ISSN 1754-3657 (Print)
ISSN 1754-3665 (Online)

CS4FN

Computer Science for Fun

Issue 27

SmartHealth

Decisions, Decisions, Decisions...

Smart pills

Computing clever care

Racist healthcare algorithms

www.cs4fn.org/smarthealth

Queen Mary
University of London

Smart health

The trouble with healthcare is that it's becoming ever more expensive: new drugs, new treatments, more patients, the ever-increasing time needed with experts.

We want everyone to get the care they need, but the costs are growing. Perhaps computer scientists can help? Research groups worldwide are exploring ways to create intelligent programs that can support patients at home, helping monitor them and make decisions about what to do.

For example, say you are on powerful drugs to manage a long term illness: should you have the vaccine? Can you have a baby? Is a flare up of your disease about to hit you and how can you avoid it? Is the new ache a side effect of the drugs? Do you need to change medicines? Do you need to see a specialist? If artificial intelligences can help support patients then the doctors and nurses can spend more time with those who need it, hospitals can save on expensive drugs that aren't working, and patients can have better lives. But what kind of technology can deliver this sort of service? In this issue, we explore one particular way being developed on the EPSRC funded PAMBAYESIAN project at Queen Mary University of London, based on an area of computing called Bayesian networks, that might just be the answer.

Image by Gerd Altmann from Pixabay

So, so tired...

Fatigue is a problem that people with a variety of long-term diseases can also suffer from.

This isn't just normal tiredness, but something much, much worse: so bad that it is a struggle to do anything at all, destroying any chance of a normal life. Doctors can often do little to help beyond managing the underlying disease, then hope the fatigue sorts itself out. Sometimes fatigue can stay with the person long, long after. Maha Albarrak, for her PhD, is exploring how computer technology might help people cope. Her first step is to interview those suffering to find out what kind of help they really need. Then she will work closely with volunteers to come up with solutions that solve the problems that matter.

Image by Małgorzata Tomczak from Pixabay

Here...

Amy Dowse wondered if an app might help people suffering with anxiety.

One way to overcome panic attacks is a mindfulness technique where you focus on the here and now - your surroundings rather than your internal feelings. For her university MSc project, she created an app to help people do this, called Here. It prompts you to look for coloured objects in the real world then use them to build a picture in the app. For example, you look at the colour of the clothes that people around you are wearing and try to fully dress a figure on the app using what you see.

Image by Zaccaria Boschetti from Pixabay

Smart tablets to swallow

by Paul Curzon, Queen Mary University of London

The first ever smart pill has been approved for use. It's like any other pill except that this one has a sensor inside it and it comes with a tracking device patch you wear to make sure you take it.

A big problem with medicine is remembering to take it. It's common for people to be unsure whether they did take today's tablet or not. Getting it wrong regularly can make a difference to how quickly you recover from illness. Many medicines are also very, very expensive. Mass-produced electronics, on the other hand, are cheap. So could the smart pill be a new, potentially useful, solution? The pill contains a sensor that is triggered when the pill dissolves and the sensor meets your stomach acids. When it does, the patch you wear detects its signal and sends a message to your phone to record the fact. The specially made sensor itself is harmless and safe to swallow. Your phone's app can then, if you allow it, tell your doctor so that they know whether you are taking the pills correctly or not.

The patch you wear detects the pill's signal and tells your phone you took it.

Smart pills could also be invaluable for medical researchers. In medical trials of new drugs, knowing whether patients took the pills correctly is important but difficult to know. If a large number of patients don't, that could be a reason why the drugs appeared less effective than expected. Smart pills could allow researchers to better work out how regularly a drug needs to be taken to still work.

More futuristically still, such pills may form part of a future health artificial intelligence system that is personalised to you. It would collect data about you and your condition from a wide range of sensors recording anything relevant: from whether you've taken pills to how active you've been, your heart rate, blood pressure and so on: in fact anything useful that can be sensed. Then, using big data techniques to crunch all that data about you, it will tailor your treatment. For example, such a system may be better able to work out how a drug affects you personally, and so be better able to match doses to your body. It may be able to give you personalised advice about what to eat and drink, even predicting when your condition could be about to get better or worse. This could make a massive difference to life for those with long term illnesses like rheumatoid arthritis or multiple sclerosis, where symptoms flare up and die away unpredictably. It could also help the doctors who currently must find the right drug and dose for each person by trial and error.

Computing in future could be looking after your health personally, as long as you are willing to wear it both inside and out.

Image by Myriams-Fotos from Pixabay

What are the chances of that?

The church minister's hobby and clever machines

by Norman Fenton, Queen Mary University of London

Image by Free-Photos from Pixabay

The hobby of a church minister over 250 years ago is helping computers make clever decisions.

Thomas Bayes was an English church minister who died in 1761. His hobby was a type of maths that today we call probability and statistics, though his writings were never really recognised during his own lifetime. So, how is the hobby of this 18th century church minister driving computers to become smarter than ever? His work is now being used in applications as varied as: helping to diagnose and treat various diseases; deciding whether a suspect's DNA was at a crime scene; accurately recommending which books and films we will like; setting insurance premiums for rare events; filtering out spam emails; and more.

How likely is that?

Bayes was interested in calculating how likely things were to happen (their probability) and particularly things that cannot be observed directly. Suppose, for example, you want to know the probability that you have an infectious virus, something you can't just tell by looking. Perhaps you're going to a concert of your favourite band - one for which you've already paid a lot of money. So you need to know you are not infected. If recent data shows that the virus currently affects one in 200 of the population, then it is reasonable to start with the assumption that the probability YOU have the virus is one in 200 (we call this the 'prior probability'). Another way of saying that is that the prior probability is 0.5 per cent.

A better estimate

However, you can get a much better estimate of how likely it is that you have the virus if you can gather more evidence of your personal situation. With a virus you can get tested. If the test was always correct, then you would know for certain. Tests are never perfect though. Let's suppose that for every 100 people taking the test, two will test positive when they actually do NOT have the virus. Scientists call this the false positive rate: here two per cent. You take the test and it is positive. You can use this information to get a better idea of the likelihood you have the virus.

