

Causing trouble

The language of science can't distinguish between cause and effect. Solving this problem could put research on firm foundations, reports **Ciarán Gilligan-Lee**

IN THE mid-1990s, an algorithm trained on hospital admission data made a surprising prediction. It said that people who presented with pneumonia were more likely to survive if they also had asthma. This flew in the face of all medical knowledge, which said that asthmatic patients were at increased risk from the disease. Yet the data gathered from multiple hospitals was indisputable: if you had asthma, your chances were better. What was going on?

It turned out that the algorithm had missed a crucial piece of the puzzle. Doctors treating pneumonia patients with asthma were passing them straight to the intensive care unit, where the aggressive treatment significantly reduced their risk of dying from pneumonia. It was a case of cause and effect being hopelessly entangled. Fortunately, no changes were rolled out on the basis of the algorithm.

Unweaving the true connection between cause and effect is crucial for modern-day science. It underpins everything from the development of medication to the design of infrastructure and even our understanding of the laws of physics. But for well over a century, scientists have lacked the tools to get it right. Not only has the difference between cause and effect often been impossible to work out from data alone, but we have struggled to reliably distinguish causal links from coincidence.

Now, mathematical work could fix that for good, giving science the causal language that it desperately needs. This has far-ranging applications in our data-rich age, from drug discovery to medical diagnosis, and may be


the essential tool to resolve this fatal flaw.

A mantra most scientists can recite in their sleep is that correlation doesn't imply causation. A simple example illustrates why. Data from seaside towns tells us that the more ice creams are sold on a day, the more bathers are attacked by sharks. Does this mean that ice cream vendors should be shut down in the interests of public safety? Probably not. A more sensible conclusion is that the two trends are likely to be consequences of an underlying third factor: more people on the beach. In that case, the rise in ice cream sales and shark attacks would both be caused by the rise in beachgoers, but only correlated to each other.


What's going on?

This analysis seems simple enough. The trouble is that the data alone can't point us in the right direction. We need some external knowledge – in this case, that a surge in people enjoying the beach on a hot day can adequately explain both trends – to correctly distinguish correlation from causation. As the data at hand gets more complicated and less familiar, however, our ability to distinguish between the two falls short.

These subtleties were lost on some of the early pioneers of statistics. One notable offender was Karl Pearson, an English mathematician and prominent eugenicist of the early 1900s. Pearson believed the mathematics of correlation was the true grammar of science, with causation being only a special case of correlation, rather than



MICHAEL HADDAD



a separate analytical concept. The statistical tools he developed remain part of the bedrock of scientific practice, and are taught in every undergraduate statistics class. As a result, for over a century, many scientific discoveries have been based on flimsy correlation rather than firm causation. This has implications far beyond the seaside. Data and correlation can tell you which of two treatments led patients to recover faster, but not why. They also can't tell you how to make treatments better, or even what to prescribe a given individual.

"If you want to actually cure a disease, or make it less likely someone gets a disease, you need to have a causal understanding," says Jonas Peters at the University of Copenhagen in Denmark. The importance of understanding causality can't be overstated, says Elias Bareinboim at Columbia University in New York. "I don't think there is any way of doing science without causality," he says. "It is the code running the system."

At the same time, science is poorly equipped to deal with questions of cause and effect. Since Galileo, modern science has been communicated using the language of algebra and equations. Physicists can write an equation describing the relationship between atmospheric pressure and the reading on a barometer, but this equation says nothing about whether it is pressure that causes the barometer reading or vice versa. The language of algebra is completely agnostic to the question of which came first.

In the early 1990s, dissatisfied with this state of affairs, Judea Pearl at the University of California, Los Angeles, set out to give science the causal language it desperately needed. His solution was to introduce a new mathematical language of "doing", allowing us to distinguish between cause and effect. If I "do" by intervening to force pressure to change, then the reading on the barometer will shift. But if I "do" a change in the barometer reading, the pressure doesn't alter as a consequence. Intervening on the cause will change the effect, but any intervention on the effect won't change the cause.

