# FIVE PHOTONS

## REMARKABLE JOURNEYS OF LIGHT ACROSS SPACE AND TIME

JAMES GEACH

REAKTION BOOKS

*For Sophie and Lucy,*
*just starting their journey.*

Published by
REAKTION BOOKS LTD
Unit 32, Waterside
44–48 Wharf Road
London N1 7UX, UK
www.reaktionbooks.co.uk

First published 2018


Printed and bound in Great Britain by
TJ International, Padstow, Cornwall

A catalogue record for this book is available from the British Library

ISBN 978 1 78023 991 0

# CONTENTS

ONE **WHAT IS LIGHT?** *7*

TWO **OLD LIGHT** *32*

THREE **STARLIGHT** *67*

FOUR **DARK ENERGY'S IMPRINT** *100*

FIVE **BLACK HOLE BEACONS** *129*

SIX **RADIO FROM THE COSMIC DAWN** *152*

**EPILOGUE** *174*

FURTHER READING *177*

ACKNOWLEDGEMENTS *179*

INDEX *180*

ONE

# WHAT IS LIGHT?

When I got up this morning it was still dark outside. Groggy with sleep, I negotiated my way downstairs, instinctively treading the steps and running my hand down the wall. At the bottom of the stairs I felt for and flipped a switch. Almost instantly, light raced out from two bulbs on the ceiling. Within a few billionths of a second that light started smashing into all the various things in the room: the wood laminate floor, the yucca plant in the corner, the sofa, the television, my daughters' toys scattered here and there.

Some of the light hitting all that domestic paraphernalia bounced back towards me. Then, about ten nanoseconds after it left the bulb, a small fraction of that reflected light was intercepted by two small apertures in my head, passing through a couple of slightly squashy, transparent and rather usefully shaped blobs of biological tissue on the way. This subtly altered the direction of the light's path, focusing it on certain photosensitive cells at the back of my eyes and triggering a response that was sent down my optic nerve to my brain through a bioelectrical stimulus. Quickly decoding this stream of information, my visual cortex rendered the scene for me. I could see the room.

Our experience of light is simple and intimate: we can see it. Light interacts with objects in our environment in different ways, and this allows us to distinguish, say, an antique oak table from a plastic chair, and polished metal from fur. Colour is maybe the most obvious example of this interaction. The bulbs in my living room emit 'white light', which is a mixture of every colour of the rainbow. In physics we'd call it a broad spectrum light source. But the things I see around me that are illuminated by that white light have different colours and shades. So how do we go from a broad spectrum source like a light bulb to the huge range of colours we see in the environment illuminated by that source?

The answer is in the way different materials absorb and reflect the incident light. If an object appears blue then it is absorbing all the colours in the white light *except* for blue, which is reflected. If an object appears totally black, then it is reflecting none of the light hitting it. White surfaces reflect all of the incident light.

We are so familiar with the stuff, but what *is* light? How does it travel from one spot to another? What does it actually mean for light to have colour, or to be reflected or absorbed by an object? Like many simple questions about the world we live in, these run deep.

In physics, when we first start learning about light and its properties, we usually begin with the so-called 'laws' of optics. These are the rules of how light travels through and between different media, such as air, glass or water. We think about light 'rays' that travel in straight lines from their origin as they traverse space. The behaviour of those rays can be described using some fairly simple rules. I'll give you an example. Take a flat mirror. A ray of light hitting the mirror will be reflected off the silvered surface, bouncing back. This we know pretty well. But the reflected ray doesn't bounce off the mirror randomly, it bounces off at an angle exactly the same as the incoming ray relative to the 'normal', which is an imaginary plane perpendicular to the reflective surface. This is called the law of reflection.

There are other effects that we encounter every day. Take a light ray passing from one medium to another, for example. If you have ever put a straw in a glass of water, you'll have noticed that the straw appears to be bent, or disjointed, where it enters the water. Of course, the straw itself has not actually broken, but your eye perceives it as so because light rays from the submerged part have to travel through the water, then into the air towards your eyes. Compare this journey to that of the light rays coming from the top of the straw sticking out of the water: they only have to travel through the air to get to your eyes.

