GRESHAM COLLEGE

7 JANUARY 2020

CAN MATHS SAVE THE WHALES AND CURE CANCER?

PROFESSOR CHRIS BUDD

"*On the planet Earth, man had always assumed that he was more intelligent than dolphins because he had achieved so much - the wheel, New York, wars and so on - whilst all the dolphins had ever done was muck about in the water having a good time. But conversely, the dolphins had always believed that they were far more intelligent than man - for precisely the same reasons.*"

Douglas Adams, The Hitchhiker's Guide to the Galaxy

1. Introduction

If were to ask anyone what are the great challenges for human kind of the 21st Century, it is very likely that curing cancer, and saving the environment would be close to top of most people's lists. I have already given a talk about maths and the climate (see November 2018), so here is a talk in which we see how mathematics plays a very important role in the fight against cancer. However, not wanting to ignore the question of the impact of mathematics on questions related to the environment I have decided to also talk about the ever popular issue of saving the whales. I admit that both of these are shameless plugs for the effectiveness of mathematics. Rather than being the useless subject, for which it is often portrayed, mathematics can be seen in this lecture to making a real difference to two questions of major concern. I would also hope that everyone would agree with me that these are both positive things. The world will undoubtedly be a better place if either of these two big challenges can be solved. Personally I hope that those in power with access to funding would want to support any initiative which cures cancer, and saves the whales. And this certainly includes support for the peaceful applications of mathematics.

Whilst saving the whales and curing cancer are clearly very important issues in their own right, they are only two of the problems for which any possible solution involves the applications of mathematics. Clearly mathematics on its own will not solve any practical problem. The main requirement of mathematics is that it should be self consistent within the statements of its axioms, and it should lead, by rigorous arguments, to amazing theorems, such as Pythagoras' Theorem, which represent eternal truths. To apply mathematics to the real world, mathematicians must work with scientists/engineers from other disciplines, to turn a real life problem into mathematics, and to then use (often very sophisticated) mathematics, to solve the resulting equations. We call this process mathematical modelling. In this lecture I will explain the basic ideas behind mathematical modelling, and will try to explain its power and its limitations. We will then look at how by using mathematical modelling we can both save the whales and help to cure cancer. We will even see a really nice application of Pythagoras' Theorem in the whale saving department.

2. The basic ideas behind mathematical modelling

2.1 Overview

Mathematical modelling is the process of describing a real world problem in mathematical terms, usually in the form of equations, and then using these equations both to help understand the original problem, and also to discover new features about the problem which could not have been predicted without the use of mathematics. Modelling both lies at the heart of much of our understanding of the world, and it allows engineers to design the technology of the future. With modelling we can travel to the edge of the universe, peer into the heart of the atom, and understand the future of our climate and the effects that climate change will have on our lives.

We are all very familiar with one application of mathematical modelling, namely the weather forecast. A modern weather forecast is based on the following steps

- Start with the laws of physics
- Encode these as (differential) equations, in particular the Navier-Stokes equations.
- Take data from satellites and weather stations to determine today's weather accurately.
- Using this as an initial condition, (using a super computer) solve the equations for 24 hours to give us the weather tomorrow.
- Continuously update the forecast by assimilating in data
- Present the results in a way that all can understand.

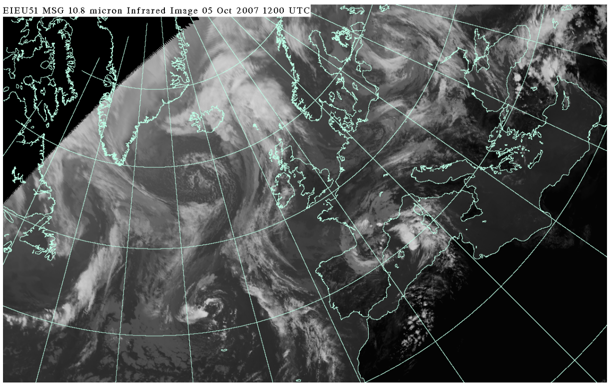

Despite rumours to the contrary this process works, and works well. At least for short term weather forecasting. This process is a special case of the more general process of mathematical modelling which can be described simply as:

1. Identify the problem i.e. talk to the people involved
2. Clarify the science
3. Formulate the science mathematically
4. Solve the mathematics possibly using a computer
5. Draw conclusions
6. Explain your results.

As well as weather forecasting, this process is used to design aeroplanes, design cars, design drugs, control the electricity supply network and, as we saw in my recent October 2019 lecture on computing, to establish the cause of the 1987 Kings Cross Fire.

2.2 Constructing a simulation

Strictly speaking what we have described above is a *simulation.* The difference between a simulation and a model is that in a simulation we are concerned to try to get all of the details as right as possible so that the conclusions are as accurate as possible. Using such simulations, we can for example, determine in advance whether a bridge will stay up after it has been built. We can also test the bridge to destruction without ever having to build it in the first place simply by varying the parameters in the computer simulation. Another important use of simulation is in the training of pilots in aircraft simulators, which are designed to be as close to reality as possible. Using these a pilot can be trained to fly an aircraft and to deal with dangerous situations, long before they have to enter the cockpit. One of my favourite examples of a simulator is the computer programme used by the line managers to control the operations of the London Underground. Simulators are also used (in a careful form to maximise speed at the expense of high accuracy) at the heart of many video games. Two examples of such, shown below, are simulators of the underground and of a factory.

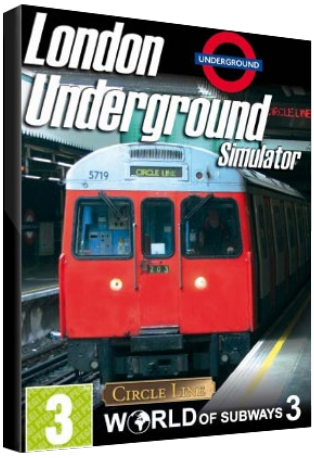

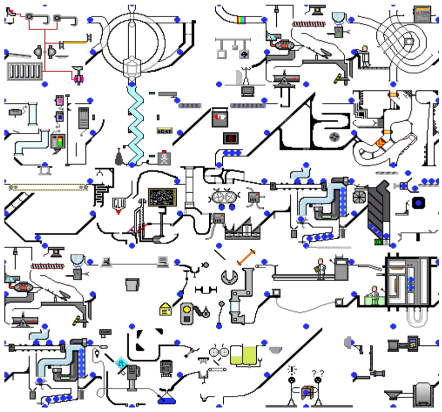

Whilst simulators are very useful they have big disadvantages. The need for high accuracy means that the equations are usually far too hard to solve analytically. Instead, they must (often) be solved by using large super-computers. The bigger the computer the better. These simulations often take a long time, consume a lot of energy, and produce vast amounts of data. So much data in fact that it is often hard to work out what is important and what is irrelevant. Furthermore it is hard to use the simulations to do 'what if' experiments as they take so long to run and are expensive.

