



Australian Government

Chief Scientist

DR ALAN FINKEL AO

2018 Samuel Alexander Lecture

The next generation

Monday 13th August 2018

**Wesley College
MELBOURNE**

I have promised you a lecture in the tradition of Samuel Alexander.

On reflection, I might want to lower the bar – because Samuel Alexander excelled in pretty much everything.

His genius encompassed almost every conceivable branch of scholarship. And he was charming, handsome and popular to boot.

In short: he was an excellent advertisement for his old school, Wesley College.

But his greatest gift was undoubtedly his philosophy.

Samuel Alexander's phenomenal brain could Hoover up the strands from thousands of years of human thought – and knit them into a vision of stupendous breadth.

Language. Logic. The art of persuasion.

The very pinnacle of human intelligence.

And heights that only an exceptional human brain could ever reach.

Or at least, so we think.

The company IBM has other ideas.

It likes to pit humans against technology in tests of intelligence.

And it's been picking off, one by one, the skills that humans have only mastered over tens of thousands of years.

We said that robots couldn't play chess. That's strategy. That's logic. That's human.

IBM's Deep Blue beat grandmaster Garry Kasparov in 2004.

We said that robots couldn't interpret complex questions and sift through multiple lines of evidence to come up with a meaningful answer. That's language and reasoning. That's human.

IBM Watson beat the humans at the quiz show *Jeopardy!* in 2011.

What's next on the intelligence scale?

Only the most Wesleyan challenge of all – the challenge at which Samuel Alexander excelled: *debating*.

And here it is: IBM's Project Debater.¹

Now I watched that video you just saw, and I was greatly impressed, not only by the fact that IBM chose to make Project Debater look like a monolith from the movie 2001, but by the arguments Project Debater articulated.

But I'm a scientist: when it comes to debating, I'm not an expert.

I needed a qualified adjudicator.

And as good luck would have it, I've got one who usually takes my telephone calls. Eventually.

And that's my son, Victor, another graduate of Wesley College, and one half of the winning team at the 2011 World Debating Championships.

¹ See IBM Project Debater in action at <https://www.youtube.com/watch?v=0Y8NwmzKsbc>.

I can tell you, as his father, that this college put enormous effort into developing his teenage instinct to answer back. In a civilised way.

Victor was drilled in the discipline of debating: matter, manner, method; principle, practice, performance. Again, and again, and again.

He carried those skills to Monash University... to the World Championships... to Oxford... and then to his current position at McKinsey & Company.

So I sent Victor the link to the video of Project Debater.

He called me almost immediately. And he was rattled.

It's possible that he was thinking about the long-term career prospects for people who make their money producing reasoned and persuasive arguments...

... just as I was thinking about all the time I spent as a parent ferrying him to debating practice, so that he could have those long-term career prospects...

... but for both of us, I suspect the disquiet was more fundamental.

We had walked straight into what's known as the "Uncanny Valley": where technologies are human enough to be familiar... and yet robotic enough to be creepy.

Now Victor and I are separated by a gap of about thirty years. But we have something in common: we're immigrants to the Uncanny Valley.

We both grew up in the period called the AI Winter. AI for Artificial Intelligence.

It was called winter because progress in the technologies – at least on the surface – appeared to be frozen.

We had the name "artificial intelligence": it was coined at an academic conference in 1956.

What did it mean?

Well, think about an ATM. That's automation – but it's not what we would consider to be intelligence.

An ATM is programmed to respond in a specific way to a specific instruction.

Intelligence, on the other hand, means something more akin to the human brain.

How do humans reason, and learn?

- 1) We're born with innate capability: a brain.
- 2) We take in content: knowledge.
- 3) And we have teachers: training.

With those three things – brain, knowledge, training – we find patterns and come up with rules that help us to understand how the world works, and make predictions.

AI can be broadly understood the same way.

It's basically a combination of the same three things.

- 1) A machine learning algorithm.
- 2) A large quantity of high-quality data.

3) And a training procedure.

That's how AI can be said to "learn" and "think" – or perform tasks that would otherwise require a human.

During the AI winter, no one could look at a computer and see anything even remotely comparable to a human brain.

