

# SalingerPrivacy

We know privacy inside and out.

## Submission in response to the Inquiry into Artificial Intelligence (AI) in New South Wales

Portfolio Committee No. 1 – Premier and Finance,  
Legislative Council, Parliament of New South Wales


12 October 2023

Salinger Consulting Pty Ltd

ABN 84 110 386 537

PO Box 1250, Manly NSW 1655

[www.salingerprivacy.com.au](http://www.salingerprivacy.com.au)



## Covering letter

12 October 2023

Portfolio Committee No. 1 – Premier And Finance  
Legislative Council  
Parliament of New South Wales

Dear Committee Members,

I am writing to submit the attached submission to your Inquiry into artificial intelligence (AI) in New South Wales

I have no objection to the publication of this submission.

Please do not hesitate to contact me if you would like clarification of any of these comments.

Anna Johnston  
Principal | [Salinger Privacy](#)

Salinger Consulting Pty Ltd  
ABN 84 110 386 537  
PO Box 1250, Manly NSW 1655  
[www.salingerprivacy.com.au](http://www.salingerprivacy.com.au)

# Introduction and Overview Position

---

We welcome the opportunity to make a submission to the Inquiry into Artificial intelligence (AI) in New South Wales.

Salinger Privacy's expertise is in privacy law and practice. We have operated as a specialist consultancy in NSW since 2004, and our Principal is a former Deputy Privacy Commissioner of NSW.

Our submission is focussed on whether current laws regarding AI in NSW that regulate privacy and data security are fit for purpose (paragraph (i) in your Terms of Reference), and recommendations to manage the risks, seize the opportunities, and guide the potential use of AI by government (paragraph (m) in your Terms of Reference).

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 social housing, education, policing and justice, health, and access to other government services – 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.

We submit that:

- robust and effective regulation is an enabler of innovation, not a barrier, and
- existing privacy legislation which regulates NSW public sector agencies needs strengthening in order to best manage the risks, seize the opportunities, and guide the potential use of AI by government.

This submission offers commentary on the gaps in existing legislation and regulatory approaches, focussing on privacy law and practice, and the changes that should be implemented to ensure that AI is deployed by government in a manner that respects the privacy of individuals, and protects personal information.

# 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, NSW public sector agencies 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 be deployed to make predictions, classifications, scores, recommendations, or decisions about humans.

We submit that the *Privacy and Personal Information Protection Act 1998* (PPIP Act) and *Health Records and Information Privacy Act 2002* (HRIP Act) are suitable vehicles to introduce a scheme to better regulate how those risks are managed within the NSW public sector.

We therefore urge the Committee to recommend specific reforms to the PPIP Act and the HRIP Act, to:

- amend the definition of 'personal information' to include information where an individual may be singled out from all others and acted upon, *even if their identity is not known*
- define consent in line with the proposed new definition in the federal Privacy Act
- introduce specific regulation of the use of personal information in automated decision-making that broadly aligns with impending changes in the Federal Privacy Act

- introduce a ‘fair and reasonable’ test in relation to collection, use and disclosure of personal information that broadly aligns with impending changes in the Federal Privacy Act
- introduce a requirement on NSW public sector agencies to conduct Privacy Impact Assessments (PIA) of inherently high risk activities, which should be defined to include the use of AI or automated decision-making; and
- provide the NSW Privacy Commissioner and the NSW Civil and Administrative Tribunal (NCAT) with updated enforcement tools, including an increase in the permitted compensation under s55 of the PPIP Act, an explicit power for NCAT to order algorithmic disgorgement as a remedy for breaches associated with AI systems; and the Privacy Commissioner with power to order ‘pause’ on the development of AI systems by NSW public sector agencies.

We address some of these issues in more detail below.

## The threshold definition is no longer suitable

There is a fundamental concern with the current wording of the PPIP and HRIP Acts, which are no longer fit for purpose in the digital age. Today, all privacy rights for individuals, and all obligations on NSW public sector agencies, hinge on the threshold definition of ‘personal information’.

If the data used in the development or deployment of AI does not meet the definition of ‘personal information’, it is unregulated. That means that if a business supplying an AI solution or automated decision-making system to a NSW public sector agency can successfully argue that some data is not ‘personal information’, they can collect, use, disclose and trade the data with impunity.

This argument is routinely made by companies utilising AI tools such as facial recognition, including in the context of policing. Recent cases which touch on the definition of ‘personal information’ in the context of AI are the *7-Eleven* case, the *Clearview AI* case, and the *Australian Federal Police* case.<sup>1</sup>

Personal information is currently defined in NSW privacy law as:

“information or an opinion ... about an individual whose identity is apparent or can reasonably be ascertained from the information or opinion”.<sup>2</sup>

We submit that the ‘identifiability’ test component of this definition is no longer fit for purpose.

---

<sup>1</sup> A summary of these cases and links to the full text is available at <https://www.salingerprivacy.com.au/2022/04/11/oaic-determinations-blog/>

<sup>2</sup> Section 4(1) of the PPIP Act.

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’.<sup>3</sup> The PPIP and HRIP Acts urgently require 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.<sup>4</sup> 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 Acts 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.<sup>5</sup>

To ensure the PPIP and HRIP Acts are fit to reflect the realities of the digital ecosystem, as well as to help NSW public sector agencies 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 the community.

In September 2023 the Australian Government announced its intention to introduce a Bill in 2024 to revise the definition of ‘personal information’ in the federal Privacy Act. The report stated: “Importantly, the Government considers that an individual may be reasonably identifiable where they are able to be distinguished from all others, even if their identity is not known”.

