must increase continuously, until the flow reaches a state of maximal entropy consistent with the conservation of the integral of u: a state where u is uniform. This is a heat-equation version of the second principle of thermodynamics. It was Eddington who first spoke of the entropy as an "arrow of time", with the universe always moving towards a state of maximal disorder or randomness.

Here comes an arrow of time of a different nature, also associated with the heat equation. Taking the gradient of (108) yields

$$\nabla u_t = \Delta \nabla u.$$

Next we take the inner product of both sides with ∇u and integrate over the full domain, yielding

$$\frac{d}{dt} \int \frac{1}{2} \|\nabla u\|^2 \, dx = \int \nabla u \cdot \Delta \nabla u \, dx = - \int \sum_{j,k} \left| \frac{\partial^2 u}{\partial x_j \partial x_k} \right|^2 \, dx \leq 0, \tag{121}$$

where we have applied the divergence theorem to the right-hand side of the equation, a component of ∇u at a time. The quantity

$$U = \int \frac{1}{2} \|\nabla u\|^2 \, dx \tag{122}$$

is sometimes denoted the *energy* associated with u (although, in the original context of the heat equation, with u representing the temperature, the thermal energy is the integral of u, not of its gradient.) We see that this energy decreases, providing another arrow of time. The meaning of this particular arrow is clearly related to the heat equation's tendency to suppress gradients. But there is more to this story. Consider the variational derivative of U with respect to u:

$$\frac{\delta U}{\delta u} = -\Delta u \tag{123}$$

(you may want to review your calculus of variations.) It follows that the heat equation (108) can be written in the form

$$u_t = -\frac{\delta U}{\delta u}. \tag{124}$$

In words, the heat equation can be thought of as a *descent* flow of the energy U, since the variational derivative is the infinite-dimensional version of the gradient. It is no wonder, therefore, that U should decrease with time:

$$\frac{dU}{dt} = \int \frac{\delta U}{\delta u} u_t \; dx = -\int \left(\frac{\delta U}{\delta u}\right)^2 \; dx = -\int (\Delta u)^2 \; dx \leq 0. \tag{125}$$

Puzzle: Are the two expressions we found for $\frac{dU}{dt}$, (121) and (125), equivalent?

5.5 Diffusion and Brownian motion

Now that we have seen the irreversibility associated with the heat equation, and alluded to it in terms of entropy, disorder and randomness, maybe it is time to look back at its physical origins, particularly as it relates to stochastic processes. We will first look at one-dimensional random walks, and soon generalize things to multidimensional Langevin processes and their corresponding Fokker-Planck equations.

Consider a walker on the real line that, at regular intervals Δt, moves to the left or right, with equal probabilities, a fixed distance Δx. If the walker starts walking at time $t = 0$ from $x = 0$, what is the probability P_j^n of finding it in position $x = x_j = j\Delta x$ at time $t = t_n = n\Delta t$? We can answer this question recursively: in order to be in $x = x_j$ at time t_n, the waker must have been one step to the left or right the step before. Then

$$P_j^{n+1} = \frac{P_{j-1}^n + P_{j+1}^n}{2}, \tag{126}$$

where the $\frac{1}{2}$ is the probability of stepping in each direction. We can reorganize this expression in the form

$$\frac{P_j^{n+1} - P_j^n}{\Delta t} = \frac{1}{2}\frac{\Delta x^2}{\Delta t}\frac{P_{j+1}^n - 2P_j^n + P_{j-1}^n}{\Delta x^2}, \tag{127}$$

suggestive of a finite-difference approximation to the one-dimensional heat equation.

To actually take the limit as Δx and Δt get small, we need some extra considerations. First, we need to re-interpret the discrete P_j^n in terms of a function defined on the real line. The natural choice is a probability density $\rho(x, t)$. In terms of this, P_j^n can be interpreted as the integral of $\rho(x, t_n)$ between $x_{j-\frac{1}{2}}$ and $x_{j+\frac{1}{2}}$. Then we need to take a succession of walks, with decreasing time intervals and step-sizes. Yet this needs to be done so that the quotient

$$r = \frac{1}{2}\frac{\Delta x^2}{\Delta t}$$

remains finite, or else the evolution equation will trivialize. Therefore the step-size Δx needs to scale as the square-root of the time interval Δt. With all this done, we can now the limit of (127), and obtain the one-dimensional heat equation

$$\rho_t = \nu \ \rho_{xx},$$

where the diffusivity ν is the limit of r as Δt and Δx go to zero.