How? Bayes worked out a general equation for calculating this new, more accurate probability, called the 'posterior' probability (see page 8). It is based, here, on the probability of having the virus before testing (the original, prior probability) and any new evidence, which here is the test result.

A surprising result

If we assume in our example that every person who does have the virus is certain to test positive then, plugging the numbers into Bayes' theorem, tells us there is actually a surprisingly low, one in five (i.e., 20 per cent) chance you have the virus after testing positive. "A Graphical Explanation of Bayes' theorem" (right) shows why the answer is correct. Although this is much higher than the probability of having the virus without testing (two per cent), it still means you are unlikely to have the virus despite the positive test result!

Many people find the result very surprising!

If you understand Bayes theorem, you might feel it unfair if your doctor still insists that you have the virus and must miss the trip. In fact, many people find the result very surprising; generally, doctors who do not know Bayes' theorem massively overestimate the likelihood that patients have a disease after a positive test result. But that is why Bayes' theorem is so important.

To go or not to go

Of course, no one knows which of the five concert goers are the ones infected. If all 25 ignore their doctor that means there are five people mingling in the crowd, passing on the virus, which would mean lots more people catch the virus who pass it on to lots more, who ... (see Ping pong vaccination, page 14).

Continued on page 6.

A graphical explanation of Bayes theorem

by Norman Fenton, Queen Mary University of London

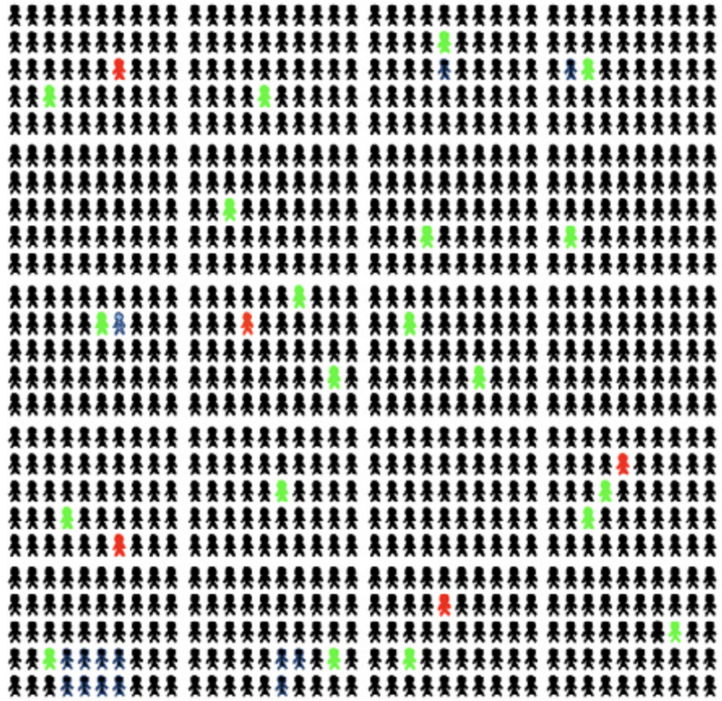

Imagine 1,000 people.

About 5 will have the virus.

The remaining 995 do not have the virus.

With a 2 per cent false positive rate about 20 (2 per cent of 995) who do NOT have the virus will test positive.

So 25 in total test positive, of whom 5 actually have the virus (assuming no false negatives).

So only one in five (20 per cent) of those who test positive actually have the virus.

What are the chances of that?

(continued from page 5)

We have seen that, with a little extra information (such as a test result), we can work out a more accurate probability and so have better information upon which to make decisions.

In practice, there are many different kinds of information that we can use to improve our estimate of the real probability. There are symptoms such as lack of taste/smell which are quite specific to the virus. Others, like a cough, are common in people with the virus but also in people with flu. There are also factors that can cause a person to have the virus in the first place such as close contact with an infected relative. So, instead of just inferring the probability of having the virus from one piece of information, like the test result, we can consider lots of interconnected data, each with its own prior probability. This is where computers come in: to do all the calculations for us.

We first need to tell the computer about what causes what. A convenient way to do this is to draw a diagram of the connections and probabilities called a 'Bayesian network' (see page 7, "A Simple Bayesian Network"). Once a computer has been given the Bayesian network, it can not only work out more accurate probabilities, but it can also use them to start making decisions for us. This is where all those applications come in. Deciding whether a suspect's DNA was at a crime scene, for example, needs the same kind of reasoning as deciding whether you have the virus.

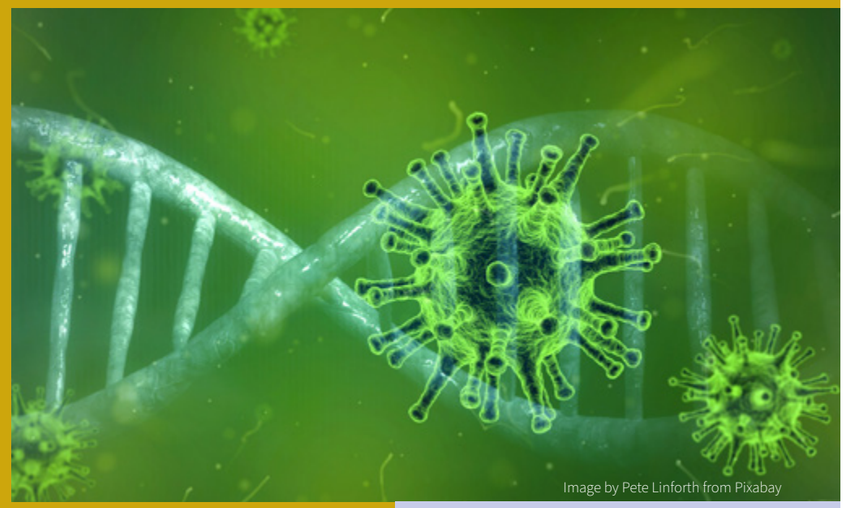

Image by Pete Linforth from Pixabay

Obviously, it is more complex to apply Bayes' theorem in realistic situations and, until quite recently, even the fastest computers weren't able to do the calculations. However, breakthroughs by computer scientists developing new algorithms mean that very complex Bayesian networks, with lots of inter-connected causes, can now be computed efficiently. Because of this, Bayesian networks can now be applied to a multitude of important problems that were previously impossible to solve. And that is why, perhaps surprisingly, the ideas of Thomas Bayes, from over 250 years ago, are showing us how to build machines that make smarter decisions when things are uncertain.

