GRESHAM COLLEGE

17 October 2017

How Maths Can Save Your Life

Professor Chris Budd OBE

Introduction

Can Maths really save your life? Of course it can! Maths has applications to many problems that are vital to human health and happiness, ranging from curing cancer to powering our mobile phones. In this talk we are going to describe on particularly important application of mathematics to saving lives, namely its applications to imaging, medical and otherwise.

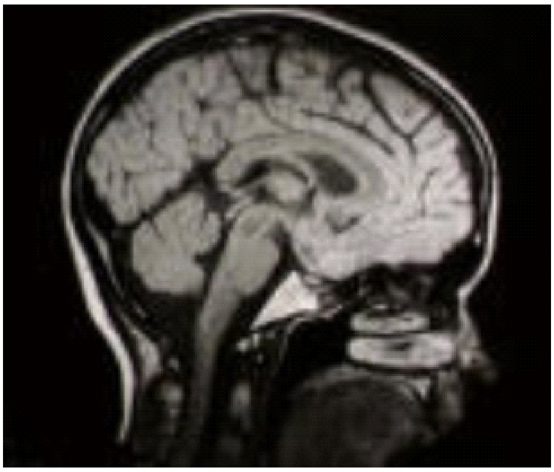

Modern medicine relies heavily on imaging methods which aim to find out what is wrong with you, without cutting you open first. The first such images used X-rays, going back to the discovery of them by Wilhelm Roentgen in 1895. By the start of the 20th Century. X-rays were being used in a variety of medical applications, and Marie Curie drove an X-ray ambulance during the First World War. Whilst X-rays were an important breakthrough, they only gave a limited view of the inside of a body. Essentially they showed you a shadow of the bones, and had limited information about soft tissue. In contrast, modern imaging methods do much better. Now only do they show a three dimensional reconstruction of the inside of the body, but they can also resolve soft tissue, such as brain matter. Advanced FMRI methods can even see what you are thinking.

Essentially these imaging methods take two forms. X-ray and ultrasound methods use a source of radiation that lies outside the body. The radiation is detected after it has passed through the body, and an image constructed from the way it has been absorbed. When X-rays are used, this process is called *computerised axial tomography* or CAT for short. Above we see a CAT scan of a head. (The word tomography comes from the Greek work *tomos* meaning "cut" or "slice".) This lecture will look at this, and other imaging processes in detail.

Other imaging methods use a source inside the body. These include *magnetic resonance imaging* (MRI and FMRI), *positron emission tomography* (PET), *single photon emission computed tomography* (SPECT), and EEG/MEG These methods have certain advantages over CAT, both in image resolution and in safety, as X-rays can easily damage soft tissue. The basic mathematics behind tomography was worked out by the mathematician Johann Radon in 1917. Much later, in the 1960s Allan McLeod Cormack, working in collaboration with Godfrey Newbold Hounsfield, developed the first practical scanning device, the celebrated EMI scanner. For this work, Cormack won the Noble Prize. Early models could only scan an object the size of a human head, but whole body scanners followed shortly after.

Medical imaging works through of a combination of very careful measurement techniques, sophisticated computer algorithms, and powerful mathematics. It is the mathematics that we will describe here. We will also show that the same mathematics has many other applications which help to save lives. These including imaging the atmosphere and improving the safety of air travel, detecting land-mines, locating oil, curing cancer, saving the whales, and slightly more frivolously, solving Sudoku puzzles.

Whilst being a very powerful application of mathematics, that saves countless lives, it is worth noting that the original work was done as an exercise in pure mathematics and it took over 50 years for it to become practical. We can speculate how many other areas of pure mathematics will have such a profound impact on people's lives in fifty more years.

Inverse Problems

Medical imaging is an example of what is called an *inverse problem*. Mathematicians distinguish between two types of problems, forward problems and inverse problems. In a forward problem you have all the information about a system and then try to predict what it will do next. A good example of this is dropping a ball onto the ground. Once you release the ball, by applying Newton's theory of gravity, you can predict (almost) exactly when it will hit the ground and when it will come to rest. An inverse problem is much harder to solve. In these types of problems you look at the evidence and try to work out what caused that evidence. For our ball problem this will amount to coming into a room, observing a ball on the ground, and then trying to deduce when it was dropped. The information is all there (in the vibrations of the floor, of the air, of the ball etc.) but it is much harder to use this to find out what has happened. Inverse problems are very common. In fact your brain has to solve an inverse problem in order to be able to make sense of what you are seeing as you read this transcript, or what you are hearing when you listen to a lecture. Forensic scientists have to solve an inverse problem when they examine a crime scene. For example working out where a bullet might have come from and what sort of gun fired it. Another example of an inverse problem is seismic prospecting for undersea oil, in which an explosion is set off at the surface of the sea and the location of the oil is deduced from measurements of the echoes of the compression waves from this explosion as they rebound from the rock underneath the ocean bed.