Let us now generalize this to arbitrary dimensions and more general walks that do not have a fixed step Δx. Instead, consider a probability distribution $\mu(y|x,t,\Delta t)$ for all possible steps y in R^n. The distribution may depend on the position x and time t from which we are walking away, and is temporarily associated with a fixed time-interval Δt. Then the probability density $\rho(x, t + \Delta t)$ is given by the following generalization of (126):

$$\rho(x,t+\Delta t) = \int \rho(x-y,t)\mu(y|x-y,t,\Delta t)\ dy, \tag{128}$$

so

$$\rho(x,t+\Delta t) - \rho(x,t) = \int \left[\rho(x-y,t)\mu(y|x-y,t,\Delta t) - \rho(x,t)\mu(y|x,\Delta t)\right]\ dy, \tag{129}$$

where we have used the fact that μ integrates to one.

We need to impose some requirement on μ as Δt gets smaller. First, we should expect the distribution of steps to be concentrated near $y = 0$. Then we can approximate the right hand-side of (129) using the first two terms in the Taylor expansion, and the left-hand size with a time derivative:

$$\Delta t \frac{\partial \rho(x,t)}{\partial t} \approx \int \left[\sum_{j=1}^{n} -y_j \frac{\partial(\rho\mu)}{\partial x_j} + \frac{1}{2}\sum_{j,k=1}^{n} y_j y_k \frac{\partial^2(\rho\mu)}{\partial x_j \partial x_k}\right]\ dy =$$
$$-\sum_{j=1}^{n} \frac{\partial\left(\rho\bar{\mu}_j\right)}{\partial x_j} + \frac{1}{2}\sum_{j,k=1}^{n} \frac{\partial^2\left(\rho c_{j,k}\right)}{\partial x_j \partial x_k},$$

where

$$\bar{\mu}(x,t,\Delta t) = \int y\ \mu(y;x,t,\Delta t)\ dy$$

is the expected value of the displacement, and

$$c_{j,k}(x,t,\Delta t) = \int y_j y_k\ \mu(y;x,t,\Delta t)\ dy$$

its covariance matrix. In order to make nontrivial contributions to the limit as $\Delta t \to 0$, both of these need to scale proportionally to Δt:

$$\bar{\mu}(x,t,\Delta t) \sim u(x,t)\ \Delta t\,, \quad c_{j,k}(x,t,\Delta t) \sim d_{j,k}(x,t)\ \Delta t,$$

in which case we obtain in the limit the Fokker-Planck equation

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho u) = \frac{1}{2}\sum_{j,k=1}^{n} \frac{\partial^2\left(d_{j,k}\rho\right)}{\partial x_j \partial x_k}. \tag{130}$$

If the *drift* u vanishes and the covariance matrix d equals twice the identity, we recover the normalized heat equation (108).

Notice the similarity between the left-hand side of (130) and the equation for conservation of mass in fluids. Indeed they are the same, though we are talking here of probability rather than fluid densities. The drift term u represents the deterministic part of the walk, and plays the same role as a fluid velocity: if, in the mean, the walker tends to go from one place to another, then the probability of finding it in the latter will increase over time. The terms on the right-hand side correspond instead to the fluctuating part of the velocity field, yielding diffusion, as we have argued in section (5.1).

5.6 Boundary conditions

We have dealt so far with the heat equation (108) in all of space. Yet more often than not we are concerned with the evolution of the solution in a bounded domain, subject to boundary conditions: the evolution of the temperature within a loaf of bread just taken from the oven and cooling in the open air, the time required for the various ingredients to mix throughout in a stew. Then we solve the heat equation in a domain Ω with initial data and with conditions on the boundary $\delta\Omega$ that are typically of one of these forms:

- Insulation:
$$\frac{\partial u}{\partial n} = 0, \tag{131}$$
where n is the direction normal to the boundary. The total heat content,
$$H = \int_\Omega u(x,t)\ dx \tag{132}$$
is constant, as follows from integrating (108) over Ω and applying the divergence theorem.

- Dirichlet:
$$u = f(x,t), \tag{133}$$
which specifies the value of the solution on the boundary.

- Neumann:
$$\frac{\partial u}{\partial n} = g(x,t), \tag{134}$$
specifying the heat flux. The insulated scenario (131) is a particular case of this. The total heat content H from (132) satisfies
$$\frac{dH}{dt} = \int_{\delta\Omega} g.$$

- Newton's law of cooling:
$$\frac{\partial u}{\partial n} = u - f(x,t), \tag{135}$$

combining Dirichlet and Neumann. Here the rate of heat loss is proportional to the difference between the temperature u at the boundary and the given temperature f of an external fluid, such as air or water.

