Journal & Proceedings of the Royal Society of New South Wales, vol. 155, part 1, 2022, pp. 76–83. ISSN 0035-9173/22/01076-08

Science and technology underpinning the digital age: past, present and future

Cathy Foley[a], **Hugh Durrant-Whyte**[b]

[a] Chief Scientist of Australia

[b] New South Wales Chief Scientist and Engineer

Moderator and Rapporteur: Ian Oppermann, Chief Data Scientist, NSW Government; Industry Professor, University of Technology, Sydney

Cathy Foley

I'm speaking from the lands of the Gadigal people of the Eora nation. It's great to be able to talk to you about how digitalisation is going to affect us; but I look at it from the past as well as the present and into the future. Let's go back and think about the past. It seems very natural to assess change in my own lifetime. I was born in the 1950s, and in that decade, Australia was filled with an enormous amount of innovation. This laid the groundwork for so much that's come since then. For example, the School of the Air started in 1951; it was cutting-edge radio education at the time. Australia's first computer, the CSIRAC, provided computing services to all the CSIRO from 1951 to 1955. It was built in the late 1940s by Trevor Pearcey, Maston Beard and Geoff Hill. In 1955, Australia benefited from mass distribution of the polio vaccine, which cut an annual number of about 1,500 cases to just 125 in the first year alone. This seems sort of reminiscent of what's happening now.

In the mid-1950s, my school was only beginning to teach senior science to girls. Today, the School of the Air relies on video conferencing and has a sophisticated studio setup in Alice Springs. We've lately seen and should remember that for many months, most Australian children have also been remotely learning via video conferencing systems.

We've learnt a lot about digital deliveries of education. A defining feature of our society today is our sheer, unbelievable connectivity: with ubiquitous smartphones, computers and wi-fi in our households and in our pockets. We all probably have a mobile phone within reach right now. With that connectivity comes a level of computing power and information access that we've just never seen before. And it's not just local or personal connectivity; but our ability to speak instantly to others around the world, regardless of the location or time zone. There was a major acceleration of remote networking solutions in response to the pandemic.

Seventy years ago, we were rolling out the polio vaccine; now we are all concerned with the COVID vaccines, which were developed in nine months compared to the usual ten years to develop a vaccine. Last night we saw the Prime Minister's Science Prize awarded to Eddie Holmes,[1] who in early 2020, when the pandemic was beginning, openly shared his knowledge of the genome sequence of

1 See Holmes E. C. (2021) The discovery, origins and evolution of SARS-CoV-2 (COVID-19). *Journal & Proceedings of the Royal Society of New South Wales* 154:161–181 [Ed.]

the virus, thus enabling laboratories to develop successful vaccines.

So where are we now? The future isn't written, obviously, and it's hard to really draw exact trends from the past and the present into the decades ahead. Yet I do see things in the areas of health improvements, connectivity and new ways to access the world around us. But there is a common thread above all — the incredible and ubiquitous impacts of digitalisation, which have given this century an entirely new paradigm — the digital world.

Throughout this conference, we are talking about a child who is born now and how digital technologies will impact their lives. I think it's important for us to give that person a name. I went digging around and realised that we had a ready-made whole family. The Jetsons. This television series aired in 1962, and it was set in 2062. That means that George Jetson was scheduled to be born in 2022; not today, but close enough. If we want to put a name to our child of the future, why not George Jetson?

Before you call me out, I do realise that the Jetsons is an old show. And of course, it's got very outdated gender roles, but it did lead me to thinking about things. So if George Jetson will come into the world next year, we know that he will live in a time of rapid change, tremendous connectivity and massive technological upheaval. So how do we train him to live in a world with technologies that don't exist yet? How do we prepare today's babies for careers that are still unnamed? And of course, when will we invent the Jetsons' flying car?

My key takeaway is that digital technologies are rapidly evolving way beyond some of the science fiction dreams that we had yesterday. So I'm going to consider future industries that are real and not make-believe. Every one of them relies on current and future technologies that require us to grow them if we want to see success.

