SalingerPrivacy

We know privacy inside and out.

Submission in response to the *Safe and Responsible AI in Australia - Discussion Paper 2023*

Australian Government, Department of Industry, Science and Resources

3 August 2023

Salinger Consulting Pty Ltd

ABN 84 110 386 537

PO Box 1250, Manly NSW 1655

www.salingerprivacy.com.au

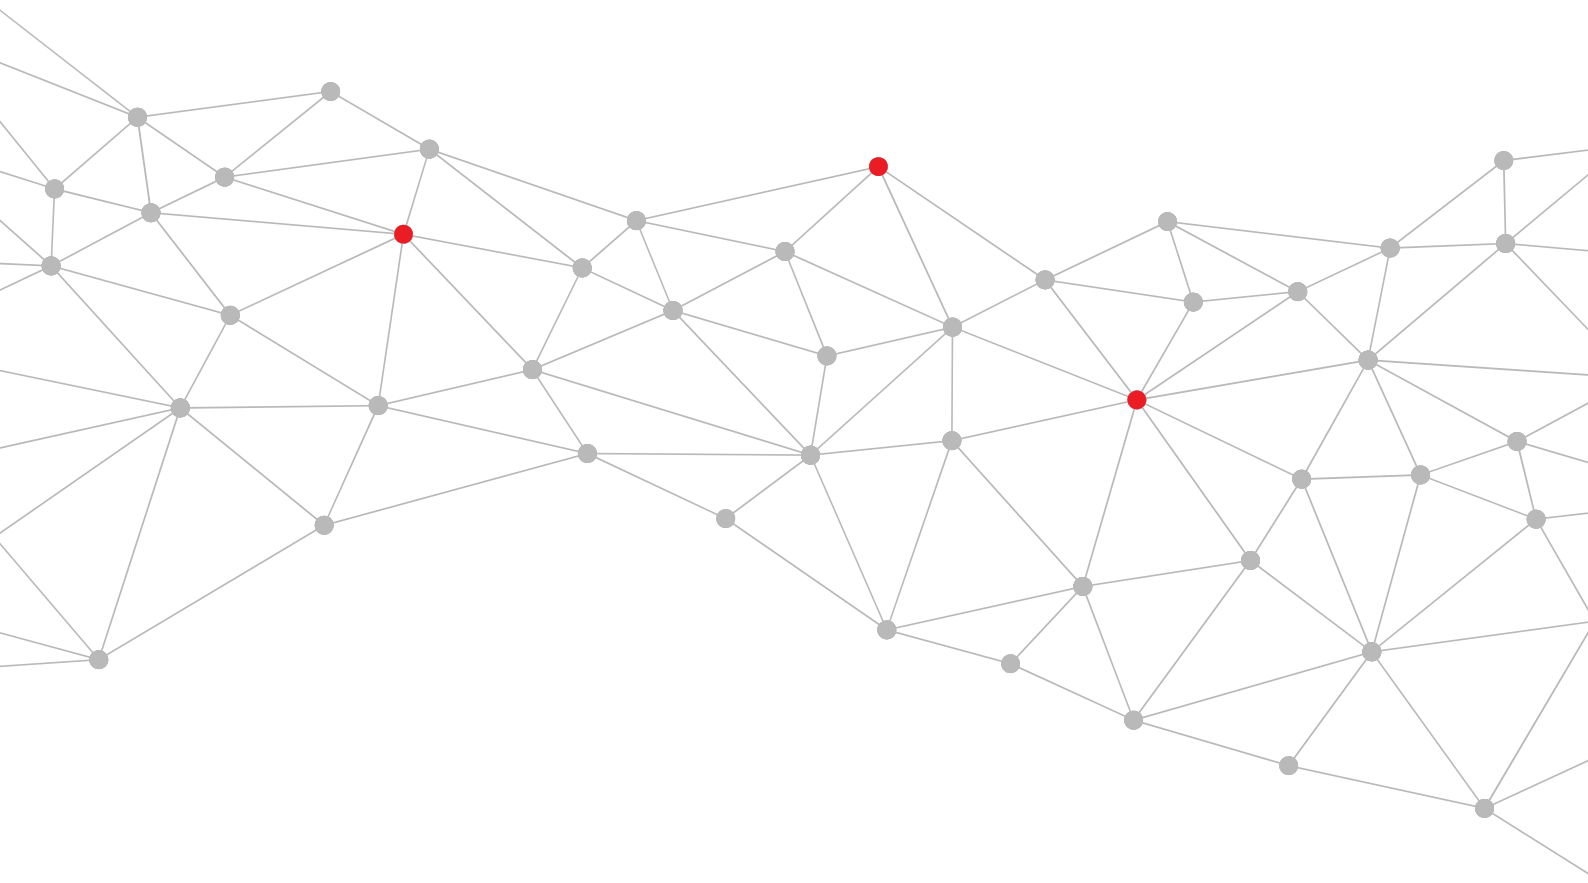

Covering letter

3 August 2023

Department of Industry, Science and Resources
By online submission

Dear Safe and Responsible AI in Australia review team,

Thank you for the opportunity to make a submission in relation to the *Safe and Responsible AI in Australia - Discussion Paper 2023*.

Please find our submission attached.

We have no objection to the publication of this submission, and no redactions are required prior to publication.

Please do not hesitate to contact me if you would like clarification of any of the comments made in this submission.

Anna Johnston
Principal | **Salinger Privacy**

Introduction and overview position

We welcome the release of the *Safe and Responsible AI in Australia - Discussion Paper 2023* (the Discussion Paper) by the Department of Industry, Science and Resources (DISR).

From setting insurance premiums to deciding who gets a home loan, from predicting the risk of a person re-offending to more accurately diagnosing disease, algorithmic systems – especially those turbo-charged by AI – have the ability to re-shape our lives. Automated decision-making systems are increasingly being used to make predictions, recommendations, or decisions vital to individuals and communities in areas such as finance, housing, social welfare, employment, education, and justice – with very real-world implications. As the use of AI and algorithmic systems increases, so too does the need for appropriate auditing, assessment, and review.

The Australian Human Rights Commission (AHRC) for example has recommended, amongst other things, mandatory human rights impact assessments before AI is deployed in administrative decision-making, a ban on the use of 'black box' algorithms by government, a moratorium on the use of facial recognition technology in law enforcement settings, and the creation of an AI Safety Commissioner. We support these and other ideas for achieving the appropriate regulatory response to the challenges posed by AI.

We submit that:

- robust and effective regulation is an enabler of innovation, not a barrier
