

Artificial Intelligence: governance and leadership

Graham Jefferson¹

NTT Communications Cloud Infrastructure Services²

0. Introduction

"Not until a machine could write a sonnet or compose a concerto because of thoughts and emotions felt, and not by the chance fall of symbols, could we say that machine equals brain - that is, not only write it but know that it had written it."

G Jefferson, "The Mind of Mechanical Man", (1949) *British Medical Journal* 1105, 1110.

"One threat in particular is worthy of further consideration: that ultraintelligent machines might lead to a future that is very different from today - we may not like it, and at that point we may not have a choice. Such considerations lead inevitably to the conclusion that we must weigh carefully, and soon, the possible consequences of AI research."

S Russell and P Norvig (Eds) "Artificial Intelligence: A Modern Approach" 3rd Edition (2007) Pearson p. 876

For decades the early promise of AI was not matched by results. There is now evidence of impressive developments in AI³ and perhaps we have reached an inflection point. Certainly, the pace at which AI is advancing is increasing.⁴ It is arguable that the rate of change is so fast that it is now almost impossible for traditional legal and regulatory models to keep up and the questionable, if not illegal, behaviour of some dominant players in the IT and AI space⁵ raise questions about how modern technology (including AI) is being employed. On the other side of the ledger are wonderful examples of how AI can be used to deliver 'social good'.⁶

NTT Communications Cloud Infrastructure Services Inc. (NTTCom CIS) is pleased to support the AHRC's investigation into what should be done to promote responsible development of AI and to respond to the main question of the White Paper "Does Australia need an organisation to take a central role in promoting responsible innovation in AI and related technology?"

This submission is in two parts: (1) a series of general observations on some of the themes raised by the White Paper; and (2) answers to some only the questions set out in the White Paper.

1. General observations

Trust and confidence in AI must be maintained

As the public becomes more aware of the risks (and supposed risks) associated with AI it is necessary that public trust and confidence in AI and all IT is maintained. NTTCom CIS sees the potential for an innovative body, a Responsible Innovation Organisation (RIO), to lead this work.

One approach to maintaining trust and confidence in AI is to make sure that 'fundamental ethical principles', 'social good' and 'public good' are the foundations upon which AI engineers build their structures. These terms are somewhat elusive and as the topic is important it demands clarity of purpose. For that reason, anchoring discussions

to more specific and widely understood principles such as universally acknowledged human rights is sound.⁷ NTTCom CIS supports the establishment of a RIO with the aims of: (1) promoting human rights through the responsible uses of AI; and (2) protecting human rights through the control of irresponsible uses of AI.

It is hard to trust something when you don't understand it and IT is notoriously complicated. The opacity of Google and Facebook's current business models and algorithms generate concern⁸ and it is inevitable that as AI evolves it will become increasingly difficult to understand. Some data protection regimes prescribe how automated decision making is explained⁹ to data subjects and already Machine Learning solutions create challenges with 'interpretability'¹⁰ or the process of explaining how the predictions from a particular algorithm or model can be explained explanation. The White Paper flags this issue in the context of bias risks but in simple terms, if it is not possible to understand the basis on which an AI algorithm operates should we be comfortable using it?¹¹ In some contexts, for example those more tightly connected with what it is that makes us a functioning society, the answer must be no. The Royal Australian and New Zealand College of Radiologists has developed a draft set of eight ethical principles governing the adoption of AI in medicine which include a requirement that the "When designing a ML system or AI tool, consideration must be given to how the decision made can be understood and explained by a discerning medical practitioner"¹²

There is a further aspect to this theme of interpretability. Many current oversight and regulatory regimes depend on notice and consent models and these regimes are materially less effective when information asymmetry is high. For the reasons set out above, AI presents real challenges on this front. How can citizens consent when they don't understand? Developing the expertise to understand these issues and craft suitable responses requires significant time and effort. The RIO acting as a centre of excellence in these matters, coordinating responses and investigations and promoting best practices would perform a valuable service to government, existing regulatory bodies and the community at large. At a minimum, information asymmetry might be managed better and the domains where 'algorithmic transparency' is needed could be identified.

Enforcement and compliance considerations must also drive the design of regulatory models. Data protection and privacy laws are ubiquitous, widely publicised and yet compliance is still patchy. Indeed, some recent public disclosures of non-compliance with privacy and data protection laws reveal that the underlying causes have little to do with ignorance about the regulatory regimes so much as resistance to compliance with those regimes.¹³ What does this mean for the promotion of human rights in connection with AI? A company (or software developer) currently willing to ignore privacy laws might also be willing to ignore human rights.

