

5 The Poisson and Laplace Equations

Until now, our focus has been very much on understanding how to differentiate and integrate functions of various types. But, with this under our belts, we can now take the next step and explore various differential equations that are written in the language of vector calculus. Our goal in this section is to find solutions to the Poisson equation and the related Laplace equation. This we will do in Section 5.2. But first we will explain why these equations underly two of the most important forces in the universe.

5.1 Gravity and Electrostatics

The first two fundamental forces to be discovered are also the simplest to describe mathematically. Newton's law of gravity states that two masses, m and M , separated by a distance r will experience a force

$$\mathbf{F}(r) = -\frac{GMm}{r^2}\hat{\mathbf{r}} \quad (5.1)$$

with G Newton's constant, a fundamental constant of nature that determines the strength of the gravitational force. Meanwhile, Coulomb's law states that two electric charges, q and Q , separated by a distance r will experience a force

$$\mathbf{F}(r) = \frac{Qq}{4\pi\epsilon_0 r^2}\hat{\mathbf{r}} \quad (5.2)$$

with the electric constant ϵ_0 a fundamental constant of nature that determines the inverse strength of the electrostatic force. The extra factor 4π reflects the fact that in the century between the Newton and Coulomb people had figured out where factors of 4π should sit in equations.

Most likely it will not have escaped your attention that these two equations are essentially the same. The only real difference is that overall minus sign which tells us that two masses always attract while two like charges repel. The question that we would like to ask is: why are the forces so similar?

Certainly it's not true that there is a deep connection between gravity and the electrostatic force, at least not one that we've uncovered to date. In particular, when masses and charges start to move, both the forces described above are replaced by something different and more complicated – general relativity in the case of gravity, the full Maxwell equations (3.7) in the case of the Coulomb force – and the equations of these theories are very different from each other. Yet, when we restrict to the simple, static set-up, the forces take the same form.

The reason for this is twofold. First, both forces are described by fields. Second, space has three dimensions. The purpose of this section is to explain this in more detail. And, for this, we need the tools of vector calculus.

5.1.1 Gauss' Law

Each of the force equations (5.1) and (5.2) contains some property that characterises the force: mass for gravity and electric charge for the electrostatic force. For our purposes, it will be useful to focus on one of the particles that carries mass m and charge q . We call this a *test particle*, meaning that we'll look at how this particle is buffeted by various forces but won't, in turn, consider its effect on any other particle. Physically, this is appropriate if $m \ll M$ and $q \ll Q$. Then it is useful to write the equation in a way that separates the properties of the test particle from the other. The force experienced by the test particle is

$$\mathbf{F}(\mathbf{x}) = m\mathbf{g}(\mathbf{x}) + q\mathbf{E}(\mathbf{x})$$

where $\mathbf{g}(\mathbf{x})$ is the gravitational and $\mathbf{E}(\mathbf{x})$ is the electric field. Clearly Newton's law is telling us that a particle of mass M sets up a gravitational field

$$\mathbf{g}(\mathbf{x}) = -\frac{GM}{r^2}\hat{\mathbf{r}} \quad (5.3)$$

while a particle with electric charge Q sets up an electric field

$$\mathbf{E}(\mathbf{x}) = \frac{Q}{4\pi\epsilon_0 r^2}\hat{\mathbf{r}} \quad (5.4)$$

So far this is just a trivial rewriting of the force laws. However, we will now reframe these force laws in the language of vector calculus. Instead of postulating the $1/r^2$ force laws (5.3) and (5.4), we will replace them by two properties of the fields from which everything else follows. Here we specify the first property; the second will be explained in Section 5.1.2.

The first property is that if you integrate the relevant field over a closed surface, then it captures the amount of "stuff" inside this surface. For the gravitational field, this stuff is mass

$$\int_S \mathbf{g} \cdot d\mathbf{S} = -4\pi GM \quad (5.5)$$

while for the electric field it is charge

$$\int_S \mathbf{E} \cdot d\mathbf{S} = \frac{Q}{\epsilon_0} \quad (5.6)$$

Again, the difference in minus sign signals the important attractive/repulsive difference between the two forces. In contrast, the factors of $4\pi G$ and $1/\epsilon_0$ are simply convention for how we characterise the strength of the fields. These two equations are known as *Gauss' law*. Or, more precisely, “Gauss' law in integrated form”. We'll see the other form below.

Examples


For concreteness, let's focus on the gravitational field. We will take a sphere of radius R and total mass M . We will require that the density of the sphere is spherically symmetric, but not necessarily constant. The spherical symmetry of the problem then ensures that the gravitational field itself is spherically symmetric, with $\mathbf{g}(\mathbf{x}) = g(r)\hat{\mathbf{r}}$. If we then integrate the gravitational field over any spherical surface S of radius $r > R$, we have

$$\int_S \mathbf{g} \cdot d\mathbf{S} = \int_S g(r) dS = 4\pi r^2 g(r)$$

where we recognise $4\pi r^2$ as the area of the sphere. From Gauss' law (5.5) we then have

$$\mathbf{g}(r) = -\frac{GM}{r^2} \hat{\mathbf{r}} \quad (5.7)$$

This reproduces Newton's force law (5.1). Note, however, that we've extended Newton's law beyond the original remit of point particles: the gravitational field (5.7) holds for any spherically symmetric distribution of mass, provided that we're outside this mass. For example, it tells us that the gravitational field of the Earth (at least assuming spherical symmetry) is indistinguishable from the gravitational field of a point-like particle with the same mass, sitting at the origin. This way of solving for the vector field is known as the *Gauss flux method*.



Another rather cute consequence of this is that, at least for spherically symmetric mass distributions, you don't feel the mass outside you. According to Gauss' law, the gravitational field at any point is determined only by what lies inside a sphere of a given radius. So if, for example, you were able to hollow out the centre of a planet (unlikely, admittedly) then anyone living there would feel no gravitational force from the mass that surrounds them.


For our second example, we turn to the electric field. Consider an infinite line of charge, with charge per unit length σ . This situation is crying out for cylindrical polar coordinates. Until now, we've always called the radial direction in cylindrical polar coordinates ρ but, for reasons that will become clear shortly, for this example alone we will instead call the radial direction r as shown in the figure. The symmetry of the problem shows that the electric field is radial so takes the form $\mathbf{E}(r) = E(r)\hat{\mathbf{r}}$. Integrating over cylinder S of radius r and length L we have

$$\int_S \mathbf{E} \cdot d\mathbf{S} = 2\pi r L E(r)$$

where there is no contribution from the end caps because $\mathbf{n} \cdot \mathbf{E} = 0$ there, with \mathbf{n} the normal vector. The total charge inside this surface is $Q = \sigma L$. From Gauss' law (5.6), we then have the electric field

$$\mathbf{E}(r) = \frac{\sigma}{2\pi\epsilon_0 r} \hat{\mathbf{r}}$$

Note that the $1/r$ behaviour arises because the symmetry of the problem ensures that the electric field lies in a plane. Said differently, the electric field from an infinite charged line is the same as we would get from a point particle in a flatland world of two dimensions.



More generally, if space were \mathbb{R}^n , then the Gauss' law equations (5.5) and (5.6) would still be the correct description of the gravitational and electric fields. Repeating the calculations above would then tell us that a point charge gives rise to an electric field

$$\mathbf{E}(r) = \frac{1}{A_{n-1}\epsilon_0 r^{n-1}} \hat{\mathbf{r}}$$

where $A_n r^n$ is the “surface area” of an n -dimensional sphere S^n of radius r . (For what it's worth, the prefactor is $A_{n-1} = 2\pi^{n/2}/\Gamma(n/2)$ where $\Gamma(x)$ is the gamma function which coincides with the factorial function $\Gamma(x) = (x-1)!$ when x is integer.) For the rest of this section, we'll keep our feet firmly in \mathbb{R}^3 .

Gauss' Law Again

There's a useful way to rewrite the Gauss' law equations (5.5) and (5.6). For the gravitational field, we introduce the *density*, or mass per unit volume, $\rho(\mathbf{x})$. Invoking the divergence theorem then, for any volume V bounded by S , we have

$$\int_V \nabla \cdot \mathbf{g} dV = \int_S \mathbf{g} \cdot d\mathbf{S} = -4\pi G M = -4\pi G \int_V \rho(\mathbf{x}) dV$$

But, rearranging, we have

$$\int_V (\nabla \cdot \mathbf{g} + 4\pi G\rho(\mathbf{x})) dV = 0$$

for any volume V . This can only hold if the integrand itself vanishes, so we must have

$$\nabla \cdot \mathbf{g} = -4\pi G\rho(\mathbf{x}) \quad (5.8)$$

This is also known as *Gauss' law* for the gravitational field, now in differential form. The equivalence with the earlier integrated form (5.5) follows, as above, from the divergence theorem.

We can apply the same manipulations to the electric field. This time we introduce the *charge density* $\rho_e(\mathbf{x})$. We then get Gauss' law in the form

$$\nabla \cdot \mathbf{E} = \frac{\rho_e(\mathbf{x})}{\epsilon_0} \quad (5.9)$$

This is the first of the Maxwell equations (3.7). (In our earlier expression, we denoted the charge density as $\rho(\mathbf{x})$. Here we've added the subscript ρ_e to distinguish it from mass density.) The manipulations that we've described above show that Gauss' law is a grown-up version of the Coulomb force law (5.2).

5.1.2 Potentials

In our examples above, we used symmetry arguments to figure out the direction in which the gravitational and electric fields are pointing. But in many situations we don't have that luxury. In that case, we need to invoke the second important property of these vector fields: they are both conservative.

Recall that, by now, we have a number of different ways to talk about conservative vector fields. Such fields are necessarily irrotational $\nabla \times \mathbf{g} = \nabla \times \mathbf{E} = 0$. Furthermore, their integral vanishes when integrated around any closed curve C ,

$$\oint_C \mathbf{g} \cdot d\mathbf{x} = \oint_C \mathbf{E} \cdot d\mathbf{x} = 0$$

You can check that both of these hold for the examples, such as the $1/r^2$ field, that we discussed above (as long as the path C avoids the singular point at the origin).

Here the key property of a conservative vector field is that it can be written in terms of an underlying scalar field,

$$\mathbf{g} = -\nabla\Phi \quad \text{and} \quad \mathbf{E} = -\nabla\phi \quad (5.10)$$

where $\Phi(\mathbf{x})$ is the gravitational potential and $\phi(\mathbf{x})$ the electrostatic potential. Note the additional minus signs in these definitions. We saw in the discussion around (1.18) that the existence of such potentials ensures that test particles experiencing these forces have a conserved energy:

$$\text{energy} = \frac{1}{2}m\dot{\mathbf{x}}^2 + m\Phi(\mathbf{x}) + q\phi(\mathbf{x})$$

Combining the differential form of the Gauss' law (5.8) and (5.9) with the existence of the potentials (5.10), we find that the gravitational and electric fields are determined, in general, by solutions to the following equations

$$\nabla^2\Phi = 4\pi G\rho(\mathbf{x}) \quad \text{and} \quad \nabla^2\phi = -\frac{\rho_e(\mathbf{x})}{\epsilon_0}$$

Equations of this type are known as the *Poisson equation*. In the special case where the “source” $\rho(\mathbf{x})$ on the right-hand side vanishes, this reduces to the *Laplace equation*, for example

$$\nabla^2\Phi = 0$$

These two equations are commonplace in mathematics and physics. Here we have derived them in the context of gravity and electrostatics, but their applications spread much further.

To give just one further example, in [Fluid Mechanics](#) the motion of the fluid is described by a velocity field $\mathbf{u}(\mathbf{x})$. If the flow is irrotational, then $\nabla \times \mathbf{u} = 0$ and the velocity can be described by a potential function $\mathbf{u} = \nabla\phi$. If, in addition, the fluid is incompressible then $\nabla \cdot \mathbf{u} = 0$ and we once again find ourselves solving the Laplace equation $\nabla^2\phi = 0$.

