


The Sydney Morning Herald

 [Print this article](#) |  [Close this window](#)

Dr Google will see you now

Paul Biegler

Published: January 2, 2017 - 10:37AM

You're whizzing along a dark, outback highway when the symptoms start.

Your mouth is so dry it's getting hard to swallow or talk, your grip on the steering wheel is weakening and that centre white line just turned into two.

You pull in at a small town hospital where the tired junior doctor tells you it's probably a virus, and then hands your case over to night staff.

But the night doc comes with a difference.

She retakes your history wearing Google Glass, and the web-enabled headgear uses voice recognition to input your symptoms into a massive data base.

Within seconds the cloud sends back a diagnosis.

It's botulism; very rare, often fatal, and you're just in time to get the antitoxin.

If it all seems a bit far-fetched the technology is, in fact, on our doorsteps.

Google Glass is alive and kicking; San Francisco health start-up [Augmedix](#) is refining the internet-browsing eyeglasses to give doctors real time access to patients' electronic health records and the web.

And Google Glass is compatible with apps, such as [Isabel](#), that can compute the likely top diagnoses from a patient's symptoms and, according to a [2016 review](#), even improve on the accuracy of clinicians.

In a world where the volume of healthcare data, including patient notes, lab tests, medications, imaging, and research articles, will soon be counted in [yottabytes](#) – that's 10 to the power of 24 Gigabytes and enough, [according to IBM](#), to fill a stack of DVDs that would stretch from Earth to Mars – it's understandable doctors could use a little help.

But the march of technology is causing frissons of nervousness in medical circles, not just about how to incorporate it into everyday practice but, ultimately, whether jobs now done by doctors could one day be taken by machines.

"There is the universe of what we know, and then there is what I know," says Herbert Chase, a physician and professor of medicine at Columbia University.

"Medical practitioners can't be expected to master the opus required to recognise all diseases," Chase says.

"In terms of knowledge, diagnosis, optimal treatment, guideline-based care, I'm pretty sure that machines are already, in some ways, much better than we are."

Chase is referring to a branch of AI that promises a tectonic shift in how medicine is practised: it's called machine learning.

Machine learning is a way of training computers to tell things apart that leaves the "learning" bit up to the computer itself; its artificial neural networks forge "knowledge" much like our own brains.

Take the question of whether a shadow on a chest X-ray is a cancer or something less sinister.

A typical machine-learning approach would feed the computer a massive database of chest X-rays with shadows that had been proven cancerous or benign.

The computer would then come to its own conclusions about what features of the X-rays robustly predicted cancer.

What's revolutionary is that, because the computer "sees" differently to a radiologist – it objectively applies statistics to millions of pixels – it could, theoretically, discover features in an X-ray not previously thought to flag cancer.

With machine learning, according to a [September editorial](#) in *The New England Journal of Medicine*, "we let the data speak for themselves".

Letting the data speak for themselves has, this year alone, delivered lung cancer prognoses with [greater accuracy than pathologists](#) and, in a study lead-authored by Google scientists [published in JAMA](#), predicted diabetic eye disease better than a panel of ophthalmologists.

And it may have saved its first life.



The speed of IBM's supercomputer known as Watson is making it possible to do diagnoses of rare cancers that would take doctors much, much longer. Photo: AP

In August [Japanese doctors reported](#) using IBM's super computer Watson to crunch through a patient's myriad genetic mutations to diagnose a rare leukaemia.

The task would have taken a person two weeks; Watson took 10 minutes.

And the stakes couldn't be higher.

A [May report](#) in the *British Medical Journal* concluded that, after heart disease and cancer, medical error is the third-highest cause of death in the US, accounting for a staggering 251,000 lives.

Wrong diagnoses make up nearly a third of all medical errors, and at least part of the problem is that human medical thinking can be derailed by cognitive biases.

"Anchoring", for example, describes our tendency to fixate on a first piece of information or idea.

"Humans get stuck on diagnoses. It's called premature closure, and you just don't expand the differential diagnosis enough. That's where many errors are made," says Chase.

Computers, by contrast, consider all possible diagnoses on a level playing field.

Data crunching might also help doctors detect illness earlier.

Chase has built an algorithm that combs patients' notes for the telltale symptoms of multiple sclerosis, such as numbness and tingling.

He [presented his findings](#) at the American Medical Informatics Association conference in November.

"We found that about 40 per cent of patients with documented MS were identified by the machine up to two years before it was recorded in their notes that they had MS," Chase says.

"They had the signs and symptoms of MS well before the diagnosis was made."

But, while diagnosis is important, Chase thinks the potential for machine learning to improve treatment is a game changer.

"There are not that many new diagnoses that have been discovered in my lifetime. The explosive side is therapy ... what is the cutting-edge therapy of common diseases?"

Chase gives the example of choosing drug A or B for high blood pressure.

Each has a given probability of effectiveness, of side effects and of interacting with the patient's other medications; the average American over 65, according to Chase, takes seven.

Do the maths and the decision tree starts to look very gnarly, so how do you work out which drug is better?

"A human can't compute that," Chase says, "this is a space where a machine can accomplish something a human can't."

A closely related space is the one in which we gauge a patient's response to treatment.

Dr Ben Glocker, a lecturer in medical image computing at London's Imperial College, is using machine learning to refine how we measure the size of brain tumours.

"It is relatively easy to see if someone has a tumour or not. But often you need to monitor progression, to see if a treatment such as a drug is successful or not," Glocker says.

"Doctors are very good at pattern recognition. What humans don't do so well is to see change," he says.