As a ray encounters a new medium, it can change direction slightly. This is called refraction. The reason the ray alters direction is because the speed of light can change in different materials. Yes, the speed of light is a fundamental constant, but that refers to light travelling in a pure vacuum. In glass, for example, light travels at about two-thirds of its maximum speed. You can see an analogous effect in water waves: a set of parallel waves approaching a shore can change direction depending on the depth of the water. Since ocean waves travel faster in deeper water, any part of the wave passing over shallow water – say, a sandbar – will slow down and lag behind, introducing a bend in the wave front. The direction of propagation of the waves is deflected.

In our straw-in-glass example, what you *perceive* as the source of the light – that is, where your brain interprets an object to be – is the apparent straight-line origin of each ray that enters your eye. Since the rays coming from below the water line get deflected as they pass from water to air, it appears that the submerged bit of the straw is slightly displaced compared to where you expect it to be.

Refraction also explains rainbows. If I hold up a chunk of glass to the sunlight, it might cast a rainbow on the wall. The

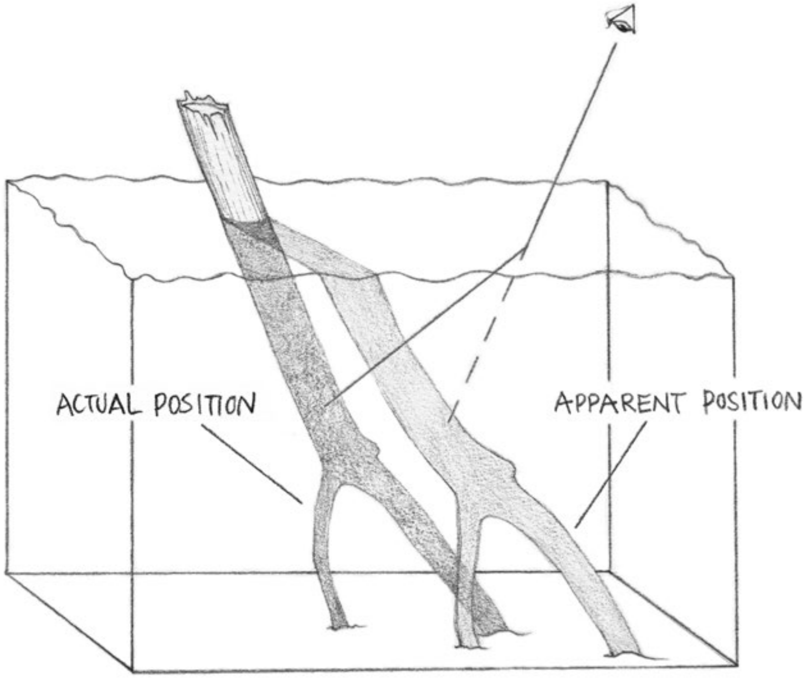


**Refraction**
*When light passes from one medium to another, its speed can change. This alters the direction of propagation of the ray, called refraction.*

refractive index of the glass, or the amount that light slows down as it passes through it, is slightly different for different colours of light. Sunlight, like our light bulb, is made up of a broad spectrum of colours, and so you can think of sunlight as a collection of different rays of light, each a different pure colour. When sunlight is refracted, the colours get dispersed into a rainbow because each ray is deflected by a slightly different angle, and the size of that angle depends on the colour.

Those were some phenomenological descriptions of the behaviour of light. If you wanted to design an optical system – say a pair of spectacles – then these rules would serve you pretty well. But they don't describe exactly *what* light is. Can we dig a little deeper? I used an ocean-wave analogy above for good reason: light is a type of wave. It's an *electromagnetic* wave.

