A JOINT PUBLICATION BY

Australian Institute of Company Directors

UTS Human Technology Institute

A Director's Guide to AI Governance

Foreword

While forms of Artificial Intelligence (AI) have been used for many years, the major development in Generative AI capabilities over recent times has prompted widespread discussion of its role in the economy and broader society.

AI, with its sophisticated pattern recognition capabilities pulled from vast datasets, has the potential to offer significant productivity and economic gains. However, alongside these benefits lie potential risks from AI system failures and/or abuse, including misuse of personal data, algorithmic discrimination and poorly controlled automated decision-making.

As stewards of organisational strategy and risk management, directors should seek to seize the opportunities and mitigate the risks of AI, with its ethical use in the interests of customers being paramount. This requires a robust governance framework that can adapt to the unique characteristics of AI systems.

Currently, research suggests that there is generally limited board oversight of AI use, with AI application often subject to inadequate controls and risk oversight. In many cases, directors and senior executives are unaware of where within the organisation's value chain AI is being used, and how. If left unaddressed, this risks significant lost opportunities and commercial, reputational and regulatory damage, with regulators and policymakers increasingly focused on regulating AI harms.

In January 2024, we saw the Australian Government commit to a range of initiatives to support the uptake of safe and responsible AI. These include consideration of the introduction of mandatory guardrails for AI deployment in high-risk settings, consideration of labelling and watermarking of AI in high-risk settings, and clarifying and strengthening existing laws to address AI harms.

Internationally, we are seeing jurisdictions attempt to walk the policy tightrope between regulating high-risk AI uses to avoid the most significant AI harms, and ensuring innovation continues to flourish by tapping into this transformational technology.

To assist boards navigate the ethical and informed use of AI, the Australian Institute of Company Directors (AICD) has partnered with the Human Technology Institute (HTI) at the University of Technology Sydney (UTS) to provide a suite of director resources.

This includes:

- '**A Director's Introduction to AI**', which lays the foundation for knowledge of AI concepts;
- '**A Director's Guide to AI Governance**', which provides practical guidance for boards' using, or wishing to deploy AI within their organisations;
- A **Concise Snapshot** of the 'Eight elements of safe and responsible AI governance'; and
- a separate **SME and NFP governance checklist** which recognises the significance of small and medium-sized enterprises to the Australian economy and the specific needs of this sector.

We hope that by applying the 'eight elements of safe and responsible AI governance', directors can guide their organisations to deploy AI systems safely and responsibly for maximum strategic and competitive advantage.

Mark Rigotti MAICD
CEO and Managing Director
Australian Institute of Company Directors

Professor Nicholas Davis MAICD
Co-Director
Human Technology Institute
University of Technology, Sydney

Contents

How to use this guide

Having considered all the boards on which you serve, select what applies to you:	What we suggest you read
☐ I know about ChatGPT, but I don't know any other types of AI ☐ I am not clear how AI is different to other technologies ☐ I am unsure about the key legal obligations applying to AI use ☐ I am not clear about the key risks or opportunities arising from AI ☐ I do not know the underlying principles of safe and responsible AI	 A Director's Introduction to AI
☐ I understand the difference between General AI and Narrow AI ☐ I understand how AI is different to other technologies, but am unclear how this impacts governance ☐ I am unsure about where AI is used within my organisation ☐ I am unsure about what questions to ask management about the governance and use of AI and how to assess the quality of management's responses	 A Director's Guide to AI Governance
☐ I am a director of a SME or NFP and do not know how to implement AI governance	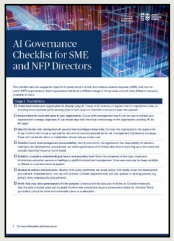 AI Governance Checklist for SME and NFP Directors

Resource purpose, audience & structure

The **purpose** of this resource is to provide practical guidance for boards and directors of organisations that are using or planning to use AI systems (as distinct from developers and distributors of AI systems).

The **primary audience** of this resource are directors of ASX300 entities who are using, or considering deploying AI.

However, recognising the significance of small and medium-sized enterprises to the Australian economy and the specific needs of this sector, we provide an AI Governance Checklist for SME and NFP Directors.

AI technology as well as AI policy and regulation is dynamic and constantly developing. This resource is not intended to 'cover the field', but to provide a suggested framework for board oversight of AI use.

The resource is structured into two sections:

- **Section 1** highlights a set of cross-cutting insights and implications for AI governance for directors.
- **Section 2** sets out eight elements of effective, safe and responsible AI governance. It also provides key questions for directors and management responses to watch out for, and provides some case studies.

As part of this Guide you can also find a separate Concise Snapshot of the 'Eight elements of safe and responsible AI governance'.