Dr Ben Glocker - Imperial College London



Doctors currently use the relatively blunt measure of a tumour's maximal diameter to see if it is shrinking or growing; Glocker's algorithms drill down into a tumour's complex structure – including dead or swollen tissue as well as active cancer cells – to give a much more detailed picture of any change after treatment.

And Glocker is applying the same fine-grained analysis to the brain scans of people after head trauma, with the aim of linking findings such as swelling and bleeding to outcomes including cognitive and memory impairment.

The research, part of the Europe-wide CENTER-TBI study, might in future, for example, help determine the optimal lay off for a head-injured rugby player.

"If you take information from the image that a radiologist can't routinely measure and add it to your prediction model you get more accurate predictions," says Glocker.

But for all these tangible advances there remains a major hurdle to implementation.

Andrew Bradley, a professor of biomedical engineering at the University of Queensland, is working with Breast Screen SA on an algorithm that combines the results of the mammogram, ultrasound, and MRI to improve breast cancer diagnosis.

"For any of the machine-learning techniques to be useful to clinicians they can't just be a black box. They have got to explain themselves to some degree, so that the clinician can look at it and go 'Oh, that's why'," Bradley says.

The big danger of the black box is that when it does silly things, nobody can see.

And in machine learning, right up there on the silly scale are things called "adversarial images".

To explain, say you take a digital image of a bus and layer it with a second image that is just pixel "noise" (imagine the old TV test pattern given a makeover by Jackson Pollock).

To you and me the new image will still look like a bus, but to a computer it could [look like an ostrich](#).

And while in every day computing these adversarial images arouse curiosity, "in the medical context they frighten the bejesus out of people", Bradley says.

That's because with just a smattering of noise a mammogram with an obvious cancer might no longer even look like a breast to a computer, and so the tumour could be missed.

"Nobody expects the mammography system to be 100 per cent accurate. There will always be false negatives and false positives," Bradley says.

"It's when the mistake is a stupid mistake that a first-year med student could pick up people get unhappy. That's when you lose any credibility you once had. It's a death blow," he says.

One antidote to the inscrutable black box is, Bradley says, a white box that makes the workings of the medical machine transparent.

Bradley is involved in a study, [led by radiologist and University of Adelaide PhD candidate Luke Oakden-Rayner](#), that shines light into the black box.

The study reviewed 48 people who had chest CT scans in 2009, half of whom died by 2014.

Just by analysing specific parts of the scans such as muscle, fat, heart and lungs a machine-learning model predicted with 65-70 per cent accuracy which patients would die over the five-year span.

Critically, the study also showed the machine's predictions could be linked back to scan changes – calcified blood vessels, emphysema, bone thinning – that a doctor could use as a "face value" check that the black box was credible.

"This is one way to make the black box white. To highlight those parts of the image that contributed to its diagnosis," Bradley says.

"Even if it gets it horribly wrong doctors could see why it made the mistake, and perhaps forgive it," he says.

Back in the consulting room, Dr Christopher Pearce, a Melbourne GP and president of the Australasian College of Health Informatics, sees a widening role for machine learning in everyday patient care.

"Say you have a patient with asthma. The system could trawl through the data, pick up subtle changes in patterns of medication usage, and predict a greater risk of hospital admission," Pearce says.

Ultimately GPs might use that information to adjust medication, arrange district nursing or schedule tests to avoid the need for hospitalisation.

Pearce doesn't see machines pinching his job anytime soon, but the zephyr of change is being felt more keenly in radiology circles.

Melbourne-based Capitol Health recently [invested \\$10 million in health start-up Enlitic](#), founded by a local data scientist Jeremy Howard, to deploy its machine-learning algorithms throughout its radiology suites.



Melbourne expatriate Jeremy Howard founded Enlitic, a medical data company in the US that's disrupting the radiology profession.

According to [its website](#), Enlitic's technology "detected lung cancer nodules in chest CT images 50 per cent more accurately than an expert panel of radiologists."

It's this kind of claim setting radiologists' nerves on edge.

Dr Mark Michalski, a radiologist and director of the Centre for Clinical Data Science at Massachusetts General Hospital, attended the November conference of the Radiological Society of North America.

"I talked with a number of people who are very concerned. The predominant question was, 'will AI subvert radiologists?' " Michalski says.

Michalski himself is involved with technology that might make his colleagues edgy.

He chairs the Medical Advisory Board of [Butterfly Network](#), a start-up that's building a smartphone-sized ultrasound that will enable non-specialists to diagnose conditions such as a bleeding ectopic pregnancy, a potential life-saver in areas with limited access to imaging services.

"The hope of this technology is to make ultrasound cheaper and more effective, and to make it possible for ultrasound to be democratised," Michalski says.

But Michalski rejects the suggestion it could replace radiologists.

"There is no solution today that allows us to have a radiologist in a box," he says.

Glocker's take is that radiology is a specialty in transition.

"At some point we probably won't need humans looking at chest X-rays," Glocker says.

"But that doesn't mean we won't need radiologists. Radiologists will perhaps become more like data scientists, using the

data from machines to make better decisions."

Pearce also sees doctors as having indispensable qualities.

"Humans have access to a different set of data. It's the soft stuff, the look of the patient, the way they talk. All data has a social context," he says.

"Machine learning is better at using the data it gets. We're better at using the data we get."

A thorny question remains, nonetheless, whether doctors practising without machine-learning assistance, including diagnostic apps such as Isabel, can maintain patient care as the technology sets new benchmarks.

"Standard of care is not a single line in the sand. The line moves over time," Bradley says.

@pbiegler

This story was found at: <http://www.smh.com.au/national/health/dr-google-will-see-you-now-20161231-gt9dr1.html>