5.2 The Poisson and Laplace Equations

In the rest of this section we will develop some methods to solve the Poisson equation. We change notation and call the potential $\psi(\mathbf{x})$ (to avoid confusion with the polar angle ϕ). We are then looking for solutions to

$$\nabla^2\psi(\mathbf{x}) = -\rho(\mathbf{x})$$

The goal is to solve for $\psi(\mathbf{x})$ given a “source” $\rho(\mathbf{x})$. As we will see, the domain in which $\psi(\mathbf{x})$ lives, together with associated boundary conditions, also plays an important role in determining $\psi(\mathbf{x})$.

The Laplace equation $\nabla^2\psi = 0$ is linear. This means that if $\psi_1(\mathbf{x})$ is a solution and $\psi_2(\mathbf{x})$ is a solution, then so too is $\psi_1(\mathbf{x}) + \psi_2(\mathbf{x})$. Any solution to the Laplace equation acts as a complementary solution to the Poisson equation. This should then be accompanied by a particular solution for a given source $\rho(\mathbf{x})$ on the right-hand side.

5.2.1 Isotropic Solutions

Both the Laplace and Poisson equations are partial differential equation. Life is generally much easier if we're asked to solve ordinary differential equations rather than partial differential equations. For the Poisson equation, this is what we get if we have some kind of symmetry, typically one aligned to some polar coordinates.

For example, if we have spherical symmetry then we can look for solutions of the form $\psi(\mathbf{x}) = \psi(r)$. Using the form of the Laplacian (3.15), Laplace equation becomes

$$\begin{aligned}\nabla^2\psi = 0 &\Rightarrow \frac{d^2\psi}{dr^2} + \frac{2}{r} \frac{d\psi}{dr} = \frac{1}{r^2} \frac{d}{dr} \left(r^2 \frac{d\psi}{dr} \right) = 0 \\ &\Rightarrow \psi(r) = \frac{A}{r} + B\end{aligned}\tag{5.11}$$

for some constants A and B . Clearly the A/r solution diverges as $r \rightarrow 0$ so we should be cautious in claiming that this solves the Laplace equation at $r = 0$. (We will shortly see that it doesn't, but it does solve a related Poisson equation.) Note that the solution A/r is relevant in gravity or in electrostatics, where $\psi(r)$ has the interpretation as the potential for a point charge.

Meanwhile, in cylindrical polar coordinates we will also denote the radial direction as r to avoid confusion with the source ρ in the Poisson equation. The Laplace equation becomes

$$\begin{aligned}\nabla^2\psi = 0 &\Rightarrow \frac{d^2\psi}{dr^2} + \frac{1}{r} \frac{d\psi}{dr} = \frac{1}{r} \frac{d}{dr} \left(r \frac{d\psi}{dr} \right) = 0 \\ &\Rightarrow \psi(r) = A \log r + B\end{aligned}\tag{5.12}$$

This again diverges at $r = 0$, this time corresponding to the entire z axis.

Note that if we ignore the z direction, as we have above, then cylindrical polar coordinates are the same thing as 2d polar coordinates, and the log form is the rotationally invariant solution to the Laplace equation in \mathbb{R}^2 . In general, in \mathbb{R}^n , the non-constant solution to the Laplace equation is $1/r^{n-2}$. The low dimensions of \mathbb{R}^2 and \mathbb{R} are special because the solution grows asymptotically as $r \rightarrow \infty$, while for \mathbb{R}^n with $n \geq 3$, the rotationally invariant solution to the Laplace equation decays to a constant asymptotically.

If $\psi(r)$ is a solution to the Laplace equation, then so too is any derivative of $\psi(r)$. For example, if we take the spherically symmetric solution $\psi(r) = 1/r$, then we can construct a new solution

$$\psi_{\text{dipole}}(\mathbf{x}) = \mathbf{d} \cdot \nabla \left(\frac{1}{r} \right) = -\frac{\mathbf{d} \cdot \mathbf{x}}{r^3}$$

for any constant vector \mathbf{d} and, again, with $r \neq 0$. This kind of solution is important in electrostatics where it arises as the large distance solution for a *dipole*, two equal and opposite charges at a fixed distance apart.

Discontinuities and Boundary Conditions

In many situations, we must specify some further data when solving the Poisson equations. Typically this is some kind of boundary condition and, in some circumstances, a requirement of continuity and smoothness on the solution.

This can be illustrated with a simple example. Suppose that we are looking for a spherically symmetric solution to:

$$\nabla^2 \psi = \begin{cases} -\rho_0 & r \leq R \\ 0 & r > R \end{cases}$$

with ρ_0 constant. We will further ask that $\psi(r=0)$ is non-singular, that $\psi(r) \rightarrow 0$ as $r \rightarrow \infty$, and that $\psi(\mathbf{x})$ and $\psi'(\mathbf{x})$ are continuous. We will now see that all of these conditions give us a unique solution.

First look inside $r \leq R$. As we mentioned above, a solution to the Poisson equation can be found by adding a complementary solution and a particular solution. Since we're looking for a spherically symmetric particular solution, we can restrict our ansatz to $\psi(r) = r^p$ for some p . It's simple to check that $\nabla^2 r^p = p(p+1)r^{p-2}$. This then gives us the general solution

$$\psi(r) = \frac{A}{r} + B - \frac{1}{6}\rho_0 r^2 \quad r \leq R$$

But now we can start killing some terms by invoking the boundary conditions. In particular, the requirement that $\psi(r)$ is non-singular at $r = 0$ tells us that we must have $A = 0$. Meanwhile, outside $r > R$ the most general solution is

$$\psi(r) = \frac{C}{r} + D$$




Figure 17. The plot of $\Phi = -4\pi G\psi$ on the left, with the radius $R = 1$ the cross over point. This is more apparent in the gravitational field $g = -\Phi'$ shown on the right.

Now we must have $D = 0$ if $\psi(r) \rightarrow 0$ as $r \rightarrow \infty$. To finish, we must patch these two solutions at $r = R$, invoking continuity

$$\psi(r = R) = B - \frac{1}{6}\rho_0 R^2 = \frac{C}{R}$$

and smoothness

$$\psi'(r = R) = -\frac{1}{3}\rho_0 R = -\frac{C}{R^2}$$

These determine our last two unknown constants, B and C . Putting this together, we have a unique solution

$$\psi(r) = \begin{cases} \frac{1}{6}\rho_0(3R^2 - r^2) & r \leq R \\ \frac{1}{3}\rho_0 R^3/r & r > R \end{cases}$$

This example has application for the gravitational potential $\Phi = -4\pi G\psi$ of a planet of radius R and density ρ_0 . The plot of Φ is shown on the left of Figure 17; the plot of the gravitational field $g = -d\Phi/dr$ is on the right, where we see a linear increase inside the planet, before we get to the more familiar $1/r^2$ fall-off.

5.2.2 Some General Results

So far our solutions to the Poisson equation take place in \mathbb{R}^3 . (Or, more precisely, $\mathbb{R}^3 - \{0, 0\}$ for the $1/r$ solution (5.11) and $\mathbb{R}^3 - \mathbb{R}$ for the $\log r$ solution (5.12).) In general, we may want to solve the Poisson or Laplace equations $\nabla^2\psi = -\rho$ in some bounded region V . In that case, we must specify boundary conditions on ∂V .

There are two common boundary conditions:

- Dirichlet condition: We fix $\psi(\mathbf{x}) = f(\mathbf{x})$ for some specific $f(\mathbf{x})$ on ∂V .

- Neumann condition: We fix $\mathbf{n} \cdot \nabla \psi(\mathbf{x}) = g(\mathbf{x})$ for some specific $g(\mathbf{x})$ on ∂V , where \mathbf{n} is the outwardly pointing normal of ∂V .

The Neumann boundary condition is sometimes specified using the slightly peculiar notation $\partial\psi/\partial\mathbf{n} := \mathbf{n} \cdot \nabla\psi$. Or even, sometimes, $\partial\psi/\partial n$. We have the following statement of uniqueness:

Claim: Consider the Poisson equation on a bounded region V , with either Dirichlet or Neumann boundary conditions specified on each boundary ∂V . If a solution exists, then it is unique. (In the case of Neumann boundary conditions everywhere, the solution is only unique up to a constant.)

Proof: Let $\psi_1(\mathbf{x})$ and $\psi_2(\mathbf{x})$ both satisfy the Poisson equation with the specified boundary conditions. Then $\psi(\mathbf{x}) = \psi_1 - \psi_2$ obeys $\nabla^2\psi = 0$ and either $\psi = 0$ or $\mathbf{n} \cdot \nabla\psi = 0$ on ∂V . Then consider

$$\int_V \nabla \cdot (\psi \nabla \psi) dV = \int_V (\nabla \psi \cdot \nabla \psi + \psi \nabla^2 \psi) dV = \int_V |\nabla \psi|^2 dV$$

But by the divergence theorem, we have

$$\int_V \nabla \cdot (\psi \nabla \psi) dV = \int_{\partial V} \psi \nabla \psi \cdot d\mathbf{S} = \int_{\partial V} \psi (\mathbf{n} \cdot \nabla \psi) dS = 0$$

where either Dirichlet or Neumann boundary conditions set the boundary term to zero. Because $|\nabla\psi|^2 \geq 0$, the integral can only vanish if $\nabla\psi = 0$ everywhere in V , so ψ must be constant. If Dirichlet boundary conditions are imposed anywhere, then that constant must be zero. \square

This result means that if we can find any solution – say an isotropic solution, or perhaps a separable solution of the form $\psi(\mathbf{x}) = \Phi(r)Y(\theta)$ – then this must be the unique solution. By considering the limit of large spheres, it is also possible to extend the proof to solutions on \mathbb{R}^3 , with the boundary condition $\psi(\mathbf{x}) \rightarrow 0$ suitably quickly as $r \rightarrow \infty$.

Note, however, that this doesn't necessarily tell us that a solution exists. For example, suppose that we wish to solve the Poisson equation $\nabla^2\psi = \rho(\mathbf{x})$ with a fixed Neumann boundary condition $\mathbf{n} \cdot \nabla\psi = g(\mathbf{x})$ on ∂V . Then there can only be a solution provided that there is a particular relationship between ρ and g ,

$$\int_V \nabla^2\psi dV = \int_{\partial V} \nabla\psi \cdot d\mathbf{S} \iff \int_V \rho dV = \int_S g dS$$

In other situations, there may well be other requirements.

If the region V has several boundaries, it's quite possible to specify a different type of boundary condition on each, and the uniqueness statement still holds. This kind of problem arises in electromagnetism where you solve for the electric field in the presence of a bunch of “conductors” (for now, conductors just means a chunk of metal). The electric field vanishes inside a conductor because, if it didn't, then the electric charges inside would move around until they created a counterbalancing field. So any attempt to solve for the electric field outside the conductors must take this into account by imposing certain boundary conditions on the surface of the conductor. It turns out that both Dirichlet and Neumann boundary conditions are important here. If the conductor is “grounded”, meaning that it is attached to some huge reservoir of charge like the Earth, then it sits at some fixed potential, typically $\psi = 0$. This is a Dirichlet boundary condition. In contrast, if the conductor is isolated and carries some non-vanishing charge then it will act as a source of electric field, but this field is always emitted perpendicular to the boundary. This, then, specifies $\mathbf{n} \cdot \mathbf{E} = -\mathbf{n} \cdot \nabla\psi$, giving Neumann boundary conditions. You can learn more about this in the lectures on [Electromagnetism](#).

