Protective Security in Government Conference

Speaker: Sven Bluemmel, Information Commissioner

Date: Tuesday, 14 August 2018

Artificial intelligence, Privacy and Security

Good morning everyone and thank you for welcoming me to speak at the Protective Security in Government Conference today.

I would like to start by acknowledging and paying my respects to the traditional owners of the land on which we meet - the Ngunnawal Peoples. As we share our own knowledge within this conference center today may we also pay respect to the knowledge and traditions of the Ngunnawal Peoples, and of their Elders past, present and emerging.

I am delighted to be here to speak with you as Victoria's inaugural Information Commissioner – and what an exciting time it is to be an Information Commissioner! Since the inception of the Office of the Victorian Information Commissioner, or OVIC, in September last year, we have seen an enormous amount of change in the information privacy and data protection landscape already, both locally and internationally.

As you may be aware, OVIC has combined oversight of Freedom of Information (or, FOI), privacy, and data protection, administering both the Privacy and Data Protection Act and the Freedom of Information Act in Victoria. OVIC combines the functions of the previous Office of the Commissioner for Privacy and Data Protection and the Office of the Freedom of Information Commissioner.

The establishment of OVIC signals a new direction for engagement with public sector agencies and entities to assist them with information management issues and to help drive systemic and cultural change in Victoria.

Now, we are coming up to OVIC's first birthday. Over the past 11 months I have been spending a significant amount of my time engaging with stakeholders to promote and drive cultural change in the information privacy and protective data security space.

One of the things I have noticed over this time has been the enthusiasm that organisations have toward new and emerging technologies, and the opportunities as well as the risks they pose for government in the information age.

One such technology is what brings me here to speak with you today: artificial intelligence or, as it is more commonly referred to, AI.

But before we get stuck into that, I wanted to briefly touch on some of the more philosophical ideas underpinning information in government, and information rights more broadly. In my opinion, these ideas truly underpin the importance of both information privacy, and protective data security.

Just a note on that – I am aware that 'protective data security' is a term we use in Victoria, so for anyone who is unaware, when I refer to protective data security I mean a subset of protective security that focuses on the protection of public sector data.

Information is the lifeblood of society. It is both personal and political, and the way that we as public servants handle it impacts people on an individual level, as well as society as a whole. What do I mean by this? Well, on a personal level, the ability to exercise control over our personal information is essential to the development of one's identity, sense of self and how we interact with each other and the government. On a societal level, the very functioning of government and democracy rely on the ability for information to be collected, stored, used and shared securely. Privacy and security work in tandem to ensure a free and fair modern society.

Perhaps I am biased, but I don't think that anyone would seriously argue that privacy and security are not of pivotal importance in our world today.

Information driving modern government

In carrying out their democratic and statutory functions, governments collect, handle, use and store vast quantities of data, including personal information about their citizens. Used and shared appropriately, public sector data can enable governments to make informed decisions and provide better policy and service responses to the issues of the day.

In a public-sector context, much of the data to which government has access contains personal information about individuals, so the objectives of government must be achieved while ensuring that Victorians are able to enjoy their information rights. This is precisely where part of OVIC's mandate lies, to ensure that appropriate protective data security protocols are in place for all public-sector data, and to protect the information rights of individuals. This is why we exist.

So, what has this got to do with artificial intelligence?

As I mentioned at the beginning, in my experience thus far as Information Commissioner I have noticed the enthusiasm with which organisations are running toward new and emerging technologies. One of the most exciting areas is of course, AI.

Recently I had the opportunity to travel to the Asia Pacific Privacy Authorities Conference, where I joined a Technology Working Group that specifically focuses on artificial intelligence. At this conference there was a palpable sense that information rights are at an important global juncture at the moment. How governments and regulators respond to technological and social developments in the next few years will have a large and lasting impact on the kind of society we live in.

In relation to AI, there was a clear tension between the camps of "don't stifle innovation or society will miss out on big benefits," and "the use of AI has far reaching effects on human freedoms and organisational accountability, and thus needs to be carefully regulated."

