



Australian Government

Chief Scientist

DR ALAN FINKEL AO

**Go8 Artificial Intelligence
Collaboration and Commercialisation Summit**

Opening Keynote

Thursday 31 October 2019

Monash University

Melbourne

In 1770, a Hungarian inventor named Wolfgang von Kempelen unveiled his latest creation to the Imperial Court in Vienna.

Was it an animated carnival mask? A gravity-fed water purifying machine?

No. It was a chess-playing machine that beat its human opponents with confident ease.

Dubbed 'The Turk', the machine consisted of a life-sized, human-like figure, dressed in robes and a turban, seated at a wooden cabinet that was overlaid with a chessboard.

von Kempelen made a great show of displaying the machine's inner workings. He would open the cabinet doors to reveal a whirling clockwork of densely packed wheels, cogs, and levers.

'The Turk' would then be wound up and, as described by Edgar Allan Poe, begin to "roll its eyes, as if surveying the board, move its head, and pronounce the word 'check' when necessary"...and the phrase 'checkmate' with glee.

The automaton became a global sensation, drawing huge crowds at exhibitions, and defeating human challengers such as Benjamin Franklin, Napoleon Bonaparte, Frederick the Great, and the Emperor and Empress of Russia.

The age of Artificial Intelligence had apparently begun.

But, alas, the Turk turned out to be a hoax – a sequential roster of human chess masters would hide inside the cabinet, controlling The Turk's movement through a clever arrangement of magnets and strings to make it appear as if the 'machine' was outsmarting humans.

Fast forward to 2005, when Amazon borrowed the concept with the launch of the Amazon Mechanical Turk – an online marketplace where, like the chess masters of the 18th century, people, hidden from view, can be hired by companies to perform discrete tasks that computers are currently unable to do, such as identifying specific content in a video.

As they go about their tasks, the actions and decisions of these online workers are providing the world's biggest tech companies with high-quality data that is then used to train computer systems to better

recognise patterns, creating ever-more accurate algorithms, until, eventually, there will no longer be a need for human intelligence – eliminated one microtask at a time.

As A.I. becomes more and more powerful, I find myself looking for areas where we mere mortals, otherwise known as humans, have the upper hand.

It's a diminishing pool.

It took more than 200 years after 'The Turk' but, in 1997, IBM created Deep Blue, the first supercomputer to defeat a reigning world chess champion ... without the hidden human.

But at the time we humans could take solace, because Deep Blue relied on "brute force" to achieve victory, rather than analysing gameplay and visualising the possible moves.

Surely, I thought, we humans would have the upper hand in poker, a game requiring human intuition and bluffing. Most commentators believed it would, therefore, be a holdout. But that border was breached in 2017.

Then surely human intelligence would have the upper hand in the Chinese game of 'Go', which has trillions more potential moves than chess. Another breach, this time in 2016, when a program named AlphaGo, developed by UK company DeepMind, beat the world's best player.

This was *truly* artificial intelligence, where AlphaGo learned from hundreds of thousands of games played between humans until, ultimately, it worked out how to master the game.

Then, to add insult to injury, a year later an improved version was produced – AlphaZero. So smart that it didn't even bother to look at human games. Instead, AlphaZero was simply given the rules, then played hundreds of thousands of games against itself, starting as a complete novice but getting better and better every second.

36 hours after it was switched on, AlphaZero defeated its predecessor and became the 'Go' world champion.

For many years I thought that recognising faces would be the mark of our superiority. In fact, like other pundits, I used to explain to anyone who would listen why it would be so difficult for A.I. to beat us on facial recognition. But sure enough, while I was still expressing my confidence in our superiority, the threshold was reached where A.I. could recognise faces more effectively than we humans.

So, what is the next human capability that is uniquely ours and beyond the reach of A.I.? Is it art? Not really, there are programs that paint original paintings in the style of Rembrandt.

“I know”, I said to myself, “it has to be speech writing”. So I did a Google search and found that, so far, it’s not happening.

So I put it to you that this is the next frontier. We can be proud to be human because we remain solely capable of stringing thoughts together for a speech.

And it is our thoughts, our unique human ability to meditate on the known and unknown, that will be critical as we delve into the challenge of ensuring that our zeal for innovation never betrays our values.

Science often moves faster than our ability to fully grasp all of its implications, leaving a trail of moral and ethical dilemmas in its wake.