A second disadvantage is that they tend to only work, and be applied to, problems where the basic science is well understood. One reason for this is that it is a very significant amount of work (in person hours) to write and code up a simulator. You do not want to put such an investment into a system which you do not understand well.

Despite these disadvantages, simulations currently have a big advantage over the data driven models that are often used in machine learning applications. By being based on well understood physical principles simulators have a level of trustworthiness and explain ability, which is not present in a machine learning algorithm. They are thus inherently more reliable, and better able to cope with unusual circumstances. It remains to be seen how long this advantage will continue.

2.3 The process of mathematical modelling

We contrast simulation with a **mathematical model.** This is a simplification of the problem to a small system of equations, which capture the essential essence of it, and, crucially, are simple enough to allow us to make analytical calculations. A formula derived from an analytical calculation can give a clear view of the role of the parameters in that system without having to run a very large number of calculations.

Perhaps the earliest example of a mathematical model of enormous predictive powers was Newton's law of gravitation applied to the solar system. Rather than modelling the whole system in all of its complexity, he treated the Sun and the planets as single points. This allowed him to write down the basic equations of motion of the whole of the solar system.

Solving these gave enormous insights. These included the understanding of the elliptical orbits of the planets, the conservation of angular momentum, and (Kepler's third law) the simple relationship between the cube of the period of motion of a planet and the square of its orbital distance illustrated here.

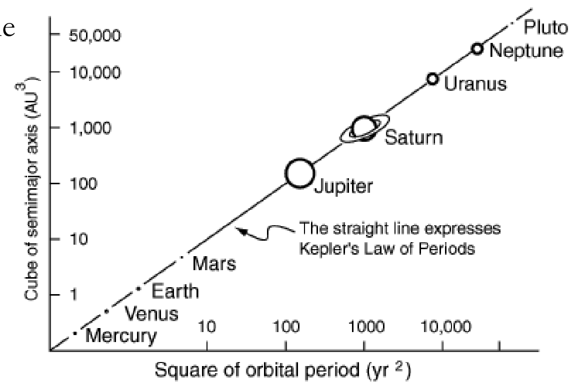

The process of simplification in constructing a model is hugely important, but also hard. A wonderful quote from Albert Einstein is:

A model should be as simple as possible, and no simpler

Einstein show strong insight here, but I'm sure that he was well aware that knowing when a model is just simple enough is very, very hard.

The 'traditional' , and often taught, approach to constructing a mathematical model is as follows:

1. Think about the problem in a mathematical way identifying all of the key ingredients
2. Write down the relevant equations, simplifying as much as possible
3. Solve the equations
4. Compare the results against data
5. If the results agree STOP
6. If not then modify the equations, such as making them more complex and including new processes
7. Repeat from 2 above.

This method of modelling is often taught in universities and at schools. However, there is one problem with this description of the modelling process. *In general it is completely wrong.* The reason that it is wrong is that it is very, very, rare to come anywhere close to writing down the right equations the first time. Indeed without looking hard at the data to start with, it is likely that the equations will not be anywhere close to the truth. The result is 'mathematical models' that might look nice are so far from the truth as to be practically useless. They are also often so simplified, that they also have no real mathematical interest either. In contrast true mathematical modelling plays close attention to the data at all stages of the process, employs computation at all times, and NEVER stops at line 5 above. A mathematical model is a living process, that if looked after well will continue to give insights into the system. Another problem with this approach is that point 4 is often very difficult. What does 'agreeing with data' really mean when it comes to a model of (say) loneliness. The best models are ones which give us excellent insight into the system which allow us to make useful future predictions. Quantitative agreement with actual data is often a bonus. Or to quote George Box [1]

All models are wrong, but some of them are useful.

In practice the modelling process (at least in my own experience of over 30 years of solving industrial problems using mathematics) is much more like the picture below:

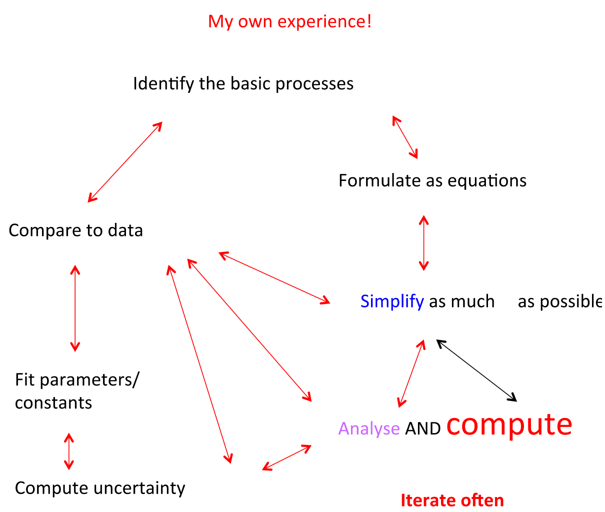

You will probably get the impression from the above that modelling is more of an art than a science. In this you would not be far wrong.

2.3 Examples of mathematical modelling

Physical Sciences and Engineering

These often start from quite well known laws, very like those used for simulators. These are then simplified. This is most often done by working out the size of different effects, and then ignoring those which are considered too small to make a big difference to the final result. This can either be done mathematically, or by experiments. Sometimes this is combined with assumptions about how the system will behave to simplify the number of possibilities that need to be studied. A good example of these forms of simplification occur in meteorology, when the effect of s0und waves (which are a perfectly valid solution of the underlying equations) are ignored as being too small. Physical models are often highly predictive, both qualitatively and quantitatively.

