Part I

Motivating examples and major applications

A *dynamical system* is a mathematical model of a system evolving in time. Most models in mathematical physics are dynamical systems. If the system has only a finite number of 'state variables', then its dynamics can be encoded in an *ordinary differential equation* (ODE), which expresses the *time derivative* of each state variable (i.e. its rate of change over time) as a function of the other state variables. For example, *celestial mechanics* concerns the evolution of a system of gravitationally interacting objects (e.g. stars and planets). In this case, the 'state variables' are vectors encoding the position and momentum of each object, and the ODE describes how the objects move and accelerate as they interact gravitationally.

However, if the system has a very large number of state variables, then it is no longer feasible to represent it with an ODE. For example, consider the flow of heat or the propagation of compression waves through a steel bar containing 10^{24} iron atoms. We *could* model this using a 10^{24}-dimensional ODE, where we explicitly track the vibrational motion of each iron atom. However, such a 'microscopic' model would be totally intractable. Furthermore, it is not necessary. The iron atoms are (mostly) immobile, and interact only with their immediate neighbours. Furthermore, nearby atoms generally have roughly the same temperature, and move in synchrony. Thus, it suffices to consider the macroscopic *temperature distribution* of the steel bar, or to study the fluctuation of a macroscopic *density field*. This temperature distribution or density field can be mathematically represented as a smooth, real-valued function over some three-dimensional domain. The flow of heat or the propagation of sound can then be described as the *evolution* of this function over time.

We now have a dynamical system where the 'state variable' is not a finite system of vectors (as in celestial mechanics), but is instead a multivariate *function*. The evolution of this function is determined by its spatial geometry – e.g. the local 'steepness' and variation of the temperature gradients between warmer and cooler regions in the bar. In other words, the *time derivative* of the function (its rate of change over time) is determined by its *spatial derivatives* (which describe its slope and curvature at each point in space). An equation that relates the different derivatives of a multivariate function in this way is a *partial differential equation* (PDE). In particular, a PDE which describes a dynamical system is called an *evolution equation*. For example, the evolution equation which describes the flow of heat through a solid is called the *heat equation*. The equation which describes compression waves is the *wave equation*.

An *equilibrium* of a dynamical system is a state which is unchanging over time; mathematically, this means that the time-derivative is equal to zero. An equlibrium of an N-dimensional evolution equation satisfies an $(N-1)$-dimensional PDE called an *equilibrium equation*. For example, the equilibrium equations corresponding to the heat equation are the *Laplace equation* and the *Poisson equation* (depending on whether or not the system is subjected to external heat input).

PDEs are thus of central importance in the thermodynamics and acoustics of continuous media (e.g. steel bars). The heat equation also describes chemical diffusion in fluids, and also the evolving probability distribution of a particle performing a random walk called *Brownian motion*. It thus finds applications everywhere from mathematical biology to mathematical finance. When diffusion or Brownian motion is combined with deterministic drift (e.g. due to prevailing wind or ocean currents) it becomes a PDE called the *Fokker–Planck equation*.

The Laplace and Poisson equations describe the equilibria of such diffusion processes. They also arise in electrostatics, where they describe the shape of an electric field in a vacuum. Finally, solutions of the two-dimensional Laplace equation are good approximations of surfaces trying to minimize their elastic potential energy – that is, soap films.

The wave equation describes the resonance of a musical instrument, the spread of ripples on a pond, seismic waves propagating through the earth's crust, and shockwaves in solar plasma. (The motion of fluids themselves is described by yet another PDE, the *Navier–Stokes equation*.) A version of the wave equation arises as a special case of Maxwell's equations of electrodynamics; this led to Maxwell's prediction of *electromagnetic waves*, which include radio, microwaves, X-rays, and visible light. When combined with a 'diffusion' term reminiscent of the heat equation, the wave equation becomes the *telegraph equation*, which describes the propagation and degradation of electrical signals travelling through a wire.

