THE LAW SOCIETY
OF NEW SOUTH WALES

Our Ref: PDL: EEap1728223

7 June 2019

Mr Jonathan Smithers
Chief Executive Officer
Law Council of Australia
DX 5719 Canberra

By email: mike.clayton@lawcouncil.asn.au

Dear Mr Smithers,

Artificial Intelligence: Australia's Ethics Framework

Thank you for the opportunity to contribute to a submission on the Discussion Paper, *Artificial Intelligence: Australia's Ethics Framework*, by CSIRO's Data61. The Law Society's Privacy and Data Law Committee and Legal Technology Committee have contributed to this submission.

The Law Society supports the steps the Australian Government is taking to ensure the emerging ethical issues associated with artificial intelligence ("AI") are properly considered. While generally supportive of the approach taken in this Discussion Paper, the Law Society considers that the principles underpinning a national AI ethics framework, particularly around privacy and data governance, require careful further consideration and development.

This submission comprises two parts:

- Part 1 responds to specific questions posed in the Discussion Paper, and
- Part 2 offers general comments for further consideration by the Department of Industry, Innovation and Science ("the Department") / CSIRO.

Part 1

Question 1: Are the principles put forward in the discussion paper the right ones? Is anything missing?

The core principles the authors have proposed should underpin the development of AI are:

1. Generates net-benefits;
2. Do no harm;
3. Regulatory and legal compliance;
4. Privacy protection;
5. Fairness;
6. Transparency and Explainability;
7. Contestability; and
8. Accountability.

The Law Society notes that at least 63 public-private initiatives have produced statements describing high-level principles, values, and other tenets to guide the ethical development, deployment, and governance of AI.[1] All statements include similar principles on transparency, equality / non-discrimination, accountability and safety. Some statements include additional principles, such as the requirement for AI to be socially beneficial and to protect human rights.

While we consider the core principles the Department / CSIRO propose broadly reflect these listings, and form a suitable basis upon which to start a discussion about when and how they should be applied, we suggest that the principles be tested for comprehensiveness. To that end, we recommend the Department / CSIRO consider including additional principles to ensure a robust ethics framework.

We suggest that the principle of 'Respect for human rights and human autonomy' should be included as an independent principle.

Other principles that warrant further exploration are:

1. human agency (sufficient information, training and education to make an informed choice);
2. human oversight (there are many possibilities for human oversight and intervention during the life cycle which require discussion);
3. accuracy, reliability and robustness of inputs, processes and outputs (including good data governance);
4. technical robustness (level of imperviousness to malware and cyberattack);
5. safety and wellbeing (what to do in the event of a problem); and
6. environmental impact and conservation of natural resources.

These principles are not specific to AI: many considerations equally apply to other advanced data analytics services that may lead to significant adverse effects on humans, whether or not these services involve humans in-the-loop decision making or automated decision making. AI may introduce more intractable issues as to explainability and autonomy, but generally, the same principles should be applied in the evaluation of all applications of advanced data analytics. For example, there are a number of statistical, data-driven tools used in criminal procedure to predict future re-offending and assess 'unacceptable risk' that are not strictly AI, but are rather 'actuarial' or 'algorithmic' instruments. These tools have been critiqued for being based on group data instead of individual data, thus challenging individualised justice.[2]

We note that the core principles reflect existing requirements under specific laws, the operation of human rights protections in Australian law, and non-mandatory ethical principles. We recommend the Department / CSIRO specifically note this overlap so that compliance with law and regulation, in addition to ethical considerations, is properly considered.

We further note the significant overlap between concerns about nurturing of citizen and consumer digital trust, or 'social licence', and the fact that often, good ethical

[1] For a non-exhaustive recent listing, see AlgorithmWatch, *AI Ethics Guidelines Global Inventory*, <https://www.rri-tools.eu/-/ai-ethics-guidelines-global-inventory>.