The actions of nation states employing AI solutions must also be considered. NTTCom CIS sees governments leading by example as an important factor in maintaining public trust and confidence. The 'moral high ground' can't be maintained when government does not meet the very standards it seeks to impose on industry and a RIO with suitable authority could play a role in ensuring transparency with government implementations of AI. We support a RIO being "an authoritative voice across government and industry".

The picture is not exclusively bleak. There are examples of formidable industry participants adopting human rights centric approaches,¹⁴ taking steps to enforce good practices against their peers,¹⁵ and even calling for regulation.¹⁶

Nevertheless, these approaches will be inevitably ad hoc and lack the objectivity that a RIO could bring to the exercise. Further, there are arguments that current 'industry led' behaviour can be partly explained by self-interest e.g. a core thesis of 'The Master Switch'¹⁷ is that regulation of emerging technology monopolies (in some cases in settlement of anti-trust action) is a part of the natural evolution of those monopolies and might even be attractive to the monopolists. An independent RIO could do much to balance all stakeholder interests.

Use of AI in surveillance technology deserves immediate attention

Innovations in technology (ubiquitous CCTV, fast networks, cheap storage and AI) make ubiquitous surveillance a reality and this raises fundamental questions for citizens - "[W]ith artificial intelligence on the verge of becoming the next force in big data, everyone should be concerned about what some have called 'surveillance capitalism.'"¹⁸

NTTCom CIS agrees generally with the view that applications of AI to surveillance raise immediate questions as the technology is both widespread and subject to weaknesses that have the potential to erode public confidence. Assumptions that existing protections under e.g. Privacy laws will be sufficient may not be sound. For example, the proposed Commonwealth Identity Matching Services Bill 2018 has attracted criticism including from the Law Council of Australia which "remains concerned that the Bill operates to create a new category of data and will erode current protections provided by the Privacy Act."¹⁹

Australian research indicates that public willingness to accept surveillance is related to perceptions of the necessity of the surveillance (i.e. threats of terrorism, questions of general safety) and to more general trust of government.²⁰ It is fair to say that trust and confidence in government is an issue in Australia at the moment and mass surveillance driven by AI has the potential to fuel fear. It should not be forgotten that Bentham's panopticon²¹ was a prison based on the theory that compliance can be improved when there is a perception of ubiquitous surveillance. Applying this approach, the potential harm is done to a society when there is a widely held belief that ubiquitous surveillance has been achieved whether that is the case or not. AI lends credence to these views and the risks must be greater when there are examples of nation states actually using mass surveillance techniques to enforce compliance.²²

NTTCom CIS is concerned that insufficient attention to the ethical and human rights issues raised by roll-out of early stage surveillance AI could undermine public confidence in AI. The White Paper talks of pilot domains for a RIO (White Paper s.5.4) and we see surveillance AI as a suitable field for this approach.

The laws and practices dealing with large data sets can be aligned to promote human rights

NTTCom CIS supports the view expressed in the White Paper that it is important to consider "ownership of big data and related concerns about the concentration of ownership among a small number of private sector entities."

If we accept the propositions that information and data will become increasingly valuable in the 21st century and that value will flow from information and data in ways not previously contemplated, then it seems strange to manage that value using legal doctrines that are no longer fit for purpose. Information itself cannot be "owned" (as opposed to the copyright, database rights and other statutory rights currently in use to protect information) and the doctrinal issues of information not having a proprietary flavour have been around for a while:

“But more generally, there is the observation that the practice of clinging to an outmoded concept – and confidentiality, as the sole basis for information protection, increasingly appears to be just that – can only be harmful in the long run. If the courts are to fashion remedies to protect information against more than broken confidences and if they are to confront the important policy issues involved, an open acknowledgement of the true basis for those remedies seems indispensable.”²³

The White Paper identifies that community concerns in this area have to date largely focused on personal information, but could this be because privacy and data protection is one area where the laws do appear to offer some recourse? Where large data sets include personal information, then privacy and data protection laws create specific rights for data subjects (the person in question) and there have been some newsworthy examples of individuals using these laws to effect.²⁴ But these rights fall short of ownership. Further, the mechanisms currently in use to provide for collection and processing of personal information are cumbersome and sometimes ineffective.²⁵ Matters become more complicated when other sector specific regimes must be considered e.g. health care legislation where medical records are involved, banking secrecy and related regimes in the case of financial services data sets.