Executive summary

ROLES & RESPONSIBILITIES

- Identify the **management and board individual/ body accountable for AI decision-making.**
- **Identify those involved in, and responsible for**, AI system procurement, development and use.
- Consider **whether decision-making processes applied by key accountable persons incorporate consideration of AI risk and opportunity.**

GOVERNANCE STRUCTURES

- Determine **which existing or new board and management governance structure** would most appropriately support AI oversight.
- **Review board and management committee charters** to determine whether and how they incorporate AI issues.
- Consider how **external experts can be leveraged within existing governance structures.**
- Consider **the nature and frequency of management reporting** to the board/relevant board committee.

PEOPLE, SKILLS & CULTURE

- Verify that management have assessed the organisation's AI **skills, capabilities and training** needs, and implement **upskilling programs** (including at the director-level).
- Discuss the **potential for AI to impact the workforce and workforce planning**.
- Consider **how AI governance structures can incorporate a diversity of perspectives, including expert views,** to aid diversity of thought and avoid 'group think'.

PRINCIPLES, POLICIES & STRATEGY

- **Require that AI is considered and, where appropriate, embedded, within the organisation's strategy.** AI use should have a clear business value – 'AI for AI's sake' should be avoided.
- Engage with management to discuss **how safe and responsible AI principles have been incorporated into relevant policies** (such as AI/ IT use, privacy, confidentiality and cyber security).
- Recognise that principles and policies need to be proactively **implemented and enforced across the supply chain.**

PRACTICES, PROCESSES & CONTROLS

- Work with management to **understand what controls are in place for AI use** (e.g. risk appetite statement and risk management framework).
- Confirm with management that there are **processes in place to assess supplier and vendor risk.**
- **Monitor and regularly review the effectiveness of controls.**

SUPPORTING INFRASTRUCTURE

- Confirm that you are **broadly aware of where, within the organisation, AI is currently being used.** Management can provide this information through an **AI inventory.**
- Verify that management is aware of, and has a **robust data governance framework** in place to manage data collected and stored by the organisation to train AI systems.
- Focus on **increasing transparency to end users** about how the organisation's AI systems use data.

STAKEHOLDER ENGAGEMENT & IMPACT ASSESSMENT

- **Identify and engage with stakeholders** to understand AI's impact and stakeholder expectations of AI use and governance.
- Confirm with management that **AI system design and assessment processes incorporate accessibility and inclusion practices.**
- Consider whether **AI-generated results/outcomes are explained to stakeholders** and whether an appeal process is available.

MONITORING, REPORTING & EVALUATION

- Confirm that a **risk-based monitoring and reporting system** for mission-critical and high-risk AI systems is in place.
- Develop and implement a **monitoring and reporting framework with metrics and outcomes** to track and measure progress.
- Consider **seeking internal and external assurance.**

SECTION 1:

AI and the governance imperative

KEY POINTS:

- The unique characteristics of AI systems (complex pattern recognition based on large and diverse datasets) mean that traditional governance approaches may not be appropriate.
- Directors should be aware of AI's unique risks and opportunities and how these require adaptations to existing governance approaches.
- Effective AI governance should be human-centred, cross-functional, adaptive and iterative.
- Directors should align investment in AI with organisational values and embed it within broader business strategy. 'AI for AI's sake' should be avoided.

1.1 WHAT IS AI?

The definition of AI adopted by the International Organisation for Standardization and the International Electrotechnical Commission ISO/IEC 22989 is:

> *An engineered system that generates outputs such as content, forecasts, recommendations or decisions for a given set of human-defined objectives.*

1.1.1 How is AI different from other technology?

AI is a special form of digital software that is particularly good at predicting outputs, optimising, classifying, inferring missing data, and generating new data.

AI systems can often outperform non-AI systems, and as a result offer significant productivity, efficiency and customer experience benefits.

AI is also more versatile and scalable than traditional software because it can be replicated and adapted to new contexts at a relatively low cost. As a result of these advantages, AI is increasingly being deployed across organisational teams and functions.

However, the differences between traditional software systems and AI systems also impact governance approaches.

Traditional software systems are built from explicit rules coded by developers, such that their behaviour is inherently more predictable and understandable (even if the software itself is complex).

By contrast, AI systems are often created by defining an objective and using historical data to create an AI model that may rely on billions of inferred connections between data points to achieve its objective. This process means that **it can be extremely challenging to replicate, explain or test an AI system's output.**

BOX 1: The role of data in AI systems

Data is the foundation of AI systems. Data, including personal information, is collected and used to train AI systems. It is both an input and an output of a deployed AI system.