Bayesian baffler

Here is another counter-intuitive result that most people get wrong. What do you estimate the probabilities are… (answers on the back page with a detailed explanation on the cs4fn website **cs4fn.org/smarthealth**).

A criminal has left their DNA at a crime scene. Only one in every 10 million people have the DNA profile found. A suspect, Fred Smith, whose DNA matches, is put on trial. There is no other evidence. The population is 10 million people. The prosecutor claims that the probability that an innocent person has the matching DNA is one in 10 million: so Fred is guilty. The defence say the correct probability is closer to one in two. Who is correct? Is Fred's guilt "beyond reasonable doubt"?

A simple Bayesian network for having a virus

by Norman Fenton, Queen Mary University of London

We can take the extra evidence for and against our having a virus and draw a Bayesian network.

For each bubble the percentages show the chance that for a random person in the population this thing is currently true. Arrows show which things can cause others. So, in the diagram, this means that 0.5 per cent of the population currently have the virus (as the probability was 1 in 200, or 0.005, and to turn a probability into a percentage you just multiply by 100); 0.4 per cent of the population have been in recent contact with an infected person; 10 per cent have a cough; 2 per cent have flu, and so on. This is all general evidence we can collect about the country as a whole.

The model also includes probabilities not shown, like the chance of a person getting the virus if they have been in recent contact with an infected person and (as we already saw) the probability of a positive test depending on whether they do, or do not, have the virus.

We then want to know about you. Do you have a cough, have you lost your sense of taste or smell, what was the result of your test, and have you been in contact with an infected relative? From this information, we can update the probabilities in the Bayesian network using Bayes' theorem to give a new probability for how likely it is that you have the virus. Computer software can do this for us, though the more complicated the Bayesian network, the longer it takes to do all the calculations.

Solving real problems

by Norman Fenton, Queen Mary University of London

The first algorithms that enabled Bayesian network models to be calculated on a computer were discovered separately by two different research groups in the late 1980s. Since then, a series of easy-to-use software packages have been developed that implement these algorithms, so that people without any knowledge of computing or statistics can easily build and run their own models.

These algorithms do 'exact' computations and can handle Bayesian networks for many different types of problems, but they can run into a barrier: when run on Bayesian networks beyond a certain size or complexity, they take too long to compute even on the world's fastest computers. However, newer algorithms – which provide good approximate calculations rather than exact ones – have made it possible to deal with much larger problems, and this is a really exciting ongoing research area.

The result, though, is that the computer can give you a personalised risk assessment of how likely it is that you have the virus based on the specific evidence about you.

You can find such a comprehensive personal COVID risk calculator, based on a Bayesian network with much more data, at **covid19.apps.agenarisk.com**

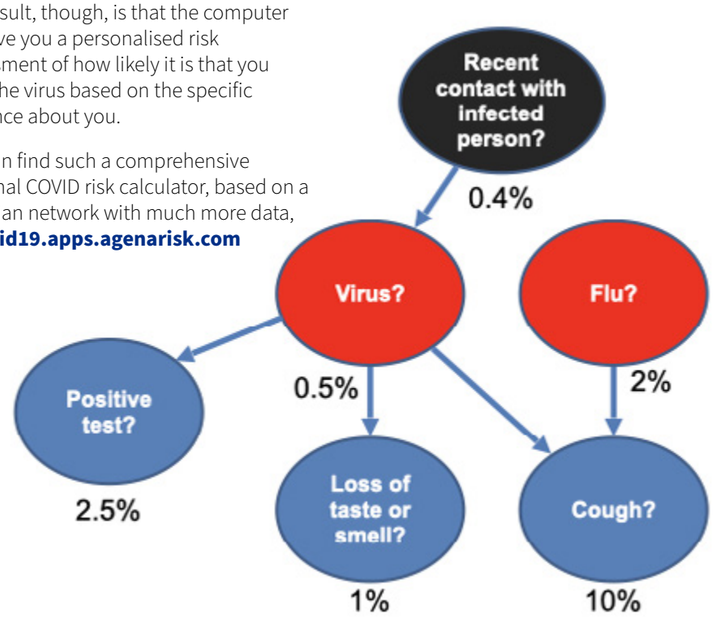

Bayes' theorem as an algorithm

by Paul Curzon, Queen Mary University of London

Thomas Bayes is famous for the theorem named after him: Bayes' theorem. It can be used in any situation where we want to calculate a more accurate probability of something given extra evidence. We will look at a version for our virus problem from page 4.

We want to know the probability that you have the virus, given that have you just tested positive. This calculated probability is called the 'posterior' probability. In this case Bayes' theorem becomes:

$$\text{Probability you have the virus after testing positive} = \frac{\text{Probability of you testing positive if you have the virus} \times \text{Probability you had the virus before knowing the test result}}{\text{Probability of a person testing positive (with or without the virus)}}$$

Image by Gerd Altmann from Pixabay

The theorem tells us that the chance that a person who tests positive actually has the virus is just the number of people **with** the virus who test positive divided by the total number of people (with or without the virus) who test positive.

The theorem can be used as the basis of an algorithm to compute the new, more accurate probability that we are after. We will assume, to make things easier to follow, that we are considering a population of a thousand people. We get the algorithm:

To calculate accurate probability that you have the virus after testing positive:

STEP 1: Calculate how many people BOTH have the virus AND test positive.

STEP 2: Calculate the number of people who will test positive (whether they have the virus or not).

