Automated Decision Making

Navigating automated decisions and Australia's evolving legal landscape

4 Dec 2025 | Insight | Australia

Article - By **Katherine Gregor, Christine Wong, Kaman Tsoi, Camille Tewari, Alex Lundie and Neeharika Palachanda**

What is ADM?

Automated decision-making (ADM) refers to the use of technology to make decisions with limited or no human intervention. ADM systems range in complexity and functionality – from traditional rule-based systems (eg, fixed-criteria loan approval processes) to more advanced models powered by complex algorithms and artificial intelligence (AI).

While ADM has existed, and been used by many companies, for decades in simpler rule-based forms, it is now receiving heightened attention due to the growing scale and sophistication of AI technologies and their ability to support ADM. As advanced ADM is increasingly embedded into critical functions across industry and government, governance and risk management frameworks will have to evolve in step.

This article looks at the lifecycle of ADM, mapping key considerations from data governance and algorithmic integrity to risk assessments and contractual frameworks. We explore how organisations can proactively assess their ADM frameworks to ensure transparency, mitigate risks, and uphold ethical decision-making principles ahead of the next tranche of privacy reforms and anticipated AI regulation.

Current legal landscape

Privacy reforms – where are we and what's coming?

Australia's privacy framework is undergoing significant reform, which is being implemented in two legislative tranches. The reform agenda includes new provisions addressing privacy implications of ADM – specifically:

- Tranche 1 of the privacy reforms (implemented on 10 December 2024)[1] introduces a requirement for privacy policies to include key details regarding ADM use. We set out further information on this below.
  This transparency obligation does not come into effect until 10 December 2026. The Office of the Australian Information

Commissioner (**OAIC**) has indicated its intention to publish more specific regulatory guidance on the topic in 2026.
- Tranche 2 of the privacy reforms is expected to contain a more sweeping set of changes than tranche 1. While details of these reforms are yet to be released, the Australian Government (Government) has indicated its support to introduce:[2]
  - a right for individuals to request meaningful information about how automated decisions with significant effect are made; and
  - a requirement for organisations to conduct privacy impact assessments for high-risk activities (potentially covering ADM activities).

For further information on the privacy reforms, see our previous article here: Navigating Australian Privacy Reform: Your guide to the changes ahead.

AI regulation

As noted above, AI often supports or enables ADM solutions. Last year the Government consulted on a proposal to introduce mandatory guardrails for AI in high risk settings[3] – and ADM that impacts on the rights of individuals or their ability to access services would likely fall into that high risk category. However, the Government has recently announced the National AI Plan[5] which steered clear of standalone AI legislation, instead confirming that the Government's regulatory approach to AI "will continue to build on Australia's robust existing legal and regulatory frameworks", in addition to other voluntary frameworks and guidance. This regulatory approach will be supported by a newly established AI Safety Institute and could have direct or indirect implications for ADM.

In September 2024, the Government released a set of Voluntary AI Safety Standards (**VAISS**), which provide guidance on good practices for the safe and responsible development and use of AI (including ADM that leverages AI). The ten key principles set out in the VAISS will assist organisations to manage risks in using AI for the purpose of ADM. In October 2025, the Government published their 'Guidance for AI Adoption' (**GfAA**).[6] The GfAA condenses and integrates the original ten principles into six practices for responsible AI adoption. It is likely to be beneficial for organisations to adopt these practices early in the roll-out of ADM to build-up organisational knowledge and governance in managing these risks.

As mentioned above, we may also start to see current laws evolve to address AI or ADM specific risks. For example, a bill[5] has recently been introduced in NSW to amend the *Work Health and Safety Act 2011* (NSW) and establishes a specific duty for organisations to ensure that digital work systems – broadly defined to include algorithms, AI, automation, online platforms and software[6] – do not create health and safety risks for workers, such as through unfair workload allocation or discriminatory practices[7]. These reforms reflect growing recognition that ADM and algorithmic management can introduce new psychosocial and physical risks[8], and require proactive governance in system design and deployment.

International comparison

ADM is subject to more stringent regulation in some jurisdictions outside Australia than what is proposed in the upcoming Australian privacy reforms. For example, since it was introduced in 2018, Article 22 of the European Union (EU) General Protection Data Protection Regulation (GDPR) imposes comprehensive obligations on organisations engaging in ADM. A number of data privacy regimes in other jurisdictions are also following suit with similar provisions.

Article 22 of the GDPR completely restricts the use of solely automated decisions (ie, those that are made without any human intervention) that produce legal or similarly significant effects for individuals, unless one or more of the specific conditions set out in Article 22 apply (eg, where conducting the ADM is a contractual necessity or if explicit consent has been obtained from affected individuals).