The selection of data, particularly its quality, quantity, and representativeness, will significantly affect the performance of AI systems.

Through the ongoing collection of data and feedback loops, the accuracy and efficiency of AI systems should improve over time.

BOX 2: What kinds of systems are usefully defined as AI?

- **Machine learning:** a broad set of models that have been trained on pre-existing data to produce useful outputs on new data.
- **Expert systems:** systems that use a knowledge base, inference engine and logic to mimic how humans make decisions.
- **Natural language systems:** models that can understand and use natural language and speech for tasks such as summarisation, translation, or content moderation.
- **Facial recognition technologies:** systems that verify a person, identify someone, or analyse personal characteristics using facial data drawn from photos or video.
- **Recommender systems:** systems that suggest products, services or information to a user based on user preferences, characteristics, or behaviour.
- **Automated decision-making systems:** systems that use data to classify, analyse and make decisions that affect people with little or no human intervention.
- **Robotic process automation:** systems that imitate human actions to automate routine tasks through existing digital interfaces.
- **Virtual agents and chatbots:** digital systems that engage with customers or employees via text or speech.
- **Generative AI:** systems that produce code, text, music, or images based on text or other inputs.
- **AI-powered robotics:** physical systems that use computer vision and machine learning models to move and execute tasks in dynamic environments.

1.1.2 Different types of AI

Box 2 provides a non-exhaustive list of systems that meet the definition of AI above.

General AI (or General Purpose AI) and **Narrow AI** are two sub-categories of AI (see Table 1).

TABLE 1: Key differences between General AI and Narrow AI

Type of AI system	Description[1]	Examples
General AI (or General Purpose AI)	An AI system that can be used for a broad range of tasks, both intended and unintended by developers. This includes Generative AI.	Text generation (e.g.GPT-4, Gemini), image generation (e.g. DALL.E, Midjourney), programming code generation (i.e. Codex).
Narrow AI	An AI system trained to deliver outputs for specialised, constrained tasks and uses to address a specific problem.	Search engines (e.g. Google, Bing), facial recognition (e.g. Apple Face ID), recommender systems (e.g. Amazon, Spotify, Netflix).

1 ISO, 2022. ISO-IEC-22989 Artificial intelligence concepts and terminology.

For more insight on What AI is and its relevance for directors, see **Chapter 1** of A Director's Introduction to AI.