It is worth considering two examples, which show the real predictive power of this process. The first comes again from celestial mechanics. After Newton had written down the laws of motion for the solar system they were found to be very accurate in predicting the paths of the then known planets. However, on 13th March 1781 the planet Uranus was discovered by William Herschel. When its orbit was plotted it was found that whilst it nearly agreed with the predictions of Newtonian mechanics, there were small discrepancies. At this stage the trust in the accuracy of the underlying model was such that even these small discrepancies caused a large amount of concern. (This would not happen with a model from biology or the social sciences!) It was postulated that there must be a cause, and one explanation was that there was another planet, which was disturbing the orbit of Uranus. Using the Newtonian model for the solar system it was possible to calculate the location of this planet. This calculation was carried out independently, by John Couch Adams at Cambridge, and by Urbain Le Verrier in Paris. Both obtained very similar answers, and in response to Le Verrier's calculations Johan Galle using the telescope in the Berlin Observatory discovered the planet in 1846. We now call this planet Neptune, and it is illustrated below with Uranus on the left.

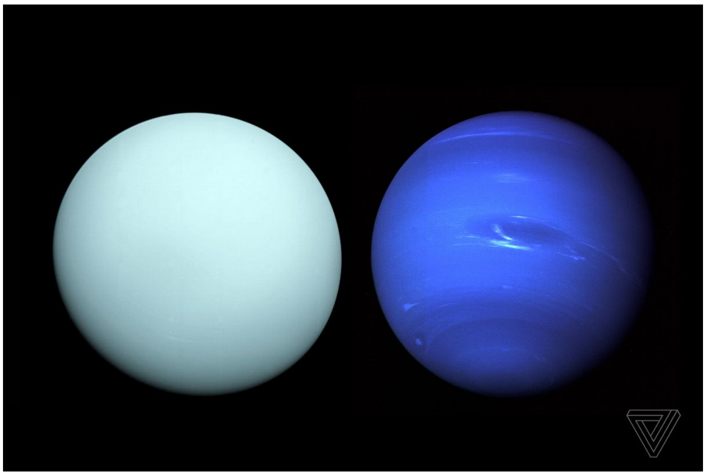

The key lesson from this example was that a mathematical model was able to *predict* something completely new, which was not built into the model in the first place. It became clear at this point that mathematical models had potentially extraordinary predictive powers.

The second example comes from climate science. I described the issues concerned with simulating the climate in a previous lecture [2]. In general climate models (or really simulations) are far too complex to be used to make analytical predictions. The best that we can do is to run them many times on different scenarios of, for example, Carbon Dioxide production and solar variability. However, a simple climate model which assumes that the energy arriving from the Sun balances the energy radiated from the Earth gives us an equation for the average temperature of the Earth T in form

$$T = \left(\frac{(1-a)\ S(t)}{\sigma e} \right)^{1/4}$$

Here S(t) is the average energy from the Sun, a is the albedo of the Earth, e is the effective emissivity of the Earth's atmosphere and sigma is Boltzman's constant. This is a true model. It is a complete simplification of a very complex system. However, it is predictive, and allows us to see how the temperature of the Earth will change if the energy from the Sun changes (varying S), the polar ice caps melt (varying a) of if Carbon Dioxide changes (varying e). It is also a good model in the sense that these predictions can be tested against experimental data.

Biological Sciences

It has taken longer for mathematical modelling to make an impact in the biological sciences. This is mainly because biological systems are inherently much more complex than physical ones (consider modelling a cell or the brain for example). It is also general harder to make, repeat, and quantitatively assess, experiments on biological systems. However, by a lot of careful study, significant progress has been made in recent years, starting with pioneering work by Alan Turing just after the Second World War. Mathematical models are now used to help understand changing animal populations, the evolution of biological patterns, the spread of disease and the functioning of the nervous system. Indeed there is now a (relatively) new subject called *Mathematical Biology* [3]. I will look at a significant example of the use of mathematical biology in the study of cancer in section 4.

Social Sciences

Modelling is both useful and important here, but is complicated by the fact that we have to take human behaviour into account. Given the complexities of the way that humans think and behave, this is generally hard, if not impossible. Personally, I am very suspicious about attempts to use mathematics model such aspects of human behaviour as love and relationships. Although have a read of [4] and decide for yourself. More usefully mathematics, particularly the mathematics of game theory, can be used to describe how economies behave, and how competition and (possibly) altruism can evolve.

The Rest

In my career I have been asked to model many things, including the reaction of homeless people to changes in government policy, microwave cooking, fish in an aquarium, organ playing, cocoa growing, and the transportation of biscuits in a factory. These I all take seriously and do my best. However, I draw the line at questions (usually posed by journalists when there is not much news around) such as what is the ideal shape for a Christmas tree, what is the perfect kiss, and what is the perfect joke. The key advice in all such cases is: avoid the question. The journalist is not looking for an answer, instead they are looking for a way to make fun of the person trying to give the answer. Avoid!

How Did the Mathematician Cross the Road?

We conclude this subsection by a short example of the use of a simple model. In particular we try to answer the question of how best to cross the road. Of course, we are all highly careful and law abiding citizens, and therefore only cross the road at a pedestrian crossing. Most of the crossings in Bath are those where you press a button and wait for the pedestrian light to go green, accompanied by a lot of beeping. I'm sure that you will all have noticed that when you press the button there is a delay D before the light goes green. The reason for this delay is to make sure that the traffic isn't blocked by a constant flow of pedestrians. But the question is: how long should the delay be? Too short and the traffic is blocked. Too long and the pedestrians give up waiting and chance it by running across the road. To answer this question, we can have a go at using a simple mathematical model for the crossing process.

We start by thinking about what the ingredients of the model. A picture always helps

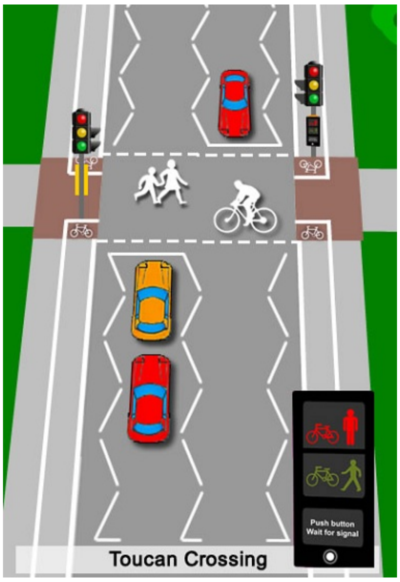

A simple set of initial assumptions are as follows:

- Pedestrians arrive at a steady rate at the crossing
- Traffic arrives at a steady rate
- It takes the pedestrians a time T to cross the road
- The ideal delay time is the one that maximises the average total flow F of both the pedestrians and the traffic.

