



**Australian Government**

**Chief Scientist**

**DR ALAN FINKEL AO**

**2018 Smallwood Oration**

***The AI will see you now***

**Thursday 16<sup>th</sup> August 2018**

**Austin Health  
MELBOURNE**

Back in 2006 I had recently sold my business, tested retirement, found it to be distasteful and was picking up some new interests in publishing, education and research governance. For reasons that are beyond simple explanation, I was invited to turn my experience as a businessman towards leading the merger of the Howard Florey Institute, the Brain Research Institute and the National Stroke Research Institute. That led to six intense months of meetings and negotiations, much of which took place at the Austin Hospital campus because this was the home of two of the three institutes.

I delighted in all the people whom I met here at the Austin during the merger negotiations and after its successful completion, so I am particularly pleased to be back on campus this evening.

Now, when I accepted Christine McDonald's invitation to deliver the Smallwood Oration I didn't anticipate that the words "healthcare" and "data" would be so interesting to so many people.

I'll have more to say on the My Health Record in the course of this speech.

But the commentary in the media has got me thinking.

I wonder how many of the people now talking about health data benefits versus privacy provisions have also sent a cheek swab to Ancestry.com or 23andMe?

They might be interested to know that 23andMe is big news at the moment in the United States.

The pharmaceutical giant GlaxoSmithKline has just acquired a \$300 million stake in the company, in exchange for exclusive access to its data.

That's the genetic information of more than 4 million people.

As one reporter put it, you don't make a fortune selling spit kits. You're selling the spit kits so you can make a fortune with the results.

Another interesting question: how many people here have a FitBit? And how many of you were given that FitBit for free by a helpful insurance provider, or perhaps your employer?

The only price you pay is sharing your information... about every step you take, every minute you sleep, every time your heart beats.

That's all.

Of course, you don't need a FitBit to track your exercise – you can do it with your smartphone.

Do you unlock it with a fingerprint? Biometric data.

And when you use your smartphone to call the tax office, or the bank, do you verify your identify just by speaking?

Your voice-print: *more* biometric data.

Then there are the symptoms you've typed into Dr Google. The loyalty card at your pharmacy. The ultrasounds you've shared on FaceBook.

That pacemaker in your chest.

And yes, it's happened: data from pacemakers have been requested by police.

In this case, the gentlemen in question was accused of burning down his house to collect the insurance.

He told the police he saw the flames, grabbed some bags, and bolted.

His nice and steady heartbeat told the police he was lying.

As they say, the heart has its secrets. It's just that now, it can spill them.

\*\*\*

I don't want to imply that the battle for privacy in healthcare is futile.

On the contrary, I want to reiterate that we all have to square up to that battle every day. We hit decision points, every day. We opt in, or opt out, every day.

In the case of the My Health Record, there are powerful arguments in favour of participation.

One. You save money. No more paying for duplicate procedures or tests.

Two. You get better care. No more dealing with multiple specialists, all treating the same body, and all relying on you, the patient, to keep the rest of them informed.

Three. You get continuity. Because at some point in your life you *will* change doctors. You will travel. Your doctor will retire. You will end up in an emergency department: you don't where, you don't know when, and you don't know for what.

I think to myself, at that moment, would I want the treating team to read my name on my driver's licence, look up my record, and see my allergies, my medical history and my current medications?

Yes, I definitely would.

And if that were to happen to my wife, or my mother, or my sons, would I want them to have My Health Records?

Yes, I definitely would.

I did hold some of the concerns that were initially flagged, regarding access to information by other government agencies, and what happens if you want to delete your records.

But I am heartened: first, by the Government's recognition of those concerns, and second, by the clear response.

The legislation will be strengthened to ensure that no record can be released to police or government agencies, for any purpose, without a court order.

Further, we have a firm commitment that if anyone wants to opt-out of the system permanently, at any time in the future, their record will be deleted for good.

But whatever we decide about the My Health Record, the big conversation about our data remains.

And it is the prelude to the even bigger conversation that follows: about the role in healthcare for artificial intelligence. AI.