- existing privacy legislation needs strengthening if any regulatory approach to AI is to succeed
- a risk-based approach is to be preferred to a 'one size fits all' approach
- a risk-based approach will only succeed if it is mandated through legislation with effective regulatory oversight, and
- useful lessons about designing a risk-based approach and the role of Algorithmic Impact Assessment (AIA) can be learned from the evolution of Privacy Impact Assessment (PIA) practice over the past two decades.

This submission offers commentary on a number of the consultation questions posed in the Discussion Paper, but in particular in relation to:

- the gaps in existing legislation and regulatory approaches, focussing on privacy law and practice, and
- how to design a risk-based approach which is meaningful, scalable and effective.

Stronger privacy regulation is essential

Privacy is interwoven with other rights. By upholding privacy, other rights and values can also be enabled or supported, such as:

- freedom of speech / expression
- freedom of association and movement
- freedom of religion
- freedom from discrimination
- the right to a fair trial
- equal access to markets and opportunities
- autonomy, free will and individual dignity.

Therefore, requiring organisations to build privacy protections into AI systems does more than just mitigate privacy law compliance risks for those organisations. It also helps to mitigate the risk of creating a range of other harms for individuals, which we refer to as privacy-related harms. To assess the risks of AI, organisations therefore need to consider a range of 'downstream' harms, rather than limiting their view of privacy harm solely to non-compliance with privacy law.

Privacy risks can arise from any AI system which is, essentially, about humans. AI systems can be developed using personal information about humans, and/r be deployed to make predictions, classifications, scores, recommendations, or decisions about humans. We submit that the Privacy Act is well placed to regulate how those risks are managed.

Further, unlike consumer law or corporations law which do not regulate the public sector, the Privacy Act is well placed to reach many organisations economy-wide – particularly if the small business exemption is to be abolished, as has recently been proposed in the current review process.

