Journal & Proceedings of the Royal Society of New South Wales, vol. 155, part 1, 2022, pp. 104–112. ISSN 0035-9173/22/010104-09

Safety and security of our digital child

Dale Lambert[a], Rory Medcalf[b]

[a] Chief of Cyber and Electronic Warfare Division, Defence Science and Technology Group, Australian Department of Defence

[b] Professor and Head of the National Security College, Crawford School of Public Policy, Australian National University

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

Dale Lambert

I'm speaking to you today from the land of the Kaurna people in Adelaide about security of our digital child. The Industrial Age is a two-tier system comprising a human domain and a physical domain in which people directly control their physical industrial machines. The Information Age expands this to a three-tiered system in which an information domain now separates the human and physical domains. In the Information Age, people issue commands to information environments and expose information environments that now directly control the physical machines.

Our digital child is born as the Information Age eclipses the Industrial Age. Society is now totally reliant on information environments. This creates an unprecedented opportunity because we are constructing a digital representation of our physical information and human domains and making those representations accessible to anyone anywhere. But it also creates a security threat because our society's physical information and human infrastructure are now totally reliant on information environments. A physical domain is now totally reliant on the information environment, and this includes our critical physical infrastructure. So, if someone controls our society's information environments, then they also control our society's physical industry.

Our information systems are also totally reliant on information environments and their algorithms, and these are susceptible to algorithmic warfare or cyber-attacks. So, if someone controls our society's information environments and algorithms, then they control our society's information. And our truth is also totally reliant on information environments. The Correspondence Theory of truth assigns truth based on correspondence with the world, but this now comes through digital images of the world that can be manipulated by image and video editors. The Coherence Theory of truth is the science truth based on coherence of opinion. But this never comes through social media that can be manipulated by fake news. So, if someone controls our society's information environments, then they control our society's truth.

What can we conclude? Well, if someone controls our information environments, they control our physical industry, they control our society's information, and they control our society's truth. In short, if someone controls our society's information environments, then they have total control of our society. Information environments are now contested. They have become the new theatre of warfare in which the participants can be nation-states, crime syndicates, ransomware syndicates, lone wolves or insiders. Human conflict now includes information warfare conducted in and through our information environments by digital ghosts within the machine.

Today's birth of our digital child corresponds with the Information Age of warfare, superseding Industrial Age warfare. By the time our digital child becomes an adult in 2039, our digital adult security will depend on automated contests conducted by digital ghosts within our machines and information environments. The security of our digital adult will critically rely on appropriate ethics and trust being embedded within the digital ghost.

So what happens as a digital adult moves beyond the Information Age? I contend that our digital adult will enter what I'm calling the Virtual Age. The computer games industry and artificial intelligence community will combine to deliver immersive technologies beyond entertainment — to become mainstream in the commerce, health, education, defence and other sectors. Why would this happen? Well, computer science has reduced the communication gulf between machines as electrical and electronic devices on the one hand and the human uses laden with rich conceptualisation, on the other hand by incrementally automating human conceptualisation within machines.

If the Virtual Age continues this application of the automation principle, then the computer gaming and AI communities will deliver us virtual people, virtual societies and virtual environments. For some of you, this might seem like a fanciful suggestion. In the Defence, Science and Technology Organisation [DSTO], we started building prototypes of such things back in 2000. In this project, someone can have a conversation with a virtual person. There is a psychological architecture underpinning the virtual person, and the virtual person can display a range of emotions. Our prototype supported agreement protocols that allowed societies to dynamically form from collections of real and virtual people, and included animated virtual environments that could represent both real and imagined worlds.

By 2063, our digital adult has reached middle age. The Information Age offered the opportunity to see something else by digitally connecting us with the wider world. But the Virtual Age goes further by offering the opportunity to be something else, by digitally experiencing real and imagined worlds. A digital middle-aged person can be their physical self or their virtual self. They can be someone else, perhaps an Indigenous Person. They can understand the mathematical curve by riding it like a rollercoaster. Or they could experience what it is like to be a DNA molecule.