STEP 3: Divide 1) by 2) to give the final answer of the probability you have the virus after testing positive.

Let's work through it with the numbers from our example. Stay calm! This is going to get hairy if you are not a computer!

What do we know? Well, actually we need another little algorithm to do Step 1:

To calculate how many people BOTH have the virus AND test positive

STEP 1a: Calculate the probability that you will test positive if you do have the virus.

STEP 1b: Calculate the probability you have the virus BEFORE knowing the test result.

STEP 1c: Multiply Answer 1a by Answer 1b by 1000 (our population). (Answer to 1c) =
(Answer to 1a) x (Answer to 1b) x 1000

This calculates the answer to Step 1 for us. We have said we have a test that is always positive if you do have the virus (in reality tests do get it wrong this way too but, to keep things simple, we will ignore that here). That means the answer needed for Step 1a is a probability of 1 (meaning it is 100 per cent certain that it gets the answer right if you have the virus).

What about Step 1b? That is the country-wide probability of having the virus we are starting with. Knowing nothing else about an individual we have said 1 in 200 people have the virus. That makes the answer needed for this step: 1 ÷ 200, so probability, 0.005

We can now calculate Step 1c: We just multiply those two numbers 1 x 0.005 and multiply that by the total number of people: 1000. This gives the answer that five people out of the 1000 have the virus and test positive.

Step 2 is a bit more tricky: it is the number of people out of our 1000 who test positive. That includes all those with the virus but ALSO those that the test wrongly says have the virus when they don't. We need to add the numbers for these two groups: those with the virus and those without.

To calculate the number of people who test positive:

STEP 2a: Calculate the number of people who have the virus AND who test positive. This is just the answer from Step 1.

STEP 2b: Calculate the number of people who do NOT have the virus AND who test positive.

STEP 2c: Add 2a and 2b together.

We have already worked out the first part (Step 2a). It is just the answer from Step 1, so we already know it is five people. Step 2b is calculated in a similar way to Step 1 as follows:

To calculate the number of people who do not have the virus AND who test positive:

STEP 2bi: Calculate the probability that you will test positive if you do NOT have the virus.

STEP 2bii: Calculate the probability you do **not** have the virus.

STEP 2biii: Multiple 2bi by 2bii and then by 1000 to give the number of people who do not have the virus but test positive.

We know the answer to Step 2bi, as we said there was a two per cent chance of the test telling you that you had the virus when you didn't. That means the answer to this step is 2 ÷ 100 = 0.02.

For Step 2bii, the probability a person does NOT have the virus, we just need to calculate the rest of the population excluding those with the virus. We said one in every 200 people have the virus. That means 199 in 200 do not have it. The answer to this step is therefore 199 ÷ 200 = 0.995.

So, to work out Step 2biii to find out the number of people who do not have the virus but test positive: we multiply our two above answers 0.02 x 0.995, then multiply this by 1000. This gives answer 19.9: so about 20 out of the 1000 people are incorrectly told they have the virus.

We can now go back to Step 2c and add the answer from Step 2a (of those correctly told they have the virus) to that from Step 2b (those told they have the virus when they do not). This is 5 + 20, so 25 people in total are given a positive result. This is the answer to Step 2.

Finally, we can work out the overall, more accurate probability (Step 3). Divide the answer from Step 1, (five people), by the answer to Step 2 (25 people), to give the final probability as 5 ÷ 25 = 0.2 or a 20 per cent chance you actually have the virus after testing positive.

Don't forget we have just made up the numbers here to show the maths. They vary from test to test, place to place and over time. Tests can also give the all-clear to people with the virus ("false negative" results).

Although no test is 100 per cent accurate, the current Covid tests can be "confirmed" with an additional test to give further evidence.

Is your healthcare algorithm racist?

by Paul Curzon, Queen Mary University of London

Image by David Mark from Pixabay

Algorithms are taking over decision making, and this is especially so in healthcare. But could the algorithms be making biased decisions? Could their decisions be racist? Yes, and such algorithms are already being used.

There is now big money to be made from healthcare software. One of the biggest areas is in intelligent algorithms that help healthcare workers make decisions. Some even completely take over the decision making. In the US, software is used widely, for example, to predict who will most benefit from interventions. The more you help a patient the more it costs. Some people may just get better without extra help, but for others it means the difference between a disability that might have been avoided or not, or even life and death. How do you tell? It matters as money is limited, so someone has to choose. You need to be able to predict outcomes with or without potential treatments. That is the kind of thing that machine learning technology is generally good at. By looking at the history of lots and lots of past patients, their treatments and what happened, these artificial intelligence programs can spot the patterns in the data and then make predictions about new patients.

This is what current commercial software does. Ziad Obermeyer, from UC Berkeley, decided to investigate how well the systems made those decisions. Working with a team combining academics and clinicians, they looked specifically at the differences between black and white patients in one widely used system. It made decisions about whether to put patients on more expensive treatment programmes. What they found was that the system had a big racial bias in the decisions it made. For patients that were equally ill, it was much more likely to recommend white patients for treatment programmes.

The program was much more likely to recommend white patients for treatment programmes.

One of the problems with machine learning approaches is it is hard to see why they make the decisions they do. They just look for patterns in data, and who knows what patterns they find to base their decisions on? The team had access to the data of a vast number of patients the algorithms had made recommendations about, the decisions made about them and the outcomes. This meant they could evaluate whether patients were treated fairly.

The data given to the algorithm specifically excluded race, supposedly to stop it making decisions on colour of skin. However, despite not having that information, that was ultimately what it was doing. How?

The team found that its decision-making was based on predicting healthcare costs rather than how ill people actually were. The greater the cost saving of putting a person on a treatment programme, the more likely it was to recommend them. At first sight, this seems reasonable, given the aim is to make best use of a limited budget. The system was totally fair in allocating treatment based on cost. However, when the team looked at how ill people were, black people had to be much sicker before they would be recommended for help. There are lots of reasons more money might be spent on white people, so skewing the system. For example, they may be more likely to seek treatment earlier or more often. Being poor means it can be harder to seek healthcare due to difficulties getting to hospital, difficulties taking time off work, etc. If more black people in the data used to train the system are poor then this will lead to them seeking help less, so less is being spent on them. The system had spotted patterns like this and that was how it was making decisions. Even though it wasn't told who was black and white, it had learnt to be biased.