Few of these problems can be solved readily in closed form; I'll pose some of these as problems (with hints) for you to work out. In all problems, $u(x,0) = u_0(x)$ within Ω.

Problem 1: Ω is the semi-infinite line $[0,\infty)$, with homogeneous Dirichlet condition

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

Hint: The heat equation preserves the oddity of solutions.

Problem 2: Ω is the semi-infinite line $[0,\infty)$, with homogeneous Neumann condition

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

Hint: The heat equation preserves the evenness of solutions.

Problem 3: Ω is the finite segment $[0,1]$, with homogeneous Dirichlet conditions

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

Hint: By expanding by oddity on the two ends, you end up building a periodic initial condition. This can be expanded in terms of sines, and we know how each sine evolves in time. *Question:* How does the solution look like for long times? Provide a simple formula.

Problem 4: Ω is the finite segment $[0,1]$, with time-independent Dirichlet conditions

$$u(0,t) = a, \quad u(1,t) = b.$$

Hint: Invoke linearity to add the solution of the previous problem and one that is time-independent and solves the non-homogeneous boundary conditions.

5.7 The backward heat equation

So far we have used the heat equation to predict the future temperature of an object that has an initial distribution $u_0(x)$. Can we do things the other way around: knowing the present temperature distribution, infer the one in the past? The problem would look identical to the one we posed before,

$$u_t = \Delta u\,, \quad u(x,0) = u_0(x)\,,$$

but now we want the solution $u(x,t)$ for $t < 0$. In order not to twist our heads too much looking backwards in time, we can make the simple change of variables $t \to -t$, which yields

$$u_t = -\Delta u\,, \quad u(x,0) = u_0(x)\,, \tag{136}$$

the initial value problem for the *backward heat equation*, and look now at positive values of t.

There are a number of intuitive reasons why we shouldn't expect this program to work out too well. If I have removed my bread loaf from the oven two hours ago, at this time it is probably as cold as it will ever be. Could I, by looking at its nearly homogeneous temperature distribution carefully, infer at which temperature it left the oven? More generally, imagine we start with a very irregular temperature field. Diffusion will tend to homogenize it very rapidly, reducing the gradients. Could we uncover, from the comparatively smooth and uniform later distribution, the details of the mess we started with? Diffusion models mixing. Given a well-mixed state, it is hard to figure out where each of its components started.

Mathematically, this translates into the statement that the initial value problem (136) is *ill-posed.* By this we mean that insignificant changes in the initial data can give rise to huge differences in the solution. In more precise mathematical terms, the solution at any fixed time t is not a continuous function of the initial data. The best way to see this is through a simple example. Consider Problem 3 from above, but for the backward heat equation (136), with sinusoidal initial data:

$$u_t = -u_{xx} \quad \text{in} \quad 0 < x < 1,$$

$$u(0,t) = u(1,t) = 0, \quad u(x,0) = u_0(x) = a\sin(2\pi n x).$$

The solution to this problem is –you worked it out for negative times in Problem 3–

$$u(x,t) = ae^{(2\pi n)^2 t}\sin(2\pi n x).$$

Now we can use this solution to play the $\epsilon - \delta$ game of continuity. With $u_0 = 0$, the solution is clearly $u(x,t) = 0$. For a given $\epsilon > 0$ and fixed time T, can we always find a δ such that, if the initial data satisfies $\|u_0(x)\| < \delta$, then $\|u(x,T)\| < \epsilon$?

The answer is negative. In the exact solution above,

$$\|u(x,T)\| = e^{(2\pi n)^2 T}\|u_0(x)\|.$$

So no matter how small δ, we can always find an integer n large enough that $\|u(x,T)\| > \epsilon$, contradicting continuity. Notice in passing that this argument is independent of the choice of a norm.

In this argument, continuity is killed by the large wavenumbers. This agrees with our intuition that it is the small-scale structure that is most rapidly averaged away by diffusion, and hence hardest to recover from the solution at a later time.