We submit that the NSW definition should be revised in line with the proposed new definition in the federal Privacy Act.<sup>6</sup>

---

<sup>3</sup> 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>

<sup>4</sup> 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](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3128593)

<sup>5</sup> 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/>

<sup>6</sup> See <https://www.ag.gov.au/rights-and-protections/publications/government-response-privacy-act-review-report>

## 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.<sup>7</sup> 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. They are designed to be 'technology neutral'.

However 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.

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. For example in 2017 it was found that an individual's sexuality could be predicted from seemingly innocuous data points on Facebook.<sup>8</sup> 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.

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.

---

<sup>7</sup> *OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data*, see: <https://www.oecd.org/sti/ieconomy/oecdguidelinesontheprivacyandtransborderflowsofpersonaldata.htm>

<sup>8</sup> '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>

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.<sup>9</sup>

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.

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.<sup>10</sup>

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.

---

<sup>9</sup> 'Demystifying De-identification,' *Salinger Privacy*, Edition 5, March 2022; see <https://www.salingerprivacy.com.au/downloads/demystifying-deid/>

<sup>10</sup> '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/>



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, social housing and access to other government 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.

We submit that the proposal included in the 2021 Discussion Paper released as part of the review of the federal Privacy Act that would see consent defined in the Act as "being voluntary, informed, current, specific, and an unambiguous indication through clear action"<sup>11</sup> presents a reasonable approach that will help prevent ongoing privacy harms caused by entities exploiting unclear consent requirements.

We recommend that the NSW Government should similarly add a definition of 'consent' into the PPIP and HRIP Act which spells out the necessary elements of a valid consent, in line with the reforms to the federal Privacy Act to be introduced in 2024.

## Privacy Impact Assessments for high-risk activities

We recommend that a requirement be introduced for all NSW public sector agencies to conduct a Privacy Impact Assessment (PIA) for all for all inherently 'high privacy risk' projects, which should be defined to include the use of AI or automated decision-making

This would mirror the legal requirement already in place for Australian Government agencies under the federal Privacy Act since 2018, which is proposed to be extended to the private sector under the impending 2024 reforms to the federal Privacy Act.

As noted by the then NSW Privacy Commissioner in her "Guide to Privacy Impact Assessments in NSW":

"A Privacy Impact Assessment ... can help you to identify and minimise privacy risks when you are starting a new project or making changes to existing initiatives. A PIA is one way to implement 'privacy by design' in your organisation's practices, and it can help you to build and demonstrate compliance with privacy laws."<sup>12</sup>

As part of mandating the completion of a PIA for all high-risk activities it is recommended that this direction be extended to include circumstances where algorithmic based systems

---

<sup>11</sup> Review of the Privacy Act (Cth) 1988, Discussion Paper 2021 p.11; available at <https://consultations.ag.gov.au/rights-and-protections/privacy-act-review-discussion-paper/>

<sup>12</sup> See <https://www.ipc.nsw.gov.au/guide-privacy-impact-assessments-nsw>

such as AI are being implemented. This may take the form of an an Algorithmic Impact Assessment (AIA).

Like a PIA, an 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.

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.<sup>13</sup> 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.

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*,<sup>14</sup> 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

---

<sup>13</sup> '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>

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

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.

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.

## Regulator powers and tools – new approaches needed

Drawing on international approaches, we also urge the Committee to recommend reforms to the PPIP Act, so that the Privacy Commissioner or NCAT has the following powers:

- the power to issue algorithmic disgorgement orders, as practiced by the Federal Trade Commission in the USA,<sup>15</sup> 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.<sup>16</sup>

We also submit that there should be an increase in the permitted compensation which can be ordered to people who have suffered demonstrable loss or harm under s55 of the PPIP Act, which has been capped at \$40,000 since 1998.

---

<sup>15</sup> The FTC has notably used the remedy of ordering algorithmic disgorgement in the Cambridge Analytica matter as well as against the company formally known as Weight Watchers; see: <https://www.protocol.com/policy/ftc-algorithm-destroy-data-privacy> and <https://www.ftc.gov/news-events/news/press-releases/2022/03/ftc-takes-action-against-company-formerly-known-weight-watchers-illegally-collecting-kids-sensitive>

<sup>16</sup> Under the General Data Protection Regulation (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.

# Further resources

---

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

- Our detailed [submission](#) to the Australian Government, Department of Industry, Science and Resources in response to its *Safe and Responsible AI in Australia - Discussion Paper 2023*
- Our detailed [submission](#) to the Australian Government, Attorney-General's Department on the Privacy Act Review Report, 2023
- 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](#)

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](#).



## About the authors

This submission has been prepared by Anna Johnston, Principal, and Justin Frank, Privacy & Technology Specialist, 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 Visiting Scholar at the Research Group on Law, Science, Technology and Society of the Faculty of Law and Criminology of the Vrije Universiteit Brussel; and a Member of the Asian Privacy Scholars Network, and
- 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.

Anna has been called upon to provide expert testimony to the European Commission as well as various Parliamentary inquiries and the Productivity Commission. 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.

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.

Justin is a specialist privacy advisor, with experience in both technology and law. He holds a Bachelor of Science, Bachelor of Laws with Honours, Master of Laws, and Graduate Diploma of Legal Practice.

## **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](http://www.salingerprivacy.com.au)