Green's Identities

The proof of the uniqueness theorem used a trick known as Green's (first) identity, namely

$$\int_V \phi \nabla^2 \psi \, dV = - \int_V \nabla \phi \cdot \nabla \psi \, dV + \int_S \phi \nabla \psi \cdot d\mathbf{S}$$

This is essentially a 3d version of integration by parts and it follows simply by applying the divergence theorem to $\phi \nabla \psi$. We used it in the above proof with $\phi = \psi$, but the more general form given above is sometimes useful, as is a related formula that follows simply by anti-symmetrisation,

$$\int_V (\phi \nabla^2 \psi - \psi \nabla^2 \phi) \, dV = \int_S (\phi \nabla \psi - \psi \nabla \phi) \cdot d\mathbf{S}$$

This is known as Green's second identity.

Harmonic Functions

Solutions to the Laplace equation

$$\nabla^2 \psi = 0$$

arise in many places in mathematics and physics. These solutions are so special that they get their own name: they are called *harmonic functions*. Here are two properties

of these functions

Claim: Suppose that ψ is harmonic in a region V that includes the solid sphere with boundary $S_R : |\mathbf{x} - \mathbf{a}| = R$. Then the value of ψ at \mathbf{a} , the centre of the sphere, is given by $\psi(\mathbf{a}) = \bar{\psi}(R)$ where

$$\bar{\psi}(R) = \frac{1}{4\pi R^2} \int_{S_R} \psi(\mathbf{x}) dS$$

is the average of ψ over S_R . This is known as the *mean value property*.

Proof: In spherical polar coordinates centred on \mathbf{a} , the area element is $dS = r^2 \sin \theta d\theta d\phi$, so

$$\bar{\psi}(r) = \frac{1}{4\pi} \int d\phi \int d\theta \sin \theta \psi(r, \theta, \phi)$$

and

$$\begin{aligned} \frac{d\bar{\psi}(R)}{dr} &= \frac{1}{4\pi} \int d\phi \int d\theta \sin \theta \frac{\partial \psi(R)}{\partial r} = \frac{1}{4\pi R^2} \int_{S_R} \frac{\partial \psi(R)}{\partial r} dS \\ &= \frac{1}{4\pi R^2} \int_{S_R} \nabla \psi \cdot d\mathbf{S} = \int_{\text{Ball}} \nabla^2 \psi dV = 0 \end{aligned}$$

But clearly $\bar{\psi}(R) \rightarrow \psi(\mathbf{a})$ as $R \rightarrow 0$ so we must have $\bar{\psi}(R) = \psi(\mathbf{a})$ for all R . □

Claim: A harmonic function can have neither a maximum nor minimum in the interior of a region V . Any maximum or minimum must lie on the boundary ∂V .

Proof: If ψ has a local maximum at \mathbf{a} in V then there exists an ϵ such that $\psi(\mathbf{x}) < \psi(\mathbf{a})$ for all $|\mathbf{x} - \mathbf{a}| < \epsilon$. But, we know that $\bar{\psi}(R) = \psi(\mathbf{a})$ and this contradicts the assumption for any $0 < R < \epsilon$. □

This is consistent with our standard analysis of maxima and minima. Usually we would compute the eigenvalues λ_i of the Hessian $\partial^2 \psi / \partial x^i \partial x^j$. For a harmonic function $\nabla^2 \psi = \partial^2 \psi / \partial x^i \partial x^i = 0$. Since the trace of the Hessian vanishes, we must have eigenvalues of opposite sign since $\sum_i \lambda_i = 0$. Hence, any stationary point must be a saddle. Note that this standard analysis is inconclusive when $\lambda_i = 0$, but the argument using the mean value property closes this loophole.

5.2.3 Integral Solutions

There is a particularly nice way to write down an expression for the general solution to the Poisson equation in \mathbb{R}^3 , with

$$\nabla^2\psi = -\rho(\mathbf{x}) \quad (\text{🐟})$$

at least for a localised source $\rho(\mathbf{x})$ that drops off suitably fast, so $\rho(\mathbf{x}) \rightarrow 0$ as $r \rightarrow \infty$.

To this end, let's look back to what is, perhaps, our simplest "solution",

$$\psi(\mathbf{x}) = \frac{\lambda}{4\pi r} \quad (5.13)$$

for some constant λ . The question we want to ask is: what equation does this actually solve?! We've seen in (5.11) that it solves the Laplace equation $\nabla^2\psi = 0$ when $r \neq 0$. But clearly something's going on at $r = 0$ because the function diverges there. In the language of physics, we would say that there is a point particle sitting at $r = 0$, carrying some mass or charge, giving rise to this potential. What is the correct mathematical way of capturing this?

To see that there must be something going on at $r = 0$, let's replay the kind of Gauss flux games that we met in Section 5.1. We integrate $\nabla^2\psi$, with ψ given by (5.13), over a volume V which we take to be a spherical region of radius R , to find

$$\int_V \nabla^2\psi \, dV = \int_S \nabla\psi \cdot d\mathbf{S} = -\lambda$$

Comparing to (🐟), we see that the function (5.13) must solve the Poisson equation with a source and this source must obey

$$\int_V \rho(\mathbf{x}) \, dV = \lambda$$

This makes sense physically, since $\int \rho dV$ is the total mass, or total charge, which does indeed determine the overall scaling λ of the potential. But what mathematical function obeys $\rho(\mathbf{x}) = 0$ for all $\mathbf{x} \neq 0$ yet, when integrated over all space, gives a non-vanishing constant λ ?

The answer is that $\rho(\mathbf{x})$ must be proportional to the 3d *Dirac delta function*,

$$\rho(\mathbf{x}) = \lambda \delta^3(\mathbf{x})$$

The Dirac delta function should be thought of as an infinitely narrow spike, located at the origin. It has the properties

$$\delta^3(\mathbf{x}) = 0 \quad \text{for } \mathbf{x} \neq 0$$

and, when integrated against any function $f(\mathbf{x})$ over any volume V that includes the origin, it gives

$$\int_V f(\mathbf{x}) \delta^3(\mathbf{x}) dV = f(\mathbf{x} = \mathbf{0})$$


The superscript in $\delta^3(\mathbf{x})$ is there to remind us that the delta function should be integrated over a 3-dimensional volume before it yields something finite. In particular, when integrated against a constant function, we get a measure of the height of the spike,

$$\int_V \delta^3(\mathbf{x}) dV = 1$$

The Dirac delta function is an example of a generalised function, also known as a distribution. And it is exactly what we need to source the solution $\psi \sim 1/r$. We learn that the function (5.13) is not a solution to the Laplace equation, but rather a solution to the Poisson equation with a delta function source

$$\nabla^2 \psi = -\lambda \delta^3(\mathbf{x}) \quad \Rightarrow \quad \psi(\mathbf{x}) = \frac{\lambda}{4\pi r} \quad (5.14)$$

With this important idea in hand, we can now do something quite spectacular: we can use it to write down an expression for a solution to the general Poisson equation.

Claim: The Poisson equation () has the integral solution

$$\psi(\mathbf{x}) = \frac{1}{4\pi} \int_{V'} \frac{\rho(\mathbf{x}')}{|\mathbf{x} - \mathbf{x}'|} dV' \quad (5.15)$$

where the integral is over a region V' parameterised by \mathbf{x}' .

Proof: First, some simple intuition behind this formula. A point particle at \mathbf{x}' gives rise to a potential of the form $\psi(\mathbf{x}) = \rho(\mathbf{x}')/4\pi|\mathbf{x} - \mathbf{x}'|$, which is just our solution (5.14), translated from the origin to point \mathbf{x}' . The integral solution (5.15) then just takes advantage of the linear nature of the Poisson equation and sums a whole bunch of these solutions.

The technology of the delta function allows us to make this precise. We can evaluate

$$\nabla^2\psi = \frac{1}{4\pi} \int_{V'} \rho(\mathbf{x}') \nabla^2 \left(\frac{1}{|\mathbf{x} - \mathbf{x}'|} \right) dV'$$

where you have to remember that ∇^2 differentiates \mathbf{x} and cares nothing for \mathbf{x}' . We then have the result

$$\nabla^2 \frac{1}{|\mathbf{x} - \mathbf{x}'|} = -4\pi\delta^3(\mathbf{x} - \mathbf{x}')$$

which is just a repeat of (5.14), but with the location of the source translated from the origin to the new point \mathbf{x}' . Using this, we can continue our proof

$$\nabla^2\psi = - \int_{V'} \rho(\mathbf{x}') \delta^3(\mathbf{x} - \mathbf{x}') dV' = -\rho(\mathbf{x})$$

which is what we wanted to show. □

The technique of first solving an equation with a delta function source and subsequently integrating to find the general solution is known as the *Green's function* approach. It is a powerful method to solve differential equations and we will meet it again in many further courses.

6 Tensors

A famously annoying definition of a tensor is:

A tensor is something whose components transform like a tensor

This becomes even more annoying when you appreciate that this is, in fact, one of the better definitions of a tensor. The purpose of this section is to explain why this definition is not as dumb as it sounds and to give some insight into what it means to be a tensor.

Very roughly speaking, tensors are generalisations of objects like vectors and matrices. In index notation, a vector has a single index while a matrix has two indices. A tensor is an object with any number of indices, something like $T_{ij\dots k}$.

However, this simplistic description hides the most important property of a tensor. Vectors, matrices and, more generally, tensors are more than just a list of numbers. Instead, those numbers should be thought of as a useful way of characterising the underlying object and, because of this, inherit some properties of that underlying object. As we will see, the key property is how the list of numbers transform under a change of basis.

We will start by explaining this in more detail, firstly with vectors and then building up to the definition of a tensor. Initially we will keep the discussion restricted to some (admittedly rather dry) mathematical formalism. Then, in Section 6.2 we will describe some physical examples.

6.1 What it Takes to Make a Tensor

Not any list of n numbers constitutes a vector in \mathbb{R}^n . Or, said more precisely, not any list of n numbers constitutes the components of a vector in \mathbb{R}^n . For example, if you write down the heights of the first three people you met this morning, that doesn't make a vector in \mathbb{R}^3 . Instead, a vector comes with certain responsibilities. In particular, the components describe an underlying object which should be independent of the choice of basis. As we now explain, that means that the components should transform in the right way under rotations.

We consider a point $\mathbf{x} \in \mathbb{R}^n$. If we wish to attach some coordinates to this point, we first need to introduce a set of basis vectors $\{\mathbf{e}_i\}$ with $i = 1, \dots, n$. We will take these to be orthonormal, meaning that $\mathbf{e}_i \cdot \mathbf{e}_j = \delta_{ij}$. Any vector can then be expressed as

$$\mathbf{x} = x_i \mathbf{e}_i \tag{6.1}$$

Usually we conflate the components $x_i = (x_1, \dots, x_n)$ with the “vector”. But, for our purposes, we should remember that these are just a useful way of representing the more abstract object \mathbf{x} . In particular, we’re entirely at liberty to take a different set of basis vectors,

$$\mathbf{e}'_i = R_{ij}\mathbf{e}_j$$

If we ask that \mathbf{e}'_i are also orthonormal, so $\mathbf{e}'_i \cdot \mathbf{e}'_j = \delta_{ij}$, then we have

$$\mathbf{e}'_i \cdot \mathbf{e}'_j = R_{ik}R_{jl}\mathbf{e}_k \cdot \mathbf{e}_l = R_{ik}R_{jk} = \delta_{ij}$$

or, in matrix notation,

$$RR^T = \mathbf{1}$$

Matrices of this kind are said to be *orthogonal*. We write $R \in O(n)$. Taking the determinant, we have $\det R = \pm 1$. Those matrices with $\det R = +1$ correspond to rotations and are said to be *special orthogonal*. We write $R \in SO(n)$. In \mathbb{R}^3 , a rotation $R \in SO(3)$ takes a right-handed orthonormal basis into another right-handed orthonormal basis. Those matrices with $\det R = -1$ correspond to a rotation together with a reflection and take a right-handed basis to a left-handed basis.

