



**Australian Government**  
**Chief Scientist**

**DR ALAN FINKEL AO**

**National Manufacturing Week**

**Opening address**

***Reversing the flow: what manufacturing can teach AI***

**14<sup>th</sup> May 2019**

**Melbourne Convention and Exhibition Centre  
MELBOURNE**

As Australia's Chief Scientist, I go to a lot of industry conferences.

I think I've spotted a general trend – and I'm sure it applies to manufacturing.

Up on stage, there are any number of people who don't work in a given industry who think they know exactly what people who do work in that industry ought to do.

And I'm certain: there are any number of people who've never set foot in a factory who want to tell manufacturers exactly what they ought to do about artificial intelligence – AI.

Now I agree that talking about AI is important – and ignoring all those people would be a terrible mistake.

But today I want to flip the script.

And instead of talking about what manufacturing needs to learn about AI, I want to talk about what AI development needs to learn from manufacturing.

And I want to encourage all of you here today to reflect on how the systems we've developed to ensure quality and safety in manufacturing can help us achieve a world of responsible AI.

But let me begin by laying out my credentials.

I come from a manufacturing family.

My father, David Finkel, was a maker of women's clothing.

He was born in an industrial town in Poland, called Bialystok, famous then and now for making two things: vodka, and carpets.

If Dad had stayed in Poland he might have followed the path my grandfather had planned for him: starting a rug-making business in another part of the country.

But this plan was shattered by the German invasion – and Dad was forced to seek refuge in Siberia instead.

Then, as soon as he could after the end of the War, Dad got on a boat, and he came to Australia – with nothing.

Or nothing, at least, in his pockets – he had courage and initiative in spades.

He also knew factories – he'd known them all his life – and he knew that manufacturing is how migrants who start with nothing can get ahead.

So that's exactly what he did.

He built a clothing business in Melbourne that at its height employed over 400 workers. And he gave many people – migrants, just like him – their start.

I admired my father and his business acumen enormously.

But I never expected to follow him into manufacturing.

When I left school, my plan was to study engineering.

I got my degree and I started my PhD on – wait for it – the electrical activity in the brains of snails.

It turns out to be extremely difficult to study the basics of what goes on in brains, even little snail brains.

I became obsessed with the need for better tools.

And eventually I came up with a design for a new kind of electronic amplifier called a voltage clamp... that you don't need to know anything about, except for the fact that it overcame a big limitation in all the existing designs.

People started asking me where they could buy my electronic amplifiers.

I realised that in order for people to buy them, I'd have to make them.

So that's what I decided to do. In 1983, at the age of thirty, I said goodbye to my research career at the Australian National University and I went with my wife to Silicon Valley.

A migrant, without suppliers, without customers, without a workspace. Everything I had was basically in my head.

I set up a company called Axon Instruments, and I went into manufacturing just like my father. Head-first.

Axon was a one man company, and the one man was me, which made it very easy to get unanimous agreement on a wages policy but very hard going in every other respect.

But I survived that first nerve-shattering year, and so did Axon.

I got that electronic amplifier onto the market, and we actually turned a profit, even though my parts alone cost as much as the retail price of my nearest competitor.

To cover direct and indirect labour and other overheads I would have to charge twice as much as the competition!

I was a novice in business, but it occurred to me that this might be a problem.

I made a panicked phone call back to Australia, and it was my step-father who picked up the phone.

“Alan”, he said to me, “is your product truly better than the competition?”

“Absolutely,” I said.

“Then charge what you need to charge, because quality is remembered, long after price is forgotten.”

That was Manufacturing 101.

But then, manufacturing 201: you’re only as good as your most recent product.

So for the next two decades I worked constantly – in my company, for my company and on my company – making new products and then making them better.

Any of you here today who have built a thriving manufacturing business, and kept it going: you have my respect.

By 2004, we employed close to 150 people, the company was still expanding, and I decided the time had come to move on. I sold the company, agreed to stick around for eighteen months as the Chief Technology Officer of the acquiring company, and then woke up on January 1st, 2006, a free agent.

I tried retirement, but it was awful.

So I went back to work, and I ended up as the Australian public’s on-call science adviser and in-house engineer.

\*\*\*

Looking back, I can match the phases of that story against the bigger trajectory of history.

My father's factory: that was Industry 2.0 at its height, the Golden Age of Capitalism; when the population was growing and so was the economy, building on the massive technology dividend from the Second World War.

My company, Axon Instruments: that was Industry 3.0, the computer age.

I founded Axon when IBM was rolling out the very first personal computer.

It was one of my first big investments: 10,000 US dollars, with, by today's standards, a miniscule 10 megabyte hard drive and just 384 kilobytes of memory.

That's about \$37,000 in Australian dollars today.

I sold the company twenty one years later, just as Apple was getting ready to launch the iPhone.

So yes, in my time as a CEO, Industry 3.0, I saw every aspect of manufacturing transformed.

