



Australian Government

Chief Scientist

DR ALAN FINKEL AO

**IEEE Sections Congress 2017: Brilliant Minds, Bright
Futures**

Keynote Address

Better than Brilliant

Saturday 12th August 2017

**International Convention Centre
SYDNEY**

When I saw the topic for this event, “Brilliant minds, bright futures”, I grew excited.

Hooray, I thought: it’s a neuroscience convention!

Well, it’s not a neuroscience convention.

If it’s possible, it’s even better: the global congress of the IEEE.

I’m proud to say that I’m a Fellow — a card-carrying member. And I’m a member because I believe in the incredible power of the human mind.

I believe it because I’ve seen it up close.

During my PhD in electrical engineering, I was fortunate to be drawn into researching the electrical activity of the brain.

Thus began my second career, as a card-carrying neuroscientist.

I was captivated by the meticulous drawings of Ramón y Cajal, the father of modern neuroscience, dating from the 1890s.

To me, he was the Leonardo da Vinci of his time: a scientist and an artist combined.

He opened my eyes to a web of such fabulous complexity that I knew I could happily spend a lifetime trying to grasp it.

In my day, we were told that we would never be able to use a microscope to observe any features smaller than 200 nanometres.

We were working to the old Abbe diffraction limit — which, at the time, had the stature of a law.

Even with a perfect lens, Ernst Abbe said, there was a fundamental limit to seeing small objects. You could never see anything smaller than half the wavelength of the illuminating light.

So we could see whole brain cells, which are several thousand nanometres wide, but not the tiny neurotransmitter capsules that fill the synapses.

We were also told that replicating human intelligence was an impossible dream.

Research funders the world over were absolutely convinced that there was no way a machine could reason, or learn, or create, like a human. Never ever, no way.

In my lifetime, I’ve seen the impossible become the everyday.

We’ve smashed the Abbe diffraction limit by much more than a factor of ten.

We’ve literally lit up our image of the brain.

And when we lit up the brain we awoke to the possibility of reverse engineering intelligence itself.

There are very few examples of genuine biomimicry: that is, technologies developed by working backwards from phenomena in the natural world.

Our aeroplanes don’t fly like birds, even if birds showed us that flying was possible.

Our submarines don't flap their tails like dolphins, even if dolphins were the proof that high speed propulsion under water could be achieved.

But unlike birds and dolphins, the human brain has become an incredible well of ideas: literally, a brains trust for humanity, with a wealth of ideas to pinch!

Our brains are neural networks, with multiple memory systems split across short-term and long-term storage. Our every thought depends on negative feedback. We have a propensity to continual learning.

All these human things are inspiring geniuses such as yourselves to develop a new generation of thinking, learning, intelligent machines.

A revolution, in a blink of history. The dawn of the artificial intelligence age.

And perhaps the twilight of humanity as we know it.

It's all there, in our brilliant minds.

Which brings me to the speech I promised you in the program: AI, automation and jobs.

By now, you probably think you've heard that speech before.

First, they came for the factory workers... the manual jobs.

Then, they came for the secretaries and the bank tellers... the process jobs.

Now, they're coming for the architects, the lawyers, the doctors... and yes, the engineers... the cognitive jobs, and the creative jobs.

The AI threat to jobs is now impossible to confine. No job, however prestigious, however complex, is safe.

So what next?

As an optimist, I'd tell you that there will be new jobs in place of the old.

As a pessimist, I'd encourage you to prepare for life on Mars.

As a transhumanist, I'd suggest some upgrades you might want to consider installing in your brain. If you can't beat the robots, be a robot.

As an economist... no, I'll let the economists explain themselves.

But optimist, pessimist or transhumanist, the message is always the same. *Embrace disruption.*

I look out at the audience today, and I realise that particular message is completely, utterly redundant.

You, the members of the IEEE, don't need to embrace disruption, you live it: you live it, breathe it and sell it.

Gram for gram, you probably pack more disruptive potential than any other gathering of human beings at this moment, anywhere on Earth!

So you don't need me to teach you your trade.

Instead, I want to throw down a challenge.

I want to challenge you, the world's best possible design team, representing the world's greatest possible supply of disruptive potential, to automate *me*.

I want us to imagine the best possible AI Chief Scientist based on the technologies either available today, or in the foreseeable future. Call it "ChiefBot".

And then, I'm going to challenge myself to make the case for employing a human.

And if I can persuade you that humans will still be required, I want to draw some conclusions about how we can sharpen humanity's competitive edge.

Not just in my job, or your jobs – but in all jobs: manual jobs, caring jobs, creative jobs.

I want to persuade you that we can *all* find ways to be better than our brilliant machines, for generations to come.

And not just better than our machines — *better off*, because of our machines. Better *together*.

I've got a lot on the line. So here goes.

This is a list of the tasks a ChiefBot would be required to do.

Deliver speeches. Easy. There are hundreds of free text to voice programs that wouldn't cost the taxpayer a cent. Naturally, I would adapt the settings to sound like David Attenborough.