Finally, an odd-looking 'complex' version of the heat equation induces wave-like evolution in the complex-valued probability fields which describe the position and momentum of subatomic particles. This *Schrödinger equation* is the starting point of quantum mechanics, one of the two most revolutionary developments in physics in the twentieth century. The other revolutionary development was relativity theory. General relativity represents spacetime as a four-dimensional manifold, whose curvature interacts with the spatiotemporal flow of mass/energy through yet another PDE: the *Einstein equation*.

Except for the Einstein and Navier–Stokes equations, all the equations we have mentioned are *linear* PDEs. This means that a sum of two or more solutions to the PDE will also be a solution. This allows us to solve linear PDEs through the *method of superposition*: we build complex solutions by adding together many simple solutions. A particularly convenient class of simple solutions are *eigenfunctions*. Thus, an enormously powerful and general method for linear PDEs is to represent the solutions using *eigenfunction expansions*. The most natural eigenfunction expansion (in Cartesian coordinates) is the *Fourier series*.

1

Heat and diffusion

The differential equations of the propagation of heat express the most general conditions, and reduce the physical questions to problems of pure analysis, and this is the proper object of theory.

Jean Joseph Fourier

1A Fourier's law

Prerequisites: Appendix A. **Recommended:** Appendix E.

1A(i) ... in one dimension

Figure 1A.1 depicts a material diffusing through a one-dimensional domain $\mathbb{X}$ (for example, $\mathbb{X} = \mathbb{R}$ or $\mathbb{X} = [0, L]$). Let $u(x, t)$ be the density of the material at the point $x \in \mathbb{X}$ at time $t > 0$. Intuitively, we expect the material to flow from regions of *greater* to *lesser* concentration. In other words, we expect the *flow* of the material at any point in space to be proportional to the *slope* of the curve $u(x, t)$ at that point. Thus, if $F(x, t)$ is the flow at the point x at time t, then we expect the following:

$$F(x, t) = -\kappa \cdot \partial_x u(x, t),$$

where $\kappa > 0$ is a constant measuring the rate of diffusion. This is an example of *Fourier's law.*

1A(ii) ... in many dimensions

Prerequisites: Appendix E.

Figure 1A.2 depicts a material diffusing through a two-dimensional domain $\mathbb{X} \subset \mathbb{R}^2$ (e.g. heat spreading through a region, ink diffusing in a bucket of water, etc.) We

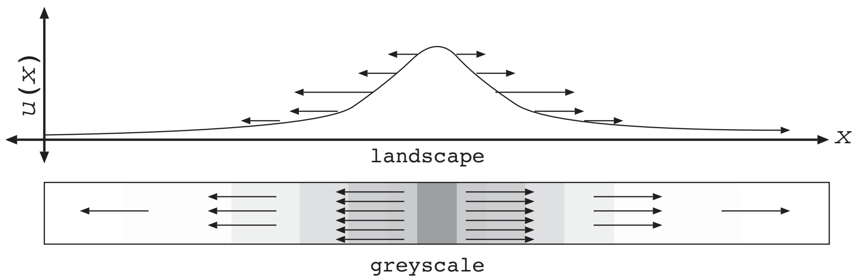

Figure 1A.1. Fourier's law of heat flow in one dimension.

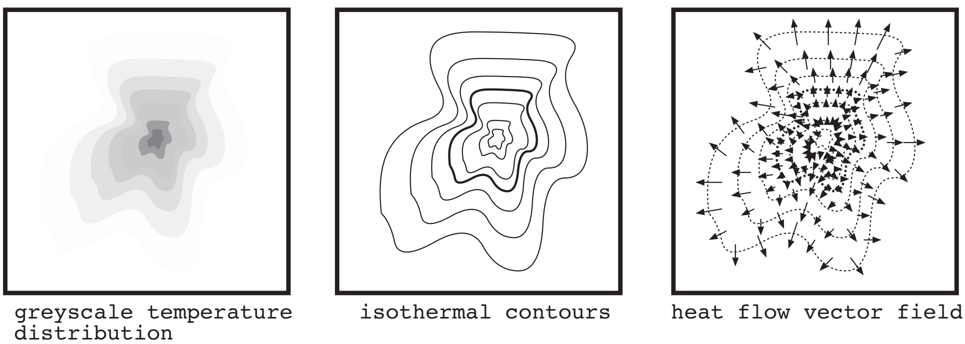

Figure 1A.2. Fourier's law of heat flow in two dimensions.

could just as easily suppose that $\mathbb{X} \subset \mathbb{R}^D$ is a D-dimensional domain. If $\mathbf{x} \in \mathbb{X}$ is a point in space, and $t \geq 0$ is a moment in time, let $u(\mathbf{x}, t)$ denote the concentration at $\mathbf{x}$ at time t. (This determines a function $u : \mathbb{X} \times \mathbb{R}_{\not -} \longrightarrow \mathbb{R}$, called a *time-varying scalar field*.)