There are a number of reasons why solving inverse problems are hard. One is that information (such as in the example of the bouncing ball) can be lost rapidly, distorted or confused with noise as time increases. This means that it can be hard to measure and to deduce what is happening. Another is that very different causes can lead to very similar measurements. As an example, a cylinder and a rectangular block can both cast a square shadow when held at the right angle to a light source. It is then impossible from just examining the shadow to work out which shape you are looking at. This has enormous practical consequences. For example in medical imaging a doctor has to know for certain that what they are looking at is a true image of the patient and not an artefact of the imaging process in which a perfectly plausible solution has constructed which fits the data but is far from the actual reality. Mathematicians call problems of this form *ill-posed.* They are traditionally very hard to solve. Indeed it has been said that solving an inverse problems is like trying to work out the precise notes of a piece of piano music, when all you can hear is the music wearing ear muffs, through a thick concrete wall, and played by a musician wearing boxing gloves. However, the reliable solution of inverse problems is vital for medical imaging, radar, remote sensing, forensic science, archaeology, geography, weather forecasting, robotics and many other areas of science and engineering. An enormous amount of work has gone into the theoretical analysis of these problems and the consequent design of careful mathematical algorithms, and associated computer software, for solving them. When people talk about the *technology* behind medical imaging they often think of scanners, large magnets and machines that make loud buzzing noises. However, the real technology behind them lies in the effective mathematical algorithms that make them work. It is this *mathematical technology* that I will describe in this lecture. For more information about inverse problems see the summary in [1] and the excellent book [2].

Milk Deliveries, Killer Sudoku and Griddler

The first example of an inverse problem that we will look at is tomography, which is the mathematics behind the CAT scanner used to produce the image of the brain above. Before delving into the depths of medical science, we will start with a simple example which illustrates the principles of tomography, and which has a very nice link to the various types of Sudoku that have become very popular recently. This example involves milk deliveries. Imagine that milk and fruit juice is delivered in bottles that are placed in trays with 9 compartments arranged as a 3×3 grid. Each compartment of the tray contains a bottle, which may contain milk, juice or be empty. The question is: which type of bottle is in which compartment?

Unfortunately, other trays are above and beneath the one we're interested in, so we can't look down on top of the tray. At first sight it would seem impossible to solve this problem. However, we can peer in through the sides and we can measure how much light is absorbed in different directions. Different types of bottle absorb different amounts of light. Careful measurements have shown that milk bottles absorb 3 units, juice bottles 2 units and empty bottles one unit. If a light beam is shone through several bottles, then this absorption adds up. If, for example, a light beam shines through a milk bottle and then a juice bottle, then 5 units are absorbed. If it passes through three empty bottles then 3 units are absorbed.

In the example below we have indicated the total amount of light absorbed when shining light through each of the rows and each of the columns.

5			
6			
4			
	6	3	6

To solve this puzzle, we must place a bottle with 1, 2 or 3 units of light absorption in each compartment, with the sum of the units in the first row equalling 5, in the second row 6 etc. The middle column contains 3 bottles and also absorbs 3 units of light. The only way this can be done is for each compartment of the middle column to contain one empty bottle absorbing one unit of light each. What about the other compartments? Unfortunately we don't have enough information (yet) to solve this puzzle. Here are two different solutions:

3	1	1
2	1	3
1	1	2

2	1	2
2	1	3
2	1	1

We are faced with a rather unusual situation for a mathematician, in that we have two perfectly plausible solutions to the same problem. This is quite unlike a problem such as 1+2 which can only have the answer of 3. Problems like the milk bottle one are examples of the ill-posed problems we described earlier and are common in situations where we are trying to extract information from an image. To find out exactly how the bottles are distributed, we need to put in a little extra information. One obvious extra thing we can measure is the light absorbed in the two diagonals of the tray. We do this and find that 6 units are absorbed in the top left to bottom right diagonal, and 3 units in the bottom left to top right diagonal. From this extra piece of information it is clear that the first solution, and not the second, corresponds to the measurements made. It can be shown with a bit of extra maths, that if we can measure the light absorbed in the rows, columns and diagonals exactly, then we can uniquely determine the arrangement of the bottles in the compartments of the tray.

This problem may seem trivial, but it is very similar to the medical imaging problem we will describe in the next section, and shows how important it is to obtain enough information about a situation to make sure that we know what is going on exactly.

If any of this looks familiar to newspaper readers, then it is. *Killer Sudoku* is an advanced version of the popular Sudoku puzzle. In Killer Sudoku, as in Sudoku, the player is asked to place the numbers 1 to 9 in a grid with each number occurring once and once only in each row and column. However, rather than giving the player some starting numbers (as in Sudoku) Killer Sudoku tells you how the numbers add up in certain combinations. This is precisely the same as the problem described above. Griddler is another puzzle in which the reader is given information on the number of occupied squares in the horizontal and vertical columns of a grid. Solving a Griddler puzzle such as the one opposite uses many of the techniques used to solve a tomography problem.