What does this mean? Electromagnetism is one of the four fundamental forces of nature, sitting alongside the force of gravity and the 'strong' and 'weak' nuclear forces. Those nuclear forces govern the structure of the nuclei of atoms, the building blocks of the material world, and they operate over extremely short distances. Gravity, as we all know, is the attractive force between any two objects in the Universe with mass, and can act over infinite range. It is the force that holds us onto the Earth, keeps the Earth in orbit around the Sun, and generally determines the overall distribution of mass in the Universe, from solar systems to clusters of galaxies. Electromagnetism is a force that acts between particles with 'charge'.

Charge is a fundamental property of subatomic particles: a particle can either be positively charged, negatively charged or have zero charge (we call that 'neutral', but we will ignore those for now). Particles with the same charge repel each other, and particles with opposite charges are attracted. A good analogy is the attraction and repulsion of the poles of a magnet. Like the force of gravity, the strength of the electromagnetic force scales with the separation of the charged particles, following an 'inverse square law': double the distance between two charged particles and the strength of the force between them drops by a factor of four. Halving the distance increases the strength of the force by a factor of four, and so on.

An example of a positively charged particle is a proton. Protons and neutrons (subatomic particles with zero charge) make up the dense nucleus of an atom. Actually, protons and neutrons are each made from groups of three subatomic particles called quarks, but we won't go down that rabbit hole here. Suffice it to say, for now, we can think of protons and neutrons as distinct particles. You may wonder, if the electromagnetic force repels particles with the same charge, how does the nucleus stay together? Shouldn't all those protons repel each other to disastrous effect? That's where the strong nuclear force comes in: it acts over very short distances to glue protons and neutrons together, and is stronger than the proton–proton repulsion on those small scales. So, overall, the nucleus of an atom can be considered a positively charged particle, where the total positive charge is set by the number of protons in the nucleus.

Surrounding the positively charged nucleus are the electrons, typically in equal number to the protons. Electrons are another type of subatomic particle, but negatively charged, and about two-thousandths of the mass of a proton. A proton carries a charge of +1, and an electron carries a charge of −1. Their charges balance. So the net charge of an atom – the sum of the charges of the electrons and protons – is zero, or neutral. By successively removing electrons from an atom, which can be done by giving an electron enough energy to break free from the attraction of the nucleus, the atom can become positively charged overall. We call this process 'ionization' and we refer to the electron-stripped atom as an ion. Conversely, negatively charged ions are formed by introducing extra electrons to an atom.

The force between charged particles plays a crucial role in nature: it binds atoms together into larger structures called molecules. The fact that you can't push your thumb through the palm of your hand is because of the 'electrostatic' bonds between molecules in your skin, muscle and bone.

A simple example of the formation of such bonds can be found in common salt. Salt is just the colloquial term for the molecule sodium chloride, which fundamentally is comprised of pairs of sodium and chlorine atoms. These atoms bond in the following way: under certain conditions, sodium can donate one of its electrons to the chlorine atom, each becoming an ion in the process. The sodium has a net positive charge because it has lost an electron, and the chlorine has a net negative charge because it has gained one. The two oppositely charged ions attract each other, and we call this an ionic bond. Countless sodium and chlorine atoms can bond together this way, arranging themselves into a regular lattice that forms salt crystals.

A different type of bond occurs when two atoms share one or more electrons. Generally, the electrons are bound to the nucleus of one atom, but when two atoms are in close proximity some of each of their retinue of electrons can be shared between them. Essentially these moonlighting electrons enjoy the attraction of both nuclei and as a result bond the two atoms together. Molecular oxygen, or $O_2$, is an example of this type of 'covalent' bond. Molecules can bind together en masse in different combinations to assemble the material world we see around us. The exact mixture and arrangement of atoms, each with different numbers of protons, neutrons and electrons, determines the properties of these materials, and it's the electromagnetic force that holds everything together.

So where do electromagnetic *waves* come into this? Every charged particle generates a 'field' around it. This field is really just a way of describing the strength of the force around a charge and the effect it would have on other charges nearby. We usually sketch the field as a set of lines radially emanating from a charged particle like spokes, stretching away to (in principle) infinite distance. Near the particle, the density of the lines is high, and this means the field is strong. As we move away from the particle the lines are more spread out, meaning the field is getting weaker. Of course, in reality the lines don't really exist, they are simply a visualization of what we call in general a 'vector field'. They describe the path of a charged 'test' particle placed somewhere within the field. For example, an electron placed near a proton would feel the force and accelerate towards the proton along a field line.