Where such ADM is permitted, the GDPR requires organisations to implement safeguards to protect individuals' rights and freedoms. These include providing meaningful information to affected individuals about the decision-making logic and offering individuals the opportunity to challenge and seek review of the decision or request human intervention.

Lifecycle of ADM

Rules and algorithms

Risk assessments

Data

Contracting with third party suppliers

Privacy policies, disclosures and consents

Ongoing monitoring

Customer feedback, complaints, objections to decisions

Rules and algorithms

There has been significant legislative and regulatory scrutiny (particularly in the context of consumer-facing digital platforms) over the risks and potential harms posed by algorithms and ADM. From a legal perspective, key risks include:

Algorithmic bias, which can result in unlawful discrimination or unfair treatment.

Consumer harms such as misleading or deceptive conduct and others.

Risks arising from the 'black box' problem.

Potentially significant liability from breaching duties and other legal obligations flowing from governance failures in the deployment of ADM tools.

To mitigate risks, the question is how to embed transparency, explainability and accountability into the design and governance of ADM models.
ADM models may be perceived as driving cost savings and efficiencies. The flip side is that when these systems go wrong, they can go systemically wrong. This means that any cost savings need to be balanced with investment in diligence and understanding how the algorithm works, including:

- Adequate testing to confirm that the algorithm works as intended, both prior to deployment and on an ongoing basis.
- Rigorous design testing and risk assessments upfront. This includes documenting decision logic, which is likely to go beyond 'chain of thought' reasoning.
- Ensuring appropriate ongoing governance and implementing structures for independent review or verification of critical ADM decision points.

Risk assessments

Safe and responsible deployment of ADM requires a structured and proactive approach to risk management. Organisations should assign clear accountability to people in relation to each element of the ADM lifecycle. Risk assessments should be conducted prior to implementing or repurposing ADM systems, particularly for high-risk use cases. For example, this may include where ADM systems make decisions which affect:

- an individual's rights under a contract, agreement or arrangement;
- an individual's access to a significant service or support, such as healthcare, insurance, housing, education or lines of credit;
- an individual's employment opportunities, including recruitment, promotion or termination; and
- a business customer's creditworthiness when applying for loans, lines of credit or trade finance.

These risks are even more acute where the decisions affect groups of potentially vulnerable individuals. Regulators have been particularly focused on ADM systems which operate unconscionably against these individuals, with significant penalties imposed.

Guidance released by the Government and regulators such as the OAIC and ASIC emphasise the need to evaluate potential harms to individuals (taking into account privacy considerations and broader community values) alongside organisational risks such as reputational damage, cybersecurity vulnerabilities and lawfulness. Where appropriate, assessments should involve cross-functional input and result in actionable strategies to mitigate or eliminate identified risks. A privacy impact assessment and a security risk assessment could be conducted together with, or as a part of, a broader risk assessment for the ADM project.

Data

The reliability and performance of ADM systems are fundamentally shaped by the quality of the input data. The risks associated with using poor quality data are well-documented, including discrimination arising from biases or inaccuracies within the dataset itself (eg historical, representation, measurement, evaluation and sampling biases) which may be replicated in outputs.

A notable instance of algorithmic bias arising from flaws in the underlying dataset is the AI-enabled recruitment tool launched by Amazon in 2014. The tool was trained on resumes submitted to Amazon over a 10-year period, and since most of the applicants in this 10-year period were male, the tool learned to favour male candidates over female candidates even though the sex of the applicants was not included in the selection information (eg, where a resume listed a candidate's role as a "women's chess club captain" or referred to an all-female college). Amazon ultimately decommissioned the tool.

Similarly, the *Mobley v Workday, Inc.*[14] class action in the United States is one of the first major legal challenges to the use of algorithmic hiring tools. The plaintiffs allege that Workday's algorithms filtered out older candidates, often without human review, in contravention of federal discrimination laws. While still ongoing, the allegations and early rulings made in the class action suggest that organisations may face legal exposure in respect of both their own practices but also the practices of vendors providing AI tools, especially where those tools are trained on biased input data or operate without human oversight.

Assessment of data quality and the identification of historical bias is a complex area and market practice in this area is still emerging. Organisations will need to carefully consider what steps need to be taken to detect and mitigate these risks.

Beyond data quality concerns, organisations must also navigate a range of other data-related risks, including as to:

- **Accuracy, de-identification and collection of personal information. Personal information contained in input or output data must be handled in accordance with privacy laws.** Amongst other things, organisations must take reasonable steps to ensure that the personal information they hold is accurate, up-to-date, complete, relevant and not misleading, having regard to the purposes for which the information is held.