As the genius of A.I. pushes the boundaries of what we can do, we are faced with increasingly complex questions about what we *should* do.

Answering these questions requires the application of ethics rather than physics. As such, it is not the province solely of scientists, but of every individual.

That is why today’s Summit is so important.

Each of us here is not simply sharing in a one-off event. We are sharing in an ongoing effort to harness the power of scientific progress for the benefit of our society, while safeguarding the ideals of our society.

The thoughts exchanged here today will go a long way to ensure that A.I. is the servant of our needs instead of the other way around.

So let me share with you some of *my* thoughts.

I believe we must pursue the tremendous possibilities of A.I., and I believe we can do so while still fostering our commitment to human values, to the good of society, and to our basic sense of right and wrong.

My belief stems from the fundamental tenets and ideals of Australia itself. It is shaped by our history, by our proven capacity to adapt to rapid changes, and by the egalitarian nature of our society.

There is a question often put to me: is Australia likely to be a leader in developing A.I. or just a follower who imports A.I.?

I believe this to be a false dichotomy.

We *are* capable technology innovators, but we have always imported more technology than we develop. That's inevitable, given our size.

However, that does not mean we have to accept a future dictated by overseas companies.

To the contrary, with smart, strategic applications we can find niches where we can excel and define our own future.

Indeed, the latest Australian Research Council review of university research performance found 11 of our universities are currently performing at world standard in the field of artificial intelligence, 11 above world standard, and 7 out of 40 well above world standard, up from only 1 university at that ranking in 2015.

This is a wonderful achievement, and testament to the strength and capacity of our university system.

But when stated by themselves, statistics can detract from the human element behind the numbers; the promise of what this can mean for people's lives.

I firmly believe that the unmatched opportunities for A.I. will only be assured in this country if it is developed with an eye to demonstrating clear benefits to individual Australians.

Focussing on what CSIRO Chief Executive Officer Dr Larry Marshall calls *A.I. for a purpose*.

Through this lens, I am personally interested in looking at fields as specific as

A.I. for medical diagnostics

A.I. for agriculture

A.I. for financial services.

Which is why, as Australia's Chief Scientist, I am currently managing work of the National Science and Technology Council on "A.I. for Manufacturing".

In a field that has always been at the forefront of progress, there is enormous potential for A.I. to shape the future of manufacturing – both the scope of what manufacturers *can* create and *how* they create it.

But for A.I.'s opportunities to be fully realised, Australian businesses and their workers will need to be adequately prepared and equipped to embrace its benefits.

And so, we look to you, our researchers and academics – the experts in A.I. development, implementation and adoption – to cultivate the necessary skillsets. Across the breadth of our universities and ultimately across the breadth of our society.

This important role of universities was very prominent last month, when M.I.T in the United States launched a brand new college for A.I.

The goal of the Schwarzman College is to "educate the bilinguals of the future". The term 'bilinguals' describes the future graduates in chemistry, politics and history who will also be skilled in the relevant techniques of modern computing; further empowering them in their discipline. A.I. everywhere, just like statistics.

It is imperative we explore this concept in Australia.

By integrating A.I. into the broader fabric of our university curricula, we can generate advances of unlimited potential in all fields, building the workforce and industries of the future.

I am counting on you to be the leaders in turning this vision into a reality and furthering the goals, aspirations, and moral principles of our society.

You are held to this benchmark precisely because you have *always* been at the forefront of our nation's proud record of upholding the

highest standards of ethics while expanding the limits of science and knowledge.

I think of IVF, which started right here, as a combined research project between Monash University and the University of Melbourne.

Building on the work of Professors Alan Trounson and Carl Wood, the first IVF baby, Louise Brown, was born in the U.K. in 1978.

Australia's first, and the world's third IVF birth, took place in 1980 here in Melbourne under the supervision of a Monash University team, as did the world's next nine IVF babies.

But think for a moment about the torment of raw emotions that early prospective mothers experienced with this procedure.

The conflicting anguish and hope.

The gnawing fear that their IVF baby might be in some way abnormal at birth, or at age 5, or 15, or 30.

Fear, magnified in their minds by the overwhelming ethical and religious debates raging across society at the time.

In most circumstances, these negative concerns would have prevented this new technology from ever being introduced.

But they were outweighed by one powerful incentive: IVF's precious gift of matchless value.