In this basic picture we are considering stationary – static – charges and fields, hence the word 'electrostatic'. What happens when the charges move?

Imagine taking an electron and wiggling it about. Like a bug trapped on a pond, the agitated electron causes the electric field around it to ripple accordingly. The electric field is no longer static – it is moving. We are now talking about electrodynamics.

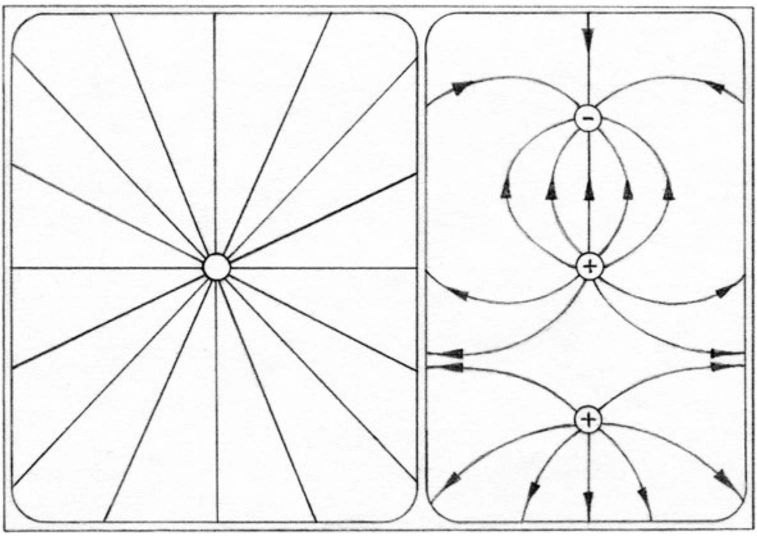


**Field lines**
*We visualize the electric field around charged particles with 'field lines', describing the strength and direction of the electrostatic force felt by other charged particles within the field.*

This is where the magnetism part of electromagnetism comes in, and it is key for the story of light.

James Clerk Maxwell (1831–1879) is considered the founder of what we now call the classical theory of electromagnetism. Maxwell, a Victorian physicist from Edinburgh, took forward the ideas of Michael Faraday (1791–1867), another pioneer of the field, and others to 'unify' electricity and magnetism, demonstrating how they are inexorably tied together. Maxwell showed that an electric field that is changing with time will induce a magnetic field, and vice versa. He also showed how the behaviour of those fields relates to charged particles.

Maxwell's achievement can be written in four elegant equations – Maxwell's Equations – that describe the 'classical' properties of the magnetic and electric fields and how they are linked. There's no need to write them down here, but the bottom line is this: Maxwell's Equations express something fundamental about the Universe. In his mathematical expression of electromagnetism, Maxwell showed that oscillations in an electric field generate an associated oscillating magnetic field, and these oscillations propagate away from their origin through 'free space' – that is, through the empty Universe – like a wave. This wave transports energy through space, which we call electromagnetic radiation. This is light.

How fast do these electromagnetic waves propagate? Using Maxwell's Equations it is possible to derive a 'wave' equation, which is a general expression for describing the properties of a wave – be it a water wave or an electromagnetic wave – in time and space. In the wave equation derived from Maxwell's work there appears a numerical constant, given the symbol *c*. It stands for *celeritas* and refers to the speed of propagation of electromagnetic waves through free space. *c* is the speed of light through a vacuum, also known as the speed limit of the Universe: two hundred and ninety-nine million, seven hundred and ninety-two thousand, four hundred and fifty-eight metres per second (299,792,458 m/s).