The evolution of the law dealing with ownership of computer software is an instructive comparison. Confusion about how software was to be treated persisted for the latter half of the 20th century before amendments to copyright legislation resolved things and it is now almost universally accepted that copyright ownership (as a literary work) is the mechanism by which legal rights in software are managed.²⁶ Database rights exist but a review would determine the extent to which AI related developments justify reform. This is a complicated and necessarily cross-disciplinary enquiry that lends itself to coordination by a suitably innovative RIO.

Few organisations have the means to collect the raw data needed to build large data sets and the way that operators in the field go about collecting, curating and cross licensing data also merits investigation. Building on its experience from the open source software movement, the Linux Foundation has created an open source data licensing agreement e.g. the Community Data Licence Agreement.²⁷ The aim is that the CDLA can be used to licence data sets including those employed in connection with AI. Currently, the CDLA is silent as to what purposes the underlying licensed data set may be used for. If the CDLA was amended, as a result of action by a RIO for example, to impose a requirement that the licensee only use the underlying data for purposes that ultimately promoted human rights (or something similar) it is possible that very quickly (1) data scientists would become aware of the human rights centric altruistic use restriction and associated issues; and (2) the altruistic use restriction would ‘spread’ to large amounts of data. Putting to one side questions of enforcement, the introduction of such an approach would help to enshrine human rights in AI and Big Data practices.

The role of government and its approach to public data sets will be influential. Why is there not a requirement that all use of public data sets be aligned to the promotion or protection of human rights? A similar approach to that described in connection with the CDLA could be taken with data controlled by the Australian Bureau of Statistics. If a human rights centric altruistic use restriction was introduced into the licence²⁸ currently used for licensing of ABS data to Secondary Distributors the Commonwealth would send a positive message about its support of human rights in the context of Big Data and AI. On the back of recent technology issues associated with the national census such a development might be well received.

The White Paper speaks of the risks of 'strategic advantage' arising from the concentration in a few hands of very large amounts of valuable data. While this is a contentious area and would benefit from detailed analysis by economists, there must be risks of anti-social monopolistic practices arising more easily when there is concentration or vertical integration of AI solutions and large data sets. One can imagine that it is easier for an operator to exercise control over a market when that operator tightly integrates development of AI services with collection and manipulation of data sets. Where the underlying processes are deeply technical and hard to understand it is reasonable to assume that community concerns will be greater. It is not clear to NTTCom CIS how ownership of data can be 'democratic' but public perception that a small number of large companies control too much information might lead to calls to force separation of (1) ownership and control of data sets (and particularly large data sets used in AI implementations) from (2) ownership and control of AI services delivered using those data sets. A similar approach was taken in connection with concentration of media ownership in Australia.²⁹ More relevant to AI, the US government anti-trust litigation against Microsoft in the late 1990s arose out of allegations that Microsoft's 'bundling' of operating systems and internet browsers conferred an unfair advantage and was used to undermine the Netscape competitor offering.³⁰

We must ensure that computer scientists and data scientists have the necessary ethical skills and training

An obvious means to promote AI solutions that incorporate human rights and other ethical considerations is to train the developers in ethics and human rights. Some branches of the sciences have strong ties to philosophy and ethics with education and practical training incorporating ethical teachings. The most obvious example is medicine and its long association with ethical practices dating back to ancient Greece and the 'Hippocratic Oath'.³¹ In general terms mathematics, computer science and data science are fields with relatively little foundation in ethics or less than is now needed.³² Including ethical training in university engineering courses and professional certification programs and even requiring AI developers to undertake some sort of continuing education with an ethical bent could be considered.

Large IT companies are taking steps to adopt guiding principles intended to ensure that AI related works are developed ethically³³ or in conformity human rights but more is probably needed. Building stronger connections between academia and industry is an approach that NTTCom CIS advocates and would support. The most effective AI implementations will solve business problems and the best new ideas are likely to come from universities and other academic institutions. Australia has a long history of making major scientific contributions, if not commercially exploiting them, and the development of world class ethical AI engineers would be yet another example of this contribution. There are already some very promising examples of this sort of work³⁴ but one challenge for Australian academic institutions is that the local IT market is small and the largest players in this space will likely be international.