We therefore urge DISR and the Minister to support specific reforms to the *Privacy Act 1988* (Cth), in particular to:

- define 'personal information' to clearly include information where an individual may be singled out and acted upon, *even if their identity is not known* (i.e. individuation)
- abolish the small business exemption
- strengthen the definition of consent, and require express consent for high impact activities
- introduce a 'fair and reasonable' test in relation to collection, use and disclosure of personal information, and
- provide the privacy regulator with additional enforcement tools.

The threshold definition is no longer suitable

There is a fundamental concern with the current wording of the Privacy Act, which is no longer fit for purpose in the digital age. Today, all privacy rights for individuals, and all obligations in organisations, hinge on the threshold definition of 'personal information'. Personal information as currently defined requires a person to be at least 'reasonably identifiable' from data, before that data will fall within the regulatory scope of the Privacy Act. However this 'identifiability' test is no longer fit for purpose.

It is our strong submission that rapid advances in technologies, including artificial intelligence and facial recognition, mean that 'not identifiable by name' is no longer an effective proxy for 'will suffer no privacy harm'.[1] The Privacy Act urgently requires updating, by *explicitly* incorporating into the threshold definition of 'personal information' the concept of *individuation*.

Individuation has been used to describe the 'singling out' of a person from a crowd – a threat to privacy, autonomy and dignity.[2] Call it 'indirect identification', call it 'singling out', call it 'distinguishing from all others', call it 'individuation' - it doesn't matter how you describe the concept. What does matter is that the wording of the definition in the Privacy Act must be clear on the face of it that what is within scope for regulation under the phrase 'personal information' includes information where an individual may be singled out and acted upon, *even if their identity is not known*.

We know the harms that can arise from individuation; and these harms are exacerbated by the use of AI and other automated decision-making systems. These harms can arise from the online tracking, profiling and targeting which forms the basis for online behavioural advertising, but also include surveillance, discrimination, behavioural engineering, and misinformation.[3]

To ensure the Privacy Act is fit to reflect the realities of the digital ecosystem, as well as meet the challenges of the future, it is critical that the definition of 'personal information' is itself fit for purpose. A strengthened statutory definition of 'personal information' will better deliver clarity for regulated entities, align with the privacy laws of our trading partners, and meet the expectations of Australians.

[1] Anna Johnston, 2020, "Individuation: Re-imagining Data Privacy Laws to Protect Against Digital Harms" (electronic). Brussels Privacy Hub. 6 (24); available at https://brusselsprivacyhub.eu/publications/wp624.html
[2] Greenleaf, Graham; Livingston, Scott (2017). "China's Personal Information Standard: The Long March to a Privacy Law". *Privacy Laws & Business International Report* (150): 25–28; available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3128593
[3] For a further discussion on the harms associated with individuation, please refer to our Blog 'Big Tech, Individuation, and why Privacy must become the Law of Everything' at https://www.salingerprivacy.com.au/2022/03/22/big-tech-blog/

Unfortunately the recommended amendments to the definition in the Attorney-General's Department's published set of proposals will *not* achieve this aim.[4] We urge further improvements before a reform Bill is finalised.

How AI systems challenge other aspects of the existing privacy law

Many privacy laws around the world are based on a set of OECD Guidelines, originally drafted in 1980.[5] The increased sophistication and availability of technologies including algorithmic systems pose new challenges to many longstanding pillars of privacy protection, including data minimisation, purpose limitation, and transparency. For example, AI systems rely on repurposing massive amounts of data, function in a way that is opaque to most people (and sometimes even to those who developed them), and can result in generation of new meanings, information or outcomes not foreseeable at the time of the original data collection.

At a fundamental level, privacy laws govern the ways that personal information can be collected and used, regardless of whether processing is done by manual or automated means.

Physical and technical limits on manual processing, computer memory and speed, and traditional programming techniques used to provide default safeguards on the scale and scope of information processing. There was only so much processing that could be done. However, the increased availability of readily available and cheap data storage alongside increasingly sophisticated algorithmic techniques means that data processing and analytics that may have been impossible at the time privacy principles were being drafted are now commonplace. The ability to process information on such a large scale and at such great speeds amplifies the possible harms that can be caused by such systems.