Just like waves on water, there is a simple way to characterize electromagnetic waves: by their wavelength or, equivalently, by their frequency. Both are related to the energy of the wave. The wavelength is simply the physical distance between two consecutive wave peaks. The frequency is the rate at which successive wave peaks pass some fixed reference point, and we measure that in units of wave cycles per second, also known as hertz (Hz). So, if the speed of a wave is constant, as it is with an electromagnetic

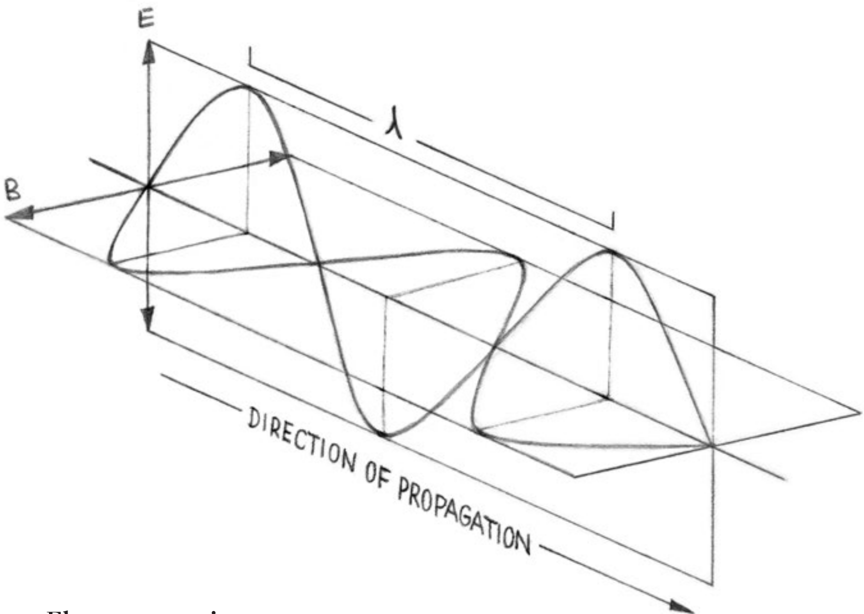


**Electromagnetic waves**
*Oscillations in an electric field (E) will cause an associated oscillating magnetic field (B) perpendicular to it. These coupled oscillations propagate through space as an electromagnetic wave – light. We characterize the wave by the distance between peaks (the wavelength) or the rate of the oscillation of the fields (the frequency).*

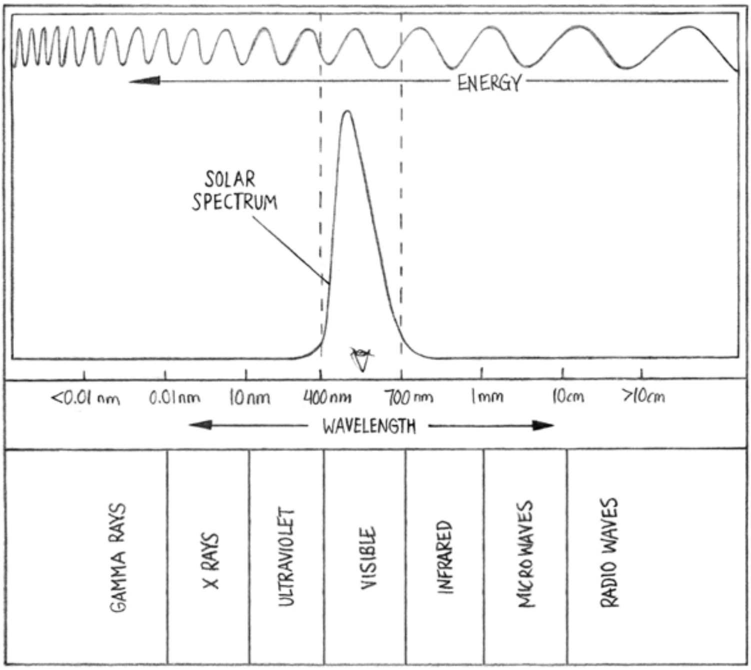


**Electromagnetic spectrum**
*Electromagnetic waves, or photons, come in a wide and continuous range of energies, characterized by their wavelength (or, equivalently, frequency). The range that humans can see roughly matches the peak range of energies of light emitted by the Sun.*

wave, then you can see that a longer wavelength corresponds to a lower frequency, and vice versa.