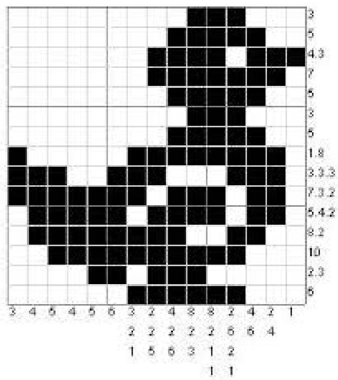

CAT and the Radon Transform

Until relatively recently, if you had something wrong with your insides, you had to be operated on to find out what it was. Any such operation carried a significant risk, especially in the case of problems with the brain. However, this is no longer the case; as we described in the introduction, doctors are able to use a whole variety of scanning techniques to look inside you in a completely safe say. A modern Computerised Axial Tomography (CAT) scanner is illustrated on the right.

In this scanner the patient lies on a bed and passes through the hole in the middle of the device. This hole contains an X-ray source, which rotates around the patient. The X-rays from this source pass through the patient and are detected on the other side. The level of intensity of the X-ray can be measured accurately and the results processed. The resulting fan of X-rays is illustrated in the following figure (with a conveniently circular patient).

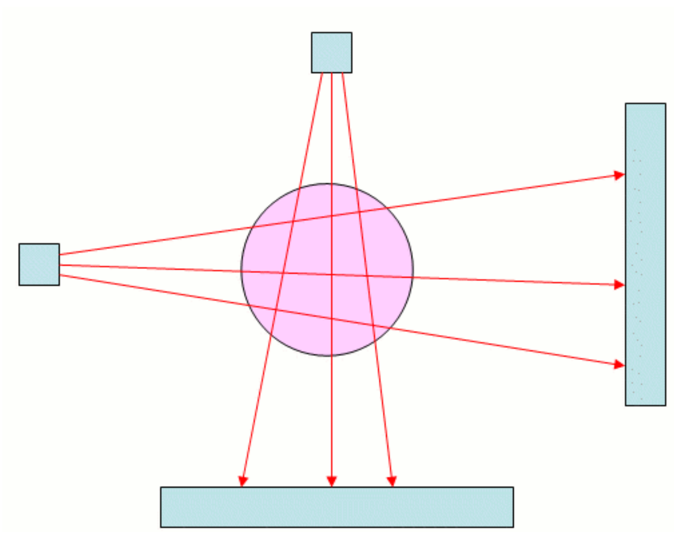

As an X-ray passes through a patient, it is attenuated so that its intensity is reduced. The degree to which this happens depends upon what material the ray passes through: its intensity is reduced more when passing through bone than when passing through muscle, an internal organ, or a tumour. A key step in reconstructing an image of the body from a set of X-ray measurements is to carefully measure exactly how different materials absorb X-rays. When an X-ray passes through a body, it does so in a straight line, and its total absorption is a combination of the amounts by which it is absorbed by the different materials that it passes through. To see how this happens, we need to use a little calculus. Imagine that the X-ray moves along a straight line and that at a distance s into the body it has an intensity $I(s)$. As s increases, so $I(s)$ decreases as the X-ray is absorbed. Now, if the X-ray travels a small distance $u(x,y)$ its intensity is reduced by a small amount δI. This reduction depends both on the intensity of the X-ray and the *optical density* $u(s)$ of the material. Significantly, the optical density tells us a lot about the properties of the body itself, and we can use this to image the inside of the body. Provided that the distance travelled is small enough, the reduction in intensity is related to the optical density by the formula

$$\delta I = -u(s)\, I(s)\, \delta s$$

Now, when the X-ray enters the body it will have intensity I_{start} and when it leaves it will have intensity I_{finish}. We can combine all of the contributions to the reduction in the intensity of the X-ray given by all of the parts of the body that it travels through. Doing this, we find that the attenuation (the reduction in the intensity) is given by

$$I_{finish} = I_{start}\, e^{-R}, \quad R = \int u(s)\, ds$$

This is the attenuation of one X-ray and it gives some information about the body. Below we see an object irradiated by several X-rays with the intensity of the rays measured on a detector. Here some X-rays pass through all of the object and are strongly absorbed so that their intensity (recorded at the centre of the detector) is low, while others pass through less of the object and are less strongly absorbed. Effectively the object casts a shadow of the X-rays and from this we can work out its basic dimensions. We illustrate this below.