A Responsible Innovation Organisation need not be a new regulatory body in order to be effective

Perceptions across business tend to be that there is too much regulation and that regulation inhibits rather than fosters innovation. Further, recent experience in Australia in connection with laws requiring retention of meta-data³⁵ and access to encrypted messages³⁶ has done little to build confidence in the ability and capacity of the various organs of government to navigate challenging policy debates involving information technology.

Credible observers have argued that AI has the potential to affect society in ways not previously imagined.³⁷ That said, it is difficult to conceive of a situation where AI is regulated in the same way as medicines and drugs or even software currently used in connection with avionics systems³⁸ i.e. there will be a regulatory system operating whereby an AI could not be released or used until it had been tested, verified or approved. The White Paper does not appear to contemplate this approach.

Australia is an island but cannot afford consider itself as such. Any effective oversight of responsible AI development must acknowledge the realities of attempting regulation in a global economy. In a purely Australian context it is likely that many domestic AI instances will actually be local versions of some larger offering designed and developed by a multi-national corporate with the dominant players most likely located in the USA and PRC³⁹. For this reason, it is encouraging that the AHRC is partnering with DFAT and working with a trans-national body like the WEF.

While NTTCom CIS supports the idea of a RIO it is not persuaded that the RIO needs to be a regulator. Rather than being a sui generis AI regulator NTTCom CIS sees the RIO acting as an interface between existing regulators, engaging with government, academia and industry, both domestically and internationally, identifying good practices and working to promote and implement them. This would necessarily be a largely educative and policy-based undertaking (as opposed a more traditional regulatory function) but the anchoring of such work in the promotion of human rights provides a solid foundation for what might otherwise be a nebulous mandate.

2. Specific responses to questions

1. What should be the main goals of government regulation in the area of artificial intelligence?

NTTCom CIS does not advocate regulation of AI per se as opposed to the effective coordination of regulatory responses to domains affected particularly by developments in AI. In that context, the main goals of government regulation of AI should be:

- Maintenance of public trust and confidence in IT generally through advancing the idea that AI implementations promote public good and advance human rights
- Aligning AI implementations to existing legal and regulatory obligations
- Promoting best practices in connection with development and implementation of AI

2. Considering how artificial intelligence is currently regulated and influenced in Australia:

(a) What existing bodies play an important role in this area?

No submission.

(b) What are the gaps in the current regulatory system?

Two trends associated with AI present challenges for Australia's regulatory regime.

The speed at which technology is developing makes it difficult for regulatory bodies to identify issues, consult and act sufficiently quickly. This challenge will only get worse as the rate of change increases and AI implementations increasingly encroach on domains affecting more and more people.

The second trend is the trans-national aspect of AI which creates challenges for domestic regulatory bodies. As trans-national corporations increasingly dominate the development of AI an appropriate response is to identify suitable partners and to work cooperatively with industry. This is a slightly different approach to historical regulatory paradigms and what is proposed is to find some way to embed corporate players in the pursuit of the RIO's objectives.

3. Would there be significant economic and/or social value for Australia in establishing a Responsible Innovation Organisation?

This depends and turns very much on: (1) what the RIO is tasked to achieve; (2) how the RIO is implemented; and (3) how AI implementations actually occur.

At a superficial level, the promotion of increased trust and confidence in AI should improve its adoption and there are numerous assessments and predictions about the economic benefits of AI. An alternative view is that regulation inhibits economic growth with the result that a RIO might restrict adoption of AI in Australia and drive away businesses that might otherwise invest in AI in Australia.

Balancing these competing considerations is art not science and CIS supports the AHRC's approach to this task in anchoring the debate in human rights.

4. Under what circumstances would a Responsible Innovation Organisation add value to your organisation directly?

If the RIO had the effect of increasing trust and confidence in AI and IT more generally it would be advantageous to NTTCom CIS. Our business assumes that end users and consumers have confidence in the products and services we build and sell. Further, the RIO could act as a 'clearing house' for industry best practices, development of standards and generally promote industry wide investment in good technology practices.

5. How should the business case for a Responsible Innovation Organisation be measured?

The business case must obviously attempt to measure the relationship between the investment in the RIO and the achievement of its ends with a particular emphasis on measuring how human rights have been advanced. Many companies use a balanced scorecard⁴⁰ to measure financial and non-financial strategic performance and the RIO might adopt this approach. Under this approach the various objectives (question 1) and stakeholder interests would form the basis of that balanced scorecard. Metrics and KPIs would be developed to measure the RIO's performance on each of the elements of the scorecard. For objectives such as promotion of trust and confidence, surveys and longitudinal studies might be the only effective metrics.⁴¹

6. If Australia had a Responsible Innovation Organisation:

(a) What should be its overarching vision and core aims?