As we increase the energy of a wave, we shorten its wavelength and increase its frequency. Imagine taking the end of a long rope and wiggling it up and down at different rates: you will create 'waves' going down the rope with different wavelengths. Your arm is like a source of waves: wiggle your arm more vigorously and you'll increase the frequency of rope-waves. In nature we encounter electromagnetic waves – radiation – with a very wide range of energies depending on the source. In fact, we refer to an unbroken 'continuum' of energies of electromagnetic radiation, like a radio transmitter you could dial to any frequency. We call it the electromagnetic spectrum.

Although this is a smooth continuum, the range of energies is so large that it has been convenient to split the electromagnetic spectrum into chunks, and we give each chunk a label. At the lowest energies we have radio waves, which are electromagnetic waves with wavelengths of a few centimetres up to a kilometre or more. Radio waves can occur naturally, particularly in astrophysics, but we have also learned to generate and manipulate these waves for practical uses, the most obvious being in communication. Oscillating electric and magnetic fields will cause charged particles within the fields to also oscillate, and so a radio wave passing through an aerial will make the electrons within it respond. This response generates an electric current that can be measured by a receiver. By encoding information in the transmission of radio waves, we can communicate things like radio and television programmes over long distances.

Moving to a higher energy we come to the microwaves, which have wavelengths of millimetres to a few centimetres. These waves can cause water molecules in food to become agitated. The reason is, again, due to the response of charged particles to electromagnetic fields. Water molecules, made of one oxygen atom and two hydrogen atoms, are called 'polar' molecules because one end of the molecule is slightly positively charged and the other end is slightly negatively charged. This imbalance, called a 'dipole', means that when they are subjected to an oscillating electromagnetic field of just the right frequency, the water molecule will rotate. This molecular rotation is a form of energy, which is dissipated through the rest of the food as thermal energy, cooking it.

Continuing our journey up the electromagnetic spectrum, after microwaves we come to infrared radiation, with wavelengths of about one millimetre down to a few thousandths of a millimetre. This is quite a wide range, so infrared radiation is split into three sectors: far-infrared, mid-infrared and near-infrared. Those prefixes refer to the difference in energy from the visible light part of the spectrum, with near-infrared light lying just beyond the reddest light we can see.

It is useful to split the infrared part of the spectrum in this way because there are rather a wide variety of sources of infrared radiation in astrophysics. Any object with a temperature greater than a few tens of degrees above absolute zero will emit infrared radiation, and we often call this 'thermal' infrared emission. Once again, the reason is down to the motion of the particles within a heated body. To have a 'temperature' implies that the atoms and molecules are agitated: jostling around with thermal energy. At a temperature of absolute zero, the particles are stationary, but turn up the heat and they start to move, at least within the shackles of their bonds. The higher the temperature, the more violent the jostling. All these moving charges cause oscillations in their electric fields, generating electromagnetic waves that propagate into space, carrying away the thermal energy. This is the infrared glow. Low-temperature objects a few tens of degrees above absolute zero will emit long-wavelength, far-infrared light, and as the temperature increases the emission moves to shorter wavelengths, through the mid-infrared and then into the near-infrared.

Then we come to the visible part of the electromagnetic spectrum. Humans, and of course other animals, have evolved to see electromagnetic radiation with a wavelength range of about four hundred to seven hundred billionths of a metre (one billionth of a metre is called a nanometre). Think about what this means for a moment. Inside your eyes are cells that can actually respond to a small set of the electromagnetic waves that criss-cross space. More than this, the cells can transfer this response to your brain, which can decode the stimulus into meaningful information: images.