  If third party datasets are used, organisations may need to understand the circumstances of the original collection of personal information (ie, data sources and compliance with privacy notice or consent requirements).

  Organisations may consider de-identifying personal information prior to use in ADM systems, however whether this is appropriate will depend on what the system needs to achieve with input data at a particular stage in the process. For example, personal information may be less important to training an AI model, compared to applying the model to a particular individual.

  Additionally, robust de-identification may be difficult, particularly where aggregated data is drawn from multiple datasets (which raises concerns as to potential re-identification through AI technologies).

- **Data minimisation.** Under the Australian Privacy Principles (APPs) (contained in Schedule 1 of the Privacy Act 1988 (Cth) (Privacy Act)), organisations must limit the collection of personal information to what is reasonably necessary and take reasonable steps to destroy or de-identify it once no longer required. Collecting excessive data not only increases privacy risks but also heightens exposure to cybersecurity threats.

- **Cybersecurity:** Strong security measures are essential to protecting training data, output data and ADM systems from unauthorised access or manipulation. Certain organisations will also need to ensure compliance with sector specific regulation around information security, such as Prudential Standard CPS 234 for APRA regulated entities.

Where organisations use ADM systems developed or operated by third parties, it is essential to address data use rights with those third parties, including confidentiality of the data, how it can be used and what safeguards are in place.

For more information on the current debate in Australia, see our article here: Australian Productivity Commission proposes text & data mining exception to copyright infringement for AI training.

Contracting with third party suppliers

Organisations may choose to outsource the design, implementation, run and/or maintenance of ADM systems to external suppliers, rather than managing these processes internally. In such cases, it will be important that contracts with third party suppliers clearly set out responsibilities, risk allocation and oversight mechanisms across the ADM system’s lifecycle.

In addition to the common areas of contention in technology and data procurement, key considerations in the ADM context include allocating roles, rights and responsibilities between the parties in respect of:

- privacy and cybersecurity controls;
- transparency, including notices and consents for individuals who will be affected by the ADM;
- governance mechanisms, including human oversight;
- management of data and know-how collected and developed during the course of the engagement, including use of such information to train or improve other ADM models or provide it to other parties;
- adversarial testing to identify any bias or other unlawful capabilities;
- third party claims, particularly in relation to IP infringements in relation to both the model and the use of any third party inputs used to train or enhance the model;
- addressing accountability for ADM outputs, including allocating responsibility for the effects of the ADM outputs and ensuring safeguards are in place to ensure compliance with applicable laws;
- capability to provide explanations for the basis on which decisions are made; and
- dispute resolution processes, including avenues for individuals to submit feedback or complaints or to contest decisions.

Case 634/21 SCHUFA Holding (Scoring) (2023) (Schufa)

The importance of clearly defining roles and responsibilities in contracts with ADM technology suppliers is underscored by the Schufa case under the GDPR in the EU. In that case, Schufa – a third party credit rating agency –argued that it was not responsible for ADM because it merely generated a credit score, while the final decision to approve or reject loan applications was made by the bank. However, the Court of Justice of the European Union (**CJEU**) rejected this argument, holding that the creation of the credit score itself constituted an automated decision with significant effects under Article 22 of the GDPR.

This ruling highlights that even upstream contributors to ADM processes may be directly subject to regulatory obligations in the EU. Accordingly, organisations outsourcing ADM functions should ensure that contracts explicitly address accountability for ADM outputs, including how decisions are generated, who is responsible for their effects, and what safeguards are in place to ensure compliance with applicable privacy laws.

Privacy policies, disclosures and consents

Organisations must comply with privacy obligations relating to notice and consent when collecting, using or disclosing personal information as part of ADM processes. However, the OAIC has cautioned against overuse of notice and consent mechanisms where this places the burden of understanding the risks of complicated information handling practices on individuals. In particular, the OAIC cautions that AI tools often use personal information as the basis for a decision in a way that is 'invisible or difficult to comprehend' and is challenging for organisations to clearly explain to individuals.[15]

As noted above, the tranche 1 reforms to the Privacy Act will require organisations to update their privacy policies in respect of ADM (from 10 December 2026 onwards). This will be implemented through the addition of new APPs 1.7 to 1.9. In anticipation of further guidance on these new APPs from the OAIC in 2026, organisations should start preparing by seeking to understand how they use ADM currently, where this involves personal information, and what effect the ADM can have on individuals.