The first technology is hydrogen; shipping our sunshine overseas as the newest clean renewable energy. Another is space: the Australian Space Agency is aiming to triple the size of the Australian space industry by 2030, and with that bring along 20,000 new jobs in artificial intelligence and machine learning. By 2024, only three years away, the number of these jobs is expected to grow by a quarter of a million. Digital technicians will be the largest workforce in Australia within the next few years, We're also seeing huge potential for the emerging quantum technology industry. I want to talk about that a bit more in a moment.

But even with those industries already underway, the future is much bigger than that. I'm hoping to throw some ideas out that came to me when I started wondering about the next decade or two.

The first idea is brain-machine interfacing: biological computers. Next are new ways of working in the virtual world, and the virtual commute. When digitally enabled home working is the norm, how will we adapt to this enormous change in our life patterns? Where do we draw the line between work and leisure? And will it be normal to work on a different continent to your company's head office? Will there be a head office? We've seen already that COVID massively accelerated the discussion that had happened before.

Next is automation. Will we not actually own cars in the future, but instead will we buy into a consortium or just rent a self-driving car that we call up on demand? And thinking about defence and security, what will wars of the future be like? Will there be a human pulling the trigger, or can our social media algorithms be weaponised?

Among all these hypotheticals, we need to consider digital ethics and social licence. These questions and all these future technologies won't go anywhere on their own. Whenever we talk about new technologies, we need to think about what I call science plus. We need to have science plus engineering. We need to have science plus user design and the user interface. Science plus the right business model, because we've seen that new digital technology often brings a whole new way of engaging economically. There's also science plus the regulation and policy settings that are needed for social licence. And, of course, marketing, which allows us to recognise what's available. And they all need to be talking to each other all of the time.

I want to talk about an area that I am championing: quantum. I think it's the next big industry for Australia. These technologies can do things such as accelerate drug and materials development for health care, enhance national security and support defence, increase productive mineral exploration and water resource management for mining and other sectors. Improve secure communications to industries like space and create optimisation processes, say for finance and logistics. And it's happening already, with noisy, intermediate scale quantum computers that are available via the cloud. We're seeing Airbus designing wings with these computers and Deutsche Bank using them to do develop transport efficiency algorithms.

But progress doesn't just happen by accident. Let's look at Moore's law, which suggests that the number of components in an integrated circuit doubles every 12 months. This observation was made by Gordon Moore in a paper in 1965.[2] It's a remarkably insightful paper. It also forecast that integrated circuits would lead to such wonders as home computers or at least terminals connected to a central computer; automatic controls for automobiles and personal portable communications equipment. Moore's observation was not driven by a particular scientific or engineering necessity, but it was a reflection that matched just how things happened.

The silicon chip industry took note of Moore's law and adopted it as a goal, a target for the whole industry to hit. As a consequence, we've seen that, if we want to get the most out of digital technologies, we usually need to give ourselves stretch targets to achieve. And this is something which I'm really observing right now as we're seeing a plethora of new quantum technologies being pushed out from research into industry and then just taking off.

To finish, things are changing rapidly and the challenge and the opportunity are there for us to take. In terms of social licence, we should make sure that any digital technologies we adopt, or have thrust upon us, are ones that really will make a difference for good.

Hugh Durrant-Whyte

Over the last 40 years, I've been working in the field of AI, artificial intelligence, and particularly its applications in robotics and autonomy. So I will reflect on the past and future of AI and what its implications could be; noting that AI receives an awful lot of

2 Moore, G. E. (1965), 'Cramming more components onto integrated circuits' *intel.com*. Electronics Magazine, 19 April.

press coverage and is often underpinned by a lot of interesting and new technologies.

It is useful to start by understanding what AI is. It has been around a very long time. Frank Rosenblatt invented the first neural network back in 1953. So AI, as a field in neural networks, has been around now for nearly 70 years. I remember reading my first book on AI: The introduction to AI by Patrick Winston,[3] who was the first director of the AI Lab at MIT. He said, and I use this quote a great deal: "What is AI? AI is anything we currently cannot do. When we know how to do it, it's called an algorithm."