There is an easy way to fix the system. Instead of including data about costs and having it use that as the basis of decision making, you can use direct measures of how ill a person is: for example, using the number of different conditions the patient is suffering from and the rule of thumb that the more complications you have, the more you will benefit from treatment. The researchers showed that if the system was trained this way instead, the racial bias disappeared. Access to healthcare became much fairer.

If we are going to allow machines to take healthcare decisions for us based on their predictions, we have to make sure we know how they make those predictions, and make sure they are fair. You should not lose the chance of the help you need just because of your ethnicity, or because you are poor. We must take care not to build racist algorithms. Just because computers aren't human doesn't mean they can't be humane.

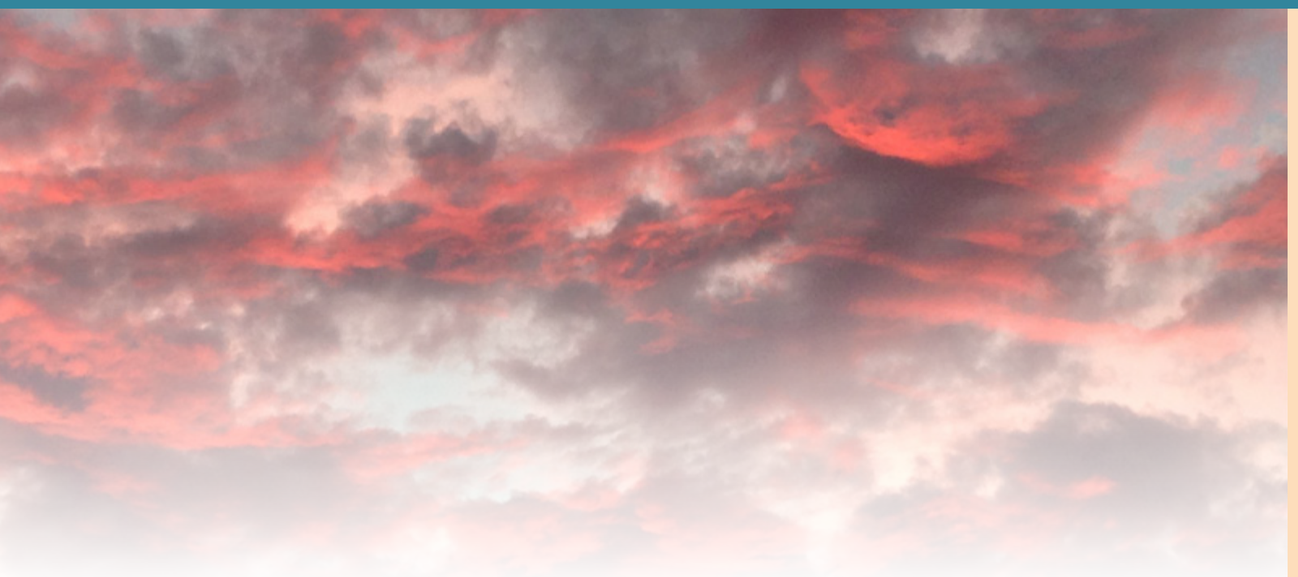

Cloudy with a chance of pain

by Paul Curzon, Queen Mary University of London

One day you may have personalised pain forecasts...

Very many people suffering from diseases like arthritis think that the weather affects the pain they feel. Many dread the coming of Autumn, for example, as they know their lives will get worse in the cold and wet weather. Others have found that trips to warmer countries have helped reduce their suffering from pain. Doctors have long been sceptical of these claims as there has been little evidence to support it, but then there have only been a few small-scale studies investigating it. It also isn't helped by the fact that different people believe different weather affects them and in different ways: some like it hot, some like it cold, some feel they suffer most when the rain comes...

A team from the University of Manchester realised that they could use crowd-sourced science, where members of the public collect data on their phones, to do a massive experiment to find out the truth. 13,000 people suffering from long-term pain took part, recording how much pain they were in every day for over a year. Their phones recorded their location and linked it to the local weather at the time. This gave the researchers millions of reports of pain to analyse against the specific weather that person was actually experiencing at the time.

So who was right: the doctors or the patients? Well, actually many of the patients were right as the weather did affect the amount of pain they personally felt. Especially problematic were days when the humidity was high, the air pressure was low, or the wind was very strong (in that order).

These results mean clinicians can now start to take the weather seriously. It may also be possible to create pain forecast programs for people affected, based on their local weather reports. It also opens up new areas to study to understand the causes of pain and so find new ways to alleviate it.

So, if you are unlucky to suffer from chronic pain, it could be your phone that, one day, might be warning you that tomorrow will be cloudy with a chance of pain.

Find out more about the "Cloudy with a Chance of Pain" project at **cloudywithachanceofpain.com**

Image by Cindy Eyler from Pixabay

Are you there yet?

by Jo Brodie, Queen Mary University of London

Plenty of people love the Weasley family's clock from the Harry Potter books and films. It shows where members of the family are at any given time. Instead of numbers giving the time, the clock face has locations where someone might be (home, school, shopping) and the many hands on the clock show the family members. The wizarding world uses magic to make their whereabouts clock work, but muggles (and squibs) can use mobile network data to build a simple version, and use Bayesian networks to improve it.

Your mobile phone is in contact with several cell towers in the mobile provider's network. When you want to send a message, it goes first to the nearest cell tower before passing through the network, finally reaching your friend's phone. As you move around, from home to school, for example, you will pass several towers. The closer you are to a tower the stronger the signal there, and the phone network uses this to estimate where you are, based on signal strength from several towers. This means that, as long as your phone is with you, it can act as a sensor for your location and track you, just like the Weasley's whereabouts clock.

