

Access Now submission to the Australian Human Rights Commission and World Economic Forum White Paper consultation on AI governance

March 2019

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

While there is no agreed-upon definition for artificial intelligence (“AI”) it is imperative that the starting point for any policy or regulatory conversation on AI should be a clear scoping exercise that includes a problem definition. In policy debates, the term AI is often used for anything between big data, algorithmic decision making, automation, machine learning and robotics.

There are several lenses through which experts examine artificial intelligence. The use of international human rights law and its well-developed standards and institutions to examine artificial intelligence systems can contribute to the conversations already happening, and provide a universal vocabulary and forums established to address power differentials. Additionally, human rights laws contribute a framework for solutions.

As one example, the role of AI in facilitating discrimination is well documented, and is one of the key issues in the debate today. In reaction to growing evidence of the risk of discriminatory harms associated with the use of machine learning systems in public and private use across many sectors, including policing, criminal justice, immigration and asylum Access Now partnered with human rights organizations and AI companies to release “The Toronto Declaration” in May 2018.¹ However, the right to non-discrimination is not the only human right implicated by AI. Because human rights are interdependent and interrelated, AI affects nearly every internationally recognized human right. **The goal of any government regulation in the area must be to as far as possible prevent and mitigate the threats posed to human rights by AI.**

There are many human rights impacted by AI.² They are largely those embodied in the three documents that form the base of international human rights law, the so-called “International

¹ “The Toronto Declaration: Protecting the right to equality and non-discrimination in machine learning systems, <https://www.accessnow.org/cms/assets/uploads/2018/05/Toronto-Declaration-D0V2.pdf>.

² Human Rights not included are: freedom from torture, right not to be enslaved, rights of detainees, right not to be imprisoned merely based on inability to fulfill a contractual obligation, rights of aliens, and the right to social security. This does not mean that AI cannot ultimately impact these rights, merely that we found no current documented violations, nor and prospective violations we believed could occur in the near future.

Bill of Human Rights.”³ This includes the Universal Declaration of Human Rights (UDHR), the International Covenant on Civil and Political Rights (ICCPR), and the International Covenant on Economic, Social and Cultural Rights (ICESCR).⁴ In our report entitled *Human Rights in the Era of Artificial Intelligence*, we outline how current AI uses violate or risk violating those rights, as well as risks posed by prospective future developments in AI.⁵ It is important to note that these violations of human rights are not necessarily unique to AI. Many already exist within the digital rights space, but the ability of AI to identify, classify, and discriminate magnifies the potential for human rights abuses in both scale and scope.

Like the human rights harms in other uses of technology that leverage data, the harms related to the use of AI often disproportionately impact marginalized populations.⁶ That can include women and children, as well as certain ethnic, racial, or religious groups, the poor, the differently abled, and members of the LGBTQ community. The long-established marginalization of these groups is reflected in the data and reproduced in outputs that entrench historic patterns.

As a society, we do not need to accept the use of AI everywhere. Indeed, there are areas where it should not be deployed, and the burden of proof should not be on society to demonstrate why that is so. Right now, it appears that AI is increasingly the default setting, and that civil society and other actors are scrambling to meet the responsibility of developing the right safeguards. In fact, it should be incumbent on government regulators and industry to make sure any developments come with sufficient protections.

Human rights establish universality

It is imperative that government regulation in this space comes from the basis of human rights, rather than ethical frameworks. Human rights are universal and binding, and are codified in a body of international law. Respecting human rights is required of both governments and companies alike, although governments have additional obligations to protect and fulfill human rights.⁷ There is an entire system of regional, international, and domestic institutions and organizations that provide well-developed frameworks for remedy and articulate the application of human rights law to changing circumstances, including technological developments. And in cases where domestic law is lacking, the moral legitimacy of human rights carries significant normative power. Human rights law can

³ Note that although there are many regional human rights systems that are more comprehensive, we mostly limited our analysis to the UN-based system in the interest of universal applicability. The exception to this is the right to data protection, which Access Now recognizes as a right and is particularly relevant in the context of AI. Further analysis of AI relating to the rights enumerated in these regional systems is merited. For example, the European Convention on Human Rights and the EU Charter of Fundamental Rights are far more comprehensive when it comes to workers’ rights, and use of AI by employers to monitor and police employee activity may violate European human rights.