What he was really trying to say is there's nothing special about AI, different from what other sides of computer science do. And yet we've heard a lot of what I would call grandstanding around AI. Here's a particular example I sometimes like from Vladimir Putin: "AI is the future for all humankind. Whoever is the leader in this sphere will be the ruler of the world." And in contrast to that, Andrew Ng, who is a professor at Stanford and a founder of some very famous AI companies, said in *Wired*, "I worry about AI superintelligence in the same way I worry about overpopulation on Mars." It will happen one day, but in such a distant future, he's no longer concerned about it.[4]

Pragmatically, when you think about AI, it is not pixie dust, it is not some kind of magic. It is basically data. It is algorithms and the way that those are put together to solve problems and applications. I think the interesting future for AI is the great applications that are being really rethought about and the kinds of new discoveries in science and the changes that will come in the lifetime of the child who is born today. I do want to emphasise the kind of difference or disparity between what the experts in the field know about AI and what I think at the moment.

A survey that was conducted at some of the most prestigious AI conferences showed how long experts in the field think it will take for AI to reach the level of human intelligence. Interestingly, more than half of the people in the field think it will not happen in 100 years. Yes, let me repeat that: most experts think AI will not match human intelligence during the lifetime of a child born now. That's quite important because although we think of AI as intelligent, it isn't yet by a long, long way. OK? I have a little cartoon, which reflects some of my conversations with the general public. In this illustration, Wally is basically saying I built an MVP, a minimum viable product, and the pointy-head boss says, Well, that's just a block of wood. And Wally says no, I call it artificial intelligence, and the pointy-head boss says, what's his middle name? And Wally says, well, it's shy like people. And the boss says it has emotions. So you get my picture here … sometimes I think we're a little bit credulous about what AI can actually do.

And I also mention Michael Jordan, who is the most highly cited researcher in AI. He has often said that it really is not a science or anything special. It's an engineering problem and we're still at the very early stages in terms of how we actually build significant AI systems. So while, as Jordan says, the science fiction discussions about AI and superintelligence are fun, they're a distraction.

3 Winston (1976) *Artificial Intelligence* Addison-Wesley.

4 Andrew Ng: Why 'Deep Learning' is a mandate for humans, not just machines, *Wired*, 2015.

There's not been enough focus on how we build large-scale machine learning systems that really work; that deliver value to us as humans and also that do not amplify inequities. There have been 'AI winters' where people who had hyped AI gradually realised that it wouldn't work. And I suspect that we're near another AI winter now. While there have been lots of predictions for what things like neural networks will do, in the last two or three years, there've been a lot of other papers on the fundamental limitations on those kinds of approaches. I think it'll be a long time before we'll see a self-driving car in the city of Sydney without a steering wheel. Probably not in my lifetime is the answer.

So what is the AI future? AI is already changing things, but it is changing things where we've got what we call weak AI, lots of data decisions which are just yes or no. Some very predictable outcomes. What we're not really good at yet is strong AI, where there is very little data, where the decisions to be made are not obvious, and where there are very high levels of uncertainty. Truthfully, we don't even really have the mathematics or the understanding, about to how to build algorithms to manage that kind of problem.

Rod Brooks, a very famous guy in the autonomy area, says just about every successful deployment of AI is used for one of two experiences. Either there's a person in the loop or the cost of failure for a blunder is very low.[5] So we see AI in areas where, say, I'm predicting what things you will buy on the internet. While that's fine, we do not see AI out there driving cars through the city of Sydney at this point. And we probably won't in the future; we still have people sitting in the driver's seat for good reasons.

So there will be impacts of AI in autonomy and automation, and there will be impacts in job replacement and elimination of work. Possibly the most important thing is that AI is beginning to revolutionise science and society; in my view, in a very, very positive way.