Beyond the bluest part of the visible spectrum we have the higher energy ultraviolet radiation, with wavelengths of tens to a few hundred nanometres. Similar to the infrared bands, this is split into 'near' and 'far' ultraviolet, with near ultraviolet light just beyond the bluest light we can see. We have all heard about the damaging effects of ultraviolet radiation from the Sun. It is damaging because of the energy of the wave; when it smashes into something – say some biological tissue – the energy being transported through space can be transferred to cellular matter. Sometimes this is to destructive effect, giving you sunburn, or worse, damaging molecules of DNA, which could go on to cause cellular mutation.

Ramping up the energy even higher, ultraviolet radiation gives way to X-ray radiation, and beyond this we encounter the gamma rays, with wavelengths of one hundredth of a nanometre and lower. This is where things get really dangerous. X-rays are useful because they can penetrate soft materials easily. Ordinarily you can't see into your hand, because visible light is absorbed by and reflected from the surface of your skin. X-rays, however, go straight through. We can use this to take images of the insides of our bodies, typically revealing the bones, which are more opaque to X-rays than skin and muscle. But this penetration can be a problem: like ultraviolet light, the high-energy X-rays can also cause damage to cells inside our bodies when their energy is deposited. Sometimes this might actually be desirable – we can use focused high-energy electromagnetic radiation to try to kill cancerous cells, for example.

On Earth, gamma rays are typically associated with radioactive elements and are simply a more extreme cousin of the X-rays. Exposure to such radiation can cause severe damage, and so sources of gamma rays must be heavily shielded by thick layers of dense material, such as lead, that can intercept as many of the rays as possible before they can do any harm.

We can imagine all these waves around us, flying this way and that in three dimensions, an ocean of oscillating electric and magnetic fields washing over us. These waves have different sources, which determine their energy. Most of them pass us by, or go straight through us, unnoticed. Some we can sense. We have harnessed some of them, transmitting, manipulating and detecting them for our benefit, be it transmitting tonight's episode of *The X Factor* or a session of radiotherapy. The point is, the waves are real; they are all over the place, travelling through space. That's what light – electromagnetic radiation – actually is. But there's a complication. There's another way of thinking about light: not as a wave but as a *particle*.

In the early twentieth century there was a revolution in our understanding of the natural world at its smallest level. You've probably heard of it: quantum mechanics. Now, quantum mechanics is a deep and complex subject, and unfortunately we don't have the time to delve too deeply into its wonders. Let's leave that for another time. The important thing is that quantum mechanics provides us with a framework for describing light and its interaction with matter that runs to a far deeper level than the 'classical' picture of electromagnetic waves.

The central principle of quantum mechanics is easy to grasp: energy, including electromagnetic radiation, comes in discrete chunks, called quanta. It is these quanta that can be thought of as particles. The idea that light is comprised of a flow of particles is not new. In the seventeenth century, Isaac Newton proposed that light was made up of 'corpuscles', or infinitesimal particles, and of course Newton was no crank. Unfortunately, this early particle model could not explain some of the observed behaviours of light, such as the patterns that are produced when light shines through small apertures. Eventually the wave model of light, championed by Newton's rivals and contemporaries, including Robert Hooke and Christiaan Huygens, became the accepted model and dominated our 'classical' thinking for centuries. It was not until the classical theory of electromagnetism itself started throwing up problems that we realized that the wave model couldn't be the end of the story.

When observations of natural phenomena cannot be explained by current theory, we have an opportunity to refine our understanding of how the world works. In science we develop theories that make predictions to be compared with observations. If these predictions don't match the observations, then the theory is refined or rejected. Most of the time these refinements are subtle, but occasionally they can be revolutionary. This is what happened at the start of the twentieth century.

One of the problems with the classical picture became known as the 'ultraviolet catastrophe'. It sounds a bit more dramatic than it was. The name refers to an issue that arose when scientists started to think about the radiation that would be emitted by a hypothetical heated cavity – a sealed box with a tiny hole drilled in it. Think of it as a kind of oven. A special property of this cavity is that its walls absorb and then re-emit all the radiation hitting them, reaching what we call 'thermal equilibrium' with the electromagnetic radiation. We call this a 'blackbody'.

Electromagnetic waves continually bounce around the walls, being absorbed and re-emitted until eventually some of them escape through the hole, to be observed.

