



Microsoft Submission to the Australian Human Rights Commission paper to Artificial Intelligence: Governance and Leadership discussion paper

Microsoft welcomes the opportunity to make a submission to the Australian Human Rights Commission (AHRC) whitepaper titled 'Artificial Intelligence: Governance and Leadership.'

Microsoft has published a book how we view the role of Artificial Intelligence in society – *The Future Computed*, and we believe that read in conjunction with this submission – provides a comprehensive look at Microsoft's approach to AI.

Artificial Intelligence or AI – a set of technologies that enable computers to perceive, learn, reason and assist in decision-making to solve problems in ways that are similar to what people do – is not new. However, progress has accelerated over the past few years in large part to the increased availability of data; growing cloud computing power and more powerful algorithms developed by AI researchers. With our lives increasingly digitized and the ubiquity of sensors, there is more data available for computers to learn from and it is this that enables them to discern the often subtle patterns. Analysing all this data requires massive computing power which is available today thanks to the efficiencies of cloud computing.

AI is in many ways, still a developing technology. Most of the progress so far has been in teaching computers to perform narrow tasks. We have a long way to go before computers adopt general intelligence.

Microsoft, like the AHRC, believes that AI has the potential to drive widespread economic and social progress. The key to attaining these benefits is to develop AI in a way that is human-centred. We aim to develop AI to augment human abilities.

For example, computers are very good at remembering things, they are good at probability, they are good at discerning patterns in data that are too subtle for people to notice. With these capabilities, computers can help people make better decisions. And this leads to faster and more profound progress in nearly every field of human endeavor.

More importantly it has the potential to help society to overcome some of its most complex and pressing problems; such as reducing poverty, improving education, eradicating diseases, addressing sustainability challenges, solving food shortages and advancing inclusion.

However, Microsoft also recognizes, that as with the great advances of the past – such as steam and electricity – AI will bring about great challenges and we will need to be thoughtful about how we

address the societal issues that arise. We will need to work together to ensure that AI is developed in a responsible manner so that people will trust it. It will require a shared understanding of the ethical and societal implications of these new technologies.

Whilst we acknowledge value in looking at the concept of responsible or ethical AI through a Human Rights lens or framework – we do have concerns with the number of separate and apparently disjointed efforts already being undertaken in Australia.

In contrast to countries like the UK and in Singapore where there is one concerted Government initiative – Australia currently has several streams of work being led by Government agencies exploring an ethical framework around AI, and we are concerned that efforts to develop an ethical framework seem to be synonymous with calls for more and new regulation.

There is the Data 61 work announced in the 2018 Budget to develop an Ethics Framework for AI. There is also work commissioned by the Chief Scientist for the Australian Council of Learned Academies (ACOLA) group under the Scanning Horizons auspice to look at the social, economic, and legal/ethical elements of AI. Standards Australia was funded in the 2018 budget to develop a standards framework for ethical AI, the Department of Industry, Innovation and Science is delivering on a co-operative research centre and chairing the inter-Departmental Committee on AI, and then there is the AHRC three year tech and human rights project of which this paper is part. In addition the Australian Competition and Consumer Commission (ACCC) paper has flagged potential regulation over the AI systems used by digital platforms to deliver content, whilst this recommendation prompted other regulatory bodies including the Australian Media and Communications Authority (ACMA) and the eSafety Commissioner to lay claim to a regulatory role over AI.

It is not clear how or whether these workstreams come together, and how the differing positions and outcomes will be reconciled.

In addition to these Government led workstreams there are an increasing number of private sector initiatives springing up all seemingly trying to achieve the same outcome, perhaps without truly understanding what the problem is that is being addressed.

Existing regulatory scheme

Australia has a strong regulatory environment with many existing regulators already having a role to play in artificial intelligence development and use. With data being central to AI, the Australian Office of the Information Commissioner and the Privacy Commissioner have existing powers and functions relevant to numerous kinds of AI applications, ranging from facial recognition, to medical diagnosis to ‘robo-advice’ in the financial services sector.

The Productivity Commission’s 2017 report on *Data Availability and Use* has resulted in a new “consumer data right”. New regulatory functions will be conferred on the Privacy Commissioner and the ACCC to realise the consumer data right whilst the Government has also established a National Data Commissioner. The remit of the National Data Commissioner’s role remains to be seen, but should be taken into account when identifying bodies relevant to the use of data in some AI applications.