AI systems can also infer information about someone, without that person ever having voluntarily provided their personal information. In 2017 it was found that an individual's sexuality could be predicted from seemingly innocuous data points on Facebook.[6] This creates challenges to the established community expectation that personal and sensitive information should be collected directly from the individual (so that the individual can exercise choice over whether or how to provide the requested information), and also raises questions about the accuracy of the information, as well as the ethics of using personal information that has been inferred from other pieces of information.

---

[4] Our analysis of the flaws in the proposed new definition of 'personal information', and how it could be fixed, is at https://www.salingerprivacy.com.au/2023/04/19/one-extra-sentence/

[5] *OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data,* see: https://www.oecd.org/sti/ieconomy/oecdguidelinesontheprotectionofprivacyandtransborderflowsofpersonaldata.htm

[6] 'Enhancing Transparency and Control when drawing data-driven Inferences about Individuals,' Daizhou Chen, Samuel P Fraiberger, Robert Moakler, and Foster Provost, *Big Data*, Volume 5, Issue 3, September 2017. See: https://www.liebertpub.com/doi/full/10.1089/big.2017.0074

Given that the OAIC has issued advice that personal information inferred about an individual from other information is considered to be 'collected' for the purposes of the Privacy Act,[7] algorithmic systems making inferences about people will need careful assessment, to ensure compliance with all Collection privacy principles, including the limitations on the collection of sensitive information, indirect collection, or collecting personal information without meeting tests such as reasonable necessity.

Further, algorithmic systems which rely for their development on the re-use of personal information originally collected for a different purpose will face significant hurdles in complying with Use principles, which typically prohibit the secondary use of personal information except in limited circumstances.

Algorithmic systems, and in particular AI, also pose challenges to some of the traditional ways to mitigate against privacy risks, such as relying on de-identification or consent. For example, de-identifying data once you have it does not resolve any of the compliance challenges faced in relation to the original collection of the data, as outlined above.

Even where it may be possible to de-identify data before it is used as a training dataset, de-identification is not a perfect solution to compliance with rules limiting secondary data use, as there will likely be a residual risk of re-identification. In fact, AI systems can actually be a tool to re-identify previously de-identified data. Furthermore, once algorithmic systems are deployed in the real world, the collection, use and disclosure of personal information will still need to be justified.

Organisations wishing to use de-identification as a risk-mitigation strategy need to ensure that they are not treating it as a privacy risk cure-all. Salinger Privacy has published an eBook, 'Demystifying De-identification' which provides an introductory guide to the techniques, benefits and limitations of de-identification.[8]

Nor will 'consent' resolve privacy risks in this context.

Taking a simplistic 'tick-box' approach to gaining permission to use or disclose someone's personal information is inappropriate, unethical, and in some cases, unlawful. In order to be a valid mechanism to authorise a collection, use or disclosure of personal information, consent must be freely given, informed and specific. The potential for unintended, or unforeseen, outcomes or inferences reduce people's ability to fully comprehend what it is they are consenting to; and indeed, can reduce the ability for organisations to understand what it is they are asking people to consent to.

[7] 'Guide to data analytics and the Australian Privacy Principles,' *The Office of the Australian Information Commissioner,* March 2018. See: https://www.oaic.gov.au/privacy/guidance-and-advice/guide-to-data-analytics-and-the-australian-privacy-principles/

[8] 'Demystifying De-identification,' *Salinger Privacy,* Edition 5, March 2022; see https://www.salingerprivacy.com.au/downloads/demystifying-deid/

This means consent is an even *less* appropriate means to authorise data flows in the context of AI than in many other contexts, as most people are not likely to understand the technology, nor be aware of the possible consequences. As a result, 'informed' and 'specific' consent can be close to impossible to achieve. For example in April 2020 the South Korean regulator, the Personal Information Protection Commission, imposed sanctions and a fine on the developer of an AI chatbot which had used customers' messages from a messaging app to train its chatbot, finding that a 'new service development' clause in the terms to log into the messaging apps did not amount to users' consent, because the description was insufficient for users to anticipate that their messages would be used to develop and operate a chatbot.[9]

Further, as the use of algorithmic systems increases, so too does the power imbalance between organisations processing data, and individuals. This poses a challenge to the requirement that consent be freely given.