And now, your factories today: Industry 4.0, the era when artificial intelligence is ascendant, coupled with rapidly accelerating progress in the Internet of Things, additive manufacturing, nanotechnology, biotechnology, materials science, energy storage, digitalisation and embedded computing.

Why do we say that we're entering a different era, with AI at its core?

Well, geologists say we can mark off a new epoch in world history if we see a universal signal – meaning it registers all over the globe – and it shows up as a distinct shift when we look back through the layers of rocks.

By analogy, we enter a new industrial era if we have a force that becomes ubiquitous, that registers in the economic indicators.

To be fair, we haven't seen a definite AI productivity spike.

But we wouldn't expect to, because we're in the learning phase, when the experiments are risky, and often, they don't go right.

For example: the world-famous "Fluffbot".

Fluffbot was a robot developed for Tesla's gigafactory in Nevada.

He had one job: to put fibreglass insulation fluff around the battery pack.

Piece of cake for a human. Seriously advanced for a machine.

And Fluffbot literally fluffed it. He couldn't pick up the fibreglass reliably. And when he did pick up the fibreglass, he couldn't find the battery. So he'd just drop it somewhere else.

Tesla concluded that he wasn't helping, and retired him.

The media loves these stories – but it would be wrong to see the failures and miss the trend.

Remember, it took us decades – decades – to see the productivity gains from developments we now understand to be transformative: such as electricity and IT.

And the trajectory in AI is clear. The individual efforts are becoming bigger and bolder – and collectively, they're surging into a wave.

Already, today, AI routes trucks.

AI makes more share trades than humans.

AI chooses the news. AI writes the news. In China, an AI even presents the news. On TV.

AI is in security cameras – an estimated 1 billion of them globally by next year.

AI is in our phones – 4 billion of them already equipped with AI assistants.

AI drives cars.

But who's impressed by cars? Think trucks. In Australia, AI drives dump-trucks on mine sites; trucks the size of two-storey houses. And AI drives the trains to the port.

On the other side of the country, at the Port of Brisbane, giant AI straddle-carriers stack and load the cargo.

And, of course, AI is seeping into every aspect of manufacturing – and manufacturing companies are buying up AI talent as fast as universities can churn it out.

Ten years ago, we'd argue about the big and abstract threat of a robot apocalypse.

Today, we're grappling with the real and present impacts of AI on our businesses, our jobs, and our children. In short, our society.

Do we want to live in a world where employees can be constantly monitored, and the least productive workers are automatically sacked?

Who should be reading our job applications and mortgage paperwork and medical scans: humans, or machines?

When is an automated driving system or production line sufficiently safe to be worthy of trust?

And how do you transition decision-making responsibility to that system over time, whilst keeping the human operators alert and engaged?

All of these questions are complicated by the massive information gap between the people who develop AI, and the people who deploy it – and the bigger gap again to the people whose lives it affects.

As consumers, we don't see the algorithms at work in our newsfeeds, or know if our job applications will be read in the first instance by a human or a machine.

And even when we do see AI in a physical form – like the SmartGates at airports that use facial recognition to verify identity – many people don't make the connection that this is AI at work.

We're still trying to find our way through an increasingly angry debate.

On the one hand, there are people who insist that AI needs to be banned – smashing the glass and pulling the emergency brake on the train of progress.

On the other hand, there are people who insist that any attempt at government control of AI is premature, that technological development and the wonders that it delivers blossom best in an unregulated free for all.

On the first path, people with scruples give up on building AI with ethics.

On the second path, we say that scruples and ethics don't count.

Either way, the unscrupulous win.

But I look at the long history of manufacturers bringing new technologies into our lives.

And I think of technologies that are inherently dangerous – like electricity and cars – that we have accepted in our lives for decades.

And I also think of technologies such as medicines, and how we have learned to minimise the adverse side effects associated with their tremendous power to heal.

We trust in our capacity to manage these technologies, not to ban them.

When you think about it, about all the things that have to go right, every time, for a safe and effective product to arrive in our hands, at a price we can afford, at the moment we want – in a country like ours, where doing business isn't cheap – that level of confidence is extraordinary.

It didn't exist at the dawn of Industry 2.0 – and it still doesn't exist in many places around the world today.

Quality is the Australian brand. Quality assurance is an Australian strength.

That says to me that there's an incredible repository of knowledge and experience in manufacturing. Right here.

And there's a lot to carry forward with AI.

\*\*\*

Let's think about how quality assurance works in manufacturing.

As a manufacturer, you understand that your practices are guided by a mesh of interlocking systems, all designed to strike the optimal balance between quality, speed and safety.

At one end of the spectrum is legislation: hard requirements, with criminal and civil penalties.

Then there are industry codes and standards: sometimes binding, sometimes voluntary; but you adhere to them because that's what your peers and your customers expect.

Next on the spectrum are the practices that you adopt internally: feedback loops to your customers, data gathering, project evaluations, employee training.