You could also have a system at home that monitors your location, so that it switches on the lights and heating as you get closer to home to welcome you back. On a typical day you might head home somewhere between 3 and 6pm (depending on after-school events) and as you leave school the connection to your phone from the tower nearest the school will weaken, but connections will strengthen with the other cell towers on your route home. But what if you appear to be heading home at 11 in the morning? Perhaps you are, or maybe actually the signal has just dropped from the tower nearest to the school so a tower nearer your home is now getting the strongest signal!

A system using Bayesian logic to determine 'near home' or 'not near home' can be trained to put things into context. Unless you are ill, it's unlikely that you'd be heading home before the afternoon so you can use these predicted timings to give a likelihood score of an event (such as you heading home). A Bayesian network takes a piece of information ('person might be nearby') and considers this in the context of previous knowledge ('and that's expected at this time of day so probably true' or 'but is unlikely to be nearby now so more information is needed').

You could also set up a similar system in a home using wifi points to predict where you are and so what you are doing. Information like that could then feed data into a personalised artificial intelligence looking after you (see page 16).

Not all magic has to be run by magic!

Image by Alistair McIntyre from Pixabay

The ping pong vaccination programming challenge

by Paul Curzon, Queen Mary University of London

Image by Sergio Pavlishko from Pixabay

Vaccination programmes work best when the majority of the population are vaccinated. One way scientists simulate the effects of disease and vaccination programmes is by using computer simulations. But what is a computer simulation?

You can visualise what a simulation is with ping pong balls bouncing around a crowd. Imagine having a large room full of people. A virus is represented by a ping pong ball, bouncing from person to person, infecting each person it touches. Each person who is hit by a ping pong ball and not already infected becomes infected. That means they toss that ping pong ball back into the crowd to infect more people, but they also toss an extra one too (and then they sit down: dead). Start with a few ping pong balls. Quickly the virus spreads everywhere and lots of people sit down (die). You have run a physical simulation of how a virus spreads!

Now start again but 'vaccinate' 80 per cent of the people first: give them a baseball cap to wear to show who is who. If those people get a ping pong ball, they just destroy it: they infect noone else. Start with the same number of ping pong balls. This time, the virus quickly dies out and only a few people sit down (die). Not only are the vaccinated people protected but they protect many of the un-protected people too who might have died.

Now (if you can program) you can write a program to do the same thing, and so simulate and explore the spread of infection, which is easier perhaps than getting a thousand people to chuck ping pong balls about. Create a grid (an array) of 1000 cells. Each represents a person. They can be infected or not. They can also be vaccinated or not. Start with five random cells (so people marked as infected). Run a series of rounds. After each round, a newly infected cell randomly chooses two others to infect. If not infected already and not vaccinated, then they become newly infected. If already infected or vaccinated, they do not pass the infection on.

You can run lots of different experiments with different conditions. For example, experiment with different proportions of people infected at the start or explore what percentage of people need to be vaccinated for the virus to quickly die out. Is 50 per cent enough? You could also change how many people one person infects, or for how long a person can infect others before dying. Perhaps they each keep causing new infections for three rounds before stopping instead of only one. In what situations does the virus infect lots of people and when does it die out quickly?

What you are doing here is computer modelling or simulating the effects of the virus in different scenarios, and that is essentially how computer scientists make the predictions that governments use to make decisions about lockdowns and mask wearing, if they are "following the science". Of course, such models are only as good as the data that goes into them, such as how many other people does each person infect. In reality, this is data provided by surveys, experimental studies, and so on.

Gadgets based on works of fiction

by Paul Curzon, Queen Mary University of London

Why might a computer scientist need to write fiction? To make sure she creates an app that people actually need.

Images from Pixabay and iStock

Writing fiction doesn't sound like the sort of skill a computer scientist might need. However, it's part of my job at the moment. Working with expert rheumatologists Amy MacBrayne and Fran Humby, I am helping a design team understand what life with rheumatoid arthritis is like, so they can design software that is actually needed and so will be used and useful.

A big problem with developing software is that programmers tend to design things for themselves. However, programmers are not like the users of their software. They have different backgrounds and needs and they have been trained to think differently. Worse, they know the system they are developing inside out, unlike its users. An important first step in a project is to do background research to understand your users. If designing an app for people with rheumatoid arthritis, you need to know a lot about the lives of such people. To design a successful product, you particularly need to understand their unfulfilled goals. What do they want to be able to do that is currently hard or impossible?

What do you do with the research? 'Personas' are a powerful next step - and this is where writing fiction comes in. Based on research, you write descriptions of lots of fictional characters (personas), each representing groups of people with similar goals. They have names, photos and realistic lives. You also write scenarios about their lives that help understand their goals. Next, you merge and narrow these personas down, dropping some, creating new ones, altering others. Your aim is to eventually end up with just one, called a primary persona. The idea is that if you design for the primary persona, you will create something that meets the goals of the groups represented by the other personas it replaced.

The primary persona (let's call her Samira) is then used throughout the design process as the person being designed for. If wondering whether some new feature or way of doing things is a good idea, the designers would ask themselves, "Would Samira actually want this? Would she be able to use it?" If they can think of her as a real person, it is much easier to make decisions than if thinking of some non-existent abstract "user" who becomes whatever each team member wants them to be. It helps stop 'feature bloat' where designers add in every great idea for a new feature they have but end up with a product so complex no one can, or wants to, use it.

As part of the Queen Mary PAMBAYESIAN project we have been talking to rheumatoid arthritis patients and their doctors to understand their needs and goals. I've then created a cast of detailed personas to represent the results. These can act as an initial set of personas to help future designers designing apps to support those with the disease.

If you thought creative writing wasn't important to a computer scientist, think again. A good persona needs to be as powerfully written and as believable as a character in a good novel. So, you should practice writing fiction as well as writing programs.

How do you solve a problem like arthritis?