Data and AI: the two are inseparable.

Let's stop for a moment to explain what we mean by AI.

\*\*\*

There are many definitions of artificial intelligence – but broadly, we're talking here about computer technology that can do tasks that ordinarily would require human intelligence.

More formally, AI is the combination of three things: machine learning algorithms, high quality data and a training procedure.

That's just like human intelligence, three things required: innate ability, access to knowledge and a good teacher.

Once trained, the AI can sort through the data to identify patterns and make predictions.

Where do we find a lot of data?

It's no secret: in healthcare.

And for decades, we've been promised extraordinary things.

AI doctors! AI nurses! AI performing surgery, reading brain scans, running hospitals!

And always, just a decade away.

The bold predictions have always failed.

And a key reason they fail is that healthcare might involve a lot of data... but it's got to be data that AI developers can actually use.

The easiest targets are narrowly defined tasks where we have high quality datasets and very structured processes.

For example: in an IVF clinic.

Doctors have to decide which of the eight or so embryos to implant.

A human doctor can assess their appearance at a handful of checkpoints.

AI can continuously watch all eight embryos as they grow, 24 hours a day, five days straight.

A human doctor might have seen several thousand embryos in the course of her career.

The AI has access to a database of tens of thousands and knows which ones were successful.

So the AI learns and does it better.

A defined task, structured process, a high quality dataset.

And wouldn't it be nice if *all* our health problems were that simple.

But of course, they're not.

How long do we spend teaching medical students to greet the patient... listen to his or her story... ask the right questions... understand what isn't said...

Doctors know what they need to know.

But AI doesn't.

Consider a study from 2015.

A team of Microsoft researchers had a dataset covering 14,000 patients with pneumonia: some who got better without medical intervention, some who were successfully treated, and some who died.

They wanted to work out if they could train an algorithm to work out which patients were high risk – who needed to go to hospital, and which patients were low risk – who were better off staying home.

Based on the dataset, the algorithm determined that – wait for it – patients with asthma were *low risk*.

It reached this startling conclusion because it didn't have the right information.

The patients with asthma in the dataset *were* less likely to die – but it wasn't because they were low risk.

It was because they knew that they were high risk, so they paid very careful attention to their breathing.

And whenever they had a problem, they got themselves to a doctor quick smart.

If the AI were running the clinic, the asthmatics would be packed off home.

So if we want AI to assist doctors with the complex decisions they have to make every day, then we need to be able to consolidate multiple sources of data.

And then we can take the next step.

We can develop AI that can draw on far more data points than any medical specialist working alone could possibly contend with.

We can make it possible for doctors to tailor treatment precisely to the genetic makeup and life experiences of the individual patient.

We can maximise the likelihood that the first drug they try is the right drug, in the right dose, at the right time.

It can be done at speed, and at scale.

And that finely calibrated care can be wrapped around every person.

That's the promise of precision medicine.

I know, because I had the privilege of commissioning a report on that very topic.<sup>1</sup>

As Chief Scientist, I serve as Executive Officer on the Commonwealth Science Council.

The Council requested the report from ACOLA, our Learned Academies.

It's a fantastic report, one of the best of its kind I've seen.

It's now being used by Minister for Health Greg Hunt and his Department as the intellectual framework for the \$500 million Genomics mission announced in the Federal Budget.

As you know, there's no big magic solution to any policy challenge – least of all, health.

But the potential here is breathtaking.

And all of it comes back to that essential foundation of trust.

Trust in the security of our data.

And trust in the tool of AI.

\*\*\*

So are we capable of extending that trust?

To earn and preserve the trust of patients?

I don't make a habit of visiting sites like RateMyDoctor.com – but I did for the purposes of this speech.

The same words crop up again and again in the five star ratings.

“Caring”.

“Really listens”.

---

<sup>1</sup> The report “The Future of Precision Medicine in Australia” (2018) is available at <https://acola.org.au/wp/pmed/>.

“Treats me like a human”.

“Seems to really know you”.

And in the zero star ratings:

‘Mechanical’.

‘Robotic’.

“On auto-pilot”.