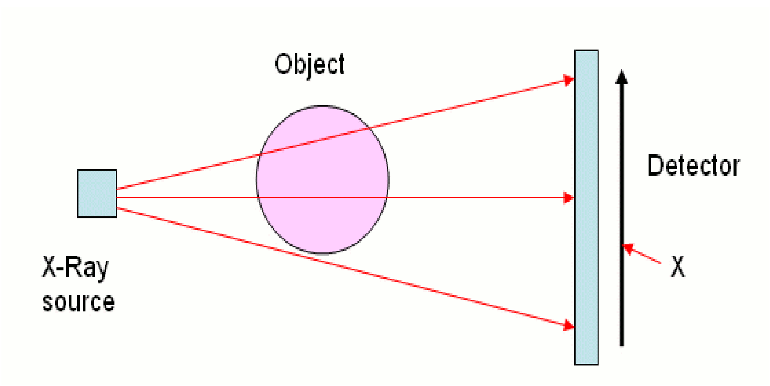

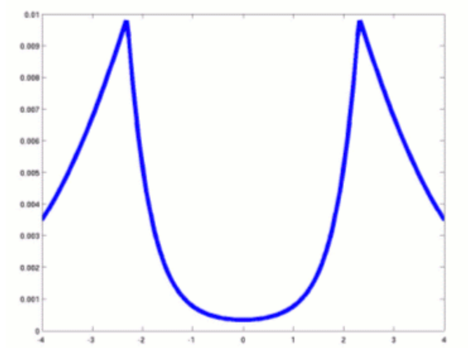

The intensity of the X-ray where it hits the detector depends on the width of object and the length of the path travelled both through the object and the air. This graph shows the intensity of the rays as they hit the detector. Rays that travel through the full width of the object have lowest intensity, as we can see from the dip in the middle of the graph. Rays that just miss the body have the highest intensity, because of all rays that are not absorbed they travel the shortest distance. This is what leads to the by the two spikes of the graph. Towards the edges the graph falls off, reflecting the fact that the corresponding rays have travelled a comparatively long distance. However, the secret to computerised axial tomography is to find out much more about the nature of the object than just its dimensions, by looking at the attenuation of as many X-rays as possible. To do this, we need to think of a number of X-rays at different angles θ and distances ρ from the centre of the object. A typical such X-ray is illustrated below.

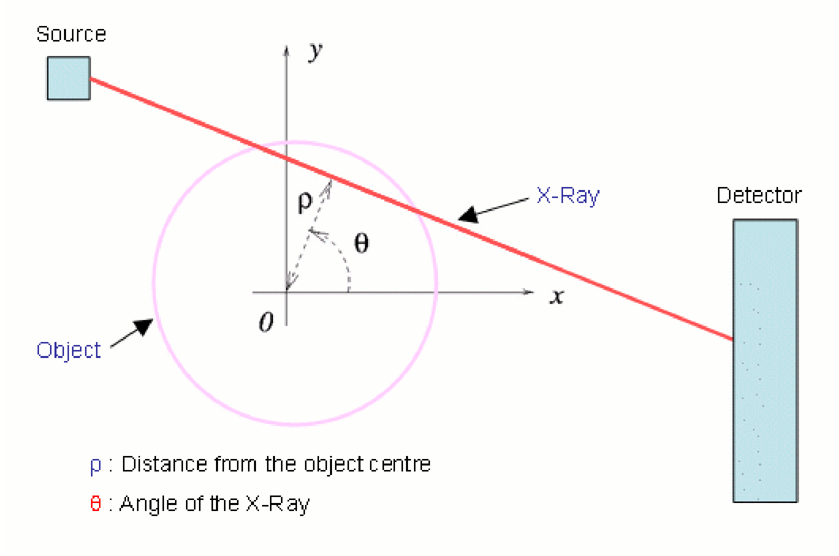

This X-Ray will pass through a series of points (x, y) at which the optical density is $u(x, y)$. Using the equation for a straight line these points are given by

$$(x, y) = (\rho\cos(\theta) - s\sin(\theta), \rho\sin(\theta) + s\cos(\theta)),$$

where s is the distance along the X-ray. In this case we now have $I_{finish} = I_{start} e^{-R(\rho,\theta)}$

where $$R(\rho,\theta) = \int u(\rho\cos(\theta) - s\sin(\theta), \rho\sin(\theta) + s\cos(\theta))\,ds$$

The function $R(\rho,\theta)$ is called the *Radon transform* of the function u(x,y). It is a map from the object to the shadow cast by that object. The larger the value of R, the more an X-ray of this particular orientation is absorbed. The Radon transformation lies at the heart of the CAT scanners and all problems in tomography. It was first studied by Johann Radon in 1917. (Radon is also famous for some very important discoveries related to the branch of mathematics called *measure theory*, which is the basis of the theory of integration.) By measuring the attenuation of the X-rays from as many angles as possible, it is possible to measure this function to a high accuracy. The big question of mathematical tomography is then the problem of *inverting* the Radon transform, in other words

Can we find the function u(x,y) if we know the function $R(\rho,\theta)$*?*

This is a classic inverse problem! In a medical scanner we need to solve this problems quickly and reliably in such a way that the results are clear to a clinician. Incidentally, this is exactly the same problem as was faced by our milk deliverer in the previous section. The short answer to this question is YES, provided that we can make enough accurate measurements. To show this rigorously is hard, and was part of Radon's great achievement. To then implement a practical algorithm is also hard, and took the combined forces of many mathematicians, engineers and computer scientists. However, a quick motivation will be given by the following example. In the two figures below we see on the left a square and on the right its Radon transform in which the large values of R are shown as darker points.