Regardless of in which camp you reside, AI is shifting the privacy and information security landscape dramatically, and we have the opportunity to shape how that landscape looks. But to do that, we need to be considering the challenges and opportunities that AI presents for both privacy and information security now. And this is why I am here today.

What is AI? Context, definitions and terminology

Artificial intelligence is certainly having a moment right now. It's attracting immense attention not just from those who work in the technology industry, but also government, the private sector, and of course, individuals. Hardly a day goes by without hearing or reading about AI and the impacts it is having on our lives.

As human beings, we have a tendency to both catastrophise and romanticise. The kind of stories we read about AI in popular media tackle both ends of this spectrum – it's either superintelligent killing machines are taking over the planet (or at the very least, taking all of our jobs), or, the benevolent machine overlords are swooping in to solve any and all of our problems. If all we go by is the media and pop culture, it seems our choices are limited to dystopia or utopia.

Today, I am going to attempt to explore the middle ground. I am going to talk about both the challenges and opportunities that AI presents to us in the information privacy and protective data security space.

Of course, I think it is important that we recognise the immense power and potential of this technology, but also that we remember that the future is not set in stone. We all have a role to play in determining what the society we want to live in looks like, and this extends to the role that AI plays in government and our lives.

First, a little bit of background.

With all of the recent hype, it might be easy to assume that AI is a new thing. But this couldn't be further from the truth. In reality, the philosophy of artificial intelligence can be dated back to as early as the 18th century. The technical concept of AI as we understand it today has existed since the early 1940s and was popularised by Alan Turing in the 1950s with the 'Turing Test,' which some of you may have heard of. The resurgence AI's popularity that is making it 'on trend' right now is due to a period of rapid technological development that has enabled those longstanding ideas to be brought to reality. This is thanks to three factors: improved algorithms, increased networked computing power, and increased ability to capture and store an unprecedented amount of data.

As an example of how times have changed, some of you in the room may be old enough to remember how expensive it used to be to store data. This has swung so far in the other direction that in many instances it's cheaper to keep all the data, forever, rather than go through the arduous task of deleting it. This is, as you can imagine, not cool for privacy.

So, we are operating in times with unprecedented computing power and access to troves of data, bringing AI from theory into reality.

Definitions/Terminology

But what does all of this mean? Before I go any further, I think it's important to clarify exactly what I mean when I say AI, as it is one of those terms than can be used to refer to many different things.

Within the context of today's session, when I refer to artificial intelligence or AI, I am using it as an umbrella term to refer to a sub-field of computer science that has the goal of creating programs to perform tasks that can be considered to be 'intelligent.' Many of these tasks have traditionally been performed by humans, but as we will explore a little further, one of the points that differentiates AI from other disruptive technologies is its potential to do things at a scope and scale far greater than that which humans could ever dream to achieve.

The umbrella term of AI can be used to describe a collection of related technologies. Most often it is used to refer to deep learning and machine learning, however it may also be used to describe techniques such as artificial neural networks, predictive analytics, natural language processing, machine vision and robotics. In reality, these techniques are often implemented together to create AI systems.

Another point worth distinguishing is the difference between narrow and general AI. Our focus today is on narrow AI, but it is important to understand the difference.

Narrow AI is deliberately intended to be competent in one specific area. What this means is, you may have built an AI system that is excellent at reviewing and interpreting mammograms for cancer at a speed and accuracy rate drastically better than any human could do, however, show it a picture of a cat and it wouldn't know what it was looking at.

As another example, in the 1980s IBM developed a computer called Deep Blue. Deep Blue can play chess at a level superior to any human being, which was a feat of huge importance in the timeline of AI development. However, while Deep Blue exhibits above-human ability in chess, ask it to play Snap and it would be completely dumbfounded. This is because it has only been trained in one specific, limited domain.

