



Australian Government

Chief Scientist

DR ALAN FINKEL AO

CREATIVE INNOVATION 2017

Human Intelligence 2.0: How will we get there?

Tuesday 14th November 2017

**Sofitel Melbourne on Collins
MELBOURNE**

Every second, about four human babies are born.

Picture just one of them.

Her brain began developing just 16 days after conception.

By birth, she had 100 billion neurons.

By the age of three, her brain will contain a thousand trillion connections.

And as she lays down those neural pathways, her parents will watch as that little baby *learns*. To gurgle. To smile. To crawl. To walk. To speak.

And then to read, to study, to question, to strive, to create, to laugh, to love, to dream.

Perhaps to be a mother and raise another human. Perhaps to design a brilliant machine...

All of that potential is bound up in that little baby.

Human Intelligence, 1.0: the original, and still the best.

Now, what's the first thing you do when you meet a new baby? Of course: whip out the iPhone and take a picture!

How many iPhones does Apple sell in a second?

Nine.

That's right: we're making new Siris at more than double the rate we're making new humans.

And Siri has siblings.

Replika, your chatbot friend, always available to talk.

Watson, from IBM, who diagnoses tumours, and talent scouts for the professional basketball leagues.

Cortana and Alexa, your personal assistants.

And countless more. We call them AIs – artificial intelligence. Like it or not, they're part of our world. From the time you snap that picture, they're going to grow up with our kids.

So they're going to have to play nice... and get along.

Now we know how to teach a child to be a civilised adult. At least in theory. We've been doing it for thousands of years.

And in all that time, the basic coding hasn't changed. Gen A is Gen Z.

But how do you develop a civilised AI – a technology that's kind to humans?

The basic coding changes, fast. The capabilities evolve incredibly quickly. And let's be honest: it's all become a little... unnerving.

Just look at the newspaper headlines!

Our AI offspring have been taking human jobs.

They've been helping the police to work out who might commit crimes in the future.

They've been implicated in more than a handful of botch-ups.

They've absorbed our biases and discriminated against vulnerable people.

They've got some of us imagining the end of the world.

And just think, right at this moment, *they could be watching your kids....* via that AI baby monitor recommended by your AI home assistant.

Okay, then: just BAN them. No more artificial intelligence! Children safe, problem solved.

You might as well ban electricity. Or the internet. Or cars. AI is that fundamental.

And AI isn't fundamental because governments imposed it, or companies wanted to sell it: it's fundamental because it *works*. It makes life better.

Hands up all the baby boomers.

All our lives, we've been told that AI was coming: just give it a decade. Make that two decades. OK, three decades.

AI was in the doldrums for fifty years, till the turn of the millennium.

All the while, our computers were evolving, progress driving progress, faster and faster every year.

Take computer memory. It's now ten trillion times cheaper than it was the day I was born.

The author Max Tegmark puts it like this. If the price of gold had fallen at the same rate, you could buy all the gold that's ever been mined for around a dollar.

That stunning progress paved the way to the AI tools we know today: tools that make it cheaper and safer to fly, tools that make it easy to take a decent photo, tools that make every one of us more productive, more perceptive and more ambitious.

And tools that make sure you never, ever forget your wedding anniversary.

That's today – think of tomorrow.

We could understand the human brain.

We could answer the fundamental questions about the origins of the universe.

We could rewire the globe for a zero-emissions world.

It would be a tragedy to rip that future away.

Alright then, let's keep the AI systems, but make them play nice... and get along.

Regulate them – like we do medicines and cars!

Good thought – but there's a fundamental difference.

It takes time to manufacture medicines and cars: all the way from research and development, through prototypes and pilots, to full-scale production.

The production is typically centralised. It's done at scale.

Even if it's done very cheaply, it produces a finite quantity of physical things, that have to be shipped and stored.

To upgrade the product, you need to remodel the production line. It all takes planning, money and time. Years.

But what if your product is an algorithm?

For near-zero cost, you can have near-infinite production. You can iterate and upgrade just by tweaking the code. And you can reach a mass audience almost instantaneously.

Twenty years ago, when AI was in the doldrums, we could afford to sit back, watch and wait. But we're not in the doldrums any more.

The organisers of this conference have challenged us to make the mental leap from Human Intelligence 1.0 to Human Intelligence 2.0.

I want to suggest to you that we are crossing a second threshold in lockstep.

The trickle of AI technologies has become an onrushing tide.