It is of particular importance for government bodies to ensure they are getting the balance right when using algorithmic systems in areas such as education, healthcare, justice and access to services. It is not appropriate to call something a 'consent-based' model, when in reality, individuals have very little opportunity to refuse or opt-out of government-run systems, especially if the consequence is not receiving a particular benefit or service. Other scenarios in which consent cannot be freely given include employee/employer relationships, tenant/landlord relationships, and in relation to access to public spaces, services, infrastructure or digital platforms.

Regulator powers and tools – new approaches needed

Drawing on international approaches, we also urge DISR to work with the Attorney-General's Department to reform the Privacy Act and/or the OAIC enabling legislation, so that the Privacy Commissioner (and/or any specific AI Safety Commissioner) has the following powers:

- the power to issue algorithmic disgorgement orders, as practiced by the Federal Trade Commission in the USA, and
- a 'veto' power over 'high privacy impact' projects where the risks to privacy cannot be mitigated satisfactorily, as enjoyed by European data protection authorities.

Under the GDPR, 'high privacy impact' projects not only require a mandatory Data Protection Impact Assessment (DPIA) to be conducted; those DPIAs must also be submitted to the relevant regulator, who then has a defined period in which they can order a pause or stop to the project.

[9] 'South Korea: The first case where the Personal Information Protection Act was applied to an AI system,' *Future of Privacy Forum,* May 2021. See: https://fpf.org/blog/south-korea-the-first-case-where-the-personal-information-protection-act-was-applied-to-an-ai-system/

Designing a risk-based regulatory approach

We welcome the general approach of the DISR, in seeking to develop a risk-based regulatory approach, rather than a 'one size fits all' approach. Laws need to remain flexible because of the different contexts in which AI systems are deployed, from banks to hospitals, and from the retail sector to military applications.

We do however strongly caution that self-regulatory models to implement such an approach will fail. A risk-based approach will only succeed if it is mandated through legislation, with effective regulatory oversight.

We also suggest that any regulatory approach must offer clarity, for both regulated entities and the Australians who seek the protection of the law. Clear guardrails allow innovation in a safe and responsible way, whereas unclear (or unenforced) requirements constrain and penalise the good actors, but create space for bad actors to profit.

In particular, we caution against any regulatory approach which only exalts companies to follow bland motherhood statements. Telling a company to 'build in privacy by design' is, on its own, meaningless.

The draft risk management approach in Attachment C of the Discussion Paper offers a good starting point. However much more work will be needed.

We suggest:

- Further clarity around what types of applications will constitute low, medium or high risk; see our further discussion below. For example, from the descriptions in the Discussion Paper we could not assess where the deployment of facial recognition technology in the retail sector would land, let alone its deployment in sporting stadiums or prisons.
- The criteria must encompass the type of technology *and* its use case. A recommendation engine suggesting what TV show to watch next poses a very different risk level to a recommendation engine being deployed in healthcare, or warfare.
- That the risk level to be described as low, medium or high should be made with reference to the *inherent* risks posed by the application, *prior to* implementation of any risk controls; otherwise organisations can make self-serving determinations that the way they will implement the technology will render it to a lower level of risk, and thus the legal obligations will not apply to them.
- The requirement for recurring training must include training for responsible decision makers on issues such as confirmation bias.
- The requirement for an 'impact assessment' must clarify: assessment in relation to what types of impacts, and for whom? (Taking the experience of privacy impact

assessment practice as an example,[10] some organisations think the assessment is simply one of ticking off their own legal compliance. Other organisations expand their gaze beyond legal compliance to consider other impacts on their organisation such as reputational damage if their conduct does not meet community expectations. However the OAIC guidance is clear that the purpose of a PIA is also to assess the privacy impacts *on affected individuals*.)