Conversely, General AI refers to a level of intelligence across multiple fields – the ability to learn and switch between contexts, and to apply the knowledge and experience of one domain to another.

We already see this distinction in the natural world. Let me give you an example:

Bees are incredibly intelligent when it comes to building beehives and constructing honeycomb. Ants are also incredibly intelligent when it comes to building a nest. However, this intelligence is only applicable to a certain domain – bees are not able to build nests, and ants are unable to build hives. This is narrow intelligence.

Humans, on the other hand, exhibit general intelligence. We have the capacity to be intelligent across a range of different areas, and we can learn intelligence in new fields through experience and observation.

Currently we have lots of examples of narrow AI, but general artificial intelligence remains in the realm of science fiction for the time being.

Lastly, something to keep in mind is the concept of artificial super intelligence. This is a purely hypothetical idea at this stage, building on the notion of general intelligence. Artificial superintelligence proposes an AI that is both general, and also far exceeds human abilities. It is also a concept that is tied up in the idea of machine sentience and consciousness – machines that develop a consciousness and emotions in the same way as a human being.

Many pop culture references to AI display it in the form of superintelligence – consider recent films such as Ex Machina and Her, or even older films such as Bicentennial Man or Terminator. This kind of portrayal contributes to the hype surrounding AI – but can also be misleading.

I want to emphasise that today, I'm not talking about superintelligent humanoid robots. Today I am talking about the AI that we are already experiencing, and the AI that will have real life implications in the immediate future – narrow AI.

Some of the techniques of AI

So, what about some of the techniques of AI that I have mentioned? While this speech is not intended to be a deep technical dive, and I am by no means a technical expert, I think it is worth introducing some of the technical ideas that we are covering today here too.

First, Machine Learning. This is a particularly prominent subfield of AI that gets a lot of attention. Machine learning is a computer science technique that uses algorithms to learn and modify itself. Through an iterative process of ingesting data, the algorithm trains itself by developing its own logic according to the data it has analysed.

Deep learning is another popular technique you may have heard of. It is a subset of machine learning, the next babushka doll in the AI family, if you will. Deep learning typically uses deep neural networks which process data through a layered approach. This means the result of each layer, or, the output, becomes the input for the next layer. The process of moving data through each layer can make it increasingly difficult to understand the decisions and inferences made at each level – resulting in what is often referred to as the 'black box' effect. This presents unique challenges to transparency and accountability, which I will get to shortly.

As a side note: the name 'neural networks' can be misleading as it conjures images of the human brain. This is not a useful analogy, as the structure of an artificial neural network, and the way that it processes information, is actually not at all like that of a human brain.

Machine learning and deep learning are extremely powerful tools that many people point to as the reason for the most recent explosion of AI. Deep learning in particular has given machines the ability to recognise spoken words almost as well as a human, has transformed computer vision, and dramatically improved machine translation. Such improvements in these areas were previously unattainable as they are far too complex to code into algorithms by hand. As we will explore though, while they are powerful, they are not immune to challenges.

What makes AI different?

Emerging technology almost always brings with it important privacy and security considerations. So what makes AI any different?

Well, AI technology brings with it the ability to drastically alter the scope and scale of how we generate, collect, use and share data. At its core, the key point of differentiation is the ability of AI to automate these processes, and to scale them far beyond human capabilities.

We need to remember that while many of our laws and standards for privacy and security are designed to be adaptable, they were formed with human agents in mind and based on the assumption that humans would be the primary handlers and controllers of data. They were not designed to contend with the computational ability of AI that does not conform to traditional ideas of information management. When we consider this, it becomes clear that AI systems do indeed present a challenge to how we apply our established law, standards, and even mentality.

Benefits of AI

Now, I promised balance to this speech, so I want to turn our attention to some of the incredible possibilities that this technology is creating.