Write speeches. Again, easy. Google has an AI system that writes poetry. A novel by a robot was shortlisted in a Japanese literary competition. And an American student has developed an AI for speeches.

That AI system draws on a database of several thousand speeches to the US Congress. The mind boggles.

I have a personal database of several hundred speeches, and many more articles and emails.

So there's scope.

Next: **scan the science landscape and identify trends.**

IBM Watson can already do it.

After it graduated by beating the humans in Jeopardy! it took on a new career as a doctor.

But Watson can do more than just scan millions of pages of scientific reports to diagnose tumours and recommend treatments.

It's a talent scout for professional teams in the NBA.

It matches guide dogs to people.

It helps wineries and airlines to maximise their profits.

It's had more careers than Barbie.

So why not add another?

Next: **serve on boards and make complex decisions.**

And ChiefBot wouldn't be the first robot to serve in that capacity. It wouldn't even be the first in this country.

For example, there's a company in Tasmania, here in Australia, that sells AI software to advise company boards on takeovers.

There's a company in Hong Kong that's gone one step further, and actually appointed a robot director.

Next: understand politics in Canberra and Washington.

If a robot can do that, it can probably build the Hyperloop.

But speak, write, research and decide — all of those functions could certainly be packaged into an immortal, impartial, apolitical form.

So, there's ChiefBot.

I admit, it's pretty good. Brilliant, in fact.

And it's only going to get better. I have to assume that ChiefBot will capture all the benefits of machine learning, high performance computing, the Internet of Things and robotics.

I also have to assume that my own data retention and processing speeds are unlikely to improve.

But let's not abandon our faith in humanity yet.

Here's my case for human.

FIRST – a simple, but obvious point. I can be better-than-brilliant because I speak fluent human.

I use words with an understanding of their nuance and impact.

And I don't just speak *in* human, I speak *as* a human.

A robot that says that science is fun is delivering a line.

A human who says that science is fun is telling you something about what it means to be alive.

And I know. I know what it's like to be a child. I know what it's like to earn a PhD, be a parent, start a company, lead a team. I know what it's like to submit a report and eagerly wait for it to be adopted.

I know what it's like to find these things both impossibly hard and endlessly rewarding.

I also know that nothing is more irritating to the humans who write to me because they care, than to receive an automated response from a system that doesn't.

A system like ChiefBot.

So better-than-brilliant means fluent in human.

SECOND POINT – I'm limited, by design.

As human beings, we are bound by certain constraints.

We are bound in time: we die. And we are bound in space: we can't be in more than one place at a time.

In practical terms, it means I can't say yes to every event.

So when I speak to an audience, I am giving them something exclusive: a chunk of my time.

It might be a good speech, or it might be a terrible speech – but I can guarantee, it will be a one-off, never-to-be-repeated, 100% robot-free delivery.

And it has value to you, as a mark of my respect.

Of course, I could *also* pre-record a message, livestream an event or post the video on YouTube. And I do take advantage of those technologies to work around my human limitations in time and space.

But it hasn't stopped people from inviting me to speak in person.

So digital Alan alone is insufficient.

Digital Alan's role is to build the market for human Alan – just as YouTube builds the market for arena spectaculars and live performances.

I see the same pattern repeated across the economy.

Thanks to technology, goods and services are cheaper, better and more accessible than ever before.

We *like* our mass-produced bread, and our on-tap lectures, and our automated FitBit advice.

But automation hasn't killed the artisan bakery. On the contrary.

Online courses haven't killed the bricks and mortar university. On the contrary.

And FitBit hasn't kill the personal trainer. On the contrary.

All of those things are booming, alongside the machine equivalents.

The robot... and the robot-free zone.

So better-than-brilliant can be limited, by design.

THIRD POINT – I can be flexible and effective in human settings.

In our world, AI are the interlopers. We are the incumbents.

It's the robots who have to make sense of us.

And let's face it: we make it difficult.

Think, for example, of a real estate negotiation.

We could rationalise it as an exchange of one economic asset for another.

But in reality, I know that there's every chance that I will raise my bid because there was a lemon tree in the backyard and it reminded me of Aunty Gerda's lemon pie.

And then... I'll withdraw from the negotiation to be strategic.

And then... I'll panic and raise my bid again.

These decisions make sense when we account for all the things that influence we humans: like sentiment, insecurity and peer pressure.

But there are so many possible permutations of Human 1, the buyer; and Human 2, the seller, that the negotiation will never follow a predictable path.

The art of the real estate agent is to anticipate, pivot and nudge.

And the human real estate agent, she's the package deal.

She can harness AI to sharpen her perceptions and overcome cognitive biases.

Then she can hit the human buttons to cajole, deflect, flatter or persuade.

That human touch is hard to replicate, and even harder to reduce to a formula and scale.

So better-than-brilliant means flexible and effective in human settings.

FOURTH – I've learned the habit of civilisation.

Let me illustrate this point by a story.