The ROI should have two complementary aims: (1) the promotion of human rights through the responsible uses of AI; and (2) the protection of human rights through the control of irresponsible uses of AI.

(b) What powers and functions should it have?

No submission.

(c) How should it be structured?

No submission.

(d) What internal and external expertise should it have at its disposal?

The peak governance body for the RIO must be diverse⁴² and this means a healthy representation of non-technical people. In addition to the usual cast of lawyers and economists one can see philosophers playing an important role. That said, the RIO must be sufficiently technically competent to analyse developments in AI and respond rapidly. Drawing on academic and industry expertise and using subcommittees and expert working groups is a simple way to manage this requirement.

(e) How should it interact with other bodies with similar responsibilities?

See general observations about the RIO acting as an interface between other bodies.

(f) How should its activities be resourced? Would it be jointly funded by government and industry? How would its independence be secured?

No submission.

(g) How should it be evaluated and monitored? How should it report its activities?

A theme in this response is that a central benefit of the RIO would be increased trust and confidence in AI. This necessitates transparency i.e. detailed evaluation and monitoring which is publicly available. If a balanced scorecard approach (see Question 5) is adopted then performance against that scorecard would also be the basis for monitoring and reporting.

¹ CEO, NTT Communications Cloud Infrastructure Services. Mr Jefferson graduated from the ANU with an LLB (Hons.) and BSc (Computer Science). Before his current role, Graham practised in-house as an IT lawyer, principally in financial services institutions, and before that he worked as a software developer including at the Commonwealth Attorney-General's Department. He was a founding member of the London Investment Bank's Technology Discussion Group, a representative of the British Bankers Association eCommerce Working Group and more recently was a member of the NSW Law Society Legal Technology Sub-committee. Graham has spoken and published widely on the intersections of IT, law and ethics and his honours thesis was an early assessment of the ethical and practical considerations arising from applications of AI to the law – see G Jefferson "A Brief History of Artificial Intelligence" (1991) 17 Computers and Law 17 (<http://www.austlii.edu.au/au/journals/ANZCompuLawJl/1991/15.pdf>)

² NTT Communications Cloud Infrastructure Services Inc. is a wholly owned subsidiary of Nippon Telegraph and Telephone Inc., (NTT) the Japanese IT and telecommunications company. We offer cloud computing infrastructure services to customers around the world and many of our customers operate technologies that fall within the notion of 'Artificial Intelligence and related technology'. NTT is itself engaged in significant research and development activities in some cases directed at advancing and implementing AI. In accordance with the UN Guiding Principles on Business and Human Rights, NTT Communications Cloud Infrastructure Services has adopted a Human Rights Charter directed to meeting our responsibilities to respect human rights and ensuring the advancement of human rights more generally.

³ Performance in chess is an interesting measure as it has been the subject of analysis of AI performance for some time or as Garry Kasparov says in the context of analysing the performance of Google's AlphaZero deep learning AI "Chess has been used as a Rosetta Stone of both human and machine cognition for over a century. AlphaZero renews the remarkable connection between an ancient board game and cutting-edge science by doing something extraordinary." (<https://deepmind.com/blog/alphazero-shedding-new-light-grand-games-chess-shogi-and-go/>)

⁴ There are numerous ways to consider this question. A simple approach is to consider the amount of information produced by humanity's collective efforts. Overwhelmingly this is now digital information. One source posits that the first 12 Exabytes of data took humanity ~300,000 years to produce and the next 12 Exabytes will be produced in less than 12 months (<https://www.sims.berkeley.edu/how-much-info/index.html>). A more thoroughgoing explanation of why this phenomenon might be so is Ray Kurzweill's 'Law of Accelerating Returns' see R Kurzweill "The Singularity is Near" (2005) Penguin especially Chapter 2 "A Theory of Technology Evolution: The Law of Accelerating Returns"

⁵ There are many but one of the most disturbing must be the 2012 'experiment' Facebook secretly conducted in which it altered the news feeds of nearly 700,000 users, showing them an abnormally low number of either positive or negative posts in order to determine whether the company could alter the emotional state of its users. See