5.8 The forced heat equation

The forced heat equation

$$u_t - \Delta u = f(x,t) \tag{137}$$

appears in the presence of heating sources or, in general, of sources of whatever the solution u models (pollution, underground water, etc.) Just like the forced wave equation, it can be solved using Duhamel's principle, which replaces the forcing by the superposition of infinitely many initial value problems, one for each time. We introduce an auxiliary function $U(x,t,s)$, satisfying the initial value problem for the unforced heat equation

$$U_t = \Delta_x U, \quad U(x,s,s) = f(x,s), \tag{138}$$

and write

$$u(x,t) = u_h(x,t) + \int_0^t U(x,t,s)\ ds, \tag{139}$$

where u_h is the unforced solution to the initial value problem

$$u_{ht} = \Delta u_h, \quad u_h(x,0) = u_0(x).$$

If the problem includes boundary conditions, one can make u_h satisfy these, and let U satisfy the corresponding homogeneous boundary conditions.

More generally, one has quite a bit of freedom on how to decompose a problem with non-homogeneous boundary data. Consider, for instance, the initial-value problem with Dirichlet boundary conditions

$$u_t - \Delta u = f(x,t)\ ,\ u(x,0) = u_0(x) \text{ for } x \in \Omega, \quad u = g(x,t) \text{ on } \delta\Omega. \tag{140}$$

One possibility is to divide u into two components:

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

where $v(x,t)$ is any function satisfying the boundary conditions, and w satisfies the forced problem with homogeneous boundary data,

$$w_t - \Delta w = f(x,t) - v_t + \Delta v\ ,\ w(x,0) = u_0(x) - v(x,0) \text{ for } x \in \Omega, \quad u = 0 \text{ on } \delta\Omega. \tag{141}$$

In other words, one can make use of the linearity of the problem to creatively trade-off among the various difficulties it may entail. This applies as well to numerical solutions, where we can build $u(x,t)$ iteratively through a sum of functions, each approximating whatever is left of the PDE and of the initial and boundary conditions after removing all the previously computed ones.

5.9 The maximum principle

Since the heat equation models diffusive phenomena, where thermal fluctuations continuously average the solution locally, we should not expect new maxima or minima to arise over time, unless they are imposed on the boundaries. A loaf of bread will never develop anywhere a temperature higher than the one it had when it had just been removed from the oven, unless it is placed in an even warmer oven later on. Clearly the same should apply, reversing the sign of the temperature, to a serving of ice-cream removed from the freezer. This is

the content of the maximum principle: the maximal and minimal values of the solution to the heat equation are to be found either on the initial data or along the boundary.

The proof is essentially the following: at an interior maximum, $u_t \geq 0$ (not necessarily zero, since u may still be growing) and $\Delta u < 0$, contradicting the heat equation. Yet we need to exclude degenerate maxima, where $\Delta u = 0$ and some higher derivatives are nonzero. For this, we need to sophisticate the argument a little bit.

Proof: Consider a domain $\Omega = \Omega_x \times [0\ T]$, where Ω_x is some fixed, bounded spatial domain, and a smooth function $v(x,t)$ satisfying

$$v_t < \Delta v. \tag{142}$$

If v achieved its maximal value in Ω at time $t = T$ and at a point $x = x_0$ in the interior of Ω_x, it would need to satisfy

$$v_t(x_0, T) \geq 0, \quad \Delta v(x_0, T) \leq 0,$$

contradicting (142). Since this applies also for all times $0 < t \leq T$, the maximal value of v is achieved either at $t = 0$ or on the boundary $\delta\Omega_x$; let's call this combination of the initial time and spatial boundary $\delta\Omega$ (notice that $\delta\Omega$ does not include the final time $t = T$.)

In order to extend this result to the heat equation, satisfying

$$u_t = \Delta u,$$

we need to create a family of v's satisfying (142) and converging to u. One possibility is to take

$$v(x,t) = u(x,t) - kt,$$

where k is a positive constant. Then

$$\max(u) \leq \max(v) + kT \leq \max_{\delta\Omega}(v) + kT \leq \max_{\delta\Omega}(u) + kT.$$

Taking the limit as $k \to 0$, we see that u also achieves its maximum on $\delta\Omega$. The argument for the minimum is the same, changing some signs and directions of the inequalities.

Uniqueness: The maximum principle provides a simple proof of uniqueness of the solution to the forced initial-boundary value problem. If two smooth functions, u and v, both satisfy

$$w_t = \Delta w + f(x,t), \quad w(x,t) = g(x,t) \text{ on } \delta\Omega,$$

then they are equal: their difference, which satisfies the unforced heat equation with homogeneous initial and boundary data, adopts its maximal and minimal values on $\delta\Omega$, where it vanishes.