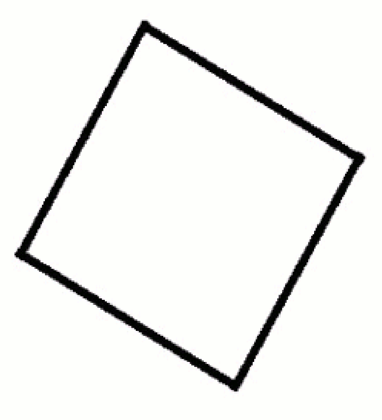

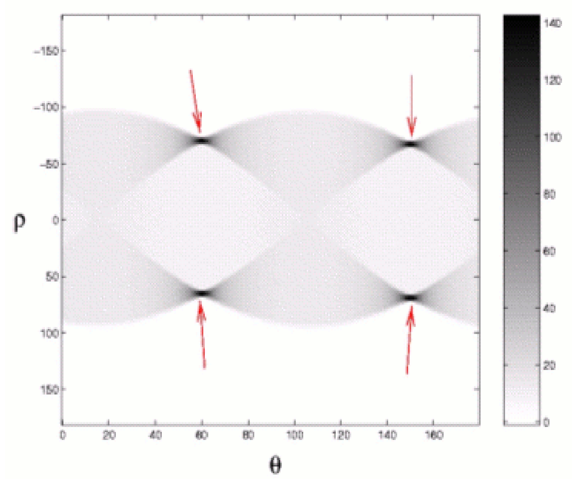

The key point to note in these two images is that the four straight lines making up the sides of the square, show up as points of high intensity (arrowed) in the Radon transform. The arrowed points give both the orientation of the lines and their distances from the centre of the square. The reasons that lines give large values for R at certain points is that an X-ray passing straight through a line is strongly absorbed, whereas one which misses it, even slightly, is hardly absorbed at all. Thus we can find the square by looking for the straight lines. Basically the Radon transform is good at finding straight lines in an image. One method for finding $u(x,y)$ called the *filtered back projection algorithm*, works (roughly) by assuming that the original image is made up of straight lines and drawing those corresponding to the high values of R. This method is fast but not particularly accurate. However, it is now possible to find u(x,y) much more accurately, reliably and quickly. The reason for this is twofold. Firstly it is due to the increases in computer power and in sensor technology. Secondly, and possibly more significantly, it is due to the development of advanced mathematical and computer algorithms (such as the *conjugate gradient method*) which are then implemented in the software of the scanning devices. Without the use of such algorithms these devices simply would not work. To see details of the various algorithms the book [3] is an excellent summary.

Saving the Bees

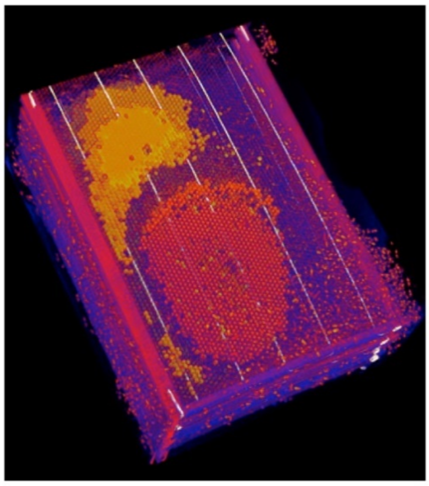

Mathematical research is constantly going on in order to improve these algorithms. This work is not only to make them faster, more accurate, and more reliable, but also to achieve the same results with fewer samples and also to deal with moving objects. Here we enter the rich and new field of *compressed sensing* [4]. A remarkable application of the use of these methods comes from efforts to save the bees. One of the major problems facing modern civilisation is the decline in the bee population. Einstein himself recognised the possible threat to us all if bees disappeared. Unfortunately, a combination of changes in land use, climate change and the effects of the Varroa mite, mean that the bee population is rapidly declining. One way to understand why this is happening is to study bees in the hive and to see how they respond to changes in climate and/or are affected by different diseases. The traditional way of doing this is to open the hive and have a look. Naturally this disturbs the bees. A better way is to study them without disturbing the hive. This can be achieved by using tomography. To do this scanners are placed around the outside of the hive and low intensity X-rays are shone through the hive. Two issues make this much harder than conventional imaging. Firstly, much lower intensities have to be used, and secondly, bees have an annoying habit of moving around whilst they are being scanned. This means that advanced algorithms need to be developed to deal with resulting lack of data. Fortunately (for the bees) these algorithms are successful and the image above shows not only the bees but also the details of the honey comb in their hive. For more detail see the work of Mark Greco in [5].

MRI and FMRI Imaging

Whilst CAT scanning is fast and reliable it has a number of disadvantages. Foremost amongst these is the way that it uses potentially harmful X-rays which should normally be avoided. Secondly, it has problems resolving fine details within the body, and certainly has trouble looking at changing physiology. Many of these problems are overcome by MRI, or *Magnetic Resonance Imaging* scanners. MRI scanners work by making the Hydrogen atoms in the abundant body water spin in a very high magnetic field. When subjected to a magnetic field, the atoms line up like tops in the direction of the field. Typically this means that they are pointing to either the patient's head or to their feet, with about half. Most of these cancel each other out, but there are a few which are not. A careful arrangement of extra *gradient magnets* alter the main magnetic field allowing slices can be taken of any part of the body Next, the MRI machine applies a carefully tuned **radio frequency pulse** directed towards the relevant part of the body. The RF pulse forces the unaligned hydrogen protons to spin in a particular direction. When the RF pulse is turned off, the protons slowly return to their natural alignment, and release the energy absorbed from the RF pulses. When they do this, they give off a signal that coils in the machine pick up and convert to an image using further careful mathematical algorithms. Many of the difficulties associated with inverse problems do not arise in MRI scanners (as the source of radiation comes from inside the body rather than outside). This means that the resulting images are often much better in an MRI scanner than in a CAT scanner. An MRI image is shown on the left below. Contrast with the CAT scan image at the start of this article.