- It should also explain who should conduct the impact assessment, when, how, and what happens next.
- Clear guidance about what the impact assessment is supposed to be assessing will also be critical. For example, concepts of necessity, legitimacy and proportionality will be key. The proposal for a new 'fair and reasonable' test in the Privacy Act will also be relevant here.
- We often suggest a cascading set of four questions, which must be asked of any project which uses data: Is it legal? Is the data fit for our purpose? *Should* we do this (aka is it ethical)? And finally – assuming the first three answers were 'yes' - How can we do this safely?
- 'Impact' or risk assessment should be more than just demonstrating the absence of bias in the data or the algorithm. Assessors should be looking for equitable impact of both harms and benefits from the system.
- An Algorithmic Impact Assessment (AIA; see further discussion below) should incorporate within it, as relevant: a Privacy Impact Assessment, a re-identification risk assessment, data bias assessment, model inversion attack testing, dataset shift testing, and cyber security assessment.
- The requirement for an impact assessment must include a requirement to provide a copy of the impact assessment report to the relevant regulator (e.g. the OAIC if there is not an AI Safety Commissioner); and the regulator must have the power to pause or stop the project.
- The requirement to conduct an impact assessment must include a requirement to publish a copy of the impact assessment report, redacted as necessary to protect data security or proprietary information.
- High risk projects should also require a feedback channel for affected individuals; an appeal mechanism including appropriate remedies; auditable code; and a 'kill switch'.
- High risk projects should also require an explanation of the data inputs and the assumptions built into the algorithmic model; a description of the data provenance (the training data, testing data, and the data used in deployment); and published tests for fairness, bias, accuracy, precision and recall.
- High risk projects should be required by law to comply with relevant standards such as 42001.

[10] See our reflections from 20 years of practice conducting PIAs in Seven tips to ensure Privacy Impact Assessments are useful; and How to implement a PIA framework.

- Activities which are *inherently high risk*, and which still pose a *residual* high risk after controls have been implemented, should instead be *prohibited by law*, unless the regulator has determined that there are overriding public benefits or public interests at stake.

Factors which we suggest should put applications in the *medium* to *high* inherent risk category should include if the AI system:

- involves a critical sector or function (such as healthcare, transport, insurance, employment, finance/credit, education, housing, welfare, benefits, taxation, political processes, defence, national security, law enforcement or the legal system)
- may have a critical impact (legal or financial effects, risk of death, damage or injury)
- will make decisions, predictions or recommendations at a large scale
- will make decisions, predictions or recommendations which have the effect of charging different people different prices for the same service; or making different offers (including excluding people from seeing any offer); or setting risk scores
- uses data inputs which include data or inferences about protected attributes (as recognised in discrimination law) or about sensitive personal information (which are subject to special restrictions under privacy law) such as gender, race/ethnicity, age, sexuality, religion, criminal record, biometrics, biometric templates, health or disability information
- will impact Indigenous communities; or
- that makes decisions, predictions or recommendations impacting particularly vulnerable groups (such as children, people with disabilities, refugees or incarcerated people).

Building a framework assessing AI risks

Building trustworthy AI: the Four D's Framework

The build of an algorithmic system comprises four stages:

- design
- data
- development, and
- deployment

By viewing algorithmic systems through the lens of these 'Four D's', Salinger Privacy has developed a framework for considering privacy and related risks in algorithmic systems, throughout the lifecycle of the entire project or program.

Building trustworthy systems will never be as simple as a checklist exercise because the specific context, purpose, data, and communities will vary from system to system. However our framework includes a range of features of trustworthy systems across the Four D's - design, data, development, and deployment - to offer a starting point for organisations to consider.

Algorithmic systems, law, and ethics

Any algorithmic system should be assessed across four key areas:

- Compliance with privacy laws
- Compliance with anti-discrimination laws (noting that in some cases differentiation may be necessary, such as to offer age-specific services to particular age groups. This is not the same as unlawful discrimination)
- Compliance with consumer protection laws[11] (e.g. prohibitions on misleading conduct), and
- FEAT issues: use existing frameworks to think about other rights and values to be protected and supported, not already included in the above.

Too often people mischaracterise something as an ethical issue which is in fact a legal issue, such as unlawful discrimination, and should be dealt with as such. Therefore we

[11] Other laws such as trade practices and product safety may also come into play, but as they have less cross-over with privacy issues we do not cover them in our analysis.

define 'unfair' to incorporate not only unlawful conduct or outcomes, but also behaviour or outcomes which are not unlawful (e.g. are not unlawful discrimination, and don't breach privacy or consumer protection law), but which raise ethical issues.

Practices such as differential pricing based on price elasticity (i.e. when a system makes a prediction about the consumer's price sensitivity, and adjusts the pricing to be shown accordingly) can impair consumer choice, lead to unlawful discrimination, or lead to treatment which could be seen as inequitable or 'unfair' but which is not unlawful. For example, a Choice investigation of Tinder showed the use of differential age-based pricing (which would be unlawful discrimination), but also pricing based on suburb of residence which was used as a proxy for income (which could be considered 'unfair', but may not constitute unlawful discrimination).[12]

As we work only in the privacy space, we have deliberately focussed this submission on matters of privacy compliance, as well as related FEAT issues. However, compliance with privacy law can also assist with compliance with anti-discrimination and consumer protection laws, to the extent that they all start with the same thing: personal information about individual customers or users.