To convey this in mathematical terms, Pearl invented a new operation to sit alongside addition, subtraction and the others. Just like the other operators, his "do operator" can manipulate variables – like the number of ice creams sold – in specific ways. Whereas addition combines the value of two or more variables, the do operator sets a variable to a specific value, irrespective of anything else.

To see why this is needed, let's head back to the seaside. If we wanted to establish the true ➤

relationship between ice cream eating and shark attacks, the scientific best practice would be to carry out a randomised control trial. This would involve randomly assigning beachgoers into two subgroups of equal size. One group would be given ice creams and the other wouldn't. Both would then be let loose in shark-infested waters, and the number of shark attacks on each group compared.

The composition of the subgroups is random, so all other potentially confounding factors, such as age, height and tastiness of flesh, are controlled for. Any remaining correlation can be explained only if there is a direct causal relationship between eating ice cream and being attacked by a shark. Pearl's do operator mathematically simulates changing the amount of ice cream someone eats, regardless of any confounding factors that would influence both the eating of ice cream and being attacked by a shark. By changing ice cream consumption alone, and keeping everything else fixed, any corresponding change in shark attacks must be due to eating ice cream, as it is the only variable that changed.

Pearl's great insight was to show that with the do operator you could effectively simulate a randomised control trial using only observational data and extract causal connections. This was a game changer, because performing real-world randomised controlled trials can be expensive and complicated, not to mention unethical. To perform a controlled trial to examine the link between pneumonia and asthma, for example, half the group would have to be infected with pneumonia.

The work won Pearl the Turing Award in 2011 – the computer science equivalent of a Nobel prize – and formed the foundations of what has come to be known as the theory of causal inference.


Besides putting science on a firmer causal footing, this mathematical framework is helping to solve problems in many disciplines, says Bareinboim, chief among which is the replication crisis that has plagued medicine and the social sciences. In the past decade, doubts have arisen about many headline-grabbing studies in these fields – from the notion that maths problems are easier for students to solve if written in a fuzzy font to the idea that willpower is a finite, exhaustible resource – because the results of their underlying experiments couldn't be replicated. In 2015, a massive replicability study in psychology found that results of 60 per cent of studies couldn't be reproduced, casting a vast shadow across the discipline.

Bareinboim believes causal inference could

help clear these problems up. In many cases, he says, the original tests were susceptible to confounding factors that the experimenters may have been unaware of, and subsequent replication attempts might have dragged new causal relationships into the mix. One classic example concerns the effect of happiness on economic decisions, which was originally measured by showing participants footage of US comedian Robin Williams. By the time the replication experiment was conducted, Williams had died, potentially skewing the participants' response. In addition, the subjects in the original study were from the US, but those in the replication one were British. By not controlling for such confounding effects, the replication study cannot legitimately comment on the original finding.

The applications extend well beyond science. "As soon as you're looking to improve decision-making, you want to understand cause and effect. Which is, if I were to do this, how would the world change?" says Suchi Saria at Johns Hopkins University in Maryland. Economists in particular were early to the

"If you don't understand the causal process, you are susceptible to bias"



PLANPICTURE/EE COENDEERS

party, realising that many of the problems they wished to solve required a causal toolkit.

Such tools could determine the effects of specific policies, such as whether an increased tax on cigarettes reduces the health impacts of smoking. For such a complicated issue, however, Pearl's mathematical tools become incredibly challenging. The relationship between smoking and health is influenced by a panoply of confounding factors, including age, sex, diet, family history, occupation and years of education. To home in on the causal connection we care about, we can look only at parts of the data where the other factors are constant. But for each confounding variable we control for, the corresponding data set gets smaller. Eventually, we are left with so little data that no robust conclusions can be drawn at all.