[2] See, for example: Jones, L & Milton, E., *Risk of Sexual Violence Protocol (RSVP): A real world study of the reliability, validity and utility of a structured professional judgement instrument in the assessment and management of sexual offenders in South East Scotland*, (2016).

decisions about applications of digital data reflect sound business judgement as to sustainability of business models. For example, there is now growing public focus on the uses of facial recognition and other automated surveillance technologies, and the secondary uses of health data and of geo-location data. As public concerns around uses of data evolve, business models that have been built upon such activities may be rendered redundant. Ethical review should include consideration of the effect of a particular application of AI upon digital trust of consumers, citizens and users.

We submit that any statement of principles is of little or no benefit unless the principles are consistently, reliably and verifiably applied in practice, whenever they should be applied.[3] This requires a clear statement of the threshold condition for determining when the principles should be applied, including a risk assessment methodology to determine at what point risks are such that a comprehensive impact assessment should be undertaken. We consider that the Framework should include an analysis that will assist organisations to understand the contexts in which risks should be initially assessed and the point at which a comprehensive impact assessment should be undertaken.

We also note that risk assessment, and mitigation in design and deployment of AI applications, requires appropriate risk management in AI-related decision making, robust internal governance systems and measures, accountability mechanisms, and appropriate relationship management of users and consumers. It is particularly important that an accountability mechanism within organisations clarifies the individuals responsible for ensuring that principles have been appropriately applied, the basis for relevant decisions made by relevant individuals, and the extent of executive oversight of those decisions.

Principle 4 – Privacy Protection

Principle 4 provides: 'Any system, including AI systems, must ensure people's private data is protected and kept confidential plus prevent data breaches which could cause reputational, psychological, financial, professional or other types of harm.'

The Discussion Paper recognises that issues related to AI ethics are closely intertwined with those that relate to data sharing. It also recognises that privacy is crucial in any discussion related to AI ethics. The Discussion Paper attempts to set out the key requirements of privacy and data sharing laws in Australia in sections 2.1.3 and 2.1.4 as well as chapter 3.

The Law Society is concerned that the Discussion Paper is not an accurate summary of privacy law in Australia and considers that a more robust understanding of privacy law, both in theory and in practice, is essential to developing an AI ethics framework. The explanation of privacy law in Australia and its application to AI is incomplete, and as a result, we consider that Principle 4 is misleading and places undue importance on 'confidentiality' and 'security' of 'private data' (further detail is set out below).

In developing an AI ethics framework, we submit that the key focus for privacy and AI should be on *privacy protection* and *adequate data governance.* We draw this

[3] See further, Luciano Floridi, *Translating Principles into Practices of Digital Ethics: Five Risks of Being Unethical*, (2019) <https://doi.org/10.1007/s13347-019-00354-x>; Personal Data Protection Commission Singapore, *A Proposed Model Artificial Intelligence Governance Framework,* (January 2019); UK Information Commissioner's Office*, An Overview of the Auditing Framework fir Artificial Intelligence and its core components,* (26 March 2019) <https://ai-auditingframework.blogspot.com/2019/03/an-overview-of-auditing-framework-for_26.html>.

framework from the European Commission's Ethics Guidelines for Trustworthy AI,[4] which covers:

a) privacy and data protection – ensuring the *lawful* collection of data initially from the user as well as data generated about the user over the course of their interaction with the system, and that the data collected about the user will not be used to unlawfully or unfairly discriminate against them; and
b) adequate data governance – ensuring the quality and integrity of the data used, its relevance in light of the domain in which the systems will be deployed, its access protocols and the capability to handle data in a manner that protects privacy.

The Law Society recommends that the Department / CSIRO engage with privacy regulators, and practitioners with specialist expertise in privacy law and practice, to assist in refining Principle 4, the general framework and the context underpinning them.

We have set out our reasoning below.

Legislative framework

The Law Society considers that the technical complexity of privacy law in Australia and privacy principles needs to be dealt with in greater depth than it is in the Discussion Paper. We recommend a more robust and technically nuanced discussion about rights and obligations under privacy law in Australia to properly address privacy as a central issue of the Discussion Paper.