We are passing from Artificial Intelligence 1.0 to 2.0.

And we are the generation to take the step.

We are the parents of Humanity 2.0 and we are the parents of AI 2.0.

We brought them both into the world.

Siblings.

And as parents, we have to take responsibility for their behaviour.

We have to raise our biological offspring and our technological offspring to play nice... and get along.

To be kind to each other, and to cherish the world. To follow in our footsteps, and to reach beyond our dreams.

Can we do it? I believe we can.

Remember, we have two things on our side – bound up in that hundred billion neurons in every new baby's head.

One, the will: nothing is more important than making AI safe and beneficial for humans.

And two, the capacity: we are curious, we are capable and we are creative.

Think about the way we train our children.

Sometimes, the best approach is to praise, and encourage with incentives.

Sometimes, to impose penalties: consequences.

Sometimes, we create safe environments for the child to try, to fail and then to get up and try again.

Sometimes, we make certain things completely off-limits: knives and guns.

Sometimes, the child learns best through formal lessons.

Sometimes, through books and stories.

And sometimes, from their peers.

By all these methods, our children grow into adults bound by manners, by morals and by laws.

I expect we will see something similar in AI: a spectrum of responses that gets more nuanced and effective over time.

On the right-hand extreme, the equivalent of knives and guns: things that we agree as a global community are simply so dangerous that they need to be managed by international treaties.

In that category we might put autonomous weaponised drones, that can select and destroy without any human decision-maker in the loop beyond establishing the rules of engagement.

Two years ago I published an article in Cosmos Magazine calling for a global accord.

In the same year, more than 3000 AI and robotics researchers signed an open letter urging the leaders of the world to take action to prevent a global arms race.

On the other end of the spectrum are tools in everyday use, such as social media platforms and smartphones. Difficult to regulate by UN Convention.

Here we might look instead for measures that empower us as consumers and citizens to make responsible choices.

Some might say that AI systems are too complicated for the average person to understand. If we don't know how algorithms work, how can we interrogate their decisions, or judge if they're fair?

But other things are complicated and important to consumers – like picking products that are manufactured in ethical ways. What do we look for? The Fairtrade stamp.

We don't have to understand the complexities of international labour law, or global markets: we just have to be aware of the issue and make a choice.

I can imagine an ethical AI stamp – call it the Asimov, in honour of the Isaac Asimov who gave us the three Robot Laws.

Imagine if companies and research labs could use the Asimov stamp if they agreed to certain ethical development standards.

And then governments and consumers could use their purchasing power to only deal with Asimov-compliant providers.

Obviously, in the case of lethal drones, it would be a woefully inadequate response.

It's got to be part of a continuum.

Between the two poles – global treaty and consumer choice – we have room to manoeuvre, through regulation. And industry has a real incentive to engage.

Think of self-driving cars.

They're coming: and we want them. Our roads will be safer, our commute will be quicker, and our wallets will be fatter with all the money we'll save on parking. Bring it on.

But do it strategically.

Today our road rules are built on the assumption of human drivers. The human behind the wheel is the human in charge.

So obviously, the rules have to change.

The rules have to evolve for *all* AI.

And they have to be enforced.

In Humanity 2.0, there are consequences for breaking the rules.

Whether it's unruly behaviour in a classroom or robbery in our streets, transgressions are met with appropriate penalties.

So too in AI 2.0. We need rules, and consequences for breaking them.

You might be surprised, but the developers and suppliers of AI 2.0 will welcome the rules.

Put yourself in the position of a CEO of an AI 2.0 company.

You don't want a total ban. You don't want a free-for-all.

You want a forward-looking regulatory regime, with clear expectations, that gives consumers confidence that you can roll out the technology safely.

Then you can go to investors with a path to market.

That means you need to engage.

You need to engage with your consumers to understand their concerns. And you need to engage with regulators to advocate for standards that lock unprincipled providers out of the market, and reward quality companies, like yours.

Let's build those quality companies in Australia.

Let's get ahead of the world and develop regulatory systems that make us a destination of choice for ethical AI developers.

Regulatory systems that make sure Humanity 2.0 and AI 2.0 play nice... and get along.

And whatever we do, let's be *ambitious*.

The time is ripe for ideas. And today, in my home town Melbourne, the platform is yours.

So be bold. Be creative. Be curious.

Play nice... and get along.

And keep in the front of your mind that hundred billion neurons.

Resolve right now: *the future is going to be great*.

Now use this conference to work out how.

THANK YOU