⁴ See <https://www.ohchr.org/EN/ProfessionalInterest/Pages/InternationalLaw.aspx> for more information

⁵ Our full report is available online at

<https://www.accessnow.org/cms/assets/uploads/2018/11/AI-and-Human-Rights.pdf>

⁶ For a more in-depth examination of how this plays out in the U.S., see *Automating Inequality* by Virginia Eubanks.

⁷ According to the UN Principles on Business and Human Rights, States must protect against human rights abuse by businesses within their jurisdiction, businesses are responsible for respecting human rights wherever they operate, victims must have access to judicial and non judicial remedy. For more information, see: https://www.ohchr.org/Documents/Publications/GuidingPrinciplesBusinessHR_EN.pdf

address some of the most egregious societal harms caused by AI, and prevent such harms from occurring in the future.

The ethics discourse has largely dominated the discussion about the social implications of AI. This is partially because AI has sparked more discussions about the interplay between human beings and machines than perhaps any previous technological development. Considering broad ethical concepts such as justice, fairness, transparency and accountability allows for valuable debate about the role of AI in our lives.⁸ And while a given use of AI that violates human rights is certainly unethical, not all unethical uses of AI constitute human rights violations. On the other hand, just because something is legal, it is not necessarily ethical. Unlike highly contextual ethics principles, human rights are universal and well-defined, and they provide for accountability and redress. Going beyond legal compliance, human rights and ethics can be mutually reinforcing. For example, a company might develop ethical AI principles such as avoiding reinforcing negative social biases. Such ethics principle should reflect the human rights of privacy and non-discrimination and should be detailed and translated into concrete measures on the company policy and product design levels. The challenges and limitations of enforcement of such principles requires domestic or international rights systems to provide for remedy should those principles be violated.⁹

Data protection ensures trust and certainty

Comprehensive data protection laws, which should apply to both the government and private sector, can go a long way in addressing some of the human rights risks posed by AI. Because data is the engine of AI, any law that mandates protection of personal data will necessarily implicate AI systems. Given the global push toward data protection legislation, this is both heartening and practical.

Consider the impact of the European Union's General Data Protection Regulation (GDPR). The GDPR is a positive framework that provides for control of a person's personal information and right to empower people to make informed decisions about how their data is used. The GDPR limits data processing to permissible purposes, with heightened protections for sensitive data. It also requires opt-in consent,¹⁰ which limits the use of personal data for training AI systems.

While the GDPR was not developed specifically for AI, it will set crucial benchmarks for the regulation of AI in Europe (complemented by other laws in the field such as the Police Directive or the ePrivacy Directive which is subject to a reform now). By setting rules and safeguards around the processing of personal data, the GDPR has the potential to directly impact the development and implementation of AI.¹¹

⁸ <https://www.considerati.com/publications/blog/marrying-ethics-human-rights-ai-scrutiny/>

⁹ See "Human Rights in the Age of Artificial Intelligence" for further discussion of the interplay between ethics and human rights in AI

¹⁰ <https://gdpr-info.eu/art-9-gdpr/>

¹¹ Access Now, Mapping regulatory proposals for artificial intelligence in Europe (2018), https://www.accessnow.org/cms/assets/uploads/2018/11/mapping_regulatory_proposals_for_AI_in_EU.pdf

Rights provided for by the GDPR, and other similar laws, offer a framework to prevent against unaccountable uses of AI that impact individual rights, while ensuring a level of control of personal data and accountability for the use of AI and ML systems.

The adoptions of the EU data protection reform was contentious: the GDPR passed in the teeth of, in the words of the European Data Protection Supervisor Giovanni Butarelli, “arguably the biggest lobbying exercise in the history of the European Union.” Precisely how they will impact AI applications is also likely to be contested. Some experts advising the EU have observed that AI’s core functions call the very cornerstones of data protection and privacy into doubt.¹²

Some have suggested that data protection laws are incompatible with AI and we should make broad exceptions for its development and use. That is misguided. While it is likely true that strong data protection laws may preempt deployment of certain AI systems, companies have never been able to “innovate” without regard for potential harm. If AI systems are used to make decisions on a basis or rationale that not even their developers can fully explain, at-risk individuals—or AI “guinea pigs”—will be the first to suffer the negative consequences. Data protection rights not only provide accountability structures to mitigate harm, they also protect people against having their personal data covertly co-opted, commodified, and otherwise exploited in ways that harm others or society at large.