Now let $\vec{\mathbf{F}}(\mathbf{x}, t)$ be a D-dimensional vector describing the *flow* of the material at the point $\mathbf{x} \in \mathbb{X}$. (This determines a *time-varying vector field* $\vec{\mathbf{F}} : \mathbb{R}^D \times \mathbb{R}_{\not -} \longrightarrow \mathbb{R}^D$.)

Again, we expect the material to flow from regions of high concentration to low concentration. In other words, material should flow *down the concentration gradient*. This is expressed by *Fourier's law of heat flow*:

$$\vec{\mathbf{F}} = -\kappa \cdot \nabla u,$$

where $\kappa > 0$ is a constant measuring the rate of diffusion.

One can imagine u as describing a distribution of highly antisocial people; each person is always fleeing everyone around them and moving in the direction with the fewest people. The constant κ measures the average walking speed of these misanthropes.

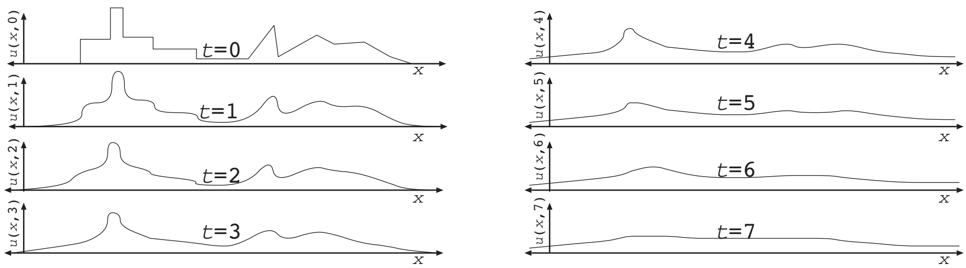

Figure 1B.1. The heat equation as 'erosion'.

1B The heat equation

Recommended: §1A.

1B(i) . . . in one dimension

Prerequisites: §1A(i).

Consider a material diffusing through a one-dimensional domain $\mathbb{X}$ (for example, $\mathbb{X} = \mathbb{R}$ or $\mathbb{X} = [0, L]$). Let $u(x, t)$ be the density of the material at the location $x \in \mathbb{X}$ at time $t \in \mathbb{R}_{\neq}$, and let $F(x, t)$ be the flux of the material at the location x and time t. Consider the derivative $\partial_x F(x, t)$. If $\partial_x F(x, t) > 0$, this means that the flow is *diverging*[1] at this point in space, so the material there is spreading farther apart. Hence, we expect the concentration at this point to *decrease*. Conversely, if $\partial_x F(x, t) < 0$, then the flow is *converging* at this point in space, so the material there is crowding closer together, and we expect the concentration to *increase*. To be succinct: the concentration of material will *increase* in regions where F converges and *decrease* in regions where F diverges. The equation describing this is given by

$$\partial_t u(x, t) = -\partial_x F(x, t).$$

If we combine this with Fourier's law, however, we get:

$$\partial_t u(x, t) = \kappa \cdot \partial_x \partial_x u(x, t),$$

which yields the *one-dimensional heat equation*:

$$\boxed{\partial_t u(x, t) = \kappa \cdot \partial_x^2 u(x, t).}$$

Heuristically speaking, if we imagine $u(x, t)$ as the height of some one-dimensional 'landscape', then the heat equation causes this landscape to be 'eroded', as if it were subjected to thousands of years of wind and rain (see Figure 1B.1).

[1] See Appendix E(ii), p. 562, for an explanation of why we say the flow is 'diverging' here.