But with opportunity comes threat. The threat of the Virtual Age depends on what we want to protect and secure. How should we balance our digital person's life between the real and virtual worlds, including, for example, the threat of virtual addiction? What virtual truths should we countenance when we can create virtual worlds in which conventional physics and psychological social norms need no longer apply? And what are the rights and status of virtual people? These are just a few of the many issues that will arise.

Rory Medcalf

I'm joining you from Canberra, from the traditional lands of the Ngunnawal people. I want to complement those great remarks by Dale and look at, I guess, a broader picture of what the security environment could look like for the so-called digital child growing up in Australia in the years and the decades ahead.

I'll start briefly with that fundamental question: What is security? It's one of those words that we all think we know what it means — but that we've all got different conceptualisations of. If you go back to its very roots, it's really about a state of mind. Yes, it's about physical protection, but there is no such thing as absolute security. In fact, if you take the Latin origin of security, it literally means without care — no worries. The security of next generations requires achieving a kind of world view or a perspective where people can engage confidently with risk — they cannot achieve absolute security. That will apply both to individuals and our society, indeed Australia as a nation state.

Let's look at the horizon of risk for Australians over the next few decades. It's very easy to be gloomy about this, when you look at the horizon of risk that we see right now — everything from pandemic through great power challenges, China's use of its coercive power, the use and misuse of technology, the continuing risk of terrorism, threats to social cohesion and of course, the overarching threats and risks of the impacts of climate change.

We shouldn't be complacent about any of this, but we should also bear in mind that the last 20 to 30 years have probably been the anomaly. You know, really, the generations growing up in the late 20th century and the beginning of this century had — I hate to say it — almost a long holiday from the historical traumas that most earlier Australians experienced.

Remember the experience of the world wars in the first half of the 20th century? Remember the shadow of the Cold War for much of the second half? So, what is the horizon of risk for our digital child for the next few decades and moving to mid-century? And what are some of the opportunities for society and governments in mitigating that risk? Just a few things to get you thinking about.

First, a lot of the risks we can already see on the horizon of the next ten years are going to be very, very influential in shaping our security environment for much of the rest of the century. There is the question of how states behave in a very competitive international environment. The tensions around, particularly the way that China is using its growing power in our region, the Indo-Pacific, and globally. How is the United States in particular responding? But how will Australia and other countries in our region — India, Japan, Indonesia and others—respond to these tensions? The risk of coercion, military force, even war, but also the more prevalent day-to-day risks in a competitive environment that Dale spoke about — the use of technology by states for strategic advantage, the use of investment and critical infrastructure for strategic advantage.

We're going to see generations growing up with that shadow to come to terms with and much more direct engagement, I think, in the idea of national security than we've seen for many years. There's also the domestic dimension. Even if, internationally, states succeed in managing their differences without confrontation or war, there will be ongoing threats to our sovereignty and our economy. What about the security picture at home?

We've already heard quite a lot about engagement with the information environment, with digital technologies. I think there will be a loss of innocence; that new generations will automatically recognise that their connection to the economy and to the information ecosystem is going to be a source of risk and security anxiety. But hopefully, government and society can engender a new maturity in engaging with that risk so that individuals grow up with a very strong sense of awareness about protecting their privacy, protecting their political freedoms and protecting their engagement with democratic institutions.

There will be risks, I think, to Australia's social cohesion. We've got to remember what a grand experiment a multicultural, federated Australia actually was in the history of our region and the world. The challenge there will be a tension between individuals wanting to simply get on with their lives, as opposed to individuals recognising the need to engage more actively with the political process, to be engaged in society, in politics, to protect those democratic institutions that really have allowed so much individual freedom to flourish in Australia.