When I was a child, the pride of Australian computing was a machine called CSIRAC. A child today would probably struggle to recognise that it was, in fact, a computer.

It filled a room. It had no display screen. It would spit out long strips of paper with punched holes that had to be fed into *another* machine to be translated to text.

But very few children like me ever saw that computer, or *any* computer, because they were incredibly expensive and extremely rare.

By the time my sons Victor and Alex arrived, computers were becoming standard features in family homes.

And arguments about computers were *also* becoming standard features in family homes.

Now, for the first time, the everyday consumer could actually see from year to year that computers were getting better.

They were still a long way from closing the gap.

And so the AI Winter dragged on.

It's been clear for some time now that the ice was beginning to crack.

And Project Debater is just one of the green shoots showing the whole world: spring has arrived.

Let's think about how that might look and feel from the perspective of a child born today.

Say that you're the son or daughter of fairly typical millennials. Your parents have a home assistant, perhaps Amazon's Alexa. It was there when you were brought home from the hospital.

So you grow up in a world where it's perfectly normal for adults to talk to machines. And it's perfectly normal for the machines to talk back.

As you grow, you join in. At your command, Alexa can play songs, or practice spelling, or suggest new games. You might consider her to be your friend.

But perhaps your parents like to think of themselves as early adopters: ahead of the curve. They might be interested in a product such as the Aristotle, announced by Mattel last year.

Here is Aristotle's mission, as explained in Mattel's media release:

"Raising kids can be hectic. Our goal is to provide parents with a platform that simplifies parenting, while helping them to nurture, teach, and protect their young ones."

It begins with Aristotle in Baby mode.

What could be more hectic than a crying baby?

Aristotle has it sorted.

Quote: *“Aristotle can automatically play a lullaby and turn lights on dim when it hears a sleeping baby begin to cry.”*

And more: *“E-commerce functionality tied directly to key retail partners to automatically reorder baby consumables.” “Hands-free way for parents to keep track and monitor sleeping, changing and feeding developments.”*

Then, when Aristotle determines the child is ready, we progress to Aristotle Toddler.

“Uses audio, visual and tactile learning methods for ABCs, 123s, first words, sing-alongs and story-time.”

And then to Aristotle Kid: *“Homework helper, entertainment unit, and playmate.”*

Now if you've got an Aristotle at home, your parents are probably attentive to the AI enhancing your learning at school.

They might be excited by the possibilities of facial recognition in the classroom.

Consider a system being trialled in China, and widely reported in Western media.

It's called the “Intelligent Classroom Behaviour Management System”.

According to the reports, there's a camera mounted on a blackboard behind the teacher. It scans the room every thirty seconds logging facial expressions and behaviour, such as reading, writing, raising a hand, or leaning on the desk.

And if this sounds far-fetched, it's worth noting that similar platforms are already marketed in Australia.

They're used in industry to monitor the concentration levels of people operating heavy equipment, or to make sure people are actually paying attention in corporate training.

In the trial site in China, they seem to be making a difference.

A quote from a student: *“Previously when I had classes that I didn't like very much, I would be lazy and maybe take a nap on the desk or flick through other textbooks. But I don't dare be distracted since the cameras were installed in the classrooms. It's like a pair of mystery eyes are constantly watching me”.*

The dawn of the AI Spring, indeed.

But take note. As in all things technology, the promises made by the developers of these platforms should never be taken at face value.

We will continue to see stories in the media about the bugs and limitations.

Last month it was Amazon's facial recognition technology. A civil liberties group in the United States scanned the images of 535 members of Congress looking for matches against a database of known criminals. It generated 28 matches.

And we all know the frustration of talking to voice-activated assistants that seem incapable of recognising Australian as an acceptable form of English.

What matters is the trend.

We will see more and more products from more and more companies, relentlessly combing through every aspect of being a parent that strikes us as messy, or dull, or simply inconvenient... and presenting their answer.

The question for all of us is how to respond.

There are many things about the concept of automated child-raising that makes people cringe – but we can distil it down to two.

Worry Number One, the fear that children soothed to sleep by AI lullabies will be incapable of meaningful relationships with human beings.

If *you* grew up with an AI companion who was always happy, always helpful, always obliging... then you might get very good at giving orders.