The Discussion Paper does not address the development of lawful and ethical AI within the framework of Commonwealth, State and Territory privacy legislation. The *Privacy Act 1988* (Cth) ("Privacy Act") and the Australian Privacy Principle ("APPs") set out in the Privacy Act are not the only privacy laws in Australia that will apply to organisations dealing with AI. The Discussion Paper fails to mention and discuss State and Territory privacy laws which may apply to, for example, deployment of surveillance and tracking technologies, secondary uses of geo-location and other data about communications derived from uses of communications services that apply to publicly funded bodies (such as universities), public sector health service providers such as public hospitals, and private sector health apps and health service providers, when employing AI systems.

We also note that the European Union General Data Protection Regulation ("GDPR") may apply to organisations in Australia employing AI systems that have an establishment in the EU, or offer their goods or services to, or monitor the behaviour of, people in the EU.[5]

Consistent terminology

Principle 4 uses the term 'private data'. While the terms 'personal information' and 'personal data' are sometimes used interchangeably, the term *'private* data' does not reflect privacy law in Australia or overseas. Personal information does not need to be

[4] *Ethics Guidelines for Trustworthy AI*, Independent High-Level Expert Group on Artificial Intelligence set up by the European Commission, (8 April 2019) <https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai> 17.
[5] *EU General Data Protection Regulation (GDPR)*: Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation), OJ 2016 L 119/1, art 3.

'private' for it to be protected under privacy laws. The Law Society recommends that the correct terminology of 'personal information' be used in future consultations and policy documents produced by CSIRO / the Department, to ensure organisations employing AI systems have certainty about their compliance requirements.

The Discussion Paper also uses the word 'sensitive' to describe personal data and the term 'sensitive data'. This can be confusing as 'sensitive information' is a subset of personal information and is defined under the Privacy Act.[6] Higher standards apply under the Privacy Act when sensitive information is collected, used or disclosed. Using the term 'sensitive data' could confuse organisations about whether these higher standards apply in the AI context. If CSIRO / the Department intends to set out different obligations for sensitive information used in AI systems, 'sensitive information' should be properly defined. Otherwise, we consider the words 'sensitive' and 'sensitive data' should be used sparingly.

Consent process is not fundamental to protecting privacy

Section 3.1 of the Discussion Paper states that 'protecting the consent process is fundamental to protecting privacy'. The Discussion Paper does not discuss why and how it arrives at this conclusion. This results in the risk assessment framework in section 7.2 of the Discussion Paper placing an over-reliance on 'consent'.

There is often a misconception that organisations require consent from the individual to collect, use and disclose their personal information. This misconception is compounded by the fact that many organisations require users to 'agree' to their privacy policy in the registration process.

Although there may be some overlap in their contents, the requirements for privacy policies pursuant to APP 1 and collection notices pursuant to APP 5 are two separate requirements under the Privacy Act. Neither APP 1 nor APP 5 require a user to consent to a privacy policy or collection notice.

Consent is not the sole mechanism by which the collection, use or disclosure of personal information may be lawfully authorised under the Privacy Act. In fact, consent is often the exception for collection, use or disclosure of personal information under the Privacy Act. The Privacy Act currently requires that individuals provide consent when their personal information is collected in only limited circumstances, including:

a) the use or disclosure of personal information for a secondary purpose (APP 6.1(a));
b) the collection of sensitive information (APP 3.3(a));
c) the collection of personal information by an agency from someone other than the individual (i.e. an individual must consent for an agency to disclose their personal information to another agency) (APP 3.6(a)(i));
d) the use or disclosure of personal information or sensitive information for direct marketing purposes (APP 7.3(b) and APP 7.4); and
e) the disclosure of personal information to an overseas recipient (APP 8.2(b)).

The Law Society therefore recommends that future consultations or policy documents produced by CSIRO / the Department emphasise the requirement for the lawful collection, use and disclosure of personal information. This should include discussion of the restrictions or limitations on collection, use and disclosure of personal information as set out in the relevant privacy laws. Future documents on

[6] *Privacy Act 1988* (Cth) s 6.

this topic should also focus on how these restrictions and limitations apply to personal information collected initially from the user, as well as data generated about the user over the course of their interaction with the AI system.