And finally, there are measures designed to inform consumers about what products do and how they are made, so that they can give their dollars to the companies that line up best with their values.

When you first go into business, you think these things are constraining.



In time, you realise that good regulation is a CEO's best friend.

It's the way you get permission from the community to be in the game.

Once you know the rules, and you know you comply with them, you can get the backing from investors, and play to win.

It means it's good business to do the right thing.

That's what we should develop around AI: not one Law of AI, but a spectrum of approaches – legal, financial, and cultural – all working together.

I've been thinking in particular about the consumer end of the spectrum.

If you're in the market for an AI baby monitor, or you're a business thinking about installing AI security cameras in your warehouse, how do you know if the product and the company that created it are trustworthy?

You could read the hundred-page disclaimer – but you won't.

Maybe, if you're a government department with a big procurement budget, you can put more resources into due diligence.

But what, exactly, are you trying to find out?

How do you know if the AI has been trained on a quality data-set?

How, for example, would you know that an AI used for targeting job ads to the best candidates isn't biased?

How can you be confident that the system you installed last week will still be properly supported in two years' time?

And even if you do have your own idea of good practice, how do the AI developers come to understand your expectations?

I was turning this problem over in my mind.

And I thought about the efforts that Australian industry has made in recent years to clean up the supply chain – in partnerships with many activists in the community.

A consumer can't tell if a T-shirt has been produced with slave labour, or if the grower of their coffee beans was paid a fair price.

But they can look for the Fair Trade mark, which tells them that the product complies with a certain minimum standard.

Then I thought about my own experience many years ago, taking my company along the journey to becoming ISO 9000 certified.

We needed ISO 9000 certification in order to be able to market a new product that inserted an electrode ten centimetres deep into human brains, during neurosurgery to treat the symptoms of Parkinson's disease.

As we discovered, the beauty of the ISO standards is that they give you a process for achieving quality by design, not by testing and rejecting.

They force you to bake high expectations into your business practices, and they keep you honest by a combination of internal and external audits.

At Axon, we maintained these exacting design and business practices for our non-medical products too, because they made us a better company and gave us a commercial edge.

So imagine if we could do the same with AI: develop a standard and certification system for quality, safety and ethics.

In the past, I've outlined one possible model for consumer products such as digital assistants, which I've called the Turing Certificate – in honour of the legendary Alan Turing.

But mine is just one of many ideas in this area.

I've just come back from the United States, where I met with the chief scientific advisor to the President, Kelvin Droegemeier.

His office is taking the lead on an Executive Order signed by the President in February.

It commits the federal government to leadership on AI governance – in its own practice, and in the standards it applies to others.

That includes the development of technical standards for reliable, robust, trustworthy, secure, portable, and interoperable AI systems, in consultation with industry – a process now underway.

The message is clear: America wants a rule-book, and they want Americans to write it.

Over in Europe, the European Commission has just kicked off a large scale pilot of its Ethics Guidelines for Trustworthy AI.

It's a set of seven principles, supported by a list of practical questions that you as a CEO need to consider, whether you're a developer or a purchaser.

For example:

- Did you put in place ways to measure whether your system is making an unacceptable amount of inaccurate predictions?
- How are you verifying that your data sets have not been compromised or hacked?

The idea of the pilot is to test the set of questions, to ensure that the guidelines can actually be embedded in practice.

Here in Australia, CSIRO's Data61 is now consulting on our own AI Ethics Framework, commissioned by the government in last year's federal budget.

The discussion paper is out there, you've got until the end of this month to make a submission.

And there will be other calls for your input, on multiple frameworks, as we get down to work on that spectrum of rules.

So, why should Australian manufacturers pay attention?

First, because it's very much in your interests to opt in.

Imagine if consumers who currently think of all things AI as an impenetrable fog had some capacity to distinguish between the good and the bad.

How much easier would it be to win support for the AI tools you want to adopt, if you could point to a rigorous external standard?

In particular, how much easier would it be to do business with big customers who will be willing to pay a premium for quality – like governments?

We know that Australian manufacturers compete on quality, safety and ethics – so let's get behind a scheme that makes those qualities count.

And second, if it's in your interests to opt in, then it's also in your interests to get involved in the standards development process – today.

You've got the experience with quality assurance approaches that work.

You know that we're strengthened by good regulation.

You can bring your perspective to best practice requirements for AI.

\*\*\*

It's still going to be a decade of tricky decisions.

And everyone here will be making them.

Am I glad that I'm a failed retiree turned public servant these days, instead of a CEO?

You bet. It's nice not to be responsible every minute for the future of the company and its employees.

But even to this day my analysis and advice is informed by my experience as a manufacturer.

The reality is that we know more than we think we know.

So, from one proud son of a manufacturing family, to the manufacturing family here today, enjoy the conference.

And, in the closing salutation of my generation,

***May the Force be with you.***