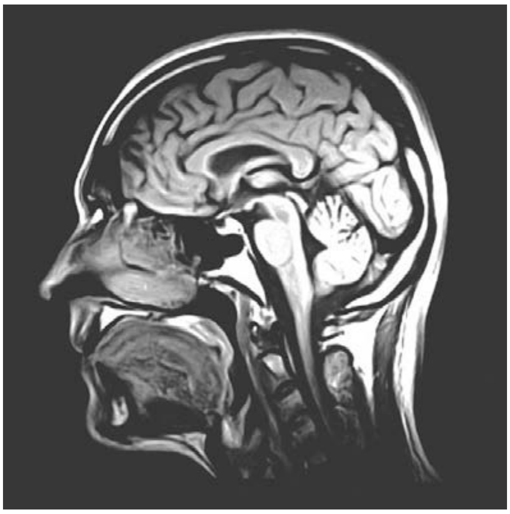

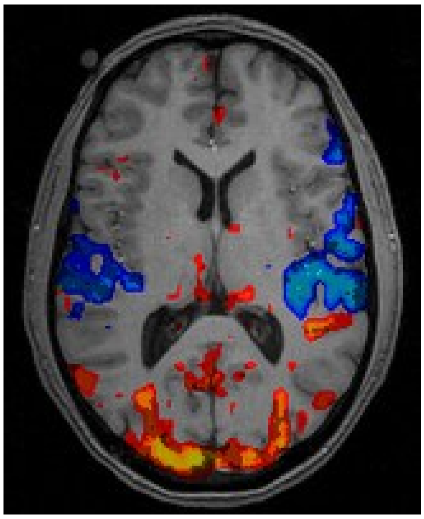

On the right we see a further MRI image obtained using the FMRI (functional MRI) process. In this the MRI process measures the amount of Oxygen and the consequent changes in blood flow in the different parts of the brain. As blood flow is related to brain activity, FMRI can see what parts of the brain are active at any one moment, and can thus, to a certain extent, see what you are thinking. See [6] for more detail.

Using Ultra-sound and MRI to Cure Cancer.

A lovely application of imaging technology (in reverse) leads to a direct treatment for cancer.
One of the most commonly used forms of medical imaging is ultra sound, which are sound waves of high frequency (much higher than we can hear). Sound is a compression wave through tissue and its speed $c(x)$ depends upon the density of the material. The sound intensity $u(x,t)$ then satisfies the *wave equation*

$$u_{tt} = c^2(x)\nabla^2 u$$

In ultra sound imaging a sound source emits a concentrated beam of sound from a transducer, which travels through a medium, such as human flesh. It is then reflected and absorbed by this medium before being received. By measuring the pattern of the received signals and solving an inverse problem involving the wave equation above, it is then possible to find the sound speed in different parts of the medium and then to work out the material density. By varying the location of the transducer a very complete picture of the material can then be built up. In the picture we see how this is done when ultrasound is used to image a foetus in the womb. Ultrasound is used here because sound waves are far less damaging than X-rays, and the procedure is far easier to implement than MRI.

An almost exactly similar process is used in seismic imaging in which the sound waves are provided by the explosion of a cylinder of compressed air. Similar methods are used in sonar (either from dolphins or man-made).

As with all inverse problems there is a trade-off between resolution and an accurate solution of the problem, the time and effort that we need to use to solve it and (especially in the case of imaging a foetus) the intensity of the ultra sound waves used. (Basically the more intense the signal the better that detail can be seen, but the higher the chance of damaging the infant). If too high an intensity of ultrasound is used then the result can be damage to the soft tissue. Usually this is a bad thing, but in one instance it can be useful. That is when the soft tissue concerned is a cancerous tumour. Basically, if a high intensity ultrasound signal can be carefully focused then it can destroy a tumour by burning away the tissue. Much as in the way that a magnifying glass can be used to start a fire by focusing the rays of the sun. In particular the source out the ultrasound can be a transducer outside the body, and it can pass through healthy body tissue before being focused on a tumour, for example in a bone marrow. With an appropriate choice of ultrasound frequency the point of focus can be as small as a grain of rice. Done appropriately this can destroy the tumour without damaging the surrounding tissue. So far, so good. But a problem arises when you try to focus the beam. Essentially the presence of bone and the unknown properties of the materials inside the body, make it very hard to focus the ultrasound beam in advance. To our rescue comes MRI imaging. By looking at the change in the *phase* of the MRI signal it is possible to tell the *temperature* of the tissue. In the MR-HIFU process (Magnetic Resonance – High Intensity Focused Ultrasound) the ultrasound is directed at the tumour and the point of focus is monitored using MRI imaging by looking at the temperature of the tissue. The point of focus can then be changed till it is on the tumour and then the intensity of the ultrasound increased to destroy the tumour. This method of destroying a tumour, and thus helping to cure cancer, is still under clinical trial, but holds promise as a new weapon for doctors in the fight to save lives.