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

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

The Office of the Australian Information Commissioner as well as state-level Information Commissioner offices are critical in establishing channels of communication between government bodies who process data and individuals’ rights to access.

Additionally, independent government bodies such as the Australian Human Rights Commission or the Commonwealth Ombudsman ensure that the interests and freedoms of individuals are protected when dealing with government agencies. The ACCC also plays an important role in protecting consumer rights which is an ever-growing responsibility in the increasingly digital marketplace. The further development of the consumer data right (CDR), which at the moment is quite restricted only to data access and portability, could be essential in shaping broader data protection rules, leading to meaningful protections and safeguards for human rights in AI technologies.¹³

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

¹² EDPS Ethics Advisory Group, Toward a Digital Ethics, Jan 2018, at 7: “The right to data protection may have so far appeared to be the key to regulating a digitised society. However, in light of recent technological developments, such a right appears insufficient to understand and address all the ethical challenges brought about by digital technologies....the tensions and frequent incompatibility of core concepts and principles of data protection with the epistemic paradigm of big data suggest limits to the GDPR even prior to its application.” Available at https://edps.europa.eu/sites/edp/files/publication/18-01-25_eag_report_en.pdf

¹³ <https://www.accc.gov.au/focus-areas/consumer-data-right-cdr-0>

Arguably the biggest gap is that there is no affirmative right to privacy in Australia. Nor do Australians have the ability to file a lawsuit against an individual, entity, or government for a violation of privacy. And while there is federal legislation related to privacy, it is full of loopholes and is woefully inadequate to the unique risks to privacy of the digital age.

The *Privacy Act* regulates the collection and use of personal information via a set of “Privacy Principles” that apply to most federal government agencies, as well as businesses and nonprofits with annual turnover of over \$3 million, with some exceptions.¹⁴ It defines personal information as “information or an opinion, whether true or not, and whether recorded in a material form or not, about an identified individual, or an individual who is reasonably identifiable.”¹⁵

The Privacy Principles include a number of laudable provisions, such as requiring entities to inform individuals why their personal information is being collected, how it will be used, and to whom it will be disclosed, allowing individuals the option of not identifying themselves or using a pseudonym (with exceptions), mandating the collection of personal information “only if necessary,” allowing individuals to request access to their personal information, as well as request the information be corrected if inaccurate.

Despite this, there are a number of important shortcomings of the Privacy Act. First, it does not address the ability of AI to easily re-identify “anonymized” information. Information in anonymized datasets does not qualify as personal information according to the definition, and therefore is not subject to any of the protections stipulated by the Privacy Principles. Second, there is no requirement for entities to obtain consent prior to collecting personal information. This removes the right of individuals to decide whether or not they are comfortable with the entity’s disclosed use of their data. Third, it does not apply to state or territory governments, although some states have passed their own privacy legislation. Fourth, it does not address the privacy threats of surveillance. Although the Privacy Act covers the Australian Federal Police and the Border and CrimTrac, it does not cover most law enforcement and intelligence agencies, which arguably the entities most likely to commit major breaches of privacy. Further, where it does apply to law enforcement agencies there are many exemptions and carve outs that allows for near limitless data collection. And finally, the Privacy Act only provides for limited civil redress via a complaints mechanism overseen by an Information Commissioner.¹⁶ These gaps make the Privacy Act nearly useless in protecting Australians against the privacy and data protection risks posed by AI.

The Data Sharing and Release Act

One new piece of legislation would erode the already weak Privacy Act even further. Whenever you interact with the government, data are created about your activities. As part

¹⁴ The following businesses are subject to the privacy act regardless of their size: private sector health providers, businesses that sell or purchase information, credit reporting bodies, and government contractors.