Which organisations must comply?	Any organisation that has arranged for a computer program, using personal information, to make, or do a thing that is substantially and directly related to making, a decision that could reasonably be expected to significantly affect the rights or interests of an individual. It may not be straightforward to determine whether your organisation meets the legislative threshold – this can often depend on the extent of automation of a system and level of human involvement.

What must be covered in the privacy policy?	• The kinds of personal information used in the operation of those computer programs. • The kinds of decisions made solely by the operation of those programs. • The kinds of decisions for which those programs do something substantially and directly related to making the decision.

These transparency requirements may prompt individuals to seek further clarification regarding how their personal information is used within ADM systems. Accordingly, organisations should be prepared to respond to such inquiries with clear and accessible explanations and further information as required. While provision of this further information is not yet a positive obligation under the Privacy Act, many organisations will look to accommodate these requests as part of meeting broader transparency standards.

Penalties for breach of the Privacy Act can be significant, with the maximum penalty for a serious or repeated interference with privacy being $50 million or more.

On 8 October 2025, the Federal Court of Australia ordered the first civil penalty ($5.8 million) under the Privacy Act against Australian Clinical Labs. Refer to our article here for further details. This signals a clear shift toward stronger enforcement by regulators, moving beyond compliance notices or warnings to imposing substantial financial penalties.

Customer feedback, complaints, objections to decisions

While not currently required under Australia law, it is considered best practice for organisations to establish processes that allow individuals to understand and, where appropriate, challenge the use or outcomes of ADM systems. To this effect, the Commonwealth Ombudsman's Better Practice Guide on ADM suggests that organisations should consider publishing plain-language transparency statements that explain how ADM systems they use operate, what decisions they make and how individuals can seek review. ADM systems should also be capable of generating a 'statement of reason'[19] (which typically, among other things, lists findings on material questions of fact and includes a probative assessment or weighing of evidence), of each decision, particularly where the outcome significantly affects an individual's rights or interests. This reflects leading global standards (most notably, under Article 22 of the EU GDPR) as well as Implementation Practices 2.2 and 4.2 of the GfAA, which encourages mechanisms for individuals to contest AI-driven decisions. Additionally, tranche 2 Privacy Act reforms are anticipated to introduce a right for individuals to request meaningful information about how significant automated decisions are made.

Case C-203/22 Dun & Bradstreet Australia (2025)

The anticipated introduction of a right for individuals to request meaningful information about how significant automated decisions are made aligns with global developments, including recent case law under the GDPR.

In Case C-203/22, the CJEU considered a challenge by an individual seeking to understand why their mobile phone contract application was rejected following an automated credit assessment by Dun & Bradstreet Australia (**D&B**), an organisation that operates in the business information and credit scoring sector. While D&B argued that disclosing the logic behind the decision would compromise its trade secrets, the Court held that while organisations are not required to disclose complex algorithms, they must provide sufficient insight into the ADM process to allow individuals to understand and contest the outcome. This case reinforces the principle that transparency and contestability are central to responsible ADM.

Ongoing monitoring

Ongoing monitoring is key to the lifecycle of ADM systems and should be designed to ensure that the relevant ADM system remains lawful, fair, accurate and operating as intended. Some examples of measures which may be appropriate include:

- **referral pathways** – ADM systems should be designed such that outputs or cases that fall outside normal parameters are automatically referred to a human decision-maker for review;
- **spot checks and audits** – ADM systems should involve regular human review of a sample of automated decisions to detect errors, biases or unintended outcomes. These can be random spot checks or targeted reviews as well as independent expert audits;
- **version control** – where changes are made following monitoring, organisations should maintain detailed records of system versions, changes and the rationale for these updates so that these decisions can be traced and explained; and
- **clear accountability** – monitoring and updating of ADM systems should be assigned to specific roles or teams.

Conclusion

Organisations that embed transparency, robust data governance, rigorous testing and defined accountability into their ADM design and procurement will be better placed to meet evolving regulatory expectations. Practical priorities for organisations adopting ADM processes or systems should include strengthening monitoring and human review mechanisms, clarifying contractual accountability with suppliers and designing accessible processes for individuals to contest significant automated outcomes. Taking such actions will assist with reducing legal exposure and encourage the trust of customers and the community as ADM inevitably becomes more complex and deeply integrated into business operations.

Footnotes

Legal Notice

The contents of this publication are for reference purposes only and may not be current as at the date of accessing this publication. They do not constitute legal advice and should not be relied upon as such. Specific legal advice about your specific circumstances should always be sought separately before taking any action based on this publication.