A role for AIAs

Evaluating the privacy risk of algorithmic systems is not just a matter of being able to test the maths behind the algorithm. It is also about understanding the legal landscape, community expectations, and social context in which an algorithm will be developed and used, applied, and interpreted. Assessing privacy risk in particular also requires consideration of potential privacy issues beyond data breaches, complying with legal obligations, or more traditional conceptualisations of the right to privacy.

Like a Privacy Impact Assessment (PIA), an Algorithmic Impact Assessment (AIA) may initially seem like extra red tape. However, in conducting an AIA, organisations can decrease their reputational and financial risk. Further, good governance and information management practices are beneficial beyond just privacy compliance requirements. AIAs facilitate good governance of technical systems and encourage organisations to better understand and manage their use of data and decision-making. Organisations that take the extra step to integrate an AIA into the design, development, and deployment of their algorithmic systems are not only demonstrating their commitment to the rights and wellbeing of their clients or customers by applying sensible risk mitigation, but also position themselves as leaders in data governance.

[12] 'Tinder charges older people more,' *Choice,* August 2020. See https://www.choice.com.au/electronics-and-technology/internet/using-online-services/articles/tinder-plus-costs-more-if-youre-older Note that even here, postcode may operate as a proxy for race, in which case it may also lead to unlawful indirect discrimination.

Assessing algorithmic systems through an AIA is an important step in identifying risks early, to avoid or mitigate unintended outcomes which can have profound impact on people's lives.

For example, without the kind of due diligence contained in the process of conducting an AIA, algorithmic systems can unintentionally exacerbate bias, and in some cases even result in unlawful discrimination. Organisations wishing to implement algorithmic systems need to consider anti-discrimination statutes which prohibit discrimination on the basis of 'protected attributes' which include an individual's age, disability, race or ethnic origin, sex, pregnancy or marital status, gender identity and sexual orientation.[13] Ensuring robust privacy protections are in place can help mitigate against discrimination occurring as a result of inappropriate collection and use of personal information. Likewise, an AIA can assist to identify risks in relation to compliance with other laws, such as consumer protection laws, which prohibit misleading conduct. AIAs can also assist organisations to look beyond purely legal compliance requirements, to include broader unethical or unjust impacts of algorithmic systems.

While there is increasing work being done in this space, there is no universal algorithmic system auditing nor simple impact assessment checklist. This is because context matters, and 'one-size-fits-all' approaches rarely meet the needs of a specific community, dataset, use case, or geographic location. There has been difficulty establishing broadly accepted ethical standards for AI, and while some existing AI ethics guides are useful, they do not necessarily encompass the broad scope of algorithmic systems that impact humans.

Recognising that mitigating privacy harms is much more than just a compliance exercise, the Salinger Privacy approach to assessing algorithmic systems for privacy risk goes beyond the legal compliance and technical accuracy of an algorithmic system, to also examine social and ethical impacts. Our guide, *Algorithms, AI, and Automated Decisions – A guide for privacy professionals*,[14] offers privacy professionals a framework for assessing the privacy risks posed by algorithmic systems, and tools to promote the design of trustworthy systems.

Our guide encourages organisations to look beyond just legal compliance, in order to understand, identify, and mitigate against privacy-related harms. We cover concepts such as fairness, ethics, accountability, and transparency (when taken together, sometimes abbreviated to 'FEAT'), which are vital factors to consider when assessing algorithmic systems. We also encourage privacy professionals to think about how to design trustworthy systems more deeply, by looking at both risks and solutions through the lens of 'The Four D's': design, data, development and deployment.