¹⁵ <https://www.oaic.gov.au/privacy-law/>

¹⁶ DRW state of digital rights report

of its modernization efforts, the Australian government would like to be able to capitalize off of all the data it collects. The government states “better use of public sector data can help us improve government services for Australians and ensure our programs and policies are informed by evidence.”¹⁷ The Data Sharing and Release Act, which is currently in Parliament, seeks to make it easier for government agencies to share data with each other, allowing any government entity to access any and all the information the government holds about you, and also permitting the government to share data with “trusted” third parties and researchers.¹⁸ And while it is certainly good to use data analysis to create evidence-based policy, the proposed law includes a total lack of privacy and security protections. It merely states that the risks should be “appropriately managed.” It also reflects an ignorance how seemingly innocuous data points, when combined and analyzed by ML systems, can quickly reveal intimate details about people’s lives.

Currently, such use of your data potentially conflicts with the over 500 existing data secrecy and confidentiality provisions across existing Australian law. Notably, it violates the Privacy Act and the Australian Privacy Principles, which state that government agencies cannot use your administrative data for secondary purposes unrelated to providing you with the respective service. However, if passed, the Data Sharing and Release Act would override any conflicting legislation, for both government and non-governmental entities alike.¹⁹ Additionally, the bill would instate data sharing by default. There would be no ability for Australians to opt-out of having their data being shared across the government and with third parties. With this bill, the government is clearly communicating its view that your data belongs to the government because they collect it, as well as continuing carelessness in its approach to risk management of technology projects.

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

By establishing an independent Responsible Innovation Organization, Australia would open the opportunity to define the direction of AI innovation, growth, sustainability, and future social impact. Using such an organization as a stepping stone in developing strong and systemic regulation which empowers individuals is an essential piece of this. The ultimate goal should be AI that benefits humanity by contributing to a more responsible and equitable society. By inviting multi-stakeholder input, considering risks, and setting rights-respecting standards, Australia has the opportunity to be among the norm-setters on the world stage.

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

¹⁷ <https://www.pmc.gov.au/resource-centre/public-data/issues-paper-data-sharing-release-legislation>

¹⁸ <https://www.eigenmagic.com/2018/07/30/tljr-data-sharing-and-release-issues/>

¹⁹

<https://independentaustralia.net/life/life-display/-the-data-sharing-and-release-act-is-coming-for-your-data,11761>

N/A

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

N/A

6. If Australia had a Responsible Innovation Organisation...

It is difficult to hypothesize the makeup of an organization without a clear mandate. However, there are several similar initiatives globally, which could be a useful indicator of parallel efforts, on which Australia can elaborate and grow its own.

As one example, in 2018, in recognition of the growing importance of AI and its impact on the digitization of Europe, the European Commission launched a High-Level Expert Group on Artificial Intelligence to support the implementation of the European strategy on AI,²⁰ and to complement the work of the European Group on Ethics in Science and New Technologies.²¹ Access Now's European Policy Manager, Fanny Hidvégi, got selected to the expert group which is now working on ethics guidelines for trustworthy AI. The group is composed of approximately 50 independent experts who were selected from an open selection process open to representatives from academia, civil society and industry. The group was tasked with two main deliverables. First, to develop an ethics guidelines for trustworthy AI. The group will publish and present the final version of the guidelines in April 2018 but a draft version is already publicly available and was subject to a consultation.²² Secondly, the group will likely continue the work on policy recommendations but this is also subject to how the mandate of the group will be re-established after the 2019 European parliamentary elections.

While the process is still ongoing, there are some important lessons learned from the establishment and operations of the group, and also from other multistakeholder fora we have been engaging with. The composition of any group must be balanced and inclusive, and the structure must mitigate the different resources of the different stakeholders. The criteria for selection (either an open call or invitation based) should be clear, objective, transparent and public. The work of the group should be published and subject to adequate consultation and feedback. The mandate of the group should be well-defined.

If the Australian Human Rights Commission decides to proceed with or recommend the development of a Responsible Innovation Organization, we would welcome the opportunity to further engage and contribute to its development.

²⁰ <https://ec.europa.eu/digital-single-market/en/high-level-expert-group-artificial-intelligence>

²¹ <https://ec.europa.eu/research/ege/index.cfm>

²²

<https://ec.europa.eu/digital-single-market/en/news/have-your-say-european-expert-group-seeks-feedback-draft-et-hics-guidelines-trustworthy>

Prepared by the Access Now policy team, with research contributions by Lindsey Andersen.

For more information, contact:

Lucie Krahulcova

Asia Policy Analyst

Access Now | www.accessnow.org

lucie@accessnow.org