Today, there are more than 8 million babies born from IVF. They are living proof not only of the wonders of modern science, but of our ability to keep our ingenuity rooted in our values.

These babies are not just statistics – they are human beings who brought mothers, fathers, grandparents and extended families the joy of bringing a baby into the world.

They are individuals who will experience birthdays, graduations, weddings, children of their own, and who, just decades ago, would not have had a chance at life.

It was precisely because of this extraordinary and visible benefit to individuals that we were able to work our way through IVF's novel challenges.

And it is important to remember and acknowledge just how critical our university and research sectors were in solving these challenges and in shaping the multi-disciplinary framework behind IVF.

In 1982, the sizeable and extremely sensitive task of designing pioneering laws to govern IVF treatments was given to Monash University's own Professor Louis Waller, and I was saddened to hear of his recent passing.

The Waller Committee report, which carefully considered the social, ethical and legal issues arising from IVF, directly led to Victoria becoming the first state in Australia, and the first government in the world, to regulate the practice of IVF.

And establish the world's first central IVF register.

And so, IVF became an accepted, mainstream procedure.

Fast forward to the present, and it will not surprise you that artificial intelligence is contributing to improved outcomes.

In a conventional IVF procedure, embryos are assessed by human beings – otherwise known as doctors – to choose which embryos to implant to maximise the likelihood of a successful pregnancy.

A.I. is now helping to make that choice more reliable.

At the forefront, is Australian company Life Whisperer Diagnostics, which emerged from the University of Adelaide.

Its A.I. diagnostics product identifies the best embryos for implantation, with the goal of reducing multiple births and improving the pregnancy success rate.

This a perfect example of how we can utilise the brilliance of A.I. to serve human needs.

And yet, while A.I. shows us how it can be of immense service to humanity, it cannot show us how to prevent its immoral use.

That's up to us. And it requires constant vigilance.

Just this month, we learned that Google has obtained a patent to use an array of sensors and cameras to monitor home activity, with the capacity to work out the title of the book you're reading in bed.

To put this into context, what if I proposed a complete stranger coming to your door and offering you unlimited free furniture and non-stick frypans, in exchange for allowing them to camp out in your bedroom for the next two weeks and observe your, and your family's, behaviour? Would you agree?

Of course not.

We are repulsed by this prospect not because of its unfamiliarity, but because we innately feel that it violates fundamental principles we rightfully hold dear.

But Google wants to do that – not for two weeks but potentially for the rest of your life.

The idea of treating humans as objects, as data, to be studied and manipulated, rather than as cherished individuals entitled to inherent worth and dignity, stirs our deepest convictions.

It crosses a moral boundary that needlessly encourages a conflict between science and ethics, which can only damage both our scientific endeavours and our nation as a whole.

No matter how fast the pace of A.I. innovation, it must never surpass the primacy of human rights.

Much will be lost if we discard our moral compass in the name of progress.

And yet, if approached correctly, this challenge can also be a golden opportunity for Australia.

We can define our own future by being world-leaders in the field of A.I. ethics and human rights.

Showing the world how to advance the cause of scientific discovery while staying true to the ideals of a prudent and virtuous society.

Like we did for IVF.

To that end, in April this year, CSIRO's Data61 and the Commonwealth Department of Industry, Innovation and Science released a discussion paper to inform the development of an Australian A.I. Ethics Framework.

In view of the emerging technological realities of A.I., the Framework aims to formulate new protections to build public trust, as well as help guide businesses and governments to responsibly develop and use A.I. systems.

At the same time, the Human Rights Commission, under the leadership of Ed Santow, is deep diving into the difficult issue of human rights and digital technology, and I am proud to be on the advisory committee.

Of course, the Human Rights Commission, the Government and the Australian community need to hear from universities about how to use A.I. for the benefit of all Australians.

Not just from the computer science department, but also from our academic leaders in ethics, philosophy, law and business.

As a reservoir of ideas, and a touchstone of our morality, input from across our universities will be crucial as we navigate the uncharted waters of promoting A.I.'s promise, while safeguarding against its potential perils.

Working together, to help steer A.I. towards preserving and enhancing the quality of our lives, and the vigour of our ideals.

As we go forward, I hope we will always be guided by our capabilities, our conscience...and our collective human thoughts.

May the Force be with you.

Thank you.