And the risks to that social cohesion could come from foreign states seeking to interfere in political processes, particularly the Russian interference in the US elections in recent years. They could come from dissatisfied elements within our own society, the challenge to accept the notions of truth, the rise of coordinated misinformation, disinformation, political violence; terrorism is a fact of life in many countries today and needs to be managed and kept in perspective. We can't let fears of terrorism dominate our daily lives, but we do need an effective national security response.

And there are risks to our social cohesion more broadly. There is a need to protect privacy and political freedoms, but also to inculcate a greater sense of responsibility for our collective future and collective destiny. Those are going to be the kind of challenges that policymakers, but also communities, indeed parents, are going to face in preparing new generations for the challenges ahead. I'm not all gloomy about this, even if it's hard not to sound that way when you're focusing on security.

We've got to think also about the extraordinary capacity of Australia. This is a nation that is often not always mobilised as the sum of its parts. The ultimate challenge for our political class and for politically mobilised communities will be to rebuild a greater sense of common purpose in a democratic Australia. Education will be absolutely key here. As a parent myself, seeing new generations of school age thinking critically about the world, engaging with science and evidence, I think that we still have enormous potential in this country to meet these challenges. We have to go forward with our eyes open, and that will be the challenge for the digital child.

Discussion

Prof Oppermann: Two very interesting presentations and quite different perspectives on issues related to safety and security. Dale, I am going to ask you a question first: What do we mean by trust in a digital environment? Is it, for example, that we believe the system has our best interests at heart? What does trust really mean in a future digital world?

DL: For me, it means the system is respecting our intent, and this goes very close to what Rory was saying about state of mind. You really want implementation in the system to represent the state of mind that we want to have as our values in Australia. I mentioned having to embed trust and ethics within the machine. We need to do that because of the time frames at which things happen. So, for example, I implemented a system for sweeping defence that had to make decisions in two milliseconds. There is no chance that a person is going to be able to be involved in that sort of decision-making. So, it's really important that we take our concepts of trust and ethics and start embedding them inside computerised environments in order to maintain control, if you like, over the information systems that are actually controlling everything now.

Prof Oppermann: Rory, I am going to ask you the more general question of the interplay between encryption and the ability to survey or sense the world around us. What do you think the consequences are of limiting access to encryption and or the interplay between privacy and security?

RM: It's a great question, and one that really frustrates policymakers. There are tensions that we've got to navigate here as a liberal democracy. On one hand, we shouldn't have any illusions that by restricting the ability of our own security establishments to access technologies, to basically co-opt the private sector, for example, in accessing data, that somehow we're going to achieve complete privacy or complete protection of our liberal democratic values. There's a very competitive international environment. Whatever constraints we put on our own security agencies, to sometimes compromise civil liberties or compromise privacy in the interests of national security, there will be authoritarian powers out there who have absolutely no such compunctions. I always find it strange that we'll have people who understandably at one level are really concerned about surveillance by intelligence agencies operating under the rule of law in a democratic system, but at the same time, they're very happy to share pretty much their entire personal data with commercial entities that might have relationships with authoritarian states.

We need to find a new balance here, and I think that balance is going to be struck through constant political scrutiny through engagement with the political process — parliamentary committees and so forth. Being able to really challenge intelligence agencies to justify the powers that they have, but at the same time, to be very open about the trade-offs we make.

I use examples such as the dreadful terrorist attack in Christchurch some years ago by an Australian national, but also terrorist plots that are regularly being frustrated or uncovered. It's going to be very difficult to talk about protection of privacy when being able to access encrypted data would have prevented such attacks. This is going to be a constant challenge, and we're going to have to have very open conversations in the political process about it.

Prof Oppermann: Dale, in the world of security, it's often stated that people are some of the weakest points in the security of systems. So how can we help individuals, old and young, to have more skills in security through the next decades?