Treatment tailored for you

by Jo Brodie, Hamit Soyel and Paul Curzon, Queen Mary University of London

Some diseases can't be cured. Doctors and nurses just try to control the disease to stop them ruining people's lives. Perhaps smartphone apps can pull off the trick of giving patients better care while giving clinicians more time to spend with the patients who most need them? A Venn diagram is at the centre of the Queen Mary team's prototype.

What is rheumatoid arthritis?

Normally your immune system does a good job of fighting infection and keeping you healthy. But, if you have an autoimmune disease, it can also attack your healthy cells, causing inflammation and damage. Rheumatoid arthritis is like this: a painful condition that mostly affects hands, knees and feet as the person's immune system attacks their joints, making them swell painfully. It affects around 400,000 people in the UK and is more common in women than men.

People with the disease alternate between periods when it is under control and they have few symptoms, and days or weeks of painful 'flares' where it is very, very bad. During these flares it especially affects a person's ability to live a normal life. It can be hard to move around comfortably, do exercise – plus it interferes with their ability to work. It can also leave them totally reliant on family and friends just to do everyday things like dress or eat, never mind go out. This can lead to depression and puts a strain on friendships.

Treating the disease

Treatment, which can include tablets, injections, physiotherapy and sometimes surgery, slows the disease, keeping it under control for long periods. Sufferers are also given advice on lifestyle changes. This all reduces the risk of joint damage and helps people live their life more fully.

At appointments, doctors collect information to help them see how the disease is progressing. A Disease Activity Score (DAS) calculator lets them combine measurements for pain, how tender or swollen their patient's joints are and how many joints are affected. Regular blood tests keep track of the amount of inflammation and how the body is reacting to drugs. This helps them decide if they need to adjust the medication.

If it is caught early, modern medicine reduces the worst effects of the disease, helped by keeping a close eye on the Disease Activity Score as treatments may need to be repeatedly adjusted to control flares. This requires regular hospital visits which uses up scarce healthcare resources and is very time-consuming for patients. It is hampered because hospital appointments may only happen twice a year due to the number of patients. Everyone wants to give more personalised care, but hospitals just can't afford to provide it.

Supporting doctors

So, what do you do when there just aren't enough doctors to see everyone as regularly as needed to maintain their patients' wellbeing? One solution is to use remote monitoring with an app on a patient's smartphone, so involving patients more directly in their own care. They can use such apps to regularly record their own disease activity measurements, sharing the information with their doctor to save visiting the hospital.

A smart app

This is an improvement, but the measurements still require expert monitoring and can take more of the doctor's time. However, if smartphones can actually be made to be, well, smart, then they could help give advice between hospital visits and alert the hospital team, when needed, so they can step in. This might involve, for example, loading the app with background knowledge about rheumatoid arthritis, expert knowledge from lots of doctors, and creating an artificial intelligence to use this information effectively for each patient.

Images by Susanne Pälmer, Free-Photos and SnapwireSnaps from Pixabay

Hospital specialists and computer scientists at Queen Mary are developing such a prototype based on Bayesian networks as the artificial intelligence core. It involves finding out if patients and clinicians find such tools useful and acceptable (some people might find clinic visits reassuring, while some may be keener to avoid taking the time off work, for example).

Smart and patient centred

This still focusses on a clinician's view of treatment using drugs though. With a smartphone app we can perhaps do better and take the person's life into account - but how? The first step is to understand patient goals (see page 15). Patients would need to be willing to share lots of information about themselves so that the software can learn as much as possible about them. Eventually, this might be done using sensors that automatically detect information: how much pain they are in, how stiff their joints are, how much they move around, how long it takes them to get out of a chair, how much sleep they get, how often they meet others, if and when they take their medicine, and so on. Rather than just focussing on medical treatment it can then focus advice 'holistically' on the whole person.

The Queen Mary team's approach is centred around three different things: helping people with physical independence so they can move around and look after themselves; empowering them to manage their condition and general well-being themselves; and participation in the sense of helping them socialise, keep friendships and maintain family bonds.

The Bayesian network processes the information about patients and computes their predicted levels of independence, empowerment and participation, working out how good or bad things are for them at the moment. This places them in one of seven positions in a Venn diagram of the three dimensions over which areas need most attention. It then gives appropriate advice, aiming to keep all three dimensions in balance, monitoring what happens, but also alerting the hospital when necessary.

So, for example, if the Bayesian network judges independence low, participation high and empowerment low, the patient is in the Venn diagram intersection of low empowerment and low independence. Advice in the following weeks, linked to this area of the Venn diagram, would focus on things like coping with pain and stiffness, getting better sleep, as well as how to manage the disease in general.

By personalising advice and focusing on the whole person, it is hoped patients will get more appropriate care as soon as they need it, but doctors' time will also be freed up to focus on the patients who most need their help.

PhD student, Ali Fahmi, won a best student paper award at an international conference. He has been developing the Bayesian network that sits underneath the PAMBAYESIAN personalisation system.

Diagnose? Delay delivery? Decisions, decisions.

Decisions about diabetes in pregnancy

by Jo Brodie, Queen Mary University of London

In the film Minority Report, a team of psychics - who can see into the future - predict who might cause harm, allowing the police to intervene before the harm happens. It is science fiction. But smart technology is able to see into the future. It may be able to warn months in advance when a mother's body might be about to harm her unborn baby and so allow the harm to be prevented before it even happens.

Image by Daniel Nebreda from Pixabay

Image by Iuliia Bondarenko from Pixabay

Gestational diabetes (or GDM) is a type of diabetes that appears only during pregnancy. Once the baby is born it usually disappears. Although it doesn't tend to produce many symptoms it can increase the risk of complications in pregnancy so pregnant women are tested for it to avoid problems. Women who've had GDM are also at greater risk of developing Type 2 diabetes later on, joining an estimated 4 million people who have the condition in the UK.