The potential applications of artificial intelligence in healthcare are, in my opinion, some of the most exciting and positive uses of this technology. Earlier I mentioned the ability of an AI system to be able to review and translate mammograms to detect cancer at a speed and accuracy rate drastically better than any human professional. This presents enormous opportunity to detect cancer early, while enhancing the ability of our doctors to focus their expertise elsewhere.

This brings me to another point that I think is worth noting here – that artificial intelligence need not replace human intelligence. Rather, it can be used as an immensely powerful tool to automate tasks at a greater speed and accuracy than a human would be able to, and as such enhance the ability of skilled humans to redirect their focus to new and exciting areas.

Another application of AI that I particularly like, is its ability to increase the autonomy and freedoms of people with disability. Things that many of us take for granted, such as being able to listen to this speech, are everyday challenges for those with sight and hearing impairments. But AI technology is making things such as live captioning more affordable and accessible. Another example, taking things a step further, is Microsoft's Seeing AI app, which is designed for the low vision community and uses AI to recognise and narrate the world around its users. These applications of AI enable humans to live more enriched, full lives.

In the public sector, government already uses AI, but certainly stands to benefit from further adoption of these technologies. Some of the immediate beneficial opportunities for the public sector are those where AI can reduce administrative burden, increase efficiency, and help solve resource allocation problems such as answering queries, filling out, searching and drafting documents, routing requests, and translation.

In the longer term, AI has the potential to go beyond enhancing established processes and alter government operations altogether. It will be exciting to see how the public sector adapts to this evolving environment.

And the positive examples don't stop there. There are also many ways in which AI can benefit information privacy and security.

Before I get into some of the challenges that AI poses to information privacy and data security, I want to emphasise that it is not a given that AI must undermine these areas by default.

As an example, there is huge potential for AI technology to be used to assist in searching, filtering and tracking harmful images and content online. For those suffering image-based abuse, or non-consensual sharing of intimate images (colloquially referred to as 'revenge porn'), the privacy enhancing nature of this application of AI is immense. And of course, its potential application in combatting child pornography on the dark web is undeniably a win not only for privacy, but for the safety and wellbeing of children and families around the world.

There are also other instances where AI could enhance privacy rather than stifle it. For example, it may mean that fewer people will need access to raw data containing personal information in order to analyse or use it. This, in turn, could minimise the risk of privacy breaches due to human error – which, I would suggest, is still the most common factor in privacy and security breaches.

Further, AI could have the potential to change the landscape of consent and empower citizens to exercise more meaningful choice. It is not hard to imagine a system that can learn your privacy and consent preferences and apply them dynamically to all the programs and services you subscribe to or receive.

AI is also already being used in the security context. For example, there are tools that use structured and unstructured machine learning techniques to learn and model network behaviour, and then identify abnormal instances to improve threat detection. In this sense, one of the biggest strengths of machine learning in security is to understand what is 'baseline' or 'normal' for a system, and then flagging anything unusual for a human to review.

These kinds of tools mean AI technology is increasingly being used to mitigate the increased volume and sophistication of attacks. According to the 2018 Cisco Annual Cybersecurity report,

"threat researchers expect to see adversaries increase their use of encryption in 2018. To keep pace, defenders will need to incorporate more automation and advanced tools like machine learning and AI to complement threat prevention, detection and remediation."

With the assistance of automation and intelligent tools, security professionals are more equipped to overcome skills and resource gaps, enabling them to be more effective at identifying and responding to both known and emerging threats. Just some of the areas it is being used in include:

- spam and phishing filters,
- network intrusion detection and prevention,
- fraud and botnet detection,

- secure user authentication, and
- hacking incident forecasting.

These AI-based security solutions don’t have to replace the human experts, in fact, it has been found that they can be immensely helpful in proving extra horsepower and automated efficiency to existing technologies and processes.

But, unfortunately, these tools are not exclusive to the good guys – they can also be used by attackers themselves. For example, machine vision has already been used to defeat captchas. Another emerging threat to machine learning security solutions is a process called ‘data poisoning.’ If attackers can determine how an algorithm is set up or where it is drawing its training data from, they can then figure out ways to introduce misleading data, building a counter-narrative about what content or traffic is legitimate versus malicious. By doing this they can skew the algorithm, resulting in truly malicious behaviour not being recognised and therefore acted upon.