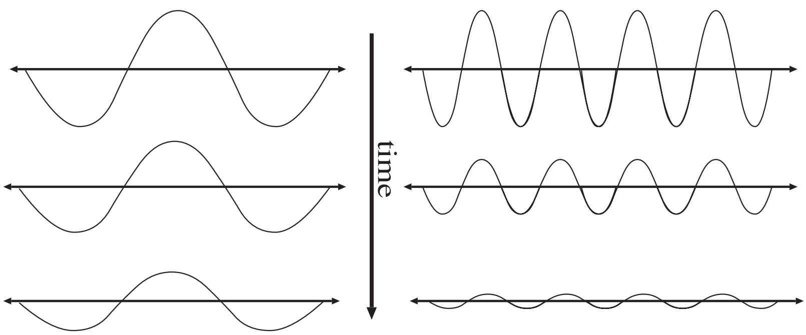

Figure 1B.2. Under the heat equation, the exponential decay of a periodic function is proportional to the square of its frequency. (a) Low frequency, slow decay; (b) high frequency, fast decay.

Example 1B.1 For simplicity we suppose $\kappa = 1$.

(a) Let $u(x, t) = e^{-9t} \cdot \sin(3x)$. Thus, u describes a spatially sinusoidal function (with spatial frequency 3) whose magnitude decays exponentially over time.
(b) The dissipating wave. More generally, let $u(x, t) = e^{-\omega^2 \cdot t} \cdot \sin(\omega \cdot x)$. Then u is a solution to the one-dimensional heat equation, and it looks like a standing wave whose amplitude decays exponentially over time (see Figure 1B.2). Note that the decay rate of the function u is proportional to the square of its frequency.
(c) The (one-dimensional) Gauss–Weierstrass kernel. Let

$$\mathcal{G}(x; t) := \frac{1}{2\sqrt{\pi t}} \exp\left(\frac{-x^2}{4t}\right).$$

Then $\mathcal{G}$ is a solution to the one-dimensional heat equation, and looks like a 'bell curve', which starts out tall and narrow, and, over time, becomes broader and flatter (Figure 1B.3). ◇

Ⓔ **Exercise 1B.1** Verify that all the functions in Examples 1B.1(a)–(c) satisfy the heat equation. ◆

All three functions in Example 1B.1 start out very tall, narrow, and pointy, and gradually become shorter, broader, and flatter. This is generally what the heat equation does; it tends to flatten things out. If u describes a physical landscape, then the heat equation describes 'erosion'.

1B(ii) . . . in many dimensions

Prerequisites: §1A(ii).

More generally, if $u : \mathbb{R}^D \times \mathbb{R}_{\not -} \longrightarrow \mathbb{R}$ is the time-varying density of some material, and $\vec{\mathbf{F}} : \mathbb{R}^D \times \mathbb{R}_{\not -} \longrightarrow \mathbb{R}$ is the flux of this material, then we would expect the

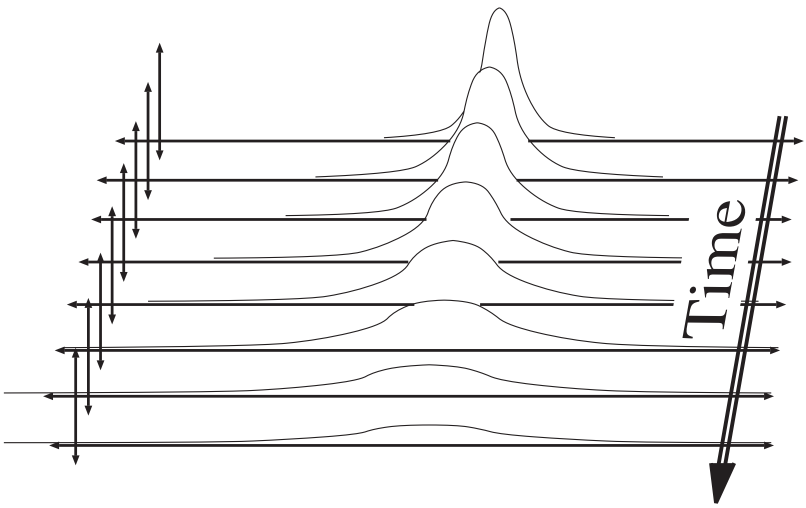

Figure 1B.3. The Gauss–Weierstrass kernel under the heat equation.