Consumer-focussed regulators may also play a role in responding to the use of AI, depending on the manner in which an application is deployed. Examples include the ACCC (general consumer protection powers and functions, and the new consumer data right functions), ASIC (consumer protection powers and functions for consumers financial services; licensing regime for providers of financial services), ACMA (focus on spam, telemarketing, interactive gambling) and the Australian eSafety Commissioner (online content regulation).

The ACCC would also play a role if AI was deployed in a manner contrary to competition law principles, for example as part of a misuse of market power resulting in substantial lessening of competition in a market.

The Therapeutic Goods Administration as Australia's health safety regulator has existing powers and functions in respect of medical devices that incorporate artificial intelligence. Software can be a medical device, particularly if it is used as a tool to diagnose disease. It is illegal to market medical devices in Australia unless they have been approved by the TGA.

In addition to Australia's own regulatory framework, there are internationally recognised bodies that play an important role in ensuring good practice(s) are maintained in the development and manufacturing of new products, including in relation to ICT. One such organisation is the International Organization for Standardization (ISO), of which Standards Australia is the national member body for Australia. In a world marked by increasingly complex supply chains, adopting common norms globally, through Standards, can be beneficial for companies and consumers.

Ensuring the responsible design and use of AI is a complex challenge: there is no single regulatory framework, or body, that can or should be responsible for AI systems in isolation. And rather than rushing to regulation, policy makers should be focused on facilitating a more comprehensive common understanding of the issues and examining the objectives of the current overlapping Government initiative in this space.

A new regulatory body for AI

With regard to the AHRC's proposal for a Responsible Innovation Organisation, Microsoft believes that this is premature and suggests that AI-based products and services are currently unregulated. Whilst it is true to say that most current standards, laws and regulations were not written specifically to account for AI, there are a wide range of regulatory schemes and laws that for example, protect the privacy and security of personal information, that govern the flow of data and how it is used, that promote fairness in the use of consumer information, or that govern decisions on credit or employment – that already apply broadly to digital products and services or their use in decision-making – whether they explicitly mention AI or not. Likewise companies and company directors have existing duties to conduct their business ethically under the Corporations Act.

In fact in Europe, the High Level Expert Group (HLEG) on Artificial Intelligenceⁱ made the point in December that the legislative status quo regarding regulation of AI is robust, stating that “no legal vacuum currently exists, as Europe already has regulation in place that applies to AI”.ⁱⁱ

The Business Software Alliance, of which Microsoft is a member, in response to the HLEG consultation, has prepared a comprehensive (albeit not necessarily exhaustive) overview of legislation already regulating AI, underlining the risk of over-regulating and creating an overlap of rules that would significantly hamper AI in the EU.

We consider this is a step that would be valuable in Australia, before rushing to create new regulations and indeed a new regulator.

As the role of AI continues to grow, it will be natural for policymakers not only to monitor its impact, but to address new questions and update laws.

One goal should be to ensure that governments work with businesses and other stakeholders to stroke the balance that is needed to maximize the potential of AI to improve lives and address new challenges as they arise.

One of the concerns around the current Government led initiatives on developing an AI framework is the lack of engagement with businesses and the tech industry.

In some public research Microsoft recently commissioned on attitudes to AI, it was found that the public expects industry involvement in the development of regulatory guidelinesⁱⁱⁱ.

We believe the most effective regulation can be achieved by providing all stakeholders with sufficient time to identify and articulate key principles guiding the development of responsible and trustworthy AI, and to implement these principles by adopting and refining best practices.

Before devising new regulations or laws, there needs to be some clarity about the fundamental issues and principles that must be addressed.

In *The Future Computed* we point to the evolution of information privacy laws in the United States and Europe as a useful model.

Today we believe policy discussion should focus on continued innovation and advancement of fundamental AI technologies, support the development and deployment of AI capabilities across different sectors, encourage outcomes that are aligned with a vision of human-centred AI, and foster the development and sharing of best practices to promote trustworthy and responsible AI.

Microsoft points to three fundamental principles that any framework should address.

1)The Importance of Data

It is likely that many near-term AI policy and regulatory issues will focus on the collection and use of data. The development of more effective AI services requires the use of data – often as much relevant data as possible.

And yet access to and use of data also involved policy issues that range from ensuring the protection of individual privacy and the safeguarding of sensitive and proprietary information to answering a range of new competition law questions. A careful and productive balancing of these objectives will require discussion and co-operation between governments, industry participants, academic researchers and civil society.