So we clearly want things from doctors that we don’t associate with machines!

There are some indications that we might in principle be accepting of at least some AI, under some conditions.

I’m a member of an organisation called the IEEE, the largest international body for engineers and technology professionals.

For two years now, the IEEE has surveyed thousands of parents in the millennial age bracket, 20 to 36, about their attitudes to AI in health.

The survey spans five countries: the United States, United Kingdom, India, China and Brazil.

This year, more than half the surveyed parents felt comfortable for their child to have a fitness tracker from infancy.

More than three in four parents, in every country, had at least some trust in AI to diagnose and recommend treatments for their sick child.

One in two parents in the US and UK said they were “likely” to use AI-powered chatbots to help them work out what to do based on their child’s symptoms.

In China and India, that figure soared to five in six.

And then, the critical question. Would you trust doctors using data from artificial intelligence to make life or death health decisions for your child?

In every country, a clear majority answered in the affirmative.

Of course, the answer you give in a survey could be very different from the answer you give when your child has a fever and your parental anxiety is spiking.



And your answer might be different again if there *wasn't* a human in the loop. If the decision came down solely to the conclusion reached by an AI.

So perhaps the best we can say about patient attitudes is: it remains to be seen.

And before we get to the point of presenting a technology to a patient, it has to pass muster with doctors, hospital authorities and regulators.

All three present challenges of their own.

Let's start with doctors.

A constant complaint of AI developers is that doctors don't perceive their technology really adds value.

IBM confronted this problem with Watson, put forward as an aid to complex decision-making in the treatment of cancers.

When Watson agreed with the human oncologists, its advice was considered unnecessary.

When Watson *disagreed*, its advice was considered wrong.

So the doctors weren't impressed.

On a deeper level, doctors know that the weight of responsibility they carry is profound.

They are right to demand a very high burden of proof before they delegate any part of their duty of care.

Next, the hospital authorities.

Here, too, some reluctance is entirely understandable.

The community's tolerance for trial and error in hospitals is very slim.

Even the trial *without* the error causes disquiet.

In January this year, Google published a research paper in Nature demonstrating the use of AI to predict what will happen when a patient is admitted to hospital.

They took the de-identified data of over 200,000 patients from two hospitals, collected over a seven year period.

That dataset contained a wealth of information about the patients: their vital signs, their medical histories, how they presented, how long they stayed.

Google's algorithm was able to predict patient outcomes far more reliably than the hospitals' existing methods.

Now this was a retrospective study. Nothing actually happened to a patient. Nothing changed about the way that hospitals were operating.

When the story broke in June, that message was helpfully boiled down by the tabloids.

From The Sun: *Google AI can now predict when a patient will DIE.*

DIE in all capital letters.

News.com gave it a name: The Google death predictor.

The implication was clear: soon enough, you'll present at a hospital, and a robot will decide if you're worth saving.

It was not, perhaps, the sensitive and nuanced approach to the issue that hospital authorities might want to see.

And what of regulatory bodies?

In April this year, we saw a major development. The American regulator, the FDA, authorised a device that provides a definitive diagnosis without the need for a clinician to confirm it.

Such approval had never been granted before.

In this instance, it was an algorithm trained to detect diabetic retinopathy from eye scans.

I know something about the rigour of the processes in the United States, through the experience of going through them.

I founded a company called Axon Instruments in San Francisco.

We made several devices, one of which was a tool to help neurosurgeons operating on patients with Parkinson's disease.

It was designed to insert an electrode connected to a high voltage power supply deep into human brains. While the patient was awake.

Quality and safety were absolutely non-negotiable.

Every part of the design, the manufacturing and the device itself was scrutinised, before it got anywhere near a human being.

Medical device makers factor that time and expense into the costs of doing business.

It might surprise you... but software developers typically don't.

For them, quality is achieved not by rigorous front-end design, development and testing, but by constant iteration in the field.

They put the product into the world and wait for the users to show them the bugs.

So the medical community's Phase IV trial is the software developer's Phase I trial.

And for them, it works: they re-code the app, update it on ten million smartphones, the user experience gets better... and nobody gets hurt.