Three examples show what's possible. Last September, DeepMind, a company in London owned by Google, announced that they'd won a competition to predict protein structure using an algorithm called Alphafold, which uses some AI techniques called reinforcement learning, deep reinforcement learning. This is basically applied statistics. It's not anything magic; pixie dust. But it is really interesting that Alphafold can predict with precision every protein in the human body. That is absolutely transformational. It will revolutionise medicine. And I don't think people in the medical community realise what this has just done. This has changed the whole way we will think about medicine, about health, about synthetic biology, everything in the future. So AI genuinely will have a transformational impact on discovery in this space.

There is also a robot system that I worked on; using machine learning, AI, to make discoveries in minerals. This is an area which people would not have considered 10 years ago — but now there's so much data out there and modelling information has been transformed. And there are systems out there like Obsidian, which our team developed, which will predict with accuracy the depth and geological mineralisation right across an entire continent. And there are

5 Brooks, R. (2021) An Inconvenient Truth About AI. AI won't surpass human intelligence anytime soon, *IEEE Spectrum*, 29 Sept.

now many companies out there doing very similar things. Again, this will transform geology and mineral discovery.

The third case shows how we use AI to understand complex human problems. This example is from my wife, who works in this area. It's about trying to understand the drivers for disengaged and vulnerable youth and trying to predict the life courses of those people to understand what impact you can have by different types of interventions. This is not a trivial problem if there are 500 factors available for predicting what a person will do in the future. Then there are 200 to 500 possible combinations of these factors. So there are more ways or models for how a person will progress through their life than there are atoms in the universe. At the moment, our standard AI techniques cannot manage that many combinations, and indeed, in this example, all the top 100 combinations produce exactly the same predictions.

So in fact we do not know what we're doing at all in this area. But nevertheless, new types of mathematics, new types of AI, are driving our ability to use data to make informed decisions about ways that we can improve the human condition. That will have a terrific impact on the lifetime of this person born today.

Let me dispel the dystopia. I think that the Googles, the Facebooks, all of these sorts of things, are a distraction at this point. What will really happen with AI is it will transform science. It will transform discovery. It will transform the way that we work in this world; all, in my view, to positive effect.

Discussion

Prof Oppermann: These were two very different perspectives on the impacts of science and technology to a future world. Cathy highlighted some technological advancements because of the quantum leap, which will change the way we do things, yet it's important to understand the limits of what's possible. It's important for us, as we think about this future environment to consider what is possible, what's likely and in fact is not possible. Understanding what's not just is as important as being able to understand the likely scenario. Hugh, let's go back to you. You said that AI will transform the way we do discovery. What do you mean by that?

HD-W: So I think the best example is, in fact, the Alphafold one, you know, people have been trying to predict protein folding for 50 years because understanding the structure of proteins tells you what they will do and how they will work, and therefore what proteins you need to tackle if you're building a vaccine, what protein to design in lots of other areas and so on. AI is now able to predict every single protein, the structure of every single protein in the human body. What's more, it can predict the structure of any protein from any sequence of amino acids. So the whole process now of discovering vaccines, of discovering mechanisms, of discovering new types of ways that life works and so on, can be tackled using these new techniques. It is transformational.

Prof Oppermann: And does it have implications for the nature of research? The nature of investigation is the future of research. An algorithm sorts through data and puts out something interesting, and then researchers say that we will take it further or … ?

HD-W: Absolutely not. I think that's part of the challenge. In the traditional way of doing science, you get a pile of data and then search through the data to confirm or deny a hypothesis you already have. Now, the whole point about discovery is, on the contrary, to really understand the area of models, e.g. the way things fold as proteins, the way geology works as geology, the way humans evolve in terms of the environment they work in and so on, and use data to understand which of those most likely explains the data. Discovery is quite different from hypothesis testing. And I think there needs to be a transformation in science to really begin to adopt these approaches. I'm not sure science is as ready for it as the machine learners are.

Prof Oppermann: Cathy, what do you think quantum really can do and what can quantum not do, what will quantum not change? What is the arc of influence of quantum knowledge?