Right to be forgotten

The Discussion Paper links the currency of consent with the absence of a 'right to be forgotten'. The 'right to be forgotten' or 'right to erasure' in the GDPR is unrelated to the issue of consent, but is connected with the right of access and correction. The 'right to be forgotten' is also not an absolute right under the GDPR and will not apply if the organisation has collected or is using the personal information lawfully and fairly and still has a current need for that information.[7] There is no 'right to be forgotten' in Australian privacy law.

APP 11.2 does adopt a similar approach, requiring an APP entity to destroy the information or to ensure that the information is de-identified if the entity no longer needs the information for any purpose for which the information may be used or disclosed, irrespective of whether the individual has requested it or not.

Protecting privacy is not just confidentiality and security

Protecting privacy is more than ensuring that personal information is ‘kept confidential’ and to ‘prevent data breaches’. The Law Society submits that privacy should not be conflated with confidentiality. Securing personal information does not necessarily mean privacy has not been breached.

Invasions of privacy could arise from unlawful collection of personal information to develop an AI system, by for example using personal information collected for an unauthorised purpose or even a secondary purpose beyond the expectation of the individual who initially provided the information. Privacy could also be invaded by subjecting individuals to automated decision-making which could significantly affect the individual. These invasions of privacy could cause harm to the individual, even in circumstances where there was no data breach.

One of the biggest challenges to protecting privacy in the development or application of AI is the use and disclosure of personal information for secondary purposes. Often, personal information collected for a different primary purpose will be re-used to develop the AI system.

Using personal information for secondary purposes is not permitted under the Privacy Act unless an exception applies under APP 6. Consent is one of those exceptions, but it may not be pragmatic to obtain consent where:

a) a large cohort of individuals’ personal information is used; and
b) in order to give informed and specific consent, the individual must know how their information will be used in the AI system.

The Law Society recommends that future consultations or policy documents produced by CSIRO / the Department set out clearly the restrictions and limitations on the use and disclosure of personal information for secondary purposes and the risks of obtaining consent in the AI context.

[7] *EU General Data Protection Regulation (GDPR)*: Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation), OJ 2016 L 119/1, art 17.

The Law Society also recommends that future documents set out the importance of the quality of data and access to data in the privacy and data protection context. Transparency of algorithms can mitigate against harmful privacy risks and visibility of the data used in automated decision making can prevent skewed data input and thus, the generation of biased datasets.

Privacy Impact Assessments may be a means by which the privacy impacts of AI can be assessed so that strategies for mitigating risks can be developed to ensure that privacy is protected before information is collected and used. Safeguards such as data minimisation and purpose limitation should also be implemented to prevent the unauthorised collection, use and disclosure of personal information. Individuals should also be able to seek access to their personal information that has been used or generated in the AI system and seek redress if they have been affected by a decision made by the AI system.

Eligible data breach

The Discussion Paper provides a description of a data breach requiring notification in section 3.2 as, 'if personal data is accessed or disclosed in any unauthorised way that may cause harm, all affected individuals must be notified'. This is description is not entirely accurate.

The Notifiable Data Breach scheme requirements are not limited to unauthorised access or disclosure; they also include loss of the information. The Notifiable Data Breach scheme applies not only to personal information, but also to credit-related information and Tax File Numbers.[8] A data breach giving rise to the obligation to notify is one where an individual is likely to suffer 'serious harm', not just harm, and the obligation is to the Office of the Australian Information Commissioner ("OAIC") as well as affected individuals.

We recommend CSIRO / the Department adopt a more nuanced approach to discussion about eligible data breach in future documents on this topic.