Under a change of basis, the vector \mathbf{x} itself doesn’t change. But its components do. We have

$$\mathbf{x} = x_i\mathbf{e}_i = x'_i\mathbf{e}'_i = x'_iR_{ij}\mathbf{e}_j$$

So the components transform under the same rotation matrix R ,

$$x_j = R_{ij}x'_i \quad \Rightarrow \quad x'_i = R_{ij}x_j \tag{6.2}$$

A *tensor* T is a generalisation of these ideas to an object with more indices. Just as the vector \mathbf{x} has an identity independent of any choice of basis, so too does the tensor T . But when measured with respect to a chosen basis $\{\mathbf{e}_i\}$, a *tensor of rank p* has components $T_{i_1\dots i_p}$. When we change the basis using (6.1), the tensor transforms as

$$T'_{i_1\dots i_p} = R_{i_1j_1}\dots R_{i_pj_p}T_{j_1\dots j_p} \tag{6.3}$$

This is known as the *tensor transformation rule*. A tensor of rank p is sometimes referred to simply as a p -tensor.

The simplest examples of tensors are very familiar. A tensor of rank 0 is just a number, or scalar, T . Here there's no requirement because a number doesn't change if you do a rotation: $T' = T$. So any single number can be said to be a tensor, although it isn't a particularly helpful designation.

A tensor of rank 1 is a vector. Here, however, it's important that the components of the vector transform as $T'_j = R_{ij}T_j$. If they don't transform in this way, then you don't have a tensor on your hands. You just have a bunch of numbers.

A tensor of rank 2 is a matrix that transforms as $T'_{ij} = R_{ik}R_{jl}T_{kl}$. Again, the transformation property is key. Just because you have an array of numbers A_{ij} , arranged in an $n \times n$ grid, doesn't mean that you have a 2-tensor. You have to check the transformation property holds. Otherwise, as with a vector, the array of numbers isn't a tensor; it's just a bunch of numbers.

What's a Tensor and What's Not?

It's worth elaborating on the definition of a tensor. For example, suppose that someone hands you a matrix, say

$$T_{ij} = \begin{pmatrix} 3 & 8 & 0 \\ 5 & -4 & 3 \\ 1 & 1 & 3 \end{pmatrix}$$

and asks you: "is this a tensor?". It's natural to answer yes. After all, it's written as T_{ij} which is the name we've given to a tensor. And it looks for all the world like a matrix. So is it a tensor? The answer is: we don't know. We haven't been given enough information². As we've stressed several times, a tensor isn't just a bunch of numbers arranged in some pattern. This sometimes goes by the name of an *array* of numbers. Instead, we only know that a given array of numbers is a tensor if it transforms as (6.3). That means that we need to firstly know what basis the array of numbers above has been measured in. And then we need to know what the array looks like when measured in other bases. Only then do we have enough information to say whether this is a tensor or not. It's a tensor only if transforms as (6.3): this transformation law is the definition of a tensor.

²When I gave these lectures on 6th March 2023, I used the example of $\begin{pmatrix} 7 \\ 0 \end{pmatrix}$ which, at least on that day, most certainly wasn't a tensor. It was a football score. A glorious, wonderful, humiliating, shameful football score.

Here's another example. In a given basis, the position of a point is given by x_i . We write this as the components of a vector

$$x_i = \begin{pmatrix} x \\ y \\ z \end{pmatrix}$$

This is a tensor. Indeed, our starting point is that the components of this simple vector transforms in the tensorial way (6.2). This is just the statement that the components of this vector transform in the familiar way under rotation.

Suppose that you now square each of these elements and decide to write them as a column vector. We'll give it a fancy name Λ_i , complete with that hanging i index,

$$\Lambda_i = \begin{pmatrix} x^2 \\ y^2 \\ z^2 \end{pmatrix}$$

That i index makes this look for all the world like it's a tensor. But it's not. We know that after a rotation, $x_i \rightarrow x'_i = R_{ij}x_j$. This means that if we do a rotation and then measure the components of the array Λ'_i we get

$$\Lambda'_i = \begin{pmatrix} (R_{11}x + R_{12}y + R_{13}z)^2 \\ (R_{21}x + R_{22}y + R_{23}z)^2 \\ (R_{31}x + R_{32}y + R_{33}z)^2 \end{pmatrix}$$

But that's most definitely not how a tensor transforms! It's not the rule (6.3) that we wanted. The upshot is that Λ_i is not a tensor and it was a little bit naughty to write it as Λ_i because it suggests that it has some property that it doesn't.

Relatedly, this explains something that you may have wondered about in school. Suppose that you're given two vectors. You know that you can take an inner product to get a scalar, or you can take the cross-product to get another vector. But what stops you from doing something much simpler, just multiplying the component of one vector with the corresponding component of another vector to get a third vector. It seems like such an obvious thing to do. But it's a bad thing to do, precisely because the thing you end up with is not a tensor. It does not transform in the way (6.2), which is how components of a vector should transform.

There is a similar story for matrices. If you have two matrices, then there’s a ridiculously complicated way to multiply them, multiplying rows with columns. Why don’t we just do something much simpler and multiply entries together component by component? You’ve probably guessed the answer by now. If we started with genuine matrices, meaning that they transform (6.3), then the object that you get if you do proper matrix multiplication will also transform as (6.3), but the simpler, stupid way to multiplying will not.

Why are we making such a big deal about this? What is so special about things that transform nicely as (6.3) under rotations? Well, there are several answers to this, depending on taste. At the most basic level, if you’re a physicist, then you might genuinely want to know how something looks in different, rotated frames of reference.

Moreover, once you realise that there’s a preferred way for things to transform — the tensor way (6.3) — this brings some extra power to the calculations, a little like dimensional analysis. Suppose that you have an equation of the form “left-hand side” = “right-hand side”. If the thing on the left is a tensor then the thing on the right better also be a tensor. And sometimes there’s not many tensors available, which limits your options for what the thing on the right can actually be. We’ll see an example of this in Section 6.1.3 when we’ll use tensors to make some scary looking integrals a little more palatable.

The discussion above is very much from a physics perspective. But what about a pure maths perspective? This gives a more formal, but arguably cleaner, definition of a tensor. We’ll explain this imminently in Section 6.1.1.

We’ll meet a number of tensors as we proceed. But there is a one that is special: this is the rank 2 tensor δ_{ij} or, equivalently, the unit matrix. Importantly, it has the same 0 and 1 entries in any basis because, under the transformation (6.3), it becomes

$$\delta'_{ij} = R_{ik}R_{jl}\delta_{kl} = \delta_{ij}$$

We will devote Section 6.1.3 to “invariant tensors” which, like δ_{ij} , take the same form in any basis.

6.1.1 Tensors as Maps

There is something a little strange about the definition of a tensor given above. We first pick a set of coordinates, and the transformation law (6.3) then requires that the tensor transforms nicely so that, ultimately, nothing depends on these coordinates. But, if that’s the case, surely there should be a definition of a tensor that doesn’t rely on coordinates at all!

There is. A tensor T of rank p is a multi-linear map that takes p vectors, $\mathbf{a}, \mathbf{b}, \dots, \mathbf{c}$ and spits out a number in \mathbb{R} ,

$$T(\mathbf{a}, \mathbf{b}, \dots, \mathbf{c}) = T_{i_1 i_2 \dots i_p} a_{i_1} b_{i_2} \dots c_{i_p} \quad (6.4)$$

Here “multi-linear” means that T is linear in each of the entries $\mathbf{a}, \mathbf{b}, \dots, \mathbf{c}$ individually. By evaluating T on all possible vectors $\mathbf{a}, \mathbf{b}, \dots, \mathbf{c}$, we get the components $T_{i_1 i_2 \dots i_p}$. The transformation rule (6.3) is simply the statement that the map T is independent of the choice of basis, and we can equally well write

$$\begin{aligned} T(\mathbf{a}, \mathbf{b}, \dots, \mathbf{c}) &= T'_{i_1 i_2 \dots i_p} a'_{i_1} b'_{i_2} \dots c'_{i_p} \\ &= (R_{i_1 j_1} R_{i_2 j_2} \dots R_{i_p j_p} T_{j_1 j_2 \dots j_p})(R_{i_1 k_1} a_{k_1})(R_{i_2 k_2} b_{k_2}) \dots (R_{i_p k_p} c_{k_p}) \\ &= T_{j_1 j_2 \dots j_p} a_{j_1} b_{j_2} \dots c_{j_p} \end{aligned}$$

which follows because $R^T R = \mathbf{1}$ or, in components, $R_{ij} R_{ik} = \delta_{jk}$. The key is that this formula takes the same form in any basis.

Tensors as Maps Between Vectors

Rather than thinking of a tensor as a map from many vectors to \mathbb{R} , you can equivalently think of it as a map from some lower-rank tensor to another. For example, in (6.4), if you don't fill in the first entry, then a rank p tensor can equally well be viewed as taking $(p - 1)$ vectors and spitting out a single vector

$$a_i = T_{i j_1 \dots j_{p-1}} b_{j_1} \dots c_{j_{p-1}}$$

This is the way that tensors typically arise in physics or applied mathematics, where the most common example is simply a rank 2 tensor, defined as a map from one vector to another

$$\mathbf{u} = T\mathbf{v} \quad \Rightarrow \quad u_i = T_{ij} v_j$$

Until now, we've simply called T a matrix but for the equation $\mathbf{u} = T\mathbf{v}$ to make sense, T must transform as a tensor (6.3). This is inherited from the transformation rules of the vectors, $u'_i = R_{ij} u_j$ and $v'_i = R_{ij} v_j$, giving

$$u'_i = T'_{ij} v'_j \quad \text{with} \quad T'_{ij} = R_{ik} R_{jl} T_{kl}$$

Written as a matrix equation, this is $T' = R T R^T$.

6.1.2 Tensor Operations

Given a bunch of tensors, there are some manipulations that leave you with another tensor. Here we describe these operations.

- We can *add* and *subtract* tensors of the same rank, so if S and T are both tensors of rank p then so too is $S + T$. We can also multiply a tensor by a constant α and it remains a tensor.
- If S is a tensor of rank p and T a tensor of rank q , then the *tensor product* $S \otimes T$ is a tensor of rank $p + q$, defined by

$$(S \otimes T)_{i_1 \dots i_p j_1 \dots j_q} = S_{i_1 \dots i_p} T_{j_1 \dots j_q}$$

You can check that the components of $(S \otimes T)$ do indeed satisfy the transformation rule (6.3). In particular, if we have p different vectors \mathbf{a} , \mathbf{b} , \dots , \mathbf{c} then we can construct a tensor

$$T = \mathbf{a} \otimes \mathbf{b} \otimes \dots \otimes \mathbf{c} \quad \text{with} \quad T_{i_1 \dots i_p} = a_{i_1} b_{i_2} \dots c_{i_p}$$

- Given a tensor T of rank p , we can construct a new tensor S of rank $(p - 2)$ by *contracting* on two indices using δ_{ij} ,

$$S_{k_1 \dots k_{p-2}} = \delta_{ij} T_{ijk_1 \dots k_{p-2}}$$

For a rank 2 tensor, the contraction is what we call the trace, $\text{Tr } T = T_{ii}$. It's a valid tensor operation because the end result is a scalar that does not transform under rotations

$$T'_{ii} = R_{ij} R_{ik} T_{jk} = \delta_{jk} T_{jk} = T_{jj}$$

The same derivation shows that higher rank tensors can also be contracted, with the additional indices unaffected by the contraction.

Combining a contraction with a tensor product gives a way to contract two different tensors together. For example, given a p -tensor P and q -tensor Q , we can form a $p + q - 2$ tensor by contracting, say, the first index on each to get $P_{ik_1 \dots k_{p-1}} Q_{il_1 \dots l_{q-1}}$. This may sound abstract, but it's very much something you've seen before: given a pair of 1-tensors \mathbf{a} and \mathbf{b} , also known as vectors, we can combine them to get a 0-tensor, also known as a number

$$\mathbf{a} \cdot \mathbf{b} = a_i b_i$$

This, of course, is just the inner-product. It is a useful operation precisely because the 0-tensor on the right-hand side is, like all 0-tensors, independent of the choice of basis that we choose to express the vectors.