CF: I don't think I can answer where it can go or where it can't. Right now, we're beginning to see the evolution of laboratory projects where we were able to control quantum states, which weren't possible before, to build machines that are able to maintain them in a stable state for periods of time. But how do we create something that can do computing for us, which allows us in some cases, to do it faster? In other cases, we can actually do computations in ways which we can't currently do, which is what they call the "quantum advantage" or "quantum supremacy." At this stage, we've only had occasional demonstrations of that ability or that idea.

So there's a wide range of views as to whether that will be possible in three years, or never. Because some people feel that a quantum computer will never work; they think that you just can't get to a point where you can control quantum states, where you can do multiple calculations at the same time — which does seem like an exhilarating capacity.

Regardless, we're now seeing the ability of new technologies to control quantum states, which allows us to do better sensing. It's allowing us to do some acceleration with even our classical computers; to do computations faster. And in theory, when we're hearing about the range of possibilities in human nature … where we have choices and can choose new pathways … it's allowing us to be more than atoms of the universe.

If we were able to get a quantum computer to operate, we would have the advantage of dealing with big questions such as climate; a major topic at the moment. Climate modelling is limited by the computation power you've got in your supercomputers. So, if we were able to do more complex modelling, that would give us the potential to have sort of a way of modelling reality in a way that might be closer to the truth because you've got to remember any model is only as good as the design of the parameters you choose.

Quite often we think that models are everything but, hopefully, when we have more powerful computers, we will be able to compute models that are closer to reality. So when you ask the question of where the boundaries are, it's impossible to say. We've seen over history that when you have a lot of people who are focused on a problem putting their brains into it together, collectively, we actually see pretty significant shifts in the hardware side of things.

Prof Oppermann: What we're trying to do today is helping create a model to help us think through the next decades: 2030 to 2050 and beyond. One question from our audience is about the bias associated with AI training on white males in Western countries … a very famous facial recognition example. Let me ask our panellists: When will we understand biased algorithms and their potential harms? When will we have enough protections and understanding of how to use data similar to the way we use electricity?

HD-W: There's a big difference between AI as an algorithm and the data you put into it. AI as an algorithm is not really biased or unbiased or anything, really. It's not ethical or unethical. It's an algorithm. But it is difficult if you take an algorithm, that say classifies things, and you try it on a set of data, which comprises only one class, then you will end up with an answer which basically classifies everything into that class … if you see what I mean.

It's the data part that really raises the ethics issues rather than the algorithm. My wife is a statistician and she looks at this and she says, well, statisticians have been dealing with bias for so long, it's not funny. You know, we need to sample the complete space in a way that's representative of the space that we wish to make decisions on. If we are to end up with a regression function, a clustering algorithm which genuinely will lead to good outcomes in the future, it's not rocket science. One thing we really need to do if we're going to address bias in AI is to make sure everyone gets trained in statistics when they're at school. Then people might appreciate the way data drives decision-making processes. For example, some people published a paper where they tried to distinguish between criminals and non-criminals, and it turned out that they had trained their method to detect whether subjects were smiling or not. Those who weren't smiling were from prisons. My point is that it's good statistics that drives accuracy more than anything else.

CF: I'm a big fan of stats and I'm really concerned that we aren't teaching it enough. We're seeing too many statistics schools being closed down in universities; that's a big flag for me. There's a group of people who are really keen to get more statistics into our education system, and I think that's absolutely the way forward. We need to understand more precisely how we're making decisions and who we are considering. Is it human nature or the way we socialise? We have this truckload of unconscious bias. To have a logical way to cut through that bias is absolutely critical.

We must add transparency. A lot of things which are being done via AI systems, behind search engines or social media, involve some decision making which doesn't give us a choice or knowledge of what parameters they are using. To work towards more transparency, we need (a) understanding, so there's education, and (b) transparency, so we know what's going on. Another one is to have ethical constraints that we as a society agree on. How are we going to manage these emerging technologies and the opportunities they offer, but also make sure that they're used in ways that allow us all to feel safe and that there isn't bias that means that some are disadvantaged.

Prof Oppermann: We've been really honoured today to hear these perspectives from Australia's Chief Scientist and New South Wales Government's Chief Scientist and Engineer.