De-identification

We consider that the discussion of de-identification and re-identification in section 3.3 is overly simplistic in the context of AI systems. De-identification, where it is often not permanent and easily reversible, does not remove the obligations to comply with the privacy laws discussed above, particularly around the lawful collection and use of personal information. De-identification should be used as a privacy enhancing tool, as discussed in the OAIC's guide on De-identification and the Privacy Act.[9]

Other legislative initiatives

The Discussion Paper outlines some government initiatives on data, such as the Consumer Data Right. Introduction of the Consumer Data Right will not change the privacy landscape, nor provide any effective privacy protections and safeguards in the AI context. This is because consumers already have a right to access their information under APP 12. The privacy obligations under the Consumer Data Right largely replicate the obligations under the Privacy Act and the Consumer Data Right merely provides another mechanism for businesses to share customer information and use it for a range of purposes, including in AI systems.

[8] *Privacy Act* s 26WE(1).

[9] <https://www.oaic.gov.au/agencies-and-organisations/guides/de-identification-and-the-privacy-act>.

However, there have been numerous other recent government initiatives that aim to enhance privacy protections, and which will have an impact in the AI context. In particular, we note the privacy-related recommendation in the Australian Competition and Consumer Commission's preliminary report on the Digital Platforms Inquiry[10] and the Attorney General's proposed Privacy Act amendments in anticipation of those recommendations.[11]

We note that despite numerous recommendations at both Federal[12] and State[13] level, there is no 'right to privacy' in Australia as mentioned at the end of section 3.3 of the Discussion Paper.

Question 2: Do the principles put forward in the discussion paper sufficiently reflect the values of the Australian public?

Values-based vs rights-based approach to AI ethics

The Discussion Paper refers to the 'values' of the Australian public in several instances but does not discuss how those values are assessed so as to determine what is important to the Australian public. The Australian public comprises many diverse groups which may place weight on different values. The Law Society considers that if a 'values-based' approach is adopted for guiding the development of AI, a consensus decision about the values of the Australian public would need to be underpinned by current empirical research.

The Law Society recommends that public consultation be undertaken in a way that would engage the members of the public who are likely to be affected by AI tools, particularly those groups likely to be affected more significantly by government activities, such as economically and socially disadvantaged groups, the disabled and chronically ill, users of public health services and senior citizens.

Instead of a values-based approach to framing AI ethics, the Law Society recommends consideration of a more objective approach to framing AI ethics, for example, an approach based on international human rights law obligations. The European Union ("EU") has taken an approach to guiding AI ethics based on fundamental rights. There is a growing body of scholarly research and case law addressing fundamental human rights, from Europe, Canada and Australia. The EU *Ethics Guidelines for Trustworthy Artificial Intelligence,* by the Independent High-Level Expert Group on Artificial Intelligence, identifies ethical principles and correlated values using an approach founded on fundamental rights (page 2 of the Guidelines). Noting Australia does not have a Human Rights Act or Charter to inform a rights-based approach to the ethical framework to underpin AI, the Law Society recommends that a number of sources could be used to inform the approach, including:

a) the seven core international human rights treaties to which Australia is a party;
b) the nine core international human rights instruments and protocols; and

[10] The Australian Competition and Consumer Commission, *Digital Platforms Inquiry*, Preliminary Report (2018).

[11] Attorney-General of Australia, 'Tougher Penalties to Keep Australians Safe Online' (Media Release, 24 March 2019).

[12] Australian Law Reform Commission, *For Your Information: Australian Privacy Law and Practice,* Report No 108 (2008); Australian Law Reform Commission, *Serious Invasions of Privacy in the Digital Era Final Report*, Report No 123 (2014).

[13] See for example, Parliament of NSW Standing Committee on Law and Justice, *Remedies for the Serious Invasion of Privacy in New South Wales,* Report No 57 (2016).

c) the fundamental rights set out in the Charter of Fundamental Rights of the EU and international human rights law.

Question 5: What other tools or support mechanisms would you need to be able to implement principles for ethical AI?

The tools the Discussion Paper proposes for implementation of ethical AI are:

1. Impact Assessments;
2. Internal or External Review;
3. Risk Assessments;
4. Best Practice Guidelines;
5. Industry Standards;
6. Collaboration;
7. Mechanisms for Monitoring and Improvement;
8. Recourse mechanisms; and
9. Consultation.