material to *increase* in regions where $\vec{\mathbf{F}}$ converges and to *decrease* in regions where $\vec{\mathbf{F}}$ diverges.[2] In other words, we have

$$\partial_t u = -\text{div}\,\vec{\mathbf{F}}.$$

If u is the density of some diffusing material (or heat), then $\vec{\mathbf{F}}$ is determined by Fourier's law, so we get the heat equation

$$\partial_t u = \kappa \cdot \text{div}\,\nabla u = \kappa \triangle u.$$

Here, $\triangle$ is the *Laplacian* operator,[3] defined as follows:

$$\boxed{\triangle u = \partial_1^2 u + \partial_2^2 u + \cdots + \partial_D^2 u}$$

Exercise 1B.2 (a) If $D = 1$ and $u : \mathbb{R} \longrightarrow \mathbb{R}$, verify that $\text{div}\,\nabla u(x) = u''(x) = \triangle u(x)$, for all $x \in \mathbb{R}$. Ⓔ

(b) If $D = 2$ and $u : \mathbb{R}^2 \longrightarrow \mathbb{R}$, verify that $\text{div}\,\nabla u(x, y) = \partial_x^2 u(x, y) + \partial_y^2 u(x, y) = \triangle u(x, y)$, for all $(x, y) \in \mathbb{R}^2$.

(c) For any $D \geq 2$ and $u : \mathbb{R}^D \longrightarrow \mathbb{R}$, verify that $\text{div}\,\nabla u(\mathbf{x}) = \triangle u(\mathbf{x})$, for all $\mathbf{x} \in \mathbb{R}^D$. ♦

By changing to the appropriate time units, we can assume $\kappa = 1$, so the heat equation becomes

$$\boxed{\partial_t u = \triangle u.}$$

[2] See Appendix E(ii), p. 562, for a review of the 'divergence' of a vector field.
[3] Sometimes the Laplacian is written as ∇^2.

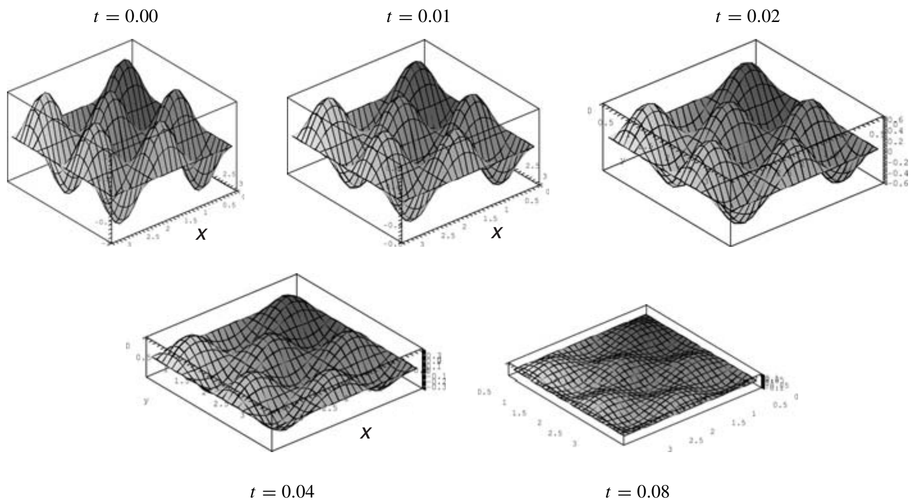

Figure 1B.4. Five snapshots of the function $u(x, y; t) = e^{-25t} \cdot \sin(3x)\sin(4y)$ from Example 1B.2.

For example,

- if $\mathbb{X} \subset \mathbb{R}$, and $x \in \mathbb{X}$, then $\triangle u(x;t) = \partial_x^2 u(x;t)$;
- if $\mathbb{X} \subset \mathbb{R}^2$, and $(x, y) \in \mathbb{X}$, then $\triangle u(x, y; t) = \partial_x^2 u(x, y; t) + \partial_y^2 u(x, y; t)$.