By a steady rate we will assume that in any time interval of length t, then a*t pedestrians will arrive and b*t cars will arrive. We will also assume that all of the pedestrians waiting at the lights cross at once when the lights turn green.

Once the lights go red again we will wait a time 1/a before a pedestrian arrives and presses the button to cross. The lights then have a delay D during which cars can cross before they go again.

Now to do some maths. We have a complete cycle of time 1/a + D+T. In this time 1 + a*D pedestrians arrive and then cross. In the same time b*(1/a +D) cars cross over the crossing.
The average total flow of pedestrians and cars is then given by:

$$F = \frac{1 + a\,D + b\left(\frac{1}{a} + D\right)}{\frac{1}{a} + T + D}$$

If we take (for example) a = ½, b = 10, T = 1 then a plot of F as a function of D is given below:

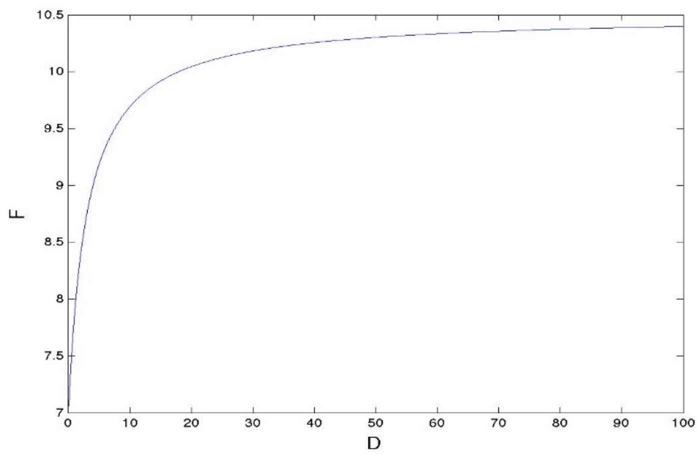

We can see from this model that the total flux *increases* as D *increases*. Indeed it rises from
F = (a+b)/(1 + aT) when D = 0, to F = (a+b) as D tends to infinity.

We conclude from this simple model that it is best to make the delay D as long as possible. However there will be an upper limit. This depends on the patience of the pedestrians. This is beyond the subject of mathematical modelling.

Readers of these notes might like to consider whether they agree with this model, the assumptions made to get it, or the conclusions that we have drawn from it. Indeed, how might the model be improved, and does it effect the conclusion? For example, is the assumption that the best delay time is on which maximises the flow, or is it better to minimise the average waiting time of the pedestrians/cars? I leave these as a matter of discussion.

2.4 Common mistakes in mathematical modelling

Don't Eat the Menu

All models, even complex simulations, are approximations of reality. They are not reality itself. This should always be born in mind when making predictions. Far better to say 'that my model for climate change predicts a rise of three degrees in temperature', than to say that there will be a three degree rise. When using a model always be aware of the limitations, uncertainties and approximations that go into them. The best models have a quantification of their uncertainty.

The Mathematical Drunkard

It is a very common mistake in mathematics (and indeed in most things) to change a problem to something that we can solve, rather than solving the original problem because it is too hard. In doing this we are behaving like the drunkard in the cartoon below.

A side effect of this is that we might end up using overly sophisticated maths (or indeed the wrong sort of maths altogether) to solve what is in effect the wrong problem.

The Ivory Tower

Another common mistake is to have such a high opinion of what we know

Nature will throw out mighty problems, but they will never reach the mathematician.
He may sit in his ivory tower waiting for the enemy with an arsenal of guns, but the enemy will never come to him. Nature does not offer her problems ready formulated. They must be dug up with pick and shovel, and he who will not soil his hand will never see them'

J. Synge. American Mathematical Monthly, 1944.

Although there are rare exceptions, mathematical modelling is a product of teamwork, with close collaboration between the mathematician, many other scientists, and the end user who needs the results.

The Curse of The Formula

The examples of models that I have shown above have led to simple formulae. But don't assume that this will always be the case, and certainly don't oversimplify the problem to get to an exact formula. Be happy and prepared to use approximation techniques if they yield useful answers. These include asymptotic and numerical methods.

2.5 Mechanisms for teaching and learning mathematical modelling

Because mathematical modelling is so important, it is essential that it is taught, and that it is taught well. Students respond very well to seeing how mathematics can have applications that they never thought were possible. It is good to see that some mathematical modelling is now taught in both A level and undergraduate mathematics courses. By far the best way to teach it is to immerse the students directly in the problem-solving environment, and to put them into teams to do this. Even better if the problems come directly from industry. One mechanism to do this, which has proved effective for over 50 years, is the *study group model* and its derivatives. In these week long workshops, industrial problems are presented on the first day, and the participants work on them in teams, to present a paper/talk on the final day. The work done during the week is then followed up with longer interactions with the company. This mechanism is very powerful, not only for solving the problems (using mathematics), but in training the participants on all sides, and introducing them to new ideas and methods. Originally developed in the UK the study group model is now used all over the world and has proved its value [5]. Below we see a typical study group team in action at a recent workshop in Ireland. Note the very full whiteboard in the background.

3 Maths saves the whales

For our first substantial example of mathematical modelling, we will return to the title of this lecture and see how maths can save the whales.

3.1 Whales are under threat

It is well known that overexploitation by the whaling industry led to serious declines in many of the world's populations of whales, although thankfully no species was brought to extinction and many are now in the process of recovering, although not all. However, many threats to whales remain. Many of these are due to changes in the whales' environment due to climate change and human activity such as fishing. Others are due to the increased pollution of the oceans. Sonar, and other noise pollution, from boats, which is used to detect fish, can disrupt the guidance system of whales causing them to get lost and beach (more of this later). A final threat to whales arises from them coming too close to shipping. This can lead to collisions between the ship and the whale, when the whale always comes off worse. Another threat arises if the ship is prospecting for oil using seismic signals. In this case the ship will detonate a small explosion and measure the echoes from this. This is a good way to find oil, but the explosions can injure any nearby whale. In this lecture we will look at how, by using some mathematical modelling, we can reduce the latter two risks.