Other Ways That Tomography Can Save Your Life

Tomography has many applications quite different from those in medicine. An interesting example comes from archaeology, where tomography was used to determine the cause of Tutankhamen's death. A CAT scan of the mummy revealed a swelling in the knee, indicating that death was the result of a massive infection. The cause of this was possibly an injury inflicted by a fall. Whether Tutankhamen was pushed or fell by accident, however, will have to remain a mystery, which even a CAT scanner cannot solve. More generally, we can apply tomography to any problem where we have information about the average of a function along a straight line. It can also be used to find evidence for straight lines in an image (such as the edge of an object). We will now describe two examples of how tomography is used to help save lives.

GPS Satellites and Flight Safety

Orbiting the Earth, are a large number of GPS satellites that are transmitting radio signals down to the ground. If you can detect the signals and find the phase difference between the signals from several different satellites, then you can work out your location with a high degree of accuracy. GPS positioning methods are very widely used by aircraft navigation systems, SATNAV devices, and hikers and their accuracy is now vital for many applications where safety is essential. (Details of how GPS systems work will be given in the next lecture in this series). However, one of the problems with this system is that variations in the *ionosphere* (the upper part of the Earth's atmosphere) can affect the radio signals and change their phase by small amounts. This phase change can lead to errors in the position given by the GPS system. These are not very large and are perfectly acceptable for navigating. However, when landing an aeroplane it is vital that its height is known to very high precision and even small GPS errors can have large consequences. Here an accurate understanding of the state of the ionosphere is essential. There are many other reasons why understanding the ionosphere is important. Chief amongst these is that fact that the ionosphere has a very significant effect on the propagation of radio waves and on communication in general. Given that we rely hugely on communication technology, understanding when or not it will work is of vital importance. Roughly speaking the effect of the ionosphere is that radio waves can either bounce off it, greatly increasing the range of a radio transmitter, or changes in the ionosphere can lead to radio interference.

Remarkably, it is possible to monitor the state of the ionosphere using tomography [7,8]. In the problem of imaging a patient we shone X-rays through their body. To image the ionosphere we use the transmissions from the GPS satellites. These form a very convenient set of "straight lines" passing through the ionosphere. The paths that they take are shown in the figure below.

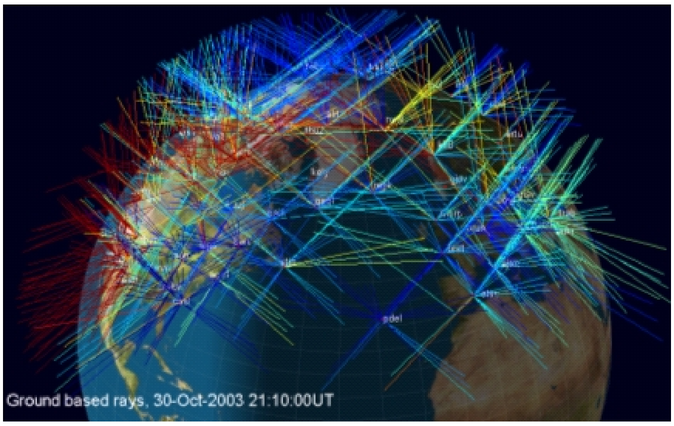

The phase of the radio waves is affected by the electron content of the atmosphere, so that the total change in the phase is proportional to the integral of the electron density along the ray path. If we can measure these phase changes, then we can estimate the electron density integrals and work out the Radon transform of the electron density. We seem to be in exactly the same situation as in the medical imaging problem and hence able to work out the electron density at any point in the atmosphere.

Well, not quite. There are two big differences between this problem and the CAT problem. Firstly, the satellites are usually moving relative to the Earth. Secondly, there are large parts of the Earth's surface where we cannot make any measurements. These include the oceans, where there are no receivers for the satellite signals, and the poles, which do not have satellites orbiting above them. Thus we have a lot less information than we had in the case of the CAT scanner. This means that we are often in the situation of the milk deliverer who couldn't distinguish between two different arrangements of milk bottles, each of which led to the same set of measurements.

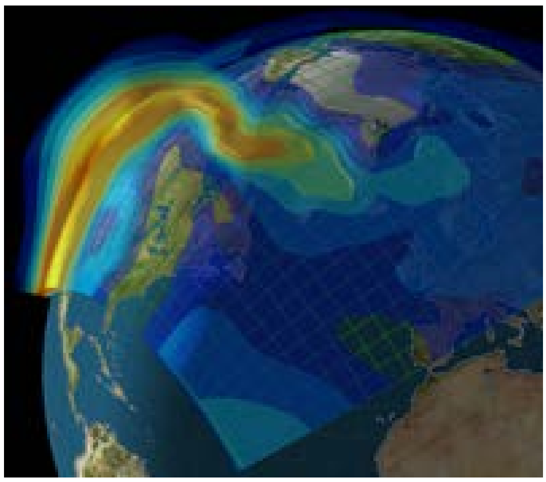

To get round this problem in the case of the ionosphere, we have to use a-priori information about the state of the ionosphere, or in other words a reasoned guess of what the solution should look like. This will allow us to reject one solution which doesn't look like this guess, and to choose the solution which looks as much like the guess as possible. Fortunately, we understand the physics of the ionosphere well enough for our reasoned guess to be pretty close to the truth. By doing this (together with some other clever refinements) it is possible to use tomography to find the state of the ionosphere. In the figure on the right we illustrate a calculation (using the MIDAS software developed at the University of Bath) of an ionospheric storm (in red) developing over the southern part of the USA. This is important as it could lead to a break down in communications over this region, and/or a loss of GPS signals, and an accurate understanding of the state of the ionosphere can be critical in this situation. More details of the methods used to produce this picture are given in [7].