DL: The short answer is education, obviously. And we have things like the Australian Signals Directorate's "Essential Eight" that people should practise, but it's more complicated than that. It's about exposing the extent to which people are vulnerable in the things they do. For example, in the previous question about encryption, people may not appreciate the fact that even when you have encrypted data, you can analyse the packet flows of the encrypted data inside the communication networks, and you stand a very good chance of understanding what they are doing.

One example is about Adversarial Machine Learning, which is where you insert data into the data stream very deliberately to cause things to appear and disappear in the outcome from the machine-learning algorithm. You can control the conclusions of the algorithm by injecting data into the system. There's all this stuff that most people probably aren't aware of. Part of the game is making those things more exposed so people understand what the vulnerabilities are. On the human side, we have a team of psychologists in our organisation who can run a bunch of instruments that will assess someone's vulnerability to these things, based on their personality types. This has been used in various organisations. We don't use it to go in and say "sack this person because they're a risk," but it has been used to give an overall profile of what the risk is like within a particular organisation. So, the results of the individuals are masked, but the overall summary of the risk of that agency is revealed.

Prof Oppermann: One of the things about cybersecurity is we don't ever arrive at a cyber-secure environment. It's an ongoing activity because the world is changing around us and it's a very dynamic environment. I want to get back to a question around trust and ethics. Does it require data and computer engineers to have a better understanding of ethics, the rule of law and other foundational values? That's the audience question, but as a follow on … how do we ensure, over the decades of the future, that we are doing appropriate things with technology, specifically thinking about defence and security?

We have heard about the countries taking moratoria on autonomous weapons, for example. These are technologies which we will not use. The ethics is very clear in a situation like that — until you have tensions with another country that won't take that moratorium. But between those black-and-white cases, how do we ensure that we're doing appropriate things with technology over the coming decades? Is it a simple matter of making sure that computer engineers have ethics and rule-of-law training before they're released on bits and algorithms? Or is it something else?

RM: Broadly, it's a self-answering question. It's an important question. I think the short answer is absolutely in technology design; whether that design is occurring in Australia or other democracies or whether it's technologies we're making use of.

Of course, ethical principles need to be raised and addressed in the design phase. But it's not enough to put the onus on design, on engineering. It's the use of technology every day that's going to require ethical decisions and that ethical sensibility needs to be instilled in policy leaders and, frankly, in ordinary citizens as we make our own decisions about using technology.

Ironically, militaries are often ahead of the curve. Ethical training and awareness in militaries is often much greater than in other parts of the policy apparatus or civil society in democracies, or certainly in the private sector, because militaries are making life-and-death decisions every day. And when they get it wrong or do it wrong, as we've seen in current scrutiny of behaviour by a very small number of special forces in Afghanistan, it becomes a major national scandal. So, ethical sensibilities have got to become mainstreamed.

DL: Would I leave it up to the computer engineers? Absolutely not. I'm a huge fan of multidisciplinary approaches to things, and when it comes to something like embedding ethics in machines, you need to understand ethics — and ethics is not like any other discipline. Ethics isn't just one thing that people understand. There are different schools of thought. There are Aristotelian ethics, there are consequentialist ethics, there are Kantian ethics.

So, part of the challenge here is not just how you code this stuff up — it's also what kind of ethics that you want to have in your system. This is quite a serious question, and I think it's a question that should be addressed globally — a bit like we're doing with climate change at the moment, where we're trying to get the world community to share an understanding. As we move forward in time — as our digital child grows up — this is something we also want the world community to embrace so that we have a coherent policy approach across the nations.

Prof Oppermann: Another question has come in, which asks you: "Is there a difference between trust and trustworthiness, when it comes to digital systems?"

DL: Yes, probably. Trustworthiness is a term that's increasingly used in the military context. It's really asking: Is this equipment going to do what I expect it to do? Part of that is about trying to protect your equipment from cyber-attacks and things of that ilk. Trust, for me, is a broader issue. As I talked about before, it goes very much to Rory's concept of security as a state of mind. I see that as a much more human and intent-driven activity.