That's never going to be viable for medical-grade AI.

\*\*\*

These challenges explain why the adoption of AI in medicine is difficult.

But they are not a justification for rejecting it.

They underscore the need for everyone involved in healthcare to stop, every time they come to a decision point, and reflect.

Recently, I spoke at a conference organised by the Australian Human Rights Commission on human rights and technology.

I said that the ultimate question for all Australians when it comes to AI is easy to express but difficult to answer. The question is:

*What kind of society do we want to be?*

We could adapt the same critical question for everyone we trust with our health.

*What kind of doctor do I want to be?*

*What kind of care do I want to deliver?*

If we keep coming back to these same questions, then we will always put the wellbeing of our fellow humans first.

And then we can work to earn their trust.

\*\*\*

I've been reflecting on that word, "trust".

Like most people, I've never stopped to ask myself precisely *why* I trust in doctors.

I *do* trust in doctors.

That's not to say that I just switch off my brain. I still take responsibility for my decisions.

But I certainly don't leave the clinic, then race to my computer to trawl through several hundred peer-reviewed papers before I take the pills.

I'd waste my life looking after my health. Which would defeat the purpose of staying healthy. So I have to trust.

Why do I give *my* trust to doctors, in particular?

I'd say it comes down to three things.

First, I've had enough experience in education and research to trust in their qualifications.

Second, I trust in the rigour of Australian regulations.

And third, I can't separate my concept of a doctor from the Hippocratic Oath.

It's imprinted on my brain: *doctors have ethics*.

Could I think about building trust in artificial intelligence in the same way?

Yes, I could.

First, qualifications. I could expect that any company developing AI for use in medicine has been thoroughly vetted, so that both the company and their products are quality by design.

And just as I expect the clinics who employ doctors to check their CVs, yes, I absolutely expect the clinics who bring AI into medicine to scrutinise every one of the company claims.

Second, regulation. I could expect the medical profession to work closely with regulators to ensure that AI is harnessed the right way.

Regulations to ensure that when a decision affects a patient, there will be a human in the loop. That human will be accountable. That process will be documented and evaluated.

And I absolutely expect that clinics will talk to each other, keep improving their protocols, and build technology expertise at every level.

And then third, the Hippocratic Oath. We could come up with an equivalent for AI in medicine.

Now I wouldn't be the first person to propose a kind of Hippocratic Oath for artificial intelligence. Microsoft has floated the idea, along with many others.

But I had never read the original Hippocratic Oath, as composed somewhere in the third century B.C.

So I did. And the elements are fascinating.

You could almost write a best-selling sequel to the Da Vinci Code, called the Hippocratic Oath.

Here are the critical sentences.

- First, do no harm.
- Second, I will use treatment to help the sick according to my ability and judgment.
- And third, whatsoever I see or hear in the course of my profession, if it be what should not be published abroad, I will never divulge, holding such things to be holy secrets.

You see?

- First, AI will do no harm.
- Second, AI will only be used to do what the evidence confirms it can.
- And third, AI will never share our personal information without our permission.

We can all hold these three simple concepts in our heads, and even better – keep them always at the forefront of our minds.

And best of all – it would be the same basic compact between doctors and patients that has endured for thousands of years.

That compact has resonated across all cultures and generations.

Yes, to me, it still has meaning in the age of medical AI.

\*\*\*

And now let me turn to you in the audience – in particular, those of you in the medical profession.

For you these questions are real, and pressing.

If you haven't already confronted them, yes, you will be asked in the next five years, or 10 years, to integrate AI into your concept of care.

You will be asked to put your trust in AI.

And you will be asked to tell a patient that *they* can trust in AI.

I can't give you the answer – but I hope that you will think first of the question.

*What kind of doctor do I want to be?*

Then, I hope, you will remember and carry forward all the wisdom that doctors have earned through hundreds of years.

I know we can live better, and longer, with AI.

Managed by doctors.

You have the wisdom to bring it about.

And with humans like you, we can be a society that never forgets what it really means to care.

**THANK YOU**