3.2 How many whales are there?

One area in which maths can make a significant difference, is the seemingly simple one of telling us how many whales there are in the first place. Good conservation and management require an understanding of the status of the whale populations. A key component of this is, an estimate of both the present population, and also the changes in the population resulting from the impact of different types of threat. Simple, you say, let's go out and count them! However, estimating the abundance of animals that are not particularly common, and which also spend most of their time below the surface is difficult. To estimate the population numbers it is therefore necessary to use statistical methods and a good way to do this mathematically is to use the *mark and recapture technique*. In this method a (small) portion of the population is either captured, marked, and released or is observed for significant features (such as fin types for whales). Some time later, another portion is captured/observed and the number of originally marked/identified individuals within this group is counted. In a well-mixed population, the number of identified individuals in the second group is proportional to the number of marked individuals in the whole population. Hence an estimate of the total population size can be obtained by dividing the number of identified individuals by the proportion of marked individuals in the second sample. For example, if 100 originals are originally identified, and in the second sample there are 5% of these, then the total population is estimated to be 20*100=2000. The accuracy and limitations of this method, which is widely used in many applications for finding population sizes, are discussed in [6].

3.3 Ship strikes

The main threat to whales, that we will now look at, are strikes from ships. Evidence of ship strikes with whales comes from a range of sources including direct reports from the vessel involved, and examination of dead whales found floating at sea or washed up on the beach. For some populations, such as the North Atlantic right whale whose main habitat is the busy waters off the east coast of the USA and Canada, the mortality rate is particularly high compared to the overall population size. It is thought that mortality due to ship strikes may make the difference between extinction and survival for this species. The most effective way to reduce collision risk is to keep whales and ships apart. However, to do this we must know where the whales are, and in particular, how far the ship is from the whale. Mathematical modelling gives us a good way to do this.

3.4 Pythagoras' Theorem saves the whales

If anyone is asked to name a theorem then chances are that they will either mention Pythagoras' theorem or (the not unrelated) Fermat's Last Theorem. Of these Pythagoras' theorem is by far the oldest having been known for over 3000 years at least. This states that for any right-angled triangle, as illustrated below

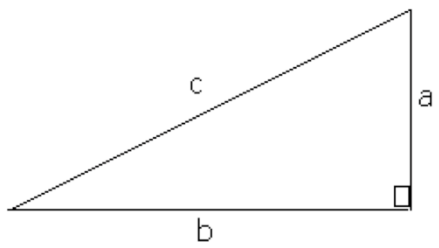

we have

$$a^2 + b^2 = c^2$$

Or in more poetic terms:

A right angled triangle opined,

My hypotenuse squared is refined,

For if anyone cares,

It's the sum of the squares,

Of my other two sides when combined.

The reason that Pythagoras' theorem is so useful to us in finding whales is that it can be used to estimate distances in general, and the distance of a whale from the ship, in particular.

To make sure that a ship does not hit a whale it has to be a certain distance from it called the 'mitigation zone'. Provided that the whale is outside the mitigation zone then it will be safe. The mitigation zone varies from ship and from whale to whale, but a distance of 400m is not unreasonable. Thus to save the whales from ship strikes, or from surveying explosions, we must make sure that we can estimate the distance from them, so that they are outside the mitigation zone.

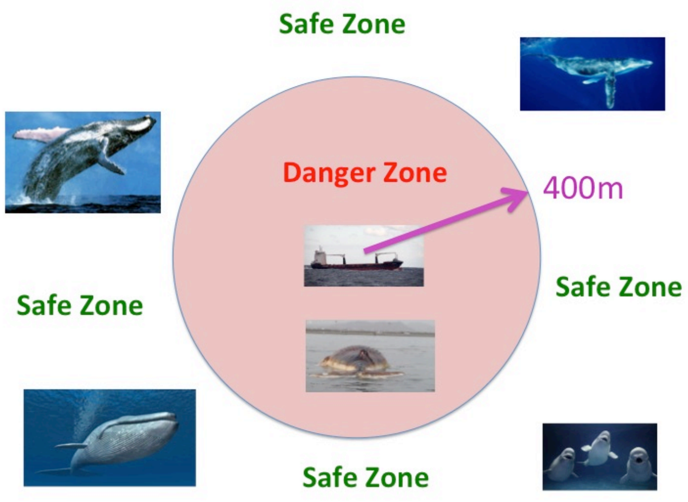

We must now work out a way of finding the distance from the whale to the ship and to do this we combine physics with mathematical modelling. There are many ways to find out how far a whale is from a ship. The simplest of these is to spot it visually, and seismology prospecting ships have to carry 'whale spotters' as a result. Other techniques such as radar are less effective due to the relatively small size of the whale and its tendency to submerge. Active sonar is often used to detect fish (and submarines) and involves sending pulses of high energy sound which reflect off nearby objects. Detecting these reflections allows the location of the fish (or submarine) to be determined. However, active sonar must be avoided when searching for whales as it damages their own very sensitive navigation system and can lead them to beach and then to die.

The best way by far to detect the whale is simply to listen to it. Whales 'sing' and we can detect the sounds that they make. A spectrograph of a typical sound is shown below

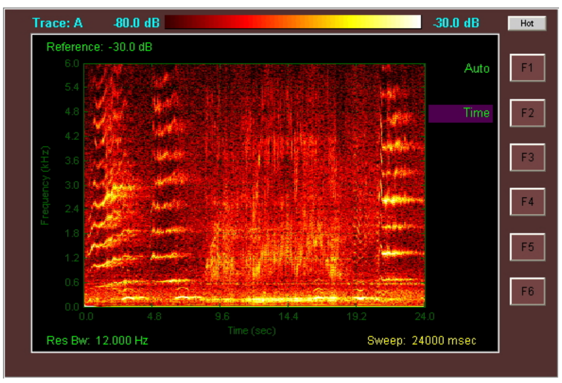

In principle by hearing when the sounds arrive, we can work out how far away the whale is.

But to do this we must use an approximate mathematical model. Our first step, as advertised above, is to look at the data. The pressure u of the detected sounds from the whale as they arrive look like this.