Diabetes happens either when someone's pancreas is unable to produce enough of a chemical called insulin, or because the body stops responding to the insulin that is produced. We need insulin to help us make use of glucose: a kind of sugar in our food that gives us energy. In Type 1 diabetes (commonly diagnosed in young people) the pancreas pretty much stops producing any insulin. In Type 2 diabetes (more commonly diagnosed in older people) the problem isn't so much the pancreas (in fact in many cases it produces even more insulin), it's that the person has become resistant to insulin. The result from either 'not enough insulin' or 'plenty of insulin but can't use it properly' is that glucose isn't able to get into our cells to fuel them. It's a bit like being unable to open the fuel cap on a car, so the driver can't fill it with petrol. This means higher levels of glucose circulate in the bloodstream and, unfortunately, high glucose can cause lots of damage to blood vessels.

During a normal pregnancy, women often become a little more insulin-resistant than usual anyway. This is an effect of pregnancy hormones from the placenta. From the point of view of the developing foetus, which is sharing a blood supply with mum, this is mostly good news as the blood arriving in the placenta is full of glucose to help the baby grow. That sounds great but if the woman becomes too insulin-resistant and there's too much glucose in her blood it can lead to accelerated growth (a very large baby) and increase the risk of complications during pregnancy and at birth. Not great for mum or baby. Doctors regularly monitor the blood glucose levels in a GDM pregnancy to keep both mother and baby in good health. Once taught, anyone can measure their own blood glucose levels using a finger-prick test and people with diabetes do this several times a day.

In-depth screening of every pregnant woman, to see if she has, or is at risk of, GDM costs money and is time-consuming, and most pregnant women will not develop GDM anyway. PAMBAYESIAN researchers at Queen Mary have developed a prototype intelligent decision-making tool, both to help doctors decide who needs further investigation, but also to help the women decide when they need additional support from their healthcare team. This will save money but also be much more flexible than the current arrangement.

The team of computer scientists and maternity experts developed a Bayesian network with information based on expert knowledge about GDM, then trained it on real (anonymised) patient data. They are now evaluating its performance and refining it. There are different decision points throughout a GDM pregnancy. First, does the person have GDM or are they at increased risk (perhaps because of a family history)? If 'yes' then the next decision is how best to care for them and whether or not to begin medical treatment or just give diet and lifestyle support. Later on in the pregnancy the woman and her doctor must consider when it's best for her to deliver her baby, then later she needs ongoing support to prevent her GDM from leading to Type 2 diabetes. Currently in early development work, it's hoped that if given blood glucose readings, the GDM Bayesian network will ultimately be able to take account of the woman's risk factors (like age, ethnicity and previous GDM) that increase her risk. It would use that information to predict how likely she is to develop the condition in this pregnancy, and suggest what should happen next.

Systems like this mean that one day your smartphone may be smart enough to help protect you and your unborn baby from future harm.

Queen Mary's Mariana Neves was keen to work on the PAMBAYESIAN project. Her background is in statistics but she was always interested in programming and found that PAMBAYESIAN not only gave her a good opportunity to use her knowledge in a real-world application, but it also let her work with a wide range of people such as various healthcare professionals, computer scientists and other statisticians. As she says, "I learn a lot working with people from different backgrounds and being part of a team that is building a tool that will improve people's lives feels very rewarding."

Image by Pexels from Pixabay

Back(page) to health

by Paul Curzon, Queen Mary University of London

Improvements in technology and decision making are transforming the way we look after our health. Here are some more interesting ideas to keep people alive and well.

The future is in your poo

You've heard of telling a person's future from reading their tea leaves. Scientists believe an effective way of seeing a town's future may be in the poo. By looking for infection in the waste at sewerage works it's possible to get fast and accurate local knowledge of where infection rates are high and where low to feed into decision making tools.

Health advice: Stay in the toilet, Stay safe. Help the NHS.

Virtually breaking quarantine

The game, World of Warcraft, a multi-user dungeon game, helped virologists understand how people might behave in pandemics. The game's developers released a plague that could be passed between avatars. The game's contaminated area was quarantined. Rather than dying out, the virus escaped - because people broke into the quarantined areas to gawk, then left taking the virus with them.

Health advice: Your avatar should obey quarantine rules too!

The missing bullet holes

To stay healthy in a war, avoid being hit by a bullet. In World War II, many aircraft returned badly damaged. Abraham Wald studied them to decide where better armour was needed. There were more bullet holes in the fuselage than the engines. Where would you add the armour? Abraham added it where there were no bullet holes. He reasoned that the lack of holes in places like engines on returning planes meant that being hit there brought the plane down. Being hit elsewhere did not kill the pilots as those planes made it home!

Health advice: Dodge bullets by making good decisions...

Cybersick of virtual reality

A problem with virtual reality is that wearing a headset can be so immersive that it makes some people actually sick. This happens if you move about when watching a 3D video that was shot from a single place. Artificial intelligence software has come to the rescue, detecting puke-inducing movement and automatically correcting the image.

Health advice: If no bucket, always keep an AI handy.

Shining light on cancers

Cancer treatments like chemotherapy and radiotherapy make patients ill. Some drugs make cancer sensitive to light, allowing tumours to be killed by painlessly shining light on them instead. Sadly, that's not easy when cancers are inside the body. A new Japanese solution is an LED chip, based on the technology used by contactless payment cards to provide power from a distance. Surgeons place it under the skin and leave it there. They glue it in place using a sticky protein from the feet of mussels. It shines low-intensity green light on the cancer, shrinking it.

Health advice: Stick a chip to your tumour

Smart sometimes means no gadgets

Being smart about health doesn't have to be high-tech or even involve drugs. Exercise, for example, can be as effective helping with depression as taking medicine. Being out in nature can help too, so sometimes it's worth leaving the gadgets behind and just going for a walk to enjoy the beauty of nature.

Health advice: Walk weekly in the woods

Bayesian Baffler answer: The probability of Fred's innocence is just less than one in two so the defence are right.

cs4fn is edited by Paul Curzon and Jo Brodie of Queen Mary University of London. Summer 2021.Thanks to Sue White for proof reading. EPSRC supported this issue through the PAMBAYESIAN research grant (EP/P009964/1): see **www.pambayesian.org** Cover Image by Myriams-Fotos from Pixabay.