Prof Oppermann: The last audience question is a statement that there is a fundamental risk that encrypted data which is stolen today will eventually be decrypted with tomorrow's quantum computers and the harms will come tomorrow. Do you see this as an issue? And if so, what should we do about this today?

RM: I think that that's a reasonable question. What, how and when will we actually see quantum encryption and quantum decryption realised? It's been much promised, but it's still some way off. But just because data is encrypted doesn't necessarily make it more important. There's an enormous amount of unencrypted material out there, or open-source information out there, that is potentially incredibly useful to a future adversary or malign actor if they can aggregate it and make sense of it, particularly with AI and machine learning. So, for example, my university some years ago had a major cyber breach that was publicly reported. Other universities, corporations, government agencies, individuals, have their data taken all the time. And many of us will say, Well, so what? It doesn't matter. Who cares if you know whether it's a cybercriminal or whether it's someone sitting in an office in a government building in Shanghai reading everything that's on my screen. Well, it's going to be useful somehow, some way. So we must get much better at that; not only because of general digital hygiene, but understanding that other governments or organisations building a complete life picture of you as an individual can be particularly dangerous in the long term. We should all remember that our own data is incredibly valuable. It's who we are. It's not only our privacy, but it's our future careers, for young people who may wish to work in government, for example, or have careers in security agencies or politics. Every piece of your digital footprint now could be used against you in future.

DL: The idea that tomorrow's quantum computers are going to completely overcome the encryption situation isn't quite right. There's a thing called Shor's algorithm, which sits in a class which is a bounded quantum polynomial — that sits higher than polynomial, but it doesn't get you to the full exponential non-polynomial (NP) complexity. I mention this is for the computer scientists who are listening.

I remember a really interesting conversation with someone from a foreign intelligence agency about their fingerprints remaining in the digital world so that they can at some point be uncovered for something they've been doing, maybe in 10, 20, or 30 years' time. It's quite a worry for them because they might do something at a time when they can't be detected. But if it's being recorded through some sort of digital mechanism, they might well be uncovered at some point in the future. It's a really interesting consideration and hopefully, it might stop some people doing some nefarious things because they'll have second thoughts about it.

Prof Oppermann: One of the points you touched on was with the use of fingerprints, and historically we've used biometrics for a whole lot of different sorts of authentication purposes. If they lose their significance as individual identifiers, then we have to rethink a lot of traditional ways of identifying and securing. Let me ask you the unfair questions: In 2050, what's no longer an issue, and what are we really focused on?

RM: The second part is much easier to answer than the first. I would never go so far as to say that the kind of terrorism that we've been dealing with — quite obsessed with over the past 20 years — will no longer be an issue, but I do think that we're on a pathway to putting it in context; to recognising that terrorism is essentially a criminal activity, politically motivated violence. I think though that the risk of terrorism and violent extremism is going to be background noise in our national security debate permanently. We need to deal with bigger issues, such as Australia's resilience and sovereignty in a pretty contested region.

Relatively speaking, Australia will probably be a less powerful country than we are now. That's a very sobering thought. There is a need to draw up the connections between national resilience and the security of our energy supplies, the sustainability of our society, the environmental sustainability to connect all of that with the idea of national defence and our security in a competitive world. That's going to be the big question.

I guess the good news is that the barriers between security and economics will break down. The bad news is it's going to be very hard to turn that into a practical policy agenda that the government can operationalise. So it's going to be tough, but I think citizen engagement in national security will be a big part of the solution.

DL: I'm hoping that we won't have to worry about information warfare by that time because we will have had some sort of stabilisation and an international agreement around what's reasonable and what's not reasonable in terms of opportunity in that time frame. I talked about three different ages: The industrial age, information age, and virtual age. It's easy to underestimate how important the virtual age will be because it basically lets human ingenuity off the leash. You're no longer restricted by a lot of the physical constraints and things that we've all grown up with. This is an opportunity space and it's really up to us how imaginative we can be in exploiting that opportunity.