The Law Society recommends that additional tools be included to encourage adherence to agreed standards, with potential coercive impacts. These tools include:

1. legislative and regulatory requirements;
2. standards, such as those prepared by the joint committee of the International Organisation for Standardisation ("ISO") and International Electrotechnical Commission ("IEC") (Joint Technical Committee 1);
3. codes of conduct and codes of ethics;
4. accreditation (which would require ongoing education, training and certification to provide services and fulfil functions to a high level, particularly those providing independent advice, assessment and expertise);
5. governance frameworks (additional organisations responsible for the development of ethical AI; independent ethics committees and panels to review for example, data used to train AI); and
6. independent experts.

The Law Society made a submission in March 2019 to the Australian Human Rights Commission and the World Economic Forum's paper *Artificial Intelligence: governance and leadership* supporting the idea of a Responsible Innovation Organisation:

a) to anticipate and articulate issues in a way that empowers civil society and regulators to engage with issues that arise from data driven decision-making and technological innovation;
b) that identifies gaps in regulation and suggests how those gaps might be filled; and
c) that advises where current regulation, expectations as to good practice, or regulatory sanctions are not fit for purpose.

A Responsible Innovation Organisation could also assist with the practical implementation and support mechanisms needed to implement principles for ethical AI. An example of a Responsible Innovation Organisation is the UK's Centre for Data Ethics and Innovation ("CDEI"). he CDEI's first Work Programme and Strategy states that it will analyse and anticipate the opportunities and risks posed by data-driven technology and will put forward practical and evidence-based advice to address them. It states its functions as:

a) Analyse and Anticipate: will convene communities and expertise to provide an overview and insight of opportunities and risks, and review existing regulatory and governance frameworks to identify gaps. It will also carry out thematic projects to enable CDEI to explore live or urgent issues, drawing together lessons from existing work and setting out how they should be taken forward.
b) Reviews: will identify and articulate best practice for the responsible use of data driven technology within specific sectors or for specific applications of technology. They will consider any gaps in governance and make recommendations to the government, as well as advice to regulators, creators and users of data-driven technology as to how those gaps should be addressed.[14]

Question 6: Are there already best-practice models that you know of in related fields that can serve as a template to follow in the practical application of ethical AI?

The Law Society would recommend the EU's Ethics Guidelines for Trustworthy AI as a useful resource. It contains a Pilot Version AI Assessment List which is practical and useful.

Question 7: Are there additional ethical issues related to AI that have not been raised in the discussion paper? What are they and why are they important?

The Discussion Paper sets out as the first core principle that the AI system must generate benefits for people that are greater than the costs. There is little discussion, however, about the benefits to an individual versus the benefits to a community. It is possible that an AI system may hold potential benefit to the community but is detrimental to the individual and vice versa. There is some limited discussion on this issue in section 6.1.3 relating to automated vehicles, but we submit this deserves further discussion as an important ethical issue.

Part 2

The Law Society offers the following general comments on the Discussion Paper.

Role of ethicists and other specialists in developing an AI ethics framework

The Law Society considers that a convening a multi-disciplinary group that includes ethicists, social scientists, data scientists, privacy specialists and lawyers, would provide depth to specialist areas and would greatly benefit the development of an AI ethics framework.

Significant and insignificant risk

The Law Society submits that the table on page 64 of the Discussion Paper contains some contradictions in relation to privacy. While the far right column acknowledges that small numbers of individuals affected could be viewed as a major or critical consequence, the left hand columns note that use of the personal information of a small number of individuals would be a minor or insignificant risk.

We consider that making such a determination on the number of individuals is dangerous. An insignificant or minor risk could only be said to arise if a small number of individuals suffered no consequences or minimal inconvenience as a result of the non-consented use. We recommend the concepts of significant and insignificant risk be revisited in a more considered way.

[14] <https://www.gov.uk/government/groups/centre-for-data-ethics-and-innovation-cdei>.

International initiatives to build a global framework

Numerous initiatives are currently underway by other government, semi-government and private entities, including various health departments, the Australian Government Digital Transformation Agency and the Australian Human Rights Commission, on the ethical and regulatory issues that arise in relation to AI. We recommend interested organisations throughout Australia co-operate and collaborate on an AI ethics framework.