The Quotient Rule

In practice, it's not hard to recognise a tensor when you see one. In any setting, they're usually just objects with a bunch of i and j indices, each of which clearly transforms as a vector. If in doubt, you can just check explicitly how the thing transforms. (There are cases where this check is needed. In later [courses](#), you'll meet an object called the Levi-Civita connection Γ_{jk}^i which looks for all the world like a tensor but turns out, on closer inspection, to be something more subtle.)

There is a more formal way to say this. Let $T_{i_1 \dots i_{p+q}}$ be a bunch of numbers that you think might comprise a tensor of rank $p+q$ in some coordinate basis. If $T_{i_1 \dots i_{p+q}}$ are indeed the components of a tensor then you can feed it a rank q tensor $u_{j_1 \dots j_q}$ and it will spit back a rank p tensor

$$v_{i_1 \dots i_p} = T_{i_1 \dots i_p j_1 \dots j_q} u_{j_1 \dots j_q} \quad (6.5)$$

There is a converse to this statement. If for every tensor $u_{j_1 \dots j_q}$, the output $v_{i_1 \dots i_p}$ defined in (6.5) is a tensor, then $T_{i_1 \dots i_p j_1 \dots j_q}$ are the components of a tensor. This is called the *quotient rule*.

It is straightforward, if a little fiddly, to prove the quotient rule. It's sufficient to restrict attention to tensors u formed from the tensor product of vectors $u_{j_1 \dots j_q} = c_{j_1} \dots d_{j_q}$. Then, by assumption, $v_{i_1 \dots i_p} = T_{i_1 \dots i_p j_1 \dots j_q} u_{j_1 \dots j_q}$ is a tensor. If we then contract with p further vectors $\mathbf{a}, \dots, \mathbf{b}$ then $v_{i_1 \dots i_p} a_{i_1} \dots b_{i_p} = T_{i_1 \dots i_p j_1 \dots j_q} a_{i_1} \dots b_{i_p} c_{j_1} \dots d_{j_q}$ is necessarily a scalar. This is then enough to ensure the correct transformation rule (6.3) for the components $T_{i_1 \dots i_p j_1 \dots j_q}$.

Symmetry and Anti-Symmetry

The symmetrisation properties of tensors are worthy of comment. A tensor that obeys

$$T_{ijp\dots q} = \pm T_{jip\dots q}$$

is said to be *symmetric* (for $+$) or *anti-symmetric* (for $-$) in the indices i and j . If a tensor is (anti)-symmetric in one coordinate system then it is (anti)-symmetric in any coordinate system

$$T'_{ijp\dots q} = R_{ik} R_{jl} R_{pr} \dots R_{qs} T_{klr\dots s} = \pm R_{ik} R_{jl} R_{pr} \dots R_{qs} T_{lkr\dots s} = \pm T'_{jip\dots q}$$

A tensor that is (anti)-symmetric in all pairs of indices is said to be *totally (anti)-symmetric*. Note that for tensors in \mathbb{R}^n , there are no totally anti-symmetric tensors of rank $p > n$ because at least one of the indices must take the same value and so the tensor necessarily vanishes. A totally anti-symmetric tensor of rank p in \mathbb{R}^n has $\binom{n}{p}$ independent components.

Let's now restrict our attention to \mathbb{R}^3 . A tensor of rank 2 is our new fancy name for a 3×3 matrix T_{ij} . In general, it has 9 independent components. We can always decompose it into the symmetric and anti-symmetric pieces

$$S_{ij} = \frac{1}{2}(T_{ij} + T_{ji}) \quad \text{and} \quad A_{ij} = \frac{1}{2}(T_{ij} - T_{ji})$$

which have 6 and 3 independent components respectively. Our discussion above shows that S and A are each, themselves, tensors. In fact, the symmetric piece can be decomposed further,

$$S_{ij} = P_{ij} + \frac{Q}{3}\delta_{ij}$$

where $Q = S_{ii}$ is the trace of S and carries a single degree of freedom, while P_{ij} is the traceless part of S and carries 5. The importance of this decomposition is that A , P and Q are individually tensors. In contrast, if you were to take, say, the upper-left-hand component of the original matrix T_{ij} then that doesn't form a tensor.

In \mathbb{R}^3 , we can also rewrite an anti-symmetric matrix in terms of a vector,

$$A_{ij} = \epsilon_{ijk}B_k \quad \iff \quad B_k = \frac{1}{2}\epsilon_{ijk}A_{ij}$$

The upshot is that in any 3×3 matrix can be decomposed as

$$T_{ij} = P_{ij} + \epsilon_{ijk}B_k + \frac{1}{3}\delta_{ij}Q \tag{6.6}$$

where $P_{ii} = 0$.

6.1.3 Invariant Tensors

There are two important invariant tensors in \mathbb{R}^n .

- We've met the first already: it is the rank 2 tensor δ_{ij} . As we noted previously, this is invariant because

$$\delta'_{ij} = R_{ik}R_{jl}\delta_{kl} = \delta_{ij}$$

Note that δ_{ij} is invariant under any $R \in O(n)$.

- The rank n totally anti-symmetric tensor $\epsilon_{i_1 \dots i_n}$. This is defined by $\epsilon_{12 \dots n} = +1$. If you swap any two indices you get a minus sign. In particular, if any two indices are repeated, the epsilon symbol vanishes. This is invariant because

$$\epsilon'_{i_1 \dots i_n} = R_{i_1 j_1} \dots R_{i_n j_n} \epsilon_{j_1 \dots j_n} = \det R \epsilon_{i_1 \dots i_n} = \epsilon_{i_1 \dots i_n}$$

Note that the epsilon symbol is only invariant under $R \in SO(n)$ but it is not invariant under $R \in O(n)$ with $\det R = -1$. It picks up a minus sign under reflections. The invariance of ϵ_{ijk} in \mathbb{R}^3 is the reason why the cross-product $(\mathbf{a} \times \mathbf{b})_i = \epsilon_{ijk} a_j b_k$ is itself a vector. Or, said differently, why the triple product $\mathbf{a} \cdot (\mathbf{b} \times \mathbf{c}) = \epsilon_{ijk} a_i b_j c_k$ is independent of the choice of basis.

In general, a tensor is said to be invariant under a given rotation R if

$$T'_{i_1 \dots i_n} = R_{i_1 j_1} \dots R_{i_n j_n} T_{j_1 \dots j_n} = T_{i_1 \dots i_n}$$

A tensor that is invariant under all rotations R is said to be *isotropic*. Obviously all tensors of rank 0 are isotropic. What about higher rank tensors?

Claim: The only non-zero isotropic tensors in \mathbb{R}^3 of rank $p = 1, 2$ or 3 are $T_{ij} = \alpha \delta_{ij}$ and $T_{ijk} = \beta \epsilon_{ijk}$ with α and β constant. In particular, there are no isotropic tensors of rank 1 (essentially because a vector always points in a preferred direction).

Proof: The idea is simply to look at how tensors transform under a bunch of specific rotations by π or $\pi/2$ about certain axes.

For example, consider a tensor of rank 1, so that

$$T'_i = R_{ij} T_j \quad \text{with} \quad R_{ij} = \begin{pmatrix} -1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & +1 \end{pmatrix} \quad (6.7)$$

Requiring $T'_i = T_i$ gives $T_1 = T_2 = 0$. Clearly a similar argument, using a different R , also gives $T_3 = 0$.

For a tensor of rank 2, consider the transformation

$$T'_{ij} = \tilde{R}_{ik} \tilde{R}_{jl} T_{kl} \quad \text{with} \quad \tilde{R}_{ij} = \begin{pmatrix} 0 & 1 & 0 \\ -1 & 0 & 0 \\ 0 & 0 & +1 \end{pmatrix} \quad (6.8)$$

which is a rotation by $\pi/2$ about the z -axis. The rotation gives $T'_{13} = T_{23}$ and $T'_{23} = -T_{13}$ so if $T'_{ij} = T_{ij}$, we must have $T_{13} = T_{23} = 0$. Meanwhile $T'_{11} = T_{22}$. Similar arguments tell us that all off-diagonal elements must vanish and all diagonal elements must be equal: $T_{11} = T_{22} = T_{33} = \alpha$ for some α . Hence $T_{ij} = \alpha \delta_{ij}$.

Finally, for a rank 3 tensor we have

$$T'_{ijk} = R_{il}R_{jp}R_{kq}T_{lpq}$$

If we pick R given in (6.7), then we find $T'_{133} = -T_{133}$ and $T'_{111} = -T_{111}$. Similar arguments show that an isotropic tensor must have $T_{ijk} = 0$ unless i, j and k are all distinct. Meanwhile, if we pick $R = \tilde{R}$ given in (6.8), then we get $T'_{123} = -T_{213}$. We end up with the result we wanted: T_{ijk} is isotropic if and only if $T_{ijk} = \beta\epsilon_{ijk}$ for some constant β . \square

Although we won't prove it here, all other isotropic tensors can be formed from δ_{ij} and ϵ_{ijk} . For example, the only isotropic 4-tensor in \mathbb{R}^3 is

$$T_{ijkl} = \alpha\delta_{ij}\delta_{kl} + \beta\delta_{ik}\delta_{jl} + \gamma\delta_{il}\delta_{jk}$$

with α, β and γ constants. You could try to cook up something involving ϵ_{ijk} but it doesn't give anything new. In particular, $\epsilon_{ijk}\epsilon_{ilp} = \delta_{jl}\delta_{kp} - \delta_{jp}\delta_{kl}$.

There is also an analogous result in \mathbb{R}^n : all isotropic tensors can be constructed from the symmetric 2-tensor δ_{ij} and the totally anti-symmetric n -tensor $\epsilon_{i_1\dots i_n}$.

Invariant Integrals

It is sometimes possible to use invariance properties to immediately write down the index structure of an integral, without doing the hard work of evaluating everything term by term. Suppose that we have some integral of the form

$$T_{ij\dots k} = \int_V f(r)x_ix_j\dots x_k dV$$

with $r = |\mathbf{x}|$. Then under a rotation, we have

$$T'_{ij\dots k} = R_{ip}R_{jq}\dots R_{kr}T_{pq\dots r} = \int_V f(r)x'_ix'_j\dots x'_k dV$$

with, as usual, $x'_i = R_{ij}x_j$. But if we now change the integration variables to x' , both $r = |\mathbf{x}| = |\mathbf{x}'|$ and $dV = dV'$ are invariant. (The latter because the Jacobian is $\det R = 1$). If the domain of integration is also rotationally invariant, so $V = V'$, then the final result must itself be an invariant tensor, $T'_{ij\dots k} = T_{ij\dots k}$.

Here are some examples. First, suppose that we have a 3d integral over the interior of a sphere of radius R , given by

$$T_i = \int_V \rho(r) x_i dV \quad (6.9)$$

This must be equal to some invariant 1-tensor (i.e. a vector), but there are no such objects. In other words, we can say immediately that $T_i = 0$. You can check this straightforwardly by doing the integral in, say, spherical polar coordinates.

Things change if we look at an integral with two hanging indices,

$$T_{ij} = \int_V \rho(r) x_i x_j dV \quad (6.10)$$

(In Section 6.2, we will find integrals of this form arising when we compute the inertia tensor of a sphere.) By the argument above T_{ij} must be an isotropic tensor and hence proportional to δ_{ij} ,

$$T_{ij} = \int_V \rho(r) x_i x_j dV = \alpha \delta_{ij}$$

for some α . If we take the trace, we get

$$\int_V \rho(r) r^2 dV = 3\alpha$$

Hence,

$$T_{ij} = \frac{1}{3} \delta_{ij} \int_V \rho(r) r^2 dV = \frac{4\pi}{3} \delta_{ij} \int_0^R dr \rho(r) r^4 \quad (6.11)$$

For example, if $\rho(r) = \rho_0$ is constant, then $T_{ij} = \frac{4}{15} \pi \rho_0 R^5 \delta_{ij}$.