"Facebook sorry – almost – for secret psychological experiment on users" The Guardian 2 October 2014 (<https://www.theguardian.com/technology/2014/oct/02/facebook-sorry-secret-psychological-experiment-users>)

⁶ "Notes from the AI Frontier: Applying AI for Social Good" McKinsey Global Institute Discussion Paper December 2018

⁷ "If AI developers treat privacy as a fundamental human right rather than an ethical preference, the privacy considerations that already exist in industry norms and technical standards would be stronger." M Latonero "Governing Artificial Intelligence: upholding human rights and dignity" Data & Society (<https://datasociety.net/output/governing-artificial-intelligence/>) p.14

⁸ For further commentary on the ACCC preliminary report into Google and Facebook see R Gittins "Who pays for the digital giants' free lunch?" Sydney Morning Herald 2-3 March 2019

⁹ See the various discussions on whether under European data protection laws individuals subject to automated decision making processes are entitled to explanations of those decisions and what this means in the context of complex algorithms - L Moerel and M Storm "Automated Decisions Based on Profiling: Information, Explanation or Justification –that's the Question!" (<https://www.scl.org/articles/10247-automated-decisions-based-on-profiling-information-explanation-or-justification-that-s-the-question>); L Edwards, and M Veale, "Enslaving the Algorithm: From a 'Right to an Explanation' to a 'Right to Better Decisions'?" (2018). IEEE Security & Privacy (2018) 16(3), pp. 46-54;

¹⁰ See "Machine Learning: The Power and Promise of Computers That Learn by Example" The Royal Society (April 2017) (<https://royalsociety.org/~media/policy/projects/machine-learning/publications/machine-learning-report.pdf>) pp 93-94

¹¹ New York City Council legislation introduced in 2017 was in part directed at ensuring transparency in connection with the workings of algorithms supporting decision making in the allocation of public funds – see J Powles "New York City's bold, flawed attempt to make algorithms accountable New Yorker Magazine (<https://www.newyorker.com/tech/annals-of-technology/new-york-citys-bold-flawed-attempt-to-make-algorithms-accountable>)

¹² See "Ethical Principles for AI in Medicine" Consultation Paper issued by the Royal Australian and New Zealand College of Radiologists (<https://www.ranzcr.com/our-work/advocacy/position-statements-and-submissions/ranzcr-ethical-principles-for-ai-in-medicine-consultation>)

¹³ It will be interesting to see what conclusions the ACCC makes in its final report into Google, Facebook and Australian news and advertising but the preliminary report indicated that there are issues with compliance with existing privacy laws and especially the terms and conditions used to manage compliance with those laws. See (<https://www.accc.gov.au/media-release/accc-releases-preliminary-report-into-google-facebook-and-australian-news-and-advertising>)

¹⁴ S Pichai, "AI at Google: Our Principles," The Keyword, Google, June 7, 2018, (<https://blog.google/technology/ai/ai-principles/>)

¹⁵ For example: Apple's decision to restrict Facebook and Google's use of Apple Enterprise Developer environments following alleged breaches of Apple's approved development principles regarding privacy and tracking of customer activity https://www.washingtonpost.com/technology/2019/01/31/apple-revokes-googles-ability-use-internal-ios-apps-just-like-facebook/?noredirect=on&utm_term=.5b9c7256993f

¹⁶ The White Paper cites several examples of calls for regulation from leading IT companies. See also "Big Tech needs to act on concerns over 'surveillance capitalism'" John Chen, CEO, Blackberry, Letter to the Editor Financial Times 30 January 2019