To overcome these difficulties, Susan Athey at Stanford University in California and her colleagues have developed techniques to approximate Pearl's methods while still holding on to as much data as possible. They aren't alone. Tools of this kind are also having a big impact in healthcare, an area where understanding cause and effect can be life-saving. Knowing that a disease is highly correlated with certain symptoms, or that a drug is highly correlated with recovery, isn't enough, and basing medical decisions on such information can be dangerous.

Saria is using causal inference to create tools to help doctors make decisions by comparing the effect of different medical actions. However, working with medical data comes with challenges. "We may be reflecting back biases that are not the true underlying phenomena in nature," she says. For example, unequal access to treatment means that the US spends less money caring for black patients than for white patients. Some algorithms conclude from such data that black patients are healthier than equally sick white patients, which is patently false.

For Kira Radinsky at the Technion-Israel Institute of Technology, causal understanding is key to a more equitable health system. "If you don't understand the causal processes, you are susceptible to bias in the data," she says. "As soon as you do understand them, you can clean out the bias."

This highlights one problem that causal inference can't solve. Before Pearl's techniques can be employed, the causal relationships need to be known. Left to analyse shark attack and ice cream sales data, for example, they wouldn't be able to determine the connection between the two without knowing that an increase in beachgoers could explain both trends.



DEEPOOL BY PLAINPICTURE/ASHLEY/ISTOCK

Shark attacks may rise along with seaside ice cream sales, but are the two trends linked?

To quantitatively compute the effect of treating someone's symptoms with a certain drug, we need to know that a causal relationship between that drug and those symptoms exists. The standard approach to this is to find out from experts on the subject.

But getting this causal knowledge from experts can be difficult and takes time, says Radinsky. The approach she and her collaborators have taken to streamline this process is to mine causal relationships from medical papers that actually verified their existence through experiments. By applying this causal knowledge to drug repurposing – using existing medicines in new ways – they have already found new treatments for hypertension and diabetes.

Learning from the data

This is a fruitful and powerful approach, but not every field has a large collection of online research papers with proven causal links just waiting to be exploited. This has led researchers in other disciplines to wonder if causal relationships could be discovered from purely observational data. The age-old difficulty of distinguishing correlation from causation would seem to rule this out. Yet a new generation of researchers bold enough to investigate the problem is starting to realise that it might not be as impossible as it sounds.

One approach gaining ground involves looking for patterns that hold true regardless of circumstances. Increased atmospheric pressure always causes a barometer reading to change, for example, regardless of whether you are in London or New York, on Earth or Mars. Likewise, physicians in different hospitals or countries may differ in how they treat people, but the underlying causal relationships between diseases and symptoms don't vary. The key idea behind new work being led by Peters and others is that this consistency can act as a signature of the underlying causal process, allowing Pearl's tools to be deployed.

To put this principle to the test, he and his colleagues dived into a complex sociological question: the true causes of a country's total fertility rate. These rates vary dramatically around the world, and understanding the factors determining them could be a boon for governments seeking to support their populations. By looking for consistent patterns in data from multiple countries, Peters and his colleagues found that mortality rates of young children were important drivers of fertility rates, a finding that tallied with previous studies from around the world. "When child mortality is high, families tend to have more children, even if none of their own children have died," says Adrian Raftery, a sociologist and statistician at the University of Washington in Seattle. "This may be proactive, to try to make sure that they do have a family."

Bareinboim is very excited about the group's ability to obtain causal insights from observational data alone. "When that work came along, it was amazing," he says. Peters and his collaborators are now using the invariance principle to paint a causal picture of biosphere and atmosphere interactions, with potentially dramatic consequences for our understanding of climate change.

But like Pearson's statistical analysis over a century ago, it isn't a silver bullet. To truly disentangle cause, effect and correlation, scientists will always need extra contextual information. Without knowing how beachgoers behave, for instance, or how doctors treat people with pneumonia if they have asthma, no analysis in the world could correctly parse even the largest data set. "The problem is the data-generating process," says Athey, "not the limits of our brains." ■



Ciarán Gilligan-Lee is a senior researcher at University College London and Babylon Health. Follow him @quantumciaran