And you might be very intolerant of the human needs and foibles of the children who might otherwise be your friends.

That's Worry Number One.

Worry Number Two, the fear that children will be trapped in a web of surveillance, never free to be simply... unobserved.

Imagine an unblinking eye, trained on you from birth.

No private space.

No escape from your mistakes.

No leeway to test the rules.

It's a terrifying thought to those of us who have lived long enough to cherish the moments when it's OK to just be yourself: less than perfect.

Together, these two worries lead some to the conclusion that it's simply impossible to raise a good and happy human in the year 2018.

But I refuse to accept that view. And I refuse to pass that fatalism to my children.

The picture that I've painted thus far is the negative – or at the very least, the concerning.

For every troubling application of AI, we could just as easily find the same technology harnessed for good.

Let's go back to AI in the classroom in a fresh light.

We know that in every cohort of students, there'll be students who race ahead and students who fall behind. But we know it's important to keep children together with other children their own age.

So we respond by teaching to the middle.

And teachers do their best to give the faster learners an extra challenge, and the slower learners extra support.

That's an enormous thing to ask of a single teacher.

AI can help to shoulder that task. In the hands of a gifted teacher, it can tailor the level of the challenge to the learning needs of the individual.

And that's the way it's being used in Australian schools: not as a substitute teacher, as a teaching tool.

Or consider the needs of students with disabilities.

AI can convert spoken words into text.

AI can help children with autism to recognise facial expressions.

AI can tell a child with vision impairment that his mother has entered the room, and she's smiling.

Would we throw the opportunities for all those children away?

And would we throw away the even greater promise of AI in all our lives?

The faster and more reliable medical tests?

The robots that get you in and out of surgery, same day?

The self-driving cars that I, for one, see as a brilliant alternative to self-driving teenagers?

And in science, the fact that we can do things in days that just a decade ago would be the work of an entire PhD?

The task as I see it is not to resist artificial intelligence, but to work out how humans and technologies can play nice and get along, always prioritising the human.

And I have hope.

I have hope because I look to our history.

How many times have we been told that a new technology would destroy human relationships... only to find that it simply connects us to humans in hitherto unimagined ways?

I fell into that trap about a decade ago, when I was asked to imagine the future of the university.

No doubt about it, I said: bricks and mortar campuses were old school. The future was all online: about ten global mega-universities, on the scale of Google or Facebook.

Then we saw the reality of online classes: the poor completion rates, the disengagement of students, and the hunger for exactly what universities have been doing for hundreds of years.

Today those bricks and mortar universities are thriving, alongside the online courses.

And why are they thriving?

Why do we keep going to gyms when we can watch fitness videos on YouTube?

Why are sales of old-fashioned board games and children's books booming?

Why do we pay for barista-made coffee, in a café, when we can have a similar precisely engineered Nespresso equivalent at home for about a fifth of the price?

It comes down to the same basic fact about human beings: we thrive on human connection.

And we've always been capable of taking the benefits of a technology without seeing those benefits as in any way a substitute for actual human beings.

I look at the young people I know: that hasn't changed.

Along with our history, I look to the present.

And I see humans encountering new technologies and actually stopping to think about what's right.

Take the home assistant Alexa. Parents complained to Amazon that she was turning their children into horrible little tyrants.

Now you can program your Alexa to thank a child for saying "please" and "thank you".

A very human response to a human concern.

Then there's Aristotle. 17,000 people signed a petition asking Mattel to cancel it. The Aristotle was promptly shelved.

We have choices. We decide.

And that includes you.

Last month, when I delivered the keynote address at a conference on human rights and technology organised by the Human Rights Commission, I said that the question for all Australians boiled down to this:

What kind of society do we want to be?

We need to answer that question at the level of public policy – but even more urgently, we need to answer it in the classroom and in the home.

The future we end up with is ultimately the sum total of all our choices.

So the question should be at the forefront of our minds, every time we come to a decision point.

What kind of parent do I want to be?

What kind of teacher do I want to be?

We can give Alexa the job of purchasing nappies – but we cannot delegate our individual duty to ask and answer these questions.