Thus, as we have already seen, the one-dimensional heat equation is given by

$$\partial_t u = \partial_x^2 u,$$

and the the two-dimensional heat equation is given by

$$\boxed{\partial_t u(x, y; t) = \partial_x^2 u(x, y; t) + \partial_y^2 u(x, y; t).}$$

Example 1B.2

(a) Let $u(x, y; t) = e^{-25t} \cdot \sin(3x)\sin(4y)$. Then u is a solution to the two-dimensional heat equation, and looks like a two-dimensional 'grid' of sinusoidal hills and valleys with horizontal spacing $1/3$ and vertical spacing $1/4$. As shown in Figure 1B.4, these hills rapidly subside into a gently undulating meadow and then gradually sink into a perfectly flat landscape.

(b) The (two-dimensional) Gauss–Weierstrass kernel. Let

$$\mathcal{G}(x, y; t) := \frac{1}{4\pi t} \exp\left(\frac{-x^2 - y^2}{4t}\right).$$

Then $\mathcal{G}$ is a solution to the two-dimensional heat equation, and looks like a mountain, which begins steep and pointy and gradually 'erodes' into a broad, flat, hill.

(c) The D-dimensional Gauss–Weierstrass kernel is the function $\mathcal{G} : \mathbb{R}^D \times \mathbb{R}_+ \longrightarrow \mathbb{R}$ defined by

$$\boxed{\mathcal{G}(\mathbf{x}; t) = \frac{1}{(4\pi t)^{D/2}} \exp\left(\frac{-\|\mathbf{x}\|^2}{4t}\right).}$$

Technically speaking, $\mathcal{G}(\mathbf{x}; t)$ is a D-dimensional *symmetric normal probability distribution* with variance $\sigma = 2t$. ◇

Exercise 1B.3 Verify that all the functions in Examples 1B.2(a)–(c) satisfy the heat equation. ◆ Ⓔ

Exercise 1B.4 Prove the *Leibniz rule* for Laplacians: if $f, g : \mathbb{R}^D \longrightarrow \mathbb{R}$ are two scalar fields, and $(f \cdot g) : \mathbb{R}^D \longrightarrow \mathbb{R}$ is their product, then, for all $\mathbf{x} \in \mathbb{R}^D$, Ⓔ

$$\triangle(f \cdot g)(\mathbf{x}) = g(\mathbf{x}) \cdot \Big(\triangle f(\mathbf{x})\Big) + 2\Big(\nabla f(\mathbf{x})\Big) \bullet \Big(\nabla g(\mathbf{x})\Big) + f(\mathbf{x}) \cdot \Big(\triangle g(\mathbf{x})\Big).$$

Hint: Combine the Leibniz rules for gradients and divergences (see Propositions E.1(b) and E.2(b) in Appendix E, pp. 562 and 564). ◆

1C The Laplace equation

Prerequisites: §1B.

If the heat equation describes the erosion/diffusion of some system, then an *equilibrium* or *steady-state* of the heat equation is a scalar field $h : \mathbb{R}^D \longrightarrow \mathbb{R}$ satisfying the Laplace equation:

$$\boxed{\triangle h \equiv 0.}$$

A scalar field satisfying the Laplace equation is called a *harmonic function*.

Example 1C.1

(a) If $D = 1$, then $\triangle h(x) = \partial_x^2 h(x) = h''(x)$; thus, the *one-dimensional Laplace equation* is just

$$\boxed{h''(x) = 0.}$$

Suppose $h(x) = 3x + 4$. Then $h'(x) = 3$ and $h''(x) = 0$, so h is harmonic. More generally, the one-dimensional harmonic functions are just the *linear* functions of the form $h(x) = ax + b$ for some constants $a, b \in \mathbb{R}$.

(b) If $D = 2$, then $\triangle h(x, y) = \partial_x^2 h(x, y) + \partial_y^2 h(x, y)$, so the two-dimensional Laplace equation is given by

$$\boxed{\partial_x^2 h + \partial_y^2 h = 0,}$$