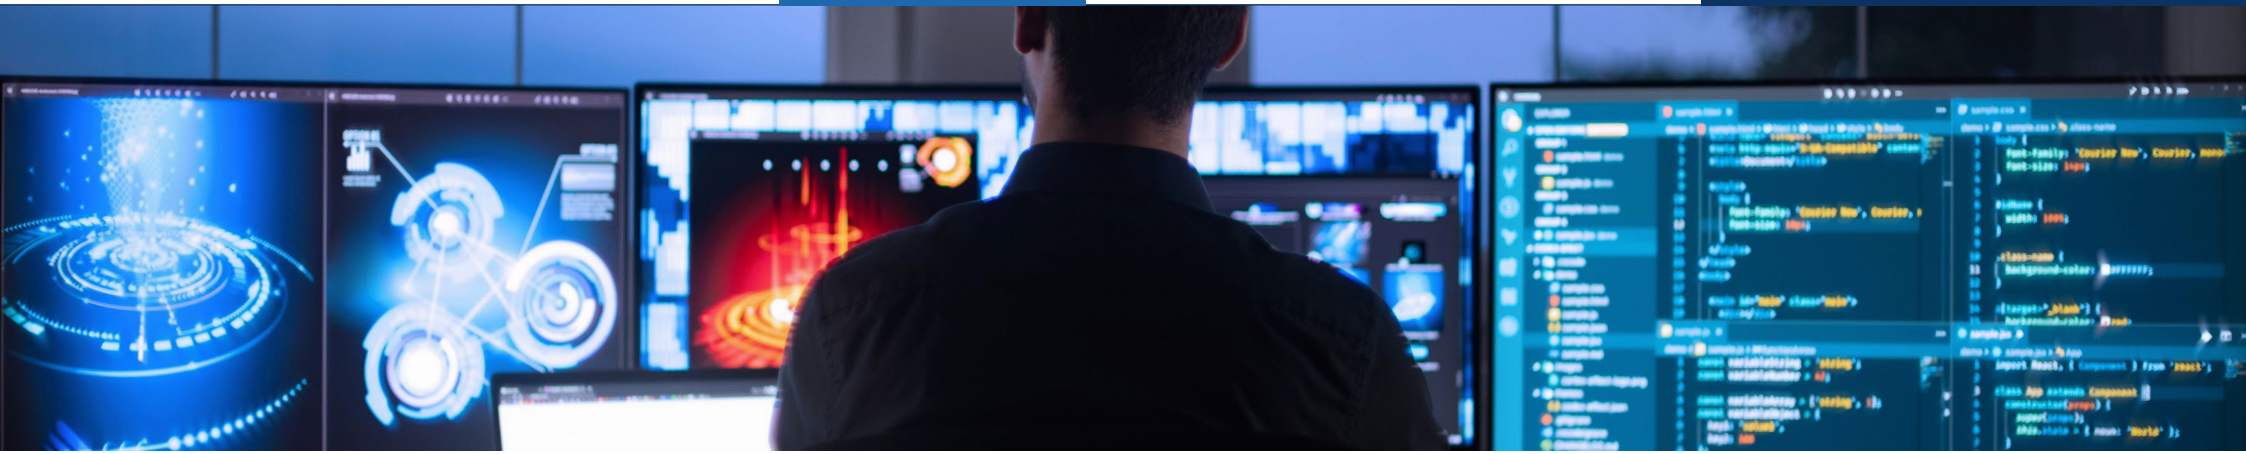

1.2 AI AND DIRECTORS' OBLIGATIONS

While stand-alone AI regulation has not yet been introduced in Australia, a range of existing laws already apply to the use of AI systems – see **Figure 1**. Some laws place obligations on the organisation, while others apply to directors and officers individually. The Australian Government has also foreshadowed further reform of these laws to apply more directly to AI use.

For more detail on existing legal obligations, as well as Australian and international regulatory developments, see **Chapter 3** of A Director's Introduction to AI.

FIGURE 1: Existing legal obligations when using AI

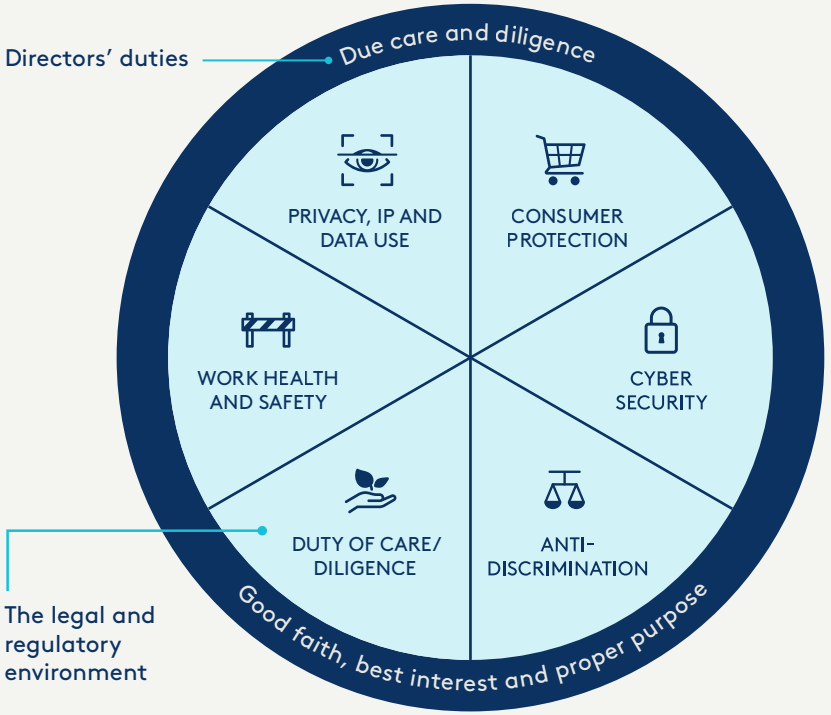

In line with their directors' duties, directors are responsible for the oversight of the organisation's strategy and risk management processes. This includes managing AI risks and opportunities.

AI risks include both AI system failures and malicious, misleading, reckless or inappropriate AI use (see the summary at Table 2). These risks can create and amplify a range of commercial, reputational and regulatory risks to organisations (see Figure 2).

TABLE 2: Key sources of AI risk for organisations

Key sources of AI risk	Examples
AI system failures – where systems create harm because they fail to perform as intended	• Poor system performance • Biased system performance • System fragility or unreliability • Security failures or vulnerabilities
Malicious, misleading, reckless, or inappropriate use – where systems are deliberately used (whether by the organisation or external parties) in a way which creates or amplifies risk of harm	• Misleading advice • Misinformation at scale • Unfair or extractive use • Opacity and lack of interpretability • Weaponisation • AI-powered cyber attacks • Fraudulent and unlawful use e.g. scams • Financial market manipulation • Excessive deployment • Deployment on vulnerable individuals

FIGURE 2: Risks to organisations from AI use