Here's a slightly more complicated example (taken from the calculation of Stokes flow around a sphere in [Fluid Mechanics](#)). Consider the surface integral over a sphere of radius R ,

$$\tilde{T}_k = a_j \int_{\mathbf{S}^2} dS_i \frac{x_i x_j x_k}{r^5}$$

This time we have a vector \mathbf{a} in the game, so it must be the case that $\tilde{T}_k = \beta a_k$ for some constant β . One way to compute β is to strip off the vector \mathbf{a} and instead look at

$$\tilde{T}_{jk} = \int_{\mathbf{S}^2} dS_i \frac{x_i x_j x_k}{r^5} = \beta \delta_{jk}$$

which now should be proportional to the invariant tensor δ_{jk} as shown, with the same coefficient β since $\tilde{T}_k = T_{jk}a_j = \beta a_k$. At this point, we again take the trace over the j and k indices to get

$$\int_{\mathbf{S}^2} dS_i \frac{x_i x_j x_j}{r^5} = 3\beta$$

But this integral is given by

$$\int_{\mathbf{S}^2} dS_i \frac{x_i}{r^3} = \int_{\mathbf{S}^2} d\mathbf{S} \cdot \frac{\mathbf{n}}{r^2} = 4\pi$$

and so we get $\beta = 4\pi/3$.

6.1.4 Tensor Fields

A tensor field over \mathbb{R}^3 is the assignment of a tensor $T_{i\dots k}(\mathbf{x})$ to every point $\mathbf{x} \in \mathbb{R}^3$. This is the generalisation of a vector field

$$\mathbf{F} : \mathbb{R}^3 \rightarrow \mathbb{R}^3$$

to a map of the kind

$$T : \mathbb{R}^3 \rightarrow \mathbb{R}^m$$

with m the number of components of the tensor. So, for example, a map that assigns a symmetric, traceless rank 2 tensor $P_{ij}(\mathbf{x})$ to every point has $m = 5$.

The tensor field $T_{i\dots k}(\mathbf{x})$ is sometimes denoted as $T_{i\dots k}(x^l)$ which is supposed to show that the field depends on all coordinates x^1, \dots, x^3 . It's not great notation because the indices as subscripts are supposed to take some definite values, while the index l in the argument is supposed to denote the whole set of indices. It's especially bad notation when combined with the summation convention and we won't adopt it here.

There is one very famous example of a tensor field. Einstein's theory of general relativity is described by a rank 2 tensor at every point in space. This is called the *metric*. The dynamics of this rank 2 tensor field describe gravity. (I've brushed something rather important under the rug here. Einstein's theory is a rank 2 tensor in *spacetime*, not just in space. Which means that the rank 2 tensor is a 4×4 matrix, rather than a 3×3 matrix.)

Before we move on, it's worth pausing to mention a slightly subtle point. Not all maps $\mathbb{R}^3 \rightarrow \mathbb{R}^3$ qualify as "vector fields". The point \mathbf{x} in the codomain \mathbb{R}^3 is a vector and so its components transform in the appropriate way under rotation. To be a vector field, the components of the map must transform under the *same* rotation. Similar comments hold for a tensor field.

To illustrate this, the electric field $\mathbf{E}(\mathbf{x})$ is an example of a vector field. If you rotate in space, and so change \mathbf{x} , then the direction \mathbf{E} also changes: the rotation acts on both the argument \mathbf{x} and the function itself \mathbf{E} .

In contrast, there are maps $\mathbb{R}^3 \rightarrow \mathbb{R}^3$ where, although the domain and codomain have the same dimension, vectors in them transform under different rotations. For example, in particle physics there exists an object called a *quark field* which, for our (admittedly, slightly dumbed down) purposes, can be thought of as a map $\mathbb{R}^3 \rightarrow \mathbb{R}^3$. This is a quantum field whose ripples are the particles that we call quarks, but these details can be safely ignored for the next couple of years of your life. We will write this field as $q_a(\mathbf{x})$ where the $a = 1, 2, 3$ label is the “colour” of the quark. If we rotate in space, then \mathbf{x} changes but the colour of the quark does not. There is then an independent rotation that acts on the codomain and rotates the colour, but leaves the point in space unchanged. Because the rotations that act on the domain and codomain are unrelated, the quark field is usually not referred to as a vector field.

Taking Derivatives

Given a tensor field, we can always construct higher rank tensors by taking derivatives. In fact, we’ve already seen a prominent example of this earlier in these lectures. There, we started with a scalar field $\phi(\mathbf{x})$ and differentiated to get the gradient $\nabla\phi$. This means that we start with a rank 0 tensor and differentiate to get a rank 1 tensor.

Strictly speaking, we didn’t previously prove that $\nabla\phi$ is a vector field. But it’s straightforward to do so. As we’ve seen above, we need to show that it transforms correctly under rotations. Any vector \mathbf{v} can be decomposed in two different ways,

$$\mathbf{v} = v^i \mathbf{e}_i = v'^i \mathbf{e}'_i$$

where $\{\mathbf{e}_i\}$ and $\{\mathbf{e}'_i\}$ are two orthonormal bases, each obeying $\mathbf{e}_i \cdot \mathbf{e}_j = \mathbf{e}'_i \cdot \mathbf{e}'_j = \delta_{ij}$, and v^i and v'^i are the two different coordinates for \mathbf{v} . If we expand \mathbf{x} in this way

$$\mathbf{x} = x_i \mathbf{e}_i = x'_i \mathbf{e}'_i \implies x_i = (\mathbf{e}_i \cdot \mathbf{e}'_j) x'_j \implies \frac{\partial x^i}{\partial x'^j} = \mathbf{e}_i \cdot \mathbf{e}'_j$$

Here $\mathbf{e}_i \cdot \mathbf{e}'_j$ is the rotation matrix that takes us from one basis to the other. Meanwhile, we can always expand one set of basis vectors in terms of the other,

$$\mathbf{e}_i = (\mathbf{e}_i \cdot \mathbf{e}'_j) \mathbf{e}'_j = \frac{\partial x^i}{\partial x'^j} \mathbf{e}'_j$$

This tells us that we could equally as well write the gradient as

$$\nabla\phi = \frac{\partial\phi}{\partial x^i} \mathbf{e}_i = \frac{\partial\phi}{\partial x^i} \frac{\partial x^i}{\partial x'^j} \mathbf{e}'_j = \frac{\partial\phi}{\partial x'^j} \mathbf{e}'_j$$

This is the expected result: if you work in a different primed basis, then you have the same definition of $\nabla\phi$, but just with primes on both \mathbf{e}'_i and $\partial/\partial x'^i$. This means that the components $\partial_i\phi$ transform correctly under a rotation, so $\nabla\phi$ is indeed a vector.

We can extend the result above to any, suitably smooth, tensor field $T(\mathbf{x})$ of rank p . We can differentiate this any number of times to get a new tensor field of rank, say, $p + q$,

$$X_{i_1\dots i_q j_1\dots j_p} = \frac{\partial}{\partial x_{i_1}} \cdots \frac{\partial}{\partial x_{i_q}} T_{j_1\dots j_p}(\mathbf{x}) \quad (6.12)$$

To verify that this is indeed a tensor, we need to check how it changes under a rotation. In a new basis, we have $x'_i = R_{ij}x_j$ (where $R_{ij} = \mathbf{e}'_i \cdot \mathbf{e}_j$ in the notation above) and so

$$\frac{\partial x'_i}{\partial x_j} = R_{ij} \quad \implies \quad \frac{\partial}{\partial x'_i} = \frac{\partial x_j}{\partial x'_i} \frac{\partial}{\partial x_j} = R_{ij} \frac{\partial}{\partial x_j}$$

which is the result we need for X in (6.12) to qualify as a tensor field.

We can implement any of the tensorial manipulations that we met previously for tensor fields. For example, if we start with a vector field $\mathbf{F}(\mathbf{x})$, we can form a rank 2 tensor field

$$T_{ij}(\mathbf{x}) = \frac{\partial F_i}{\partial x_j}$$

But we saw in (6.6) that any rank 2 tensor field can be decomposed into various pieces. There is an anti-symmetric piece

$$A_{ij}(\mathbf{x}) = \epsilon_{ijk} B_k(\mathbf{x}) \quad \text{with} \quad B_k = \frac{1}{2} \epsilon_{ijk} \frac{\partial F_i}{\partial x_j} = -\frac{1}{2} (\nabla \times \mathbf{F})_k$$

and a trace piece

$$Q = \frac{\partial F_i}{\partial x_i} = \nabla \cdot \mathbf{F}$$

and, finally, a symmetric, traceless piece

$$P_{ij}(\mathbf{x}) = \frac{1}{2} \left(\frac{\partial F_i}{\partial x_j} + \frac{\partial F_j}{\partial x_i} \right) - \frac{1}{3} \nabla \cdot \mathbf{F}$$

Obviously, the first two of these are familiar tensors (in this case a scalar and vector) from earlier sections.

6.2 Physical Examples

Our discussion above was rooted firmly in mathematics. There are many places in physics where tensors appear. Here we give a handful of examples.

6.2.1 Electric Fields in Matter

Apply an electric field \mathbf{E} to a lump of stuff. A number of things can happen.

If the lump of stuff is an insulator then the material will become *polarised*. This means that the positive electric charge will be pushed in one direction, the negative in another until the lump of stuff acts like a dipole. (This is described in some detail in Section 7 of the lectures on [Electromagnetism](#).) One might think that the resulting polarisation vector \mathbf{P} points in the same direction as the electric field \mathbf{E} , but that's too simplistic. For many lumps of stuff, the underlying crystal structure allows the electric charges to shift more freely in some directions than others. The upshot is that the relation between polarisation \mathbf{P} and applied electric field \mathbf{E} is given by

$$\mathbf{P} = \alpha \mathbf{E}$$

where α is a matrix known as the *polarisation tensor*. In a given basis, it has components α_{ij} .

There is a similar story if the lump of stuff is a conductor. This time an applied electric field gives rise to a current density \mathbf{J} . Again, the current is not necessarily parallel to the electric field. The relationship between them is now

$$\mathbf{J} = \sigma \mathbf{E}$$

This is known as *Ohm's law*. In general σ is a 3×3 matrix known as the *conductivity tensor*; in a given basis, it has components σ_{ij} .

What can we say about σ when the material is isotropic, meaning that it looks the same in all directions? In this case, no direction is any different from any other. With no preferred direction, the conductivity tensor must be proportional to an invariant tensor, so that it looks the same in all coordinate systems. What are our options?

For 3d materials, the only option is $\sigma_{ij} = \sigma \delta_{ij}$, which ensures that the current does indeed run parallel to the electric field. In this case σ is just referred to as the *conductivity*.

However, suppose that we're dealing with a thin wafer of material in which both the current and electric field are restricted to lie in a plane. This changes the story because now we're dealing with vectors in \mathbb{R}^2 rather than \mathbb{R}^3 and \mathbb{R}^2 is special because there are two invariant 2-tensors in this dimension: δ_{ij} and ϵ_{ij} . This means that the most general conductivity tensor for an isotropic 2d material takes the form

$$\sigma_{ij} = \sigma_{xx}\delta_{ij} + \sigma_{xy}\epsilon_{ij} = \begin{pmatrix} \sigma_{xx} & \sigma_{xy} \\ -\sigma_{xy} & \sigma_{xx} \end{pmatrix}$$

Here σ_{xx} is called the *longitudinal conductivity* while σ_{xy} is called the *Hall conductivity*. If $\sigma_{xy} \neq 0$ then an electric field in the x -direction induces a current in the y -direction.