A few years ago, some researchers set out to investigate the way that people interact with robots.

So they built a small robot and sent it off to patrol a nearby shopping mall.

When the robot encountered a human, it would politely ask the human to step aside.

The outcomes of this experiment are summarised in the title of the subsequent report: “Escaping from Children’s Abuse of Social Robots”.

Yes, that robot had a terrible time.

It was beaten by children with plastic bottles. It was smacked in the head with a soccer ball.

It was kicked, bullied, obstructed, and called what the researchers delicately referred to as “bad words”.

The more children in the group, the worse they behaved.

When asked why they attacked the robot, the top reason was “for fun”.

The researchers decided that the most reliable way to protect the robot was to program it to steer clear of humans shorter than 1.4 metres.

“But Alan, that’s a terrible story!” you’ll say.

Well, yes – it is.

But the point is not that the children were violent. The point is that the people *above* the 1.4 metre threshold were not violent.

The adults in the shopping centre were able to restrain whatever primitive impulse they felt to smack something smaller and weaker in the head.

The adults had absorbed the habit of living together. Civilisation.

And when you think about it, it’s an amazing thing.

Take electricity. It’s probably the most dangerous substance that we handle day to day. And yet it’s available to us at the flick of a switch.

We’ve tamed it. And that was only possible because we tamed ourselves – through layers of regulations and standards and industry codes and market incentives and cultural norms. *Human* systems.

Human systems that put technology to work for humans.

And if we want artificial intelligence for the people, of the people, and by the people, then we will need to remember what we’ve learned over thousands of years.

Better-than-brilliant means civilised.

Together, these four points suggest to me that humanity has a powerful competitive edge.

We can co-exist with our increasingly capable machines.

And there is plenty of room for all of our human talents to flourish, whether we want to build the robots, harness the robots, civilise the robots, or create a robot-free zone.

But if we want that future, we have to claim it.

People call me a techno-optimist because I believe humanity can do it.

You, the technologists and engineers, you are the torchbearers. I rely on you.

Let me suggest three things that we can do, as individuals and as a global movement, to sharpen humanity's competitive edge.

First, we can resist the temptation to let humans dumb down.

I have seen that temptation expressed in many ways.

There is the argument that we don't need advanced maths and science in schools.

And there's the suggestion that you don't actually need to learn a discipline, or any hard content, because all the information you could ever need is available on line. And Siri or Alexa can find it on your behalf.

I strongly disagree.

The evidence suggests that workers will need to make their niche in a fluid and unpredictable environment.

That suggests to me that we need to be *more* capable — not less.

Let's work backwards from the skills we need in our university graduates.

At a minimum, they need to master a discipline.

By mastering a discipline they learn how to learn.

With a discipline under their belts, they can enter a professional community with the mental toolkit to allow them to learn from their peers.

They can learn about leadership and life in a job, and then they can make new jobs for others as well as themselves.

But to master a discipline in the first place, they need to have strong foundation skills in language and maths.

With language and maths, they can tackle science and technology.

So we have to grab them young and keep them going.

Then they can enter the artificial intelligence age with the capacity to hold their own.

So skill up, don't dumb down.

Second, we can improve our fluency in human.

As engineers, we are trained to think and speak in systems.

But the essence of engineering is people: it is the capacity to design around our human limitations so that we, as a species, can transcend them.

It takes a grasp of how humans behave — and an understanding of what humans want.

We call that empathy.

It's the difference between the engineer who designs a product, and the engineer who delivers a solution.

We don't teach our students a formula for empathy. We haven't got one.

My alternative is to advise them with a single word: *respect*. Respect the intelligence of your audience. Respect their right to participate in the debate. And respect your rights and responsibilities as an expert — someone worthy of trust.

With respect, you can't go too far wrong.

Third, we can embrace... wait for it... we can embrace — regulation!

I'm serious.

Let me explain.

If someone told you that we must resist AI because it's dangerous, you would disagree.

No, you'd say: AI is a tool that we can use in either good or harmful ways.

I'd say the same is true of regulation: it's a tool.

Bad regulation is an impediment.

But so too is *no* regulation, when it results in consumer backlash or investor uncertainty.

What we need is effective regulation to give our society the confidence to experiment. Effective regulation sets the rules that allow you to contribute to our human future.

So let's be proactive in making the case.

The IEEE gets it. It's the essence of the IEEE Global Initiative launched last year for Ethically Aligned Design.

How often do you come across a paper from the tech community that begins with a statement like this: *human wellbeing is the highest virtue for a society, and human flourishing begins with conscious contemplation*.

It goes on: *Our mission is to ensure every technologist is educated, trained, and empowered to prioritise ethical considerations in the design and development of autonomous and intelligent systems.*

That is the tech community I recognise: a community of intelligent and articulate people, with a genuine desire to make the world a better place.

It's a community to which I'm proud to belong.

So skill up, show respect and regulate the smart way.

And of course, embrace disruption... but you knew that.

After all, you're better than brilliant. You're brilliantly human.

THANK YOU