Another example that I think demonstrates the two-sided nature to using AI in security is based on some research done by the cloud security firm, Cyxtera. Their researchers built a machine-learning based phishing attack generator that was trained on more than 100 million examples of particularly effective historic attacks. This was used to optimize and automatically generate scam emails and links. They established that an average phishing attack will bypass a detection system that uses AI only 0.3% of the time. But, if you fight fire with fire and use phishing attacks generated through AI, it is able to bypass the system more than 15% of the time.

That is a really significant jump and presents an undeniable risk for many companies using AI-based defence. Something to note here is that the researchers only used data and tools that were open source and available to the general public, meaning anyone with the skills and the malicious intent would be able to build similar systems.

As always, in security, we should be cognisant of the remaining security domains. The term ‘cybersecurity’ is used to describe many different things, but we must be careful to ensure that even when we’re talking about something like AI, we consider all four domains: physical, personnel, ICT, and information security, as well as proper overarching governance. Otherwise, we risk overlooking risk and making ourselves vulnerable.

For instance, earlier I used the example of ‘data poisoning’. There are many ways an attacker might be able to determine an algorithm that is being used, but we would be unwise to discount methods such as social engineering in order to gain this knowledge. As we all know, the human element is often where mistakes can happen, so having robust personnel security remains vital as we move toward increased use of AI.

On this note, recently we have seen how AI can be used to manipulate people on a large scale. I’m talking about the recent Cambridge Analytica scandal which I am sure you have all heard about to some degree. I won’t get stuck into the detail here, but an interesting offshoot of this issue are the potential security implications. As I see it, the threat landscape changes when you put a human up against a machine. If people, arguably through no fault of their own, are able to be manipulated into voting a certain way or buying a certain product as a result of the abuse of their personal information, what’s to say that this approach cannot be leveraged in a security context? Food for thought.

Privacy challenges

Now I’d like to turn our attention to some of the challenges posed to information privacy by AI.

I would like to highlight that from a practical point of view, while privacy and security are different concepts they are inextricably linked – you can’t expect to maintain overall good privacy governance if data security practices are poor.

Modern information privacy as we know it in Australia and many countries around the world is based on a core set of principles and several key concepts. This is not the first time these principles and concepts have been challenged by new technology. As an example, in the late 19th century, the portable camera was considered to be a threat to “the right to be let alone”.

But the potential of AI proposes an information landscape that is unlike anything we’ve experienced before, and prompts us to reflect on our established ideas, to consider whether the privacy status quo remains fit for purpose.

As many of you may know, the core principles that underpin much of privacy legislation around the world are based on the OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data, which were originally drafted in 1980.

Of the eight principles of these Guidelines, there are three which I consider to be particularly challenged by the growth of AI:

- First: The Collection limitation, which proposes that we should only collect personal information that is necessary

- Second: The purpose specification, which says we ought to specify the purpose of collecting the personal information, and

- Third: The use limitation, which states that the personal information collected should only be used for the purpose for which it was collected.

As you can tell, these three principles are fundamental to how we currently understand information privacy, and they are by their nature, very intertwined. Taken together, they have the goals of minimising the amount of information any one organisation holds about citizens and ensuring that the way that information is being handled is consistent with the expectations of that individual.

Now consider this: the very nature of many techniques of artificial intelligence relies on ingesting and learning from vast amounts of data – much of which could be considered to be personal information. This flies in the face of the collection limitation principle. Many organisations using AI may not fully understand how the personal information being ingested by the AI will be used – how then, can you genuinely implement the use limitation principle? And how can you communicate the purpose of collecting the personal information of individuals if its use is unclear?

Transparency and consent

This brings me to my next concern – transparency and consent.