[13] 'Using artificial intelligence to make decisions: Addressing the problem of algorithmic bias,' *Australian Human Rights Commission*, November 2020. See: https://humanrights.gov.au/our-work/rights-and-freedoms/publications/using-artificial-intelligence-make-decisions-addressing

[14] *Algorithms, AI, and Automated Decisions – A guide for privacy professionals*, Salinger Privacy, June 2021; available at https://www.salingerprivacy.com.au/downloads/algorithms-guide/

More detail can be found in our guide about:

- When an AIA will be needed
- What an AIA should assess
- How to integrate AIAs into other risk assessment frameworks to avoid duplication or gaps
- How to assess for factors such as necessity and proportionality
- Different types of bias to look out for, and
- A list of features of 'trustworthy' systems.

Further resources

For more on ways in which genuine accountability and transparency can be achieved in practice for AI and algorithmic systems, see Algorithms, AI, and Automated Decisions – A guide for privacy professionals.

For further details on other points raised in this submission, please see:

- Our detailed submission to the Attorney-General's Department on the Privacy Act Review Report, 2023
- Our analysis of the flaws in the proposed new definition of 'personal information'
- Our analysis of the proposals to introduce 'algorithmic transparency' via the Privacy Act
- Our critique of earlier attempts to manage privacy risks via 'ethical AI principles': 'The ethics of artificial intelligence: start with the law'
- Our reflections on Privacy Impact Assessment as a methodology, after two decades of practice:
    - Seven tips to ensure Privacy Impact Assessments are useful
    - How to implement a PIA framework

About the author

This submission was prepared by Anna Johnston.

Anna Johnston is founder and Principal of Salinger Privacy. Anna has served as:

- Deputy Privacy Commissioner of NSW
- Chair of the Australian Privacy Foundation, and member of its International Committee
- a founding member and Board Member of the International Association of Privacy Professionals (IAPP), Australia & New Zealand
- a member of the Advisory Board for the EU's STAR project to develop training on behalf of European Data Protection Authorities
- a Visiting Scholar at the Research Group on Law, Science, Technology and Society of the Faculty of Law and Criminology of the Vrije Universiteit Brussel
- a Member of the Asian Privacy Scholars Network
- a member of the Australian Law Reform Commission's Advisory Committee for the Inquiry into Serious Invasions of Privacy, and expert advisory group on health privacy, and
- an editorial board member of both the Privacy Law Bulletin and the Privacy Law & Policy Reporter.

Anna has been called upon to provide expert testimony to the European Commission as well as various Parliamentary inquiries and the Productivity Commission, spoken at numerous conferences, and is regularly asked to comment on privacy issues in the media. In 2018 she was recognised as an industry leader by the IAPP with the global designation of Fellow of Information Privacy (FIP).

In 2022, Anna was honoured for her 'exceptional leadership, knowledge and creativity in privacy' with the IAPP Vanguard Award, one of five privacy professionals recognised globally whose pioneering work is helping to shape the future of privacy and data protection. While her day-to-day work involves assisting clients to develop innovative approaches to privacy protection, the Vanguard award was bestowed in reflection of Anna's contributions to the privacy profession, and to the protection of privacy for the benefit of all.

Anna holds a first class honours degree in Law, a Masters of Public Policy with honours, a Graduate Certificate in Management, a Graduate Diploma of Legal Practice, and a Bachelor of Arts. She was admitted as a Solicitor of the Supreme Court of NSW in 1996. She is a Certified Information Privacy Professional, Europe (CIPP/E), and a Certified Information Privacy Manager (CIPM).

About Salinger Privacy

Established in 2004, Salinger Privacy offers privacy consulting services, specialist resources and training.

Our clients come from government, the non-profit sector and businesses across Australia. No matter what sector you are in, we believe that privacy protection is essential for your reputation. In everything we do, we aim to demystify privacy law, and offer pragmatic solutions – to help you ensure regulatory compliance, and maintain the trust of your customers.

Salinger Privacy offers specialist consulting services on privacy and data governance matters, including Privacy Impact Assessments and privacy audits, and the development of privacy-related policies and procedures. Salinger Privacy also offers a range of privacy guidance publications, eLearning and face-to-face compliance training options, and Privacy Tools such as templates and checklists.

Qualifications

The comments in this submission do not constitute legal advice, and should not be construed or relied upon as legal advice by any party. Legal professional privilege does not apply to this submission.

SalingerPrivacy

We know privacy inside and out.

Salinger Consulting Pty Ltd

ABN 84 110 386 537

PO Box 1250, Manly NSW 1655

www.salingerprivacy.com.au