Detecting Land-Mines

Anti-personnel land-mines are one of the nastiest aspects of the modern warfare. They are typically triggered by almost invisible trip-wires attached to the detonators. Any algorithm for the detection of trip-wires must work quickly and not get confused by the leaves and foliage that obscure the wire. An example of the problem that such an algorithm has to face is given in the figure below, in which some trip-wires are hidden in an artificial jungle.

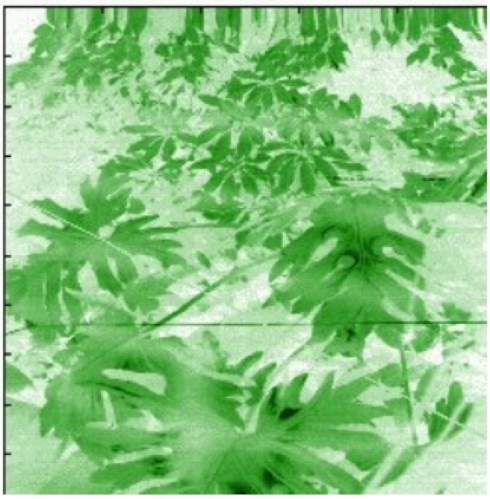

Finding trip-wires involves finding partly obscured straight lines in an image. Fortunately, just such a method exists; it is the Radon transform! For the problem of finding the trip-wires we don't need to find the inverse, instead we can apply the Radon transform directly to the image. Of course life isn't quite as simple as this for real images of trip-wires, and some extra work has to be done to detect them. In order to apply the Radon transform the image must first be pre-processed to enhance any edges. Following the application of the transform to the enhanced image a threshold must then be applied to the resulting values to distinguish between true straight lines caused by trip-wires (corresponding to large values of R) and false lines caused by short leaf stems (for which R is not quite as large). Following a sequence of calibration calculations and analytical estimates with a number of different images, it is possible to derive a fast algorithm which detects the trip-wires by first filtering the image, then applying the Radon transform, then applying a threshold and then applying the inverse Radon transform. The result of applying this method to the previous image is given below, with the three detected trip-wires are highlighted.

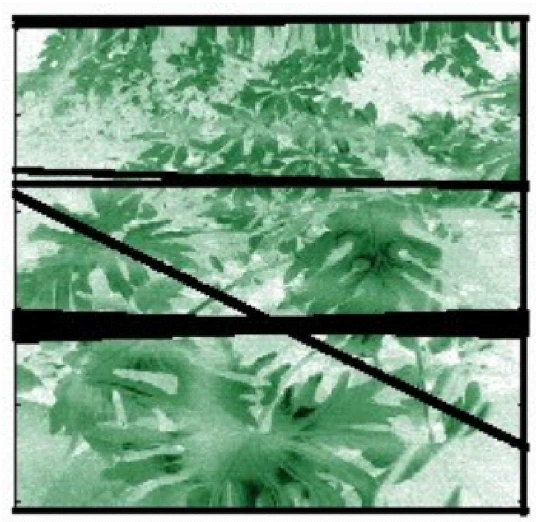

Note how the method has not only detected the trip-wires, but, from the width of the lines, an indication is given of the reliability of the calculation. This by using this algorithm we can find trip wires and hence get rid of land mines.

As advertised in the title, maths truly does save lives!

References

[1] C.J. Budd and C.N.Mitchell, *Saving lives: the mathematics of tomography*, (2010), **https://plus.maths.org/content/saving-lives-mathematics-tomography**.

[2] P.C Hansen, *Discrete inverse problems, insight and algorithms*, (2009), SIAM.

[3] J Mueller and S. Siltanen, *Linear and nonlinear inverse problems with practical applications*, (2012), SIAM

[4] S. Foucart and H. Rauhut, *A Mathematical Introduction to Compressive Sensing*, (2013), Birkhauser

[5] Mark Greco's work on bees is described in:
http://news.bbc.co.uk/earth/hi/earth_news/newsid_8682000/8682842.stm

[6] A noce account of MRI and FMRI is given in:
https://en.wikipedia.org/wiki/Functional_magnetic_resonance_imaging

[7] Mitchell, C. N. and Spencer, P. S. J., 2003. A three-dimensional time-dependent algorithm for ionospheric imaging using GPS. *Annals of Geophysics*, 46 (4), pp. 687-696.

[8] The Bath Invert Centre is described in http://www.bath.ac.uk/elec-eng/invert/