AI will have global impact and the use of AI products will not be limited by national borders. Global opportunities require a global solution. We consider that international collaboration on appropriate standards should be encouraged. Various international standards organisations are evaluating global standards for AI. Standards are usually developed through committees of experts and relevant stakeholders. The joint committee of the ISO and IEC (Joint Technical Committee 1) can enlist countries to collaborate on international standards.[15] Given the range of international initiatives, we consider that, once consensus is reached on an approach in Australia, Australia should take a more active role and partner with like-minded international organisations to develop a shared approach to AI development and the embedding of a high standard of privacy and ethical principles in AI design.

Data Governance

While the central role of data in AI is acknowledged in the Discussion Paper, we recommend that it be more fully explored in developing an AI ethics framework. Effective data governance requires a bottom-up review of processes, systems and methodologies. This is to ensure that principles are properly embedded in everyday operations so that application of principles is reliably repeatable, robust (taking due account of the range of contexts and circumstances in which AI is being deployed and the variety of decisions to be aided or powered by AI), and with appropriate transparency and reporting built in so as to ensure ongoing oversight and feedback loops enabling continuous refinement and improvement.

The Law Society considers that additional attention should be given to the nature and quality of the data used to develop and train AI. Not all data is equal, with the quality and accuracy of datasets varying greatly. This can adversely inform and impact the end result of AI. Standards, training and education need to be consistently applied and provided to improve the quality and accuracy of data, collection practices and methodologies.

We recommend that future consultation documents set out the importance of assessing the data sets being captured and analysed to determine:

a) whether there is a potential breach of an individual's privacy (for example, if the data set is so small that it is clear who the individuals are),
b) if biases are likely to be entrenched by using the data sets; and
c) whether AI is the right tool to analyse the data.

Lethal Autonomous Weapons

The Discussion Paper acknowledges the risks and benefits of the use of AI and the fact that there may be many unanticipated and latent impacts, the effects of which

[15] <https://www.iso.org/isoiec-jtc-1.html>.

may not be apparent or fully understood until a significant time after introduction. However, we note there is an immediate concern about the development and proliferation of 'robot killers', or lethal autonomous weapons, which do not require human intervention to inflict lethal harm. We consider CSIRO / the Department should further explore this use in the Framework developed for ethical AI.

Discrimination

There are various references throughout the Discussion Paper to 'unfair discrimination'. Under discrimination law at both the Federal, and State and Territory level, there is no 'unfair' or 'fair' discrimination. A person either discriminates against another person or does not. As set out above under the heading 'Consistent Terminology' the Law Society recommends that the correct terminology be used in future consultations and policy documents, to ensure organisations employing AI systems have certainty of their compliance requirements.

New Regulation

Principle 3 provides: 'The AI system must comply with all relevant international, Australian Local, State/Territory and Federal government obligations, regulations and laws'. The Discussion Paper makes little comment on the fact that new regulation may also be required.

Section 5.5 discusses the topic of medical predictions and comments that 'AI systems used in health care require close management and gold standard research before implementation.' The Law Society agrees with this comment but notes that AI is well advanced in health care, with limited ethical and regulatory guidelines. The Royal Australian and New Zealand College of Radiologists has prepared its own draft ethical principles for AI in medicine and the Therapeutic Goods Administration has some regulatory oversight, where the AI system fits the definition of 'medical device' in the *Therapeutic Goods Act 1989.* A 'medical device' is essentially limited to software involved in diagnosis, prevention and monitoring of disease and does not apply to health and lifestyle apps or software that does not meet the legislative definition (section 41BD).

We recommend that future consultation or policy documents acknowledge that new laws or regulation may be needed to guide the ethical deployment of AI.

Thank you again for seeking our feedback on this Discussion Paper. Should you have any questions in relation to this submission, please contact Adi Prigan, Policy Lawyer, on 9926 0285 or email adi.prigan@lawsociety.com.au.

Yours sincerely,

Elizabeth Espinosa

Elizabeth Espinosa,
President