I’ll pose another question: if you, as an individual, don’t fully understand how your personal information is being used and for what purpose by an AI system, how can you exercise meaningful and informed consent? The complexity around AI means that processes are often unclear to those whose information is being used, which makes true consent a challenge.

This idea of consent is intrinsically linked to transparency. If you can't be transparent about processes and decisions, it's very difficult to properly inform people of what is being done with their information, let alone gain meaningful consent.

Many of the AI techniques used today make transparency difficult. For instance, deep learning techniques, which I touched on earlier, work through a process of ingesting data through layers, in which the output of each layer becomes the input for the next layer. After several layers, the logic can become increasingly obscure to the human eye, sometimes even resulting in those who built the algorithms themselves not fully understanding how a conclusion was drawn.

When we put this in the context of decision making within government or the judicial system – if we are unable to explain the reason behind a decision, this poses significant difficulty not only for consent and the three principles I mentioned before – but consider this: who do we hold accountable for these outcomes?

This may seem a little abstract, so here are some examples:

- Determining a prison sentence for criminal offences
- Deciding who does or does not received particular government services

When these decisions are made by humans, we have established fairness and accountability processes. Now, humans are not perfect in their decision making, but if these decisions are made by opaque AI algorithms, who is held accountable for their outcomes?

Accountability + governance

While we're on the topic of accountability, this raises the next challenge: governance.

Isn't it interesting to notice how all of these challenges begin to link together?

Governance and oversight are championed in information privacy law to ensure appropriate structures are in place that prevent a power imbalance between citizens and the state. And I know that governance is also of course championed as a fundamental part of a mature approach to data security.

Good data security is about more than just ICT. You may have the most secure software and technology in place to protect your information, however, if there aren't appropriate overall controls in place to limit access to that information, such as physical and personnel security, then the information may still be compromised.

Good governance frameworks can be used to promote good design, structure and oversight of AI technologies and how they interact with privacy and security.

Personal information

Finally, one of the last points related to privacy that I want to touch on is the idea of 'personal information'. As you may have noticed, a lot of the issues I have raised today relate to how personal information is collected, used, and shared. In fact, most of the information privacy law in place around the world is predicated on this notion of what is and what is not personal information. In this sense, the definition of what constitutes 'personal information' acts as a gatekeeper to the legal protections offered to individuals – if it's not personal information, it doesn't receive privacy protections.

Generally speaking, the concept of personal information is based on the idea of identifiability – on whether or not you can reasonably ascertain the identity of an individual from that information. However, the distinction between what is and is not personal information is being increasingly blurred by new technology.

What is personal information in a world of artificial intelligence? How can we continue to use a binary distinction between what is personal and what is not, when so much of the value of AI is to bring together seemingly disparate data, and to recognise and identify patterns unseen to the human eye?

This raises important questions about whether or not distinguishing between personal information and information more broadly will continue to be legally or technically practical, and beyond that, whether it holds individuals' best interests at heart when protecting their privacy.

Human computer interaction

Now, before I wind this up, there is one other aspect of artificial intelligence that I think is fascinating, but not yet talked about all that often. And that is the relationship aspect on AI. Some of you may have heard of the notion of 'Human Computer Interaction'? It is a well-established field of study and design. And yet, some experts in the field are beginning to question whether the term 'interaction' really captures the full essence of where we're headed as AI continues to evolve.

At the Human Rights and Technology Conference put on by the Human Rights Commission several weeks ago, Distinguished Professor Genevieve Bell posed the question: what will our relationships with artificial intelligence look like?

This notion of relationships, rather than 'interaction' is striking. Research from MIT has shown that humans have a tendency to project life-like qualities onto machines. Beyond this, humans have a knack for anthropomorphising all kinds of objects – we name our cars, bikes and boats; we have conversations with Siri and Alexa; we instil personality into our houseplants and form sentimental attachments to all kinds of intimate objects of significance to us.