As an aside, it turns out that the seemingly mundane question of understanding σ_{xy} in real materials is closely tied to some of the most interesting breakthroughs in mathematics in recent decades! This is the subject of the [Quantum Hall Effect](#).

6.2.2 The Inertia Tensor

Another simple example of a tensor arises in Newtonian mechanics. A *rigid body* rotating about the origin can be modelled by some number of masses m_a at positions \mathbf{x}_a , all moving with velocity $\dot{\mathbf{x}}_a = \boldsymbol{\omega} \times \mathbf{x}_a$. Here $\boldsymbol{\omega}$ is known as the angular velocity. The angular velocity $\boldsymbol{\omega}$ is related to the angular momentum \mathbf{L} by

$$\mathbf{L} = I\boldsymbol{\omega} \tag{6.13}$$

with I the *inertia tensor*. The angular momentum does not necessarily lie parallel to the angular velocity and, correspondingly, I is in general a matrix, rather than a single number. In fact, we can easily derive an expression for the inertia tensor. The angular momentum is

$$\mathbf{L} = \sum_a m_a \mathbf{x}_a \times \dot{\mathbf{x}}_a = \sum_a m_a \mathbf{x}_a \times (\boldsymbol{\omega} \times \mathbf{x}_a) = \sum_a m_a \left(|\mathbf{x}_a|^2 \boldsymbol{\omega} - (\mathbf{x}_a \cdot \boldsymbol{\omega}) \mathbf{x}_a \right)$$

In components, $L_i = I_{ij}\omega_j$, where

$$I_{ij} = \sum_a m_a \left(|\mathbf{x}_a|^2 \delta_{ij} - (\mathbf{x}_a)_i (\mathbf{x}_a)_j \right)$$

For a continuous object with density $\rho(\mathbf{x})$, we can replace the sum with a volume integral

$$I_{ij} = \int_V \rho(\mathbf{x}) \left(|\mathbf{x}|^2 \delta_{ij} - x_i x_j \right) dV \tag{6.14}$$

So, for example, $I_{33} = \int \rho(x_1^2 + x_2^2) dV$ and $I_{12} = -\int \rho x_1 x_2 dV$.

An Example: A Sphere

For a ball of radius R and density $\rho(r)$, the inertia tensor is

$$I_{ij} = \int_V \rho(r)(r^2\delta_{ij} - x_i x_j) dV$$

The second of these terms is the integral (6.10) that we simplified in Section 6.1.3 using isotropy arguments. Using (6.11), we have

$$I_{ij} = \frac{2}{3}\delta_{ij} \int_V \rho(r)r^2 dV = \frac{8\pi}{3}\delta_{ij} \int_0^R dr \rho(r)r^4$$

For example, if $\rho(r) = \rho_0$ is constant, then $I_{ij} = \frac{8}{15}\pi\rho_0 R^5\delta_{ij} = \frac{2}{5}MR^2\delta_{ij}$ where M is the mass of the sphere.

Another Example: A Cylinder

The sphere is rather special because the inertia tensor is proportional to δ_{ij} . That's not the case more generally. Consider, for example, a solid 3d cylinder of radius a and height $2L$, with uniform density ρ . The mass is $M = 2\pi a^2 L\rho$. We align the cylinder with the z -axis and work in cylindrical polar coordinates $x = r \cos \phi$ and $y = r \sin \phi$. The components of the inertia tensor are then

$$I_{33} = \int_V \rho(x^2 + y^2) dV = \rho \int_0^{2\pi} d\phi \int_0^a dr \int_{-L}^{+L} dz r r^2 = \rho\pi L a^4$$
$$I_{11} = \int_V \rho(y^2 + z^2) dV = \rho \int_0^{2\pi} d\phi \int_0^a dr \int_{-L}^{+L} dz r(r^2 \sin^2 \phi + z^2) = \rho\pi a^2 L \left(\frac{a^2}{2} + \frac{2L^2}{3} \right)$$

By symmetry, $I_{22} = I_{11}$. For the off-diagonal elements, we have

$$I_{13} = - \int_V \rho x_1 x_3 dV = -\rho \int_0^{2\pi} d\phi \int_0^a dr \int_{-L}^L dz r^2 z \cos \phi = 0$$

where the integral vanishes due to the ϕ integration. Similarly, $I_{12} = I_{13} = 0$. We find that the inertia tensor for the cylinder is

$$I = \text{diag} \left(M \left(\frac{a^2}{4} + \frac{L^2}{3} \right), M \left(\frac{a^2}{4} + \frac{L^2}{3} \right), \frac{1}{2} M a^2 \right) \quad (6.15)$$

Note that the inertia tensor is diagonal in our chosen coordinates.

The Eigenvectors of the Inertia Tensor

The inertia tensor I defined in (6.14) has a special property: it is symmetric

$$I_{ij} = I_{ji}$$


Any symmetric matrix I can always be diagonalised by an appropriate rotation. This means that there exists an $R \in SO(n)$ such that

$$I' = RIR^T = \text{diag}(I_1, I_2, I_3)$$

Another way of saying this is that any symmetric rank 2 tensor has a basis of orthonormal eigenvectors $\{\mathbf{e}_i\}$, with I_i the corresponding eigenvalues.

In the case of the inertia tensor, the eigenvectors \mathbf{e}_1 , \mathbf{e}_2 and \mathbf{e}_3 are called the *principal axes* of the solid. It means that any object, no matter how complicated, has its own preferred set of orthonormal axes embedded within it. If the object has some symmetry, then the principal axes will always be aligned with this symmetry. This, for example, was the case for the cylinder that we computed above where aligning the cylinder with the z -axis automatically gave us a diagonal inertia tensor (6.15).

In general, it will be less obvious where the principal axes lie. For example, the figure on the right shows the asteroid Toutatis, which is notable for its lumpy shape. The principal axes are shown embedded in the asteroid.



From (6.13), the angular momentum \mathbf{L} is aligned with the angular velocity $\boldsymbol{\omega}$ only if a body spins about one of its principal axes. It turns out that, in this case, nice things happen and the body spins smoothly.

However, if \mathbf{L} and $\boldsymbol{\omega}$ are misaligned, the body exhibits more complicated tumbling, wobbling motion as it spins. You can learn all about this in the lectures on [Classical Dynamics](#). (For what it's worth, Toutatis does not spin about a principal axes.)

6.2.3 Higher Rank Tensors

You might reasonably complain that, after all that work defining tensors, the examples that we've given here are nothing more exotic than matrices, mapping one vector to another. And you would be right. However, as we get to more sophisticated theories of physics, tensors of higher rank do make an appearance. Here we don't give full details, but just say a few words to give you a flavour of things to come.

Perhaps the simplest example arises in the theory of elastic materials. These materials can be subjected to *strain*, which describes the displacement of the material at each point, and *stress*, which describes the forces acting on the material at each point. But each of these is itself a 2-tensor (strictly a tensor field). The *strain tensor* e_{ij} is a symmetric tensor that describes the way the displacement in the x^i direction varies in the x^j . The *stress tensor* σ_{ij} describes the component of the force F_i across a plane normal to x^j . These two tensors are related by

$$\sigma_{ij} = C_{ijkl}e_{kl}$$

This is the grown up version of Hooke's law. In general an elastic material is characterised by the *elasticity tensor*, also known as the *stiffness tensor*, C_{ijkl} .

Higher rank tensors also appear prominently in more advanced descriptions of geometry. In higher dimensions, the simple Gaussian curvature that we met in Section 2 isn't enough to capture all the interesting ways in which spaces can curve in different directions. Instead, it is replaced by a 4-tensor R_{ijkl} known as the Riemann curvature. In the context of physics, this 4-tensor describes the bending of space and time and is needed for the grown-up version of Newton's law of gravity.

6.3 A Unification of Integration Theorems

In this final section, we turn back to matters of mathematics. The three integral theorems that we met in Section 4 are obviously closely related. To end these lectures, we show how they can be presented in a unified framework. This requires us to introduce some novel and slightly formal ideas. These go quite a bit beyond what is usually covered in an introductory course on vector calculus, but we will meet these objects again in later courses on [Differential Geometry and General Relativity](#). View this section as a taste of things to come.

6.3.1 Integrating in Higher Dimensions

Our unified framework will give us integral theorems in any dimension \mathbb{R}^n . If you look back at Section 4, you'll notice that the divergence theorem already holds in any \mathbb{R}^n . Meanwhile, Stokes' theorem is restricted to surfaces in \mathbb{R}^3 for the very simple reason that the cross-product is only defined in \mathbb{R}^3 . This suggests that before we can extend our integral theorems to higher dimensions, we should first ask a more basic question: how do we extend the cross product to higher dimensions?

The introduction of tensors gives us a way to do this. Given two vectors \mathbf{a} and \mathbf{b} in \mathbb{R}^3 , the cross-product is

$$(\mathbf{a} \times \mathbf{b})_i = \epsilon_{ijk} a_j b_k$$

From this perspective, the reason that the cross product can only be employed in \mathbb{R}^3 is because it's only there that the ϵ_{ijk} symbol has three entries. If, in contrast, we're in \mathbb{R}^4 then we have ϵ_{ijkl} and so if we feed it two vectors \mathbf{a} and \mathbf{b} , then we find ourselves with a tensor of rank 2, $T_{ij} = \epsilon_{ijkl} a_k b_l$.

The tensors that we get from an epsilon symbol are always special, in the sense that they are totally anti-symmetric. The anti-symmetry condition doesn't impose any extra constraint on a 0-tensor ϕ or a 1-tensor a_i as these are just scalar fields and vector fields respectively. It only kicks in when we get to tensors of rank 2 or higher.

With this in mind, we can revisit the cross product. We can define the cross product in any dimension \mathbb{R}^n : it is a map that eats two vectors \mathbf{a} and \mathbf{b} and spits back an anti-symmetric $(n - 2)$ -tensor

$$(\mathbf{a} \times \mathbf{b})_{i_1 \dots i_{n-2}} = \epsilon_{i_1 \dots i_n} a_{i_{n-1}} b_{i_n}$$

The only thing that's special about \mathbb{R}^3 is that we get back another vector, rather than a higher dimensional tensor.

There is also a slightly different role played by the epsilon symbol $\epsilon_{i_1, \dots, i_n}$: it provides a map from anti-symmetric p -tensors to anti-symmetric $(n - p)$ -tensors, simply by contracting indices,

$$\epsilon : T_{i_1 \dots i_p} \mapsto \frac{1}{(n - p)!} \epsilon_{i_1 \dots i_n} T_{i_{n-p+1} \dots i_n} \quad (6.16)$$

This map goes by the fancy name of the *Hodge dual*. (Actually, it's an entirely trivial version of the Hodge dual. The proper Hodge dual is a generalisation of this idea to curved spaces.)

Our next step is to think about what this has to do with integration. Recall that earlier in these lectures we found two natural ways to integrate vector fields in \mathbb{R}^3 . The first is along a line

$$\int_C \mathbf{F} \cdot d\mathbf{x} \quad (6.17)$$

which captures the component vector field *tangent* to the line. We can perform this procedure in any dimension \mathbb{R}^n . The second operation is to integrate a vector field over a surface

$$\int_S \mathbf{F} \cdot d\mathbf{S} \quad (6.18)$$

where $d\mathbf{S}$ points in the direction normal to the surface. This integration captures the component of the vector field *normal* to the surface and only makes sense in \mathbb{R}^3 . This is because it's only in \mathbb{R}^3 that a two-dimensional surface has a unique normal. More operationally, this normal, which is buried in the definition of $d\mathbf{S}$, requires us to use the cross product. For a parameterised surface $\mathbf{x}(u, v)$, the vector area element is

$$d\mathbf{S} = \frac{\partial \mathbf{x}}{\partial u} \times \frac{\partial \mathbf{x}}{\partial v} du dv$$

or, in components,

$$dS_i = \epsilon_{ijk} \frac{\partial x^j}{\partial u} \frac{\partial x^k}{\partial v} du dv$$

Now comes a mathematical sleight of hand. Rather than thinking of (6.18) as the integral of a vector field projected normal to the surface, instead think of it as the integral of an anti-symmetric 2-tensor $F_{ij} = \epsilon_{ijk} F_k$ integrated *tangent* to the surface. We then have

$$\int_S \mathbf{F} \cdot d\mathbf{S} = \int_S F_{ij} dS_{ij} \quad \text{with} \quad dS_{ij} = \frac{1}{2} \left(\frac{\partial x^j}{\partial u} \frac{\partial x^k}{\partial v} - \frac{\partial x^j}{\partial v} \frac{\partial x^k}{\partial u} \right) du dv \quad (6.19)$$

This is the same equation as before, just with the epsilon symbol viewed as part of the integrand F_{ij} rather than as part of the measure dS_i . Note that we've retained the anti-symmetry of the area element dS_{ij} that was inherent in our original cross product definition of $d\mathbf{S}$. Strictly speaking this isn't necessary because we're contracting with anti-symmetric indices in F_{ij} , but it turns out that it's best to think of both objects F_{ij} and dS_{ij} as individually anti-symmetric.