-
- ¹⁷ T Wu "The Master Switch: the Rise and Fall of Information Empires" (2010) Atlantic Books
- ¹⁸ "Big Tech needs to act on concerns over 'surveillance capitalism'" John Chen, CEO, Blackberry, Letter to the Editor Financial Times 30 January 2019. Microsoft's General Counsel, Brad Smith, has also warned of the potentially harmful consequences of mass surveillance technologies powered by AI engines. See (<https://blogs.microsoft.com/on-the-issues/2018/12/17/six-principles-to-guide-microsofts-facial-recognition-work/>)
- ¹⁹ <https://www.lawcouncil.asn.au/media/news/review-of-the-identity-matching-services-bill-2018>
- ²⁰ See A Bunn and N Thompson "Australians accept government surveillance, for now" The Conversation 5 February 2019 (<https://theconversation.com/australians-accept-government-surveillance-for-now-110789>)
- ²¹ "What does the panopticon mean in the age of digital surveillance?" The Guardian 23 July 2015 (<https://www.theguardian.com/technology/2015/jul/23/panopticon-digital-surveillance-jeremy-bentham>)
- ²² "Does China's digital police state have echoes in the West?" The Economist May 3 2018 (<https://www.economist.com/leaders/2018/05/31/does-chinas-digital-police-state-have-echoes-in-the-west>)
- ²³ J McKeough, A Stewart and P Griffith "Intellectual Property in Australia" 3rd Ed. (2004) LexisNexis Butterworths p. 83
- ²⁴ The actions of Mr Schremm in enforcing his data protection rights against Facebook ultimately led to the collapse of the US – EU Safe Harbour regime.
- ²⁵ See (<https://www.accc.gov.au/media-release/accc-releases-preliminary-report-into-google-facebook-and-australian-news-and-advertising>)
- ²⁶ See generally the discussion in J McKeough, A Stewart and P Griffith "Intellectual Property in Australia" 3rd Ed. (2004) LexisNexis Butterworths pp. 165 - 167
- ²⁷ <https://cdla.io>
- ²⁸ Creative Commons Attribution 2.5 Australia Licence (<https://creativecommons.org/licenses/by/2.5/au/>)
- ²⁹ NTTCom CIS does not know if the ACCC review of Google and Facebook extends into this area but clearly there are considerations that would ultimately fall within the ACCC's current jurisdiction.
- ³⁰ "What the Microsoft Anti-trust case Taught Us" New York Times 18 May 2018 (<https://www.nytimes.com/2018/05/18/opinion/microsoft-antitrust-case.html>)
- ³¹ For an excellent discussion of this topic in the context of AI see "Ethical Principles for AI in Medicine" Consultation Paper issued by the Royal Australian and New Zealand College of Radiologists (<https://www.ranzcr.com/our-work/advocacy/position-statements-and-submissions/ranzcr-ethical-principles-for-ai-in-medicine-consultation>)
- ³² "Currently, we're naively assuming that the AI coders and developers have the ethical knowledge, understanding and skills to navigate the challenges that their technological innovations create". S Longstaff "How will we teach intelligent machines to behave themselves?" AFR 4 January 2019 (<https://www.afr.com/news/economy/how-will-we-teach-the-robots-to-behave-themselves-20190104-h19pi3>)
- ³³ S Pichai, "AI at Google: Our Principles," The Keyword, Google, June 7, 2018, (<https://blog.google/technology/ai/ai-principles/>)
- ³⁴ For example the ANU has recently established such a body <https://3ainstitute.cecs.anu.edu.au>
- ³⁵ 2015 amendments to the Telecommunications (interception and Access) Act 1979 (Cth)
- ³⁶ See A Bogle "'Outlandish' encryption laws leave Australian tech industry angry and confused" (<https://www.abc.net.au/news/science/2018-12-07/encryption-bill-australian-technology-industry-fuming-mad/10589962>)

³⁷ Stephen Hawking's pessimistic observations about AI were widely publicised (https://www.washingtonpost.com/news/speaking-of-science/wp/2014/12/02/stephen-hawking-just-got-an-artificial-intelligence-upgrade-but-still-thinks-it-could-bring-an-end-to-mankind/?utm_term=.e1d8a28ccea9)

³⁸ The FAA requires software used in avionics systems to be developed in ways that minimise risks of accidents and other harm see https://www.faa.gov/documentlibrary/media/advisory_circular/ac_20-115d.pdf

³⁹ Europe is struggling to keep pace with AI developments coming out of the US and the PRC – see “Tackling Europe’s gap in digital and AI” McKinsey Global Institute (<https://www.mckinsey.com/featured-insights/artificial-intelligence/tackling-europes-gap-in-digital-and-ai>)

⁴⁰ As a performance management tool the approach is versatile and can be applied in many different ways but see generally R Kaplan and D Norton “The Balanced Scorecard: measures that drive performance” Harvard Business Review July-August 2005 (<https://hbr.org/2005/07/the-balanced-scorecard-measures-that-drive-performance>)

⁴¹ Given AI has the potential to affect Australian society so widely, inclusion of a census question might one day be justified.

⁴² The merits of diversity in governance contexts is now unarguable but see also P Tetlock and D Gardner “Superforecasting: the Art and Science of Prediction” (2015) Penguin Random House and S Page “The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools and Societies” (2007) Princeton University Press