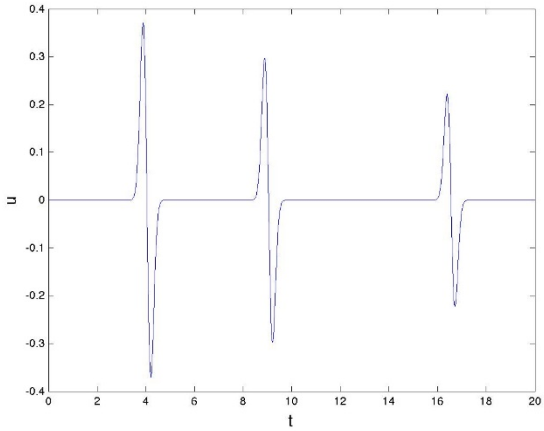

To interpret this picture, we must next consider the physical situation. This is illustrated below. The whale is a distance L from the ship. We also assume that it is just underneath the surface of the sea, and that the sea has a (constant) depth of H. Using this we can interpret the graph of the detected sounds.

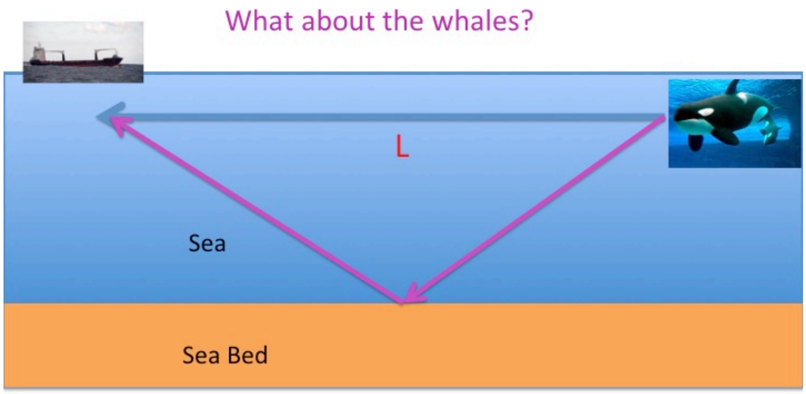

Hear **Two Sounds**. A direct one and an echo from the sea bed

We can measure the time difference between them.

What we are seeing here (at t=4) is the sound from the whale, followed shortly afterwards by the same sound reflected from the sea surface. Later (at t=8.5) we see the sound reflected from the seabed followed closely by one reflected from the sea surface. Later still (at t=16) we see the sound reflected twice from the seabed. The further they have to travel, the longer the sounds take to arrive, and the weaker they are when they do arrive.

To construct the mathematical model we start, as ever, with the laws of physics. In particular if u is the amplitude of the sound from the whale, then u satisfies the *wave equation*. This is given by the partial differential equation

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

Here c is the speed of sound in water which is approximately 1500 metres per second. It is possible to solve this equation directly, indeed this is exactly what oil prospectors do. However, the computer codes to do this are extremely complex and slow running. With a bit of mathematical approximation, we can make the process very much faster. The first thing to do is to assume that the sounds from the whale come have a constant frequency so that:

$$u(x,t) = e^{i\omega t} w(x)$$

where

$$\omega = 2\pi f$$

and f is the frequency of the sound. Substituting this into the wave equation, we obtain the Helmholtz equation

$$\nabla^2 w + \frac{\omega^2}{c^2}\, w = 0.$$

Finally, we rescale this equation by introducing a natural length scale X. This is a representative distance for any physical calculation and allows us to link the equation above to the real-life problem that we are trying to solve. For saving the whales it is reasonable to take X = 200 m which is roughly the depth of the sea and also of the mitigation zone. If we do this we get the rescaled equation

$$\nabla^2 w + k^2\, w = 0.$$

where

$$k = \frac{X\omega}{c}$$

The nature of the solutions of this equation completely depends on the size of the (wave number) k. If k is *small* the solutions look like waves on the sea (and are quite hard to compute). In contrast if k is *large* then the solutions look like rays and *travel on straight lines*. We are very used to this, because it is how we experience light travelling. (This is because light has a very high frequency).

The next step in our modelling process is to calculate k. Whales typically emit sounds at about 30Hz. This means that

$$k = \frac{200 * 2\pi * 30}{1500} \approx 25.$$

A value of k = 25 is big enough for the straight line approximation to work well for our calculations. This is a *major simplification* which makes the rest of the calculations *much easier*.

Using the straight line approximation from the model we can now work out how long it takes the different sounds from the whale to arrive.

If the whale is a distance L away from the boat then the time of arrival of a direct sound from the whale to the boat is given by

$$t_1 = \frac{L}{c}$$

If we knew when the whale emitted this sound then we could find L directly from this formula. However, we do not know this. To find the whale we must find listen for the echo of this first sound, which will arrive a bit later as it has to travel a greater distance. If we suppose that the depth of the ocean is H and the distance travelled by the sound (on its straight line path) is 2S then we can find S by drawing a couple of right angled triangles and applying Pythagoras' Theorem (as advertised earlier) as follows:

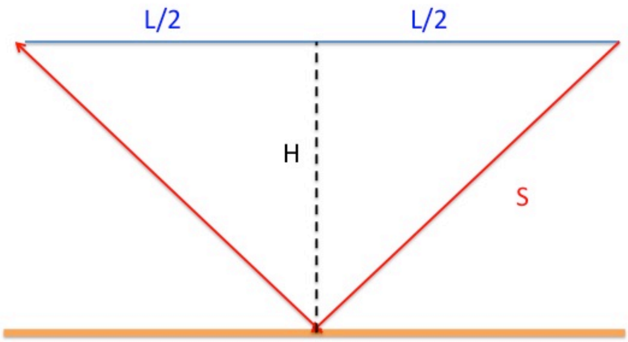

$$S^2 = H^2 + \frac{L^2}{4}$$

$$S = \sqrt{H^2 + \frac{L^2}{4}}$$

The total distance 2S is then given by:

$$2S = 2\sqrt{H^2 + \frac{L^2}{4}} = \sqrt{4H^2 + L^2}$$

The arrival time of the echo is then given by:

$$t_2 = \frac{\sqrt{4H^2 + L^2}}{C}$$

The resulting times for the pulses are then as shown in the figure below.

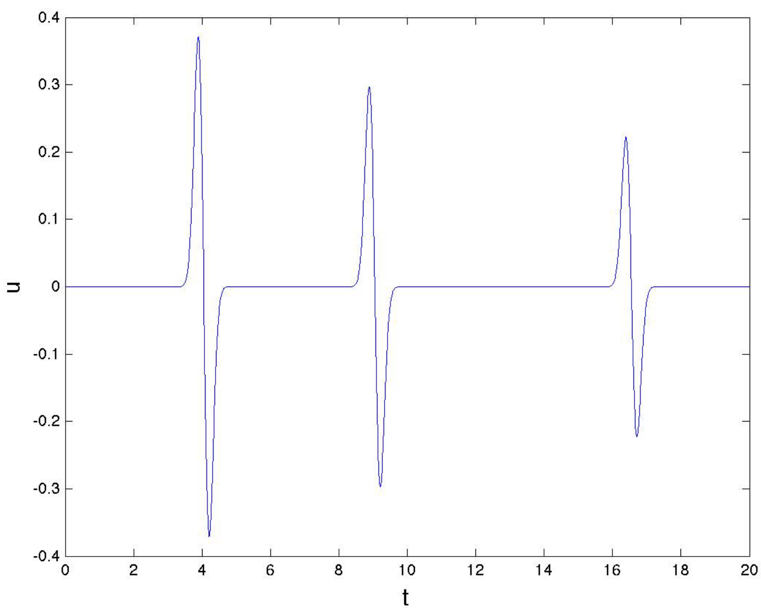

Now, whilst we can't tell when the sound of the whale was made, we can measure *the time difference* between the echo and the direct signal. This is given by:

$$\Delta = t_2 - t_1$$

So that

$$\Delta = \frac{\sqrt{4H^2 + L^2}}{C} - \frac{L}{C}$$

Now comes the clever bit. Knowing this time difference, we can work out the value of L and thus find out how far away the whale is. To do this we square the equation above to give

$$\Delta^2 + 2\Delta\frac{L}{C} + \frac{L^2}{C^2} = \frac{4H^2 + L^2}{C^2}$$

Cancelling the same terms from each side and rearranging we get

$$L = \frac{2H^2}{\Delta C} - \frac{\Delta C}{2}$$

Provided that we know H then we can find L, and hence find the whale. Here is and example of doing this when H = 100m is given below. In this we can see that if the time difference is less than 30ms then the whale is safe, but if it is longer than 30ms then the whale is in trouble. By means of this simple formula we can thus save the whales.

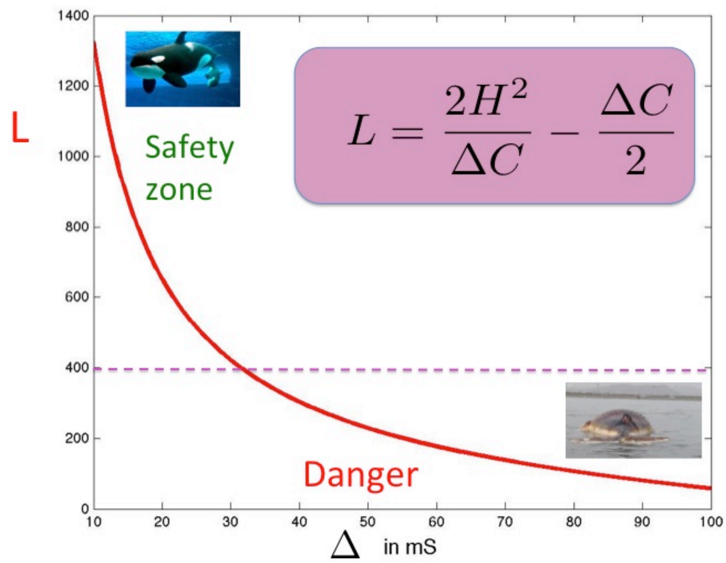

Of course as with all models we have made many simplifications and assumptions to get to this result. One of these is the assumption that we know the sea depth H. One way to find this, which is commonly used, is to employ active sonar. However, this as we have seen, can injure the whales. A much better way is to make use of the other echoes from the seabed. By listening to these it is possible not only to locate the whale, but to also find H and much more besides.

4 How maths helps to cure cancer

We will now look briefly at how a much more complex application of mathematical modelling can help in the fight against cancer, before returning in more detail, to show how the same maths that was used to save the whales can also directly cure bone cancer.

4.1 Modelling tumour growth

Cancer is the name given to diseases which are characterised by rapid, uncontrolled cell growth. There are over 200 different types of cancer, classified by the type of cell that is initially affected. Normally, our bodies form new cells only as we need them. However, when cells acquire mutations that disrupt the tightly controlled processes of cell division and death, a self-sustaining wave of cellular multiplication can occur and result in the formation of a tumour. To ensure its continued growth, a tumour must acquire a continuous supply of nutrients and the ability to export metabolic waste. It does this by recruiting new blood vessels from the nearby existing vasculature, a process known as tumor-induced angiogenesis. If the tumour grows too large then it will kill its host. Mathematical models allow us to quickly search and identify the most effective drug combinations for cancer patients. They are also deepening our understanding of how and why cancer cells often become resistant to chemotherapy drugs. Mathematical and computational modeling approaches, usually based on differential equation models, have been applied to every aspect of tumour growth from initiation, mutation acquisition and tumourigenesis to metastasis and also treatment response. An excellent summary of these is given in [7]. For example a differential equation for the rate of growth of a tumour of density c before angiogenesis is given by

$$c_t = \nabla.(D\ \nabla c) + \rho\ c\ (1 - c/K)$$

Here the first term on the right-hand side of this equation describes the dispersal of the cancer cells and is similar to the conduction of heat. The second term describes the net proliferation rate, with K being the maximum carrying capacity of the tissue for the cells. Solving this nonlinear equation allows us to estimate how rapidly the tumour will grow in its initial stages of development

Ultimately, the outcomes of all cancers in a patient are a product of tumour biology and the response to therapy. Since there is no practical means to define these variables in isolation for an individual patient, the creation of a *virtual tumour* through application of a careful mathematical model provides an immediate means to investigate the interplay of biology and response in outcomes. Such models should then be used to inform and guide a spectrum of clinical trials to improve outcomes for patients suffering from cancer. To quote [7] *'the field of mathematical oncology has the opportunity to drive forward a completely new trajectory of precision treatment strategies for cancer patients'.*

The account in [8] gives an excellent overview of the way that mathematical models for tumours and their treatments are constructed in the context of the ideas presented in section 2 above. They make the point that *'mathematicians have a unique set of skills that may be very useful, even essential, in the war on cancer. They understand complex systems both nonlinear and stochastic, they truly recognize what is required to make causal claims and understand that "reducing" a system may be the best way to handle it without any loss of essential information.*

4.2 Destroying tumours by sound

In the last section we looked at how it is possible to use maths locate a whale by listening to both the sound it makes and also the echoes of that sound. In this section I will take the maths in the opposite direction, and will describe a procedure I have been involved with that uses the same ideas to help cure cancer. When looking for the whales we were keen not to use sonar in case it damaged the whale. However, in treating cancer we want to be able to damage, indeed to destroy, a tumour, without damaging the surrounding tissues. On way to do this is to generate high intensity ultrasound (very high frequency sound of frequency around 1 MHz) in a transducer. This can then be focused on a tumour in much the same way that a magnifying glass can focus light onto a spot. As we all know from having done this (admit it) the point at the focus then starts to burn. Exactly the same happens to the tumour, which is then destroyed. It the point of focus is not on the tumour itself, then this process could cause damage to a vital organ. To avoid this, the patient is monitored throughout using magnetic resonance imaging (MRI), which can detect the changes in temperature inside the body. The whole treatment process is called MR-HIFU (Magnetic resonance High Intensity Focused Ultra Sound) , and is shown below. The patient lies on top of a table in which there is a transducer producing the ultrasound. They are then surrounded by a large ring, which contains the magnets for the MRI scan.

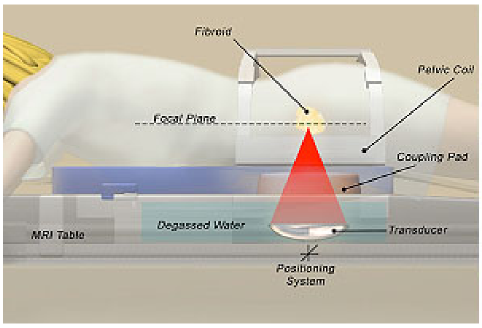

Mathematics is needed to make all stages of this treatment work. It is used to design the shape of the transducer, to determine the strength of the sound waves and the time needed for them to radiate the tumour, and to create the images in the MRI scan. This is all done by using mathematical models, and I am pleased to say my introduction to this problem came from a mathematical study group held in Toronto, Canada.

As usual in the mathematical modelling process we start with a picture. Sounds starts from a transducer, travels through layers of the body, and eventually reaches the tumour. The sounds waves are in general of low intensity (they need to be to avoid damaging the body)

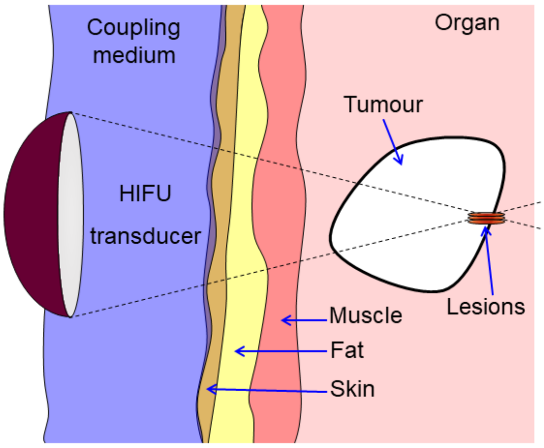

Apart from the point of maximum focus, the sound waves satisfy the same (wave) equation as the sounds from the whale. Therefore we can try to use the same analysis as worked so well in that case. (In maths we are always stealing ideas from one problem so that we can use them on another). We can calculate the wave number k in the same way as before. The difference in this case is that the frequency of the ultra sound $f = 1$ MHz, and the length scale $L = 5$ mm is much smaller (although the sound speed $c = 1500$ metres per second) is about the same. Substituting these in we get

$$k = \frac{2\pi * 10^6 * 5 * 10^{-3}}{1500} = 20.9$$

This is remarkably similar to the value of k that we used to listen to the whales. This means that we can use the same mathematics as before to calculate where the sound waves go to inside the body. In particular they will *travel in straight lines*. This simple modelling approximation can then be used both to design the shape of the transducer and to work out approximately where the point of focus will be. (It is never quite exact due to variations in the speed of sound as it travels through different parts of the body. Most importantly, by counting the number of rays that pass through any particular part of the body we can (to a good approximation) calculate the intensity $Q(x,t)$ of the sound waves at any point x and at any time t. From this we can then determine the temperature $T(x,t)$ of that part of the body by solving Pennes (partial differential) equation which is given by:

$$\rho c\, T_t = \nabla.\,(\,\kappa(x)\ \nabla T) - \omega c\,(T - T_b) + Q(x,t)$$

This equation models the heat loss both from conduction and by transport from the blood, in response to the heating Q from the sound waves. It can be solved relatively easily to find the temperature either analytically or using a numerical method. Once T has been determined, the damage to the tissue can be estimated using known physiological models.

The result is a method which can be used to destroy tumours safely and reliably. It is still under testing, but has already shown success in destroying tumours in bone cancers at the SickKids hospital in Toronto [9].

5 What next?

Given that so much of the real world can be modelled using mathematics, the potential applications of mathematical modelling are almost limitless. Watch this space. But always be aware of the limitations of any model.

6. References

[1] G. Box, 'All models are wrong but some are useful', (1976)

A discussion of this comment is given in

https://en.wikipedia.org/wiki/All_models_are_wrong

[2] C. Budd, *The mathematics of climate*, Gresham lecture, February, 2018

[3] J. Murray, Mathematical Biology, (1989), Springer.

[4] H. Fry, *The Mathematics of Love: Patterns, Proofs, and the Search for the Ultimate Equation,* (2015), TED Books

[5] M. Barons, C. Budd, M. Butchers and J. Jordan, *Study groups, what are their value?,* Mathematics Today, December, (2020).

[6] The mark and recapture method is described in

https://en.wikipedia.org/wiki/Mark_and_recapture

[7] T. Jackson, N. Komarova and K. Swanson, *Mathematical Oncology: Using Mathematics to Enable Cancer Discoveries*, (2014), The American Mathematical Monthly, 121:9, pp. 840-856

[8] V. M. Pérez-García, S. Fitzpatrick, L. A. Pérez-Romasanta, M. Pesic, P. Schucht, E. Arana, P. Sánchez-Gómez, *Applied mathematics and nonlinear sciences in the war on cancer*, Applied Mathematics and Nonlinear Sciences 1(2) (2016) 423–436

[9] An account of the uses of the MR-HIFU procedure at the SickKids hospital in Toronto is given in :

http://www.sickkids.ca/AboutSickKids/Newsroom/Past-News/2014/SickKids-doctors-destroy-bone-tumour-using-incisionless-surgery.html