This new perspective suggests a way to generalise to higher dimensions. In the line integral (6.17) we're integrating a vector field over a line. In the surface integral (6.19), we're really integrating an anti-symmetric 2-tensor over a surface. The key idea is that one can integrate a totally anti-symmetric p -tensor over a p -dimensional subspace.

Specifically, given an anti-symmetric p -tensor, the generalisation of the line integral (6.17) is the integration over a p -dimensional subspace,

$$\int_M T_{i_1 \dots i_p} dS_{i_1 \dots i_p} \quad (6.20)$$

where $\dim(M) = p$. Here $dS_{i_1 \dots i_p}$ is a higher dimensional version of the “area element” defined in (6.19).

Alternatively, the higher dimensional version of the surface integral (6.18) involves first mapping the p -tensor to an $(n - p)$ -tensor using the Hodge dual. This can subsequently be integrated over an $(n - p)$ -dimensional subspace

$$\int_{\tilde{M}} T_{i_1 \dots i_p} \epsilon_{i_1 \dots i_p j_1 \dots j_{n-p}} d\tilde{S}_{j_1 \dots j_{n-p}} \quad (6.21)$$

with $\dim(\tilde{M}) = n - p$.

In fact, we’ve already met an integral of the form (6.21) elsewhere in these lectures, since this is what we’re implicitly doing when we integrate a scalar field over a volume. In this case the “area element” is just $dS_{i_1 \dots i_n} = \frac{1}{n!} \epsilon_{i_1 \dots i_n} dV$ and the two epsilon symbols just multiply to a constant.. When actually computing a volume integral, this extra machinery is more of a distraction than a help.. But if we want to know how to think about things more generally then it’s extremely useful.

6.3.2 Differentiating Anti-Symmetric Tensors

We’ve now learned how to integrate anti-symmetric tensors. Our next step is to learn how to differentiate them. We’ve already noted in (6.12) that we can differentiate a p tensor once to get a tensor of rank $p + 1$, but in general differentiating loses the anti-symmetry property. As we now explain, there is a way to restore it so that when we differentiate a totally anti-symmetric p tensor, we end up with a totally anti-symmetric $(p + 1)$ -tensor.

For a scalar field, things are trivial. We can construct a vector field $\nabla\phi$ and this is automatically “anti-symmetric” because there’s nothing to anti-symmetrise.

If we’re given a vector field F_i , we can differentiate and then anti-symmetrise by hand. I will introduce a new symbol for “differentiation and anti-symmetrisation” and write

$$(\mathcal{D}F)_{ij} := \frac{1}{2} \left(\frac{\partial F_i}{\partial x^j} - \frac{\partial F_j}{\partial x^i} \right)$$

where the anti-symmetry is manifest on the right-hand side. I should confess that the notation $\mathcal{D}F$ is not at all standard. In subsequent courses, this object is usually viewed as something called a “differential form” and written simply as dF but the notation dF is loaded with all sorts of other connotations which are best ignored at this stage. Hence the made-up notation $\mathcal{D}F$.

In \mathbb{R}^3 , this anti-symmetric differentiation is equivalent to the curl using the Hodge map (6.16),

$$(\nabla \times \mathbf{F})_i = \epsilon_{ijk}(\mathcal{D}F)_{jk}$$

But now we can extend this definition to any anti-symmetric p -tensor. We can always differentiate and anti-symmetrise to get a $(p+1)$ -tensor defined by

$$(\mathcal{D}T)_{i_1 \dots i_{p+1}} = \frac{1}{p+1} \left(\frac{\partial T_{i_1 \dots i_p}}{\partial x^{i_{p+1}}} + p \text{ further terms} \right)$$

where the further terms involve replacing the derivative $\partial/\partial x^{i_{p+1}}$ with one of the other coordinates $\partial/\partial x^j$ so that the whole shebang is fully anti-symmetric.

Note that, with this definition of \mathcal{D} , if we differentiate twice then we take a p -tensor to a $(p+2)$ -tensor. But this $(p+2)$ -tensor always vanishes!

$$(\mathcal{D}\mathcal{D}T)_{i_1 \dots i_{p+2}} = 0$$

for any tensor T . This is because we’ll have two derivatives contracted with an epsilon and is the higher dimensional generalisation of the statements that $\nabla \times \nabla\phi = 0$ or $\nabla \cdot (\nabla \times \mathbf{F}) = 0$.

As an aside: this is actually the second time in these lectures that we’ve seen something vanish when you act twice, although you’d be forgiven for failing to notice the connection. Here our new anti-symmetric derivative obeys $\mathcal{D}^2(\text{anything}) = 0$. But we previously saw that the “boundary of a boundary” is always zero. This means that if a higher dimensional space (really a manifold) M has boundary ∂M then $\partial(\partial M) = 0$. Conceptually, these two ideas are very different but one can’t help but be struck by the similarity of the equations $\mathcal{D}^2(\text{anything}) = 0$ and $\partial^2(\text{anything}) = 0$, even though the “anything”’s are very different objects in the two formulae. It turns out that this similarity is pointing at a deep connection between the topology of spaces and the kinds of tensors that one can put on these spaces. In fancy maths words, this is the link between homology and cohomology.

Finally, we can now state the general integration theorem. Given an anti-symmetric p -tensor T , then

$$\int_M (\mathcal{D}T)_{i_1 \dots i_{p+1}} dS_{i_1 \dots i_{p+1}} = \int_{\partial M} T_{i_1 \dots i_p} dS_{i_1 \dots i_p} \quad (6.22)$$

Here $\dim(M) = p + 1$ and, therefore the boundary has $\dim(\partial M) = p$. Note that we don't use a different letter to distinguish the integration measure over these various spaces: everything is simply dS and you have to look closer at the indices to see what kind of space you're integrating over.

The equation (6.22) is a unification of all integration theorems. It contains the fundamental theorem of calculus (when $p = 0$), the divergence theorem (when $p = n - 1$) and Stokes' theorem (when $p = 1$ and $\mathbb{R}^n = \mathbb{R}^3$). Geometers refer to this generalised theorem simply as *Stokes' theorem* since that is the original result that it resembles most. The proof is simply a higher dimensional version of the proofs that we sketched previously.

There is, to put it mildly, quite a lot that I'm sweeping under the rug in the discussion above. In particular, the full Stokes' theorem does not hold only in \mathbb{R}^n but in a general curved space known as a manifold. In that context, one has to be a lot more careful about what kind of tensors we're dealing with and, as I mentioned above, Stokes' theorem should be written using a kind of anti-symmetric tensor known as a *differential form*. None of this really matters when working in flat space, but the differences become crucial when thinking about curved spaces. If you want to learn more, these topics will be covered in glorious detail in later courses on Differential Geometry or, for physicists, [General Relativity](#).

What You Really Need

Here are expressions for div, grad, curl and the Laplacian in various coordinate systems.

Cartesian: $\mathbf{x} = (x, y, z)$

$$\nabla f = \frac{\partial f}{\partial x} \hat{\mathbf{x}} + \frac{\partial f}{\partial y} \hat{\mathbf{y}} + \frac{\partial f}{\partial z} \hat{\mathbf{z}}$$

$$\nabla \cdot \mathbf{F} = \frac{\partial F_x}{\partial x} + \frac{\partial F_y}{\partial y} + \frac{\partial F_z}{\partial z}$$

$$\nabla \times \mathbf{F} = \left(\frac{\partial F_z}{\partial y} - \frac{\partial F_y}{\partial z} \right) \hat{\mathbf{x}} + \left(\frac{\partial F_x}{\partial z} - \frac{\partial F_z}{\partial x} \right) \hat{\mathbf{y}} + \left(\frac{\partial F_y}{\partial x} - \frac{\partial F_x}{\partial y} \right) \hat{\mathbf{z}}$$

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

Cylindrical Polars: $\mathbf{x} = (\rho \cos \phi, \rho \sin \phi, z)$

$$\nabla f = \frac{\partial f}{\partial \rho} \hat{\boldsymbol{\rho}} + \frac{1}{\rho} \frac{\partial f}{\partial \phi} \hat{\boldsymbol{\phi}} + \frac{\partial f}{\partial z} \hat{\mathbf{z}}$$

$$\nabla \cdot \mathbf{F} = \frac{1}{\rho} \frac{\partial(\rho F_\rho)}{\partial \rho} + \frac{1}{\rho} \frac{\partial F_\phi}{\partial \phi} + \frac{\partial F_z}{\partial z}$$

$$\nabla \times \mathbf{F} = \left(\frac{1}{\rho} \frac{\partial F_z}{\partial \phi} - \frac{\partial F_\phi}{\partial z} \right) \hat{\boldsymbol{\rho}} + \left(\frac{\partial F_\rho}{\partial z} - \frac{\partial F_z}{\partial \rho} \right) \hat{\boldsymbol{\phi}} + \frac{1}{\rho} \left(\frac{\partial(\rho F_\phi)}{\partial \rho} - \frac{\partial F_\rho}{\partial \phi} \right) \hat{\mathbf{z}}$$

$$\nabla^2 f = \frac{1}{\rho} \frac{\partial}{\partial \rho} \left(\rho \frac{\partial f}{\partial \rho} \right) + \frac{1}{\rho^2} \frac{\partial^2 f}{\partial \phi^2} + \frac{\partial^2 f}{\partial z^2}$$

Spherical Polars: $\mathbf{x} = (r \sin \theta \cos \phi, r \sin \theta \sin \phi, r \cos \theta)$

$$\nabla f = \frac{\partial f}{\partial r} \hat{\mathbf{r}} + \frac{1}{r} \frac{\partial f}{\partial \theta} \hat{\boldsymbol{\theta}} + \frac{1}{r \sin \theta} \frac{\partial f}{\partial \phi} \hat{\boldsymbol{\phi}}$$

$$\nabla \cdot \mathbf{F} = \frac{1}{r^2} \frac{\partial(r^2 F_r)}{\partial r} + \frac{1}{r \sin \theta} \frac{\partial(\sin \theta F_\theta)}{\partial \theta} + \frac{1}{r \sin \theta} \frac{\partial F_\phi}{\partial \phi}$$

$$\nabla \times \mathbf{F} = \frac{1}{r \sin \theta} \left(\frac{\partial(\sin \theta F_\phi)}{\partial \theta} - \frac{\partial F_\theta}{\partial \phi} \right) \hat{\mathbf{r}} + \frac{1}{r} \left(\frac{1}{\sin \theta} \frac{\partial F_r}{\partial \phi} - \frac{\partial(r F_\phi)}{\partial r} \right) \hat{\boldsymbol{\theta}} + \frac{1}{r} \left(\frac{\partial(r F_\theta)}{\partial r} - \frac{\partial F_r}{\partial \theta} \right) \hat{\boldsymbol{\phi}}$$

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