In the classical theory, when it is in thermal equilibrium, the cavity is filled with a series of 'standing' electromagnetic waves, which are waves in which the amplitude of the field is changing, but the positions of the peaks and troughs are fixed in space. We often imagine these waves as a set of strings inside the cavity that are attached at fixed points on opposite sides of the walls. We can vibrate these strings like a guitar, and they oscillate with different frequencies. The different frequencies are called 'modes' of oscillation. Now, the classical theory states that the average amount of energy associated with each mode is proportional to the temperature of the system, and that you can, in principle, fit an infinite number of modes of increasing frequency inside the cavity.

If the total amount of energy in the cavity was divided between the different modes of oscillation of the electromagnetic field in the way the classical theory predicted, then the spectrum of the light emerging from the hole would be expected to diverge towards higher frequencies. In other words, the theory predicted that the intensity of light being emitted by the cavity should increase with increasing frequency, such that the total integrated emission of radiation from the hole becomes infinite. This clearly wasn't the case in practice.

Although the classical theory did a reasonable job of modelling the spectrum of electromagnetic radiation escaping from the cavity at low frequencies, it went drastically wrong once you reached the frequencies corresponding to ultraviolet light. Actually, the spectrum of light coming from the cavity follows a rather particular distribution: the intensity rises to a

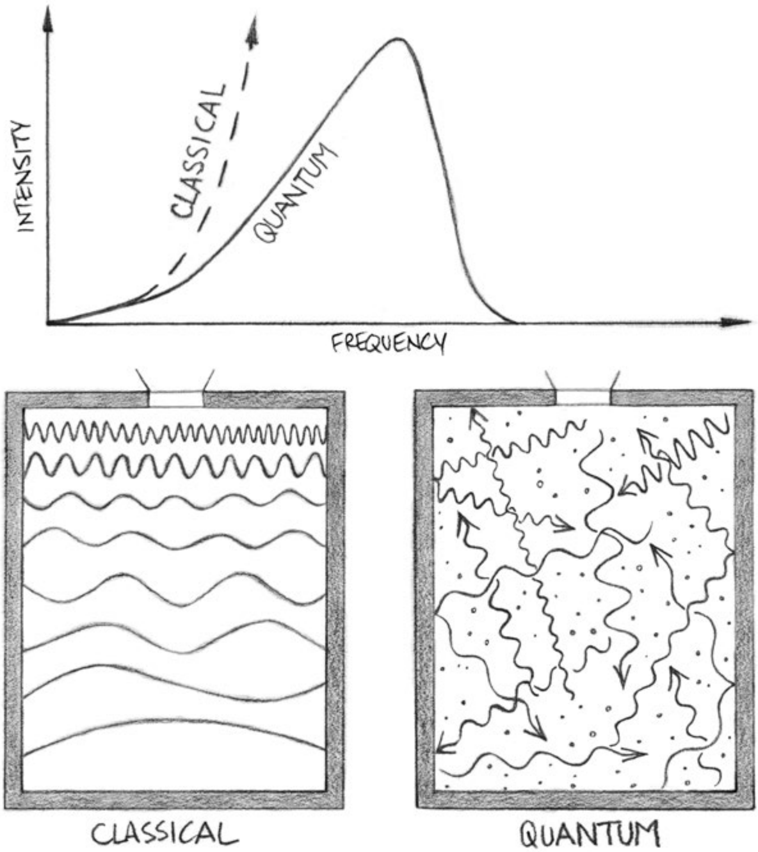


**Cavity radiation**
*A comparison of the classical and quantum views of electromagnetic radiation filling a hypothetical heated 'blackbody' cavity in thermodynamic equilibrium.*

peak at a particular frequency and drops away again at higher frequencies. The exact frequency of the peak intensity does depend on the temperature of the system, with hotter cavities emitting higher frequency (bluer) light, but the integrated, or total, emission is decidedly finite.

Max Planck (1858–1947), considered the father of quantum theory, proposed a solution right at the start of the twentieth