On one hand we believe governments should help accelerate AI advances by promoting common approaches to making data broadly available for machine learning – such as releasing public datasets. Governments can also invest in and promote methods and processes for linking and combining datasets from public and private organisations whilst preserving confidentiality, privacy and security as circumstances require.

At the same time it will be important for Governments to develop and promote effective approaches to privacy protection that take into account the type of data and the context in which it is used. Additional research into ‘de-identification’ techniques will be important.

Another important policy area involves competition law. As vast amounts of data are generated through the use of smart devices, applications and cloud-based services, there are growing concerns about the concentration of information by a relatively small number of companies. But in addition to the data that companies generate from their customers, there is publicly available data. Governments can help add to the supply of data by ensuring public data is usable by AI developers on a non-exclusive basis. There are concerns that access to unique datasets is becoming a barrier to competition, and whether sophisticated algorithms will enable rivals to effectively fix prices. Whilst these questions warrant consideration – they can most likely be addressed within the existing framework of competition law, and by the current competition regulator.

2) Promotion of Responsible and Ethical Uses of AI

In addition to addressing issues relating to data, governments have an important role to play in promoting responsible and effective uses of AI itself. This should start with the adoption of AI technologies in the public sector. While enabling more effective delivery of services for citizens, this will also provide governments with first hand experience in developing best practices to address ethical principles.

Governments also have an important role to play in funding core research to further advance AI deployment and support multidisciplinary research that focusses on studying and fostering solutions to the socio economic issues that may arise – this will also be valuable for the design of future AI laws and regulations.

3) Liability

Governments must also balance support for innovation with the need to ensure consumer safety by holding the makers of AI systems responsible for harm caused by unreasonable practices. Well-tested principles of negligence law are most appropriate for addressing these issues.

Relying on a negligence standard that is already applicable to software generally to assign responsibility

for harm caused by AI is the best way for policymakers and regulators to balance innovation and consumer safety, and promote certainty for developers and users of AI.

Transparency or Intelligibility

We felt it might be helpful to specifically address the issue of transparency or as we refer to it at Microsoft – intelligibility. Demands for explanations of the behavior of AI systems are increasing, fueled by concerns over the growing ubiquity of “black-box” technologies.

Microsoft and others have described this goal of making AI understandable to humans as fundamental to “transparency.” In policy circles, transparency represents not only the idea that people should be able to understand and monitor how AI systems behave, but also that those who use AI systems should be honest and forthcoming about when, why, and how they choose to deploy them. In AI and machine learning circles, the term “intelligibility,” often used interchangeably with the term “interpretability,” refers to the concept of making the behavior of AI systems, or components of systems, understandable to humans.

But achieving useful explanations of the behavior of AI systems and their components can be quite complex and highly dependent on a host of variables, precluding anything resembling a “one-size-fits-all” approach. It is therefore an area of intense, cutting-edge research.

Selecting an intelligibility approach therefore starts with asking questions about the people at the center of an AI system’s development and use: Who are they? What do they want to know? Why do they want to know it? Do they need to understand the overall behavior of an entire AI system, or do they require more specific understanding of a particular output or prediction? Do they need to know key characteristics of the data used to train a particular machine learning model?

Simply publishing the algorithms underlying AI systems will rarely provide meaningful transparency. With the latest AI techniques such as deep neural networks, there typically isn’t any algorithmic output that would help people understand the subtle patterns that systems find. This is why we need a more holistic approach.

Microsoft is working with the Partnership on AI and other organisations to develop best practices for enabling meaningful transparency of AI systems. This is an area that will require further research to understand how machine learning models work and to develop new techniques that provide more meaningful transparency.

Conclusion

In summary, Microsoft shares the AHRC’s optimism about the potential for AI to deliver economic and societal progress. We also acknowledge there are risks and challenges that come with it. We therefore consider it essential that we continue to convene open discussions amongst governments, industry, civil society, academic researches and all interested individuals and organisations. It is only by working together, that we can identify issues and prioritise the development of solutions that protect people without restricting future innovation.

ⁱ https://ec.europa.eu/newsroom/dae/document.cfm?doc_id=56433

ⁱⁱ Ibid, p 2

ⁱⁱⁱ Futureye's qualitative research into social expectations in relation to ethical AI (2019)'.