So, if we already have this tendency, what does it mean when we consider artificial intelligence that is built to replicate human behaviour? Many of you may have seen the Google Duplex demonstration in which a virtual assistant made a booking at a hair salon over the phone. The human on the other end of the phone had no idea that they were speaking to a machine. What was particularly striking was that the machine was programmed to include human-like filler words into their dialogue like 'umm' and 'ahh' to replicate what an actual human would sound like.

So, what is the point of me bringing this up here today?

Well, when we consider the relationship between humans and intelligent machines, it raises unique questions regarding privacy and security.

I already touched on the notion that this may create an uneven playing field when humans come up against machines in relation to security, but let's look at this idea of relationships in another way.

Consider for a moment, the idea of humans forming trusting relationships with their virtual personal assistant. Your virtual personal assistant that, through the power of the Internet of Things, is connected to your fridge, your toaster, your heating system… maybe even your garage door or your house security system. Imagine that maybe you've named your virtual assistant and have instilled upon it a personality. Imagine that the virtual assistant, through collecting and interpreting massive amounts of your data, has learnt your preferences, understands you on an intimate level, and is able to adapt itself to suit your unique wants and needs.

Now, in this scenario, it is undeniable that the human would be more likely to share personal information with this system. We would be comfortable in doing so. The potential for immensely privacy-invasive systems designed to make you feel comfortable, based on the knowledge that humans have a tendency to anthropomorphise, is immense.

We have already seen the way that human cognitive bias can be tapped into and exploited in the way many apps have been designed. As a classic example, many social media platforms are designed to replicate the addictive nature of casinos, with flashy rewards (the like button) to keep us hooked. We would be rash to think that the design of AI products is immune to this kind of design thinking.

Now consider the security risks of this scenario. In a connected world, our security is only as good as the weakest element. That smart toaster probably doesn't have robust security protections, and now it's linked to the automated security system for your house.

This may seem like a dystopian nightmare – but all of the technology is there. And this is why we need to consider the privacy and security considerations from the outset – and to be creative in how we think about challenges and their solutions when it comes to both protective security and information privacy.

Concluding remarks

I like to think of privacy as a public good, similar to clean air or safe drinking water. I also like to think about secure systems in a similar way. If you think about privacy and security as public goods, we all suffer when government information is not secure, or if we allow broad access to our personal information, just as we all suffer if some of us choose to pollute the water supply.

We need to shift the way we think about privacy from a purely individualistic or transactional mind-set, to one that is more collective and long-term. We need to shift the way we think about security away from something that is the focus of a few people in the office, and toward something that everyone takes some responsibility for.

Our aim for data protection is to work to facilitate the ongoing improvement of the data security practices of Victorian government agencies, rather than striving for perfection. The nature of data protection and the constantly evolving threat environment means that encouraging agencies to take a risk-based approach to data security will support businesses achieving their objectives, in a manner commensurate with their resources. And also note that I refer to improving practice, not just systems or technology.

Let me finish by saying that it has been a pleasure to come here and talk to you all today. I want to conclude by leaving you with three key messages:

1) Artificial intelligence is not neutral. We often hear that "technology is neither good nor bad, it's just the way you use it." But even with the best of intentions, systematic societal bias and discrimination is difficult to prevent from being built in to intelligent systems. AI is not neutral because humans are not neutral. Information management is of huge importance in this space to protect against the misuse of data.

2) AI technology has huge potential in the security space, but this technology is not exclusive to the good guys. We need to be agile in the way we approach the new opportunities, as well as the new threats.

3) As government organisations, everything we do in this space is ultimately about people. We must never lose sight of that.

Thank you for being such an engaged audience. If you are interested in reading more about artificial intelligence and privacy, we have published an issues paper which is available on our website, as well as plenty of other guidance material covering security, privacy and freedom of information. If you like to use Twitter, I encourage you to follow us @OVIC_AU, where we regularly post new resources and events as they come up.

We're reaching out, so look us up and get in touch.

Thank you.