In my mind, that means instilling in every child the fundamental principle that humans are special, to be respected and cherished.

And that will be true no matter what human-like abilities we might have to share with robots.

I've been reflecting in recent months on how we might teach that principle, on a deeper level than a prompt from Alexa to say "thank you".

And I know from my own experience: children may or may not listen to our words, but they do pay attention to our behaviour.

So they will be guided by the way they see adults bringing artificial intelligence into the classroom.

We should start with the motto that every school principal should be forced to write out on an old-fashioned blackboard at least a thousand times:

Technology for technology sake's is far worse than no technology at all.

Either we don't use it, and it gets discarded.

Or we do use it, and it becomes a distraction.

I look around this room, and I see people of different generations.

But irrespective of the year you graduated, I can promise you that when you look back to your school days, you remember not a technology, but a *teacher*.

A human.

A person who held your attention not with a screen, but with the thrill of discovering that yes, I can learn. Yes, I can get better if I practice. Yes, I have a brain and I'm going to use it.

You fell in love with learning because a teacher believed you could.

I think back to the times as a parent I observed the Wesley Big Band practicing and performing, when the musical director Mr Peter Foley oozed so much musicality, enthusiasm and joy that it infused into every student in the band and they exceeded their own dreams.

I can promise you, human beings have not evolved since you went to school.

So let's not be dazzled by the claims of AI developers. Let's show our children how responsible humans assess those claims, and make intelligent choices.

But Alan, you say, I'm just an adult. What do I know about technology?

I'm here to tell you: don't be dissuaded.

As parents, we have to grapple with things beyond our expertise all the time.

Think about all the times when your children were sick with mystery illnesses.

You took them to doctors, and then you had to decide what to do with the doctors' advice.

At those moments, you couldn't plead ignorance: life doesn't care if you failed biology in Year 10.

You still had to make a decision and take responsibility.

And you knew enough about doctors to remember the Hippocratic Oath, penned by Hippocrates in the third century BC.

That Oath told you what it means for a doctor to be worthy of your trust.

In brief and colloquialised, the first two of the three oaths are:

- I will do no harm.
- I will use treatment to help the sick according to my ability and judgement.

The third I will read out in full:

- 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.

In the modern context this basically says: Never share personal information without the patient's explicit permission.

You can think in exactly the same way about the standards you should expect for AI in the classroom.

First, do no harm. If educational AI is undermining the relationship between the teacher and the student it is doing harm.

Second, ensure that all decisions are for the good of the student. The Secretary of the Federal Department of Home Affairs, Michael Pezzullo, recently proposed a Golden Rule that says that no one should be deprived of their fundamental rights, privileges or entitlements by a computer rather than an accountable human being.

Apply the same rule in schools: when a decision affects a student, we need to insist that there will always be a human in the loop.

And third, educational AI should always respect the privacy of the student so that the relationship between the students and teacher is a trusted one.

One, two, three: you can remember those three ancient standards.

And we can all do something else as well, every one of us.

As a community, we can keep investing all our energies in human intelligence: in education.

Because nothing says that humans are special more clearly than the effort we put into the development of human talents, human passions and human skills.

There are people who will tell you that technology is a substitute for education.

If you've got a calculator, you don't need maths!

If you've got audiobooks, you don't need to read!

If you've got the internet, you don't need to teach any content, whatsoever, to little humans!

Everyone here today, your duty is to disregard the people who say those things.

For all time, education has been the path out of poverty and dependence.

Parents trapped by their circumstances have found in education their child's only route to a better life.

And they have fought for that opportunity.

That has not changed.

Technology is not a solution for ignorance.

It makes the opportunities for the educated even greater.

And it demands that all of us – and remember, we're the humans, with the choices – all of us need to have the knowledge in our heads to make those choices wisely.

So the answer is not to lower the bar in education. It's to raise it, and help all our children to clear it.

Aristotle not required.

And I am sure Samuel Alexander would agree.

And if you ever lose hope in the capacity of human beings to rescue themselves from all the problems we create...

Just remember what it's like to watch a newborn baby learn.

Human intelligence: the original, and still the best.

What kind of society do we want to be?

One that takes responsibility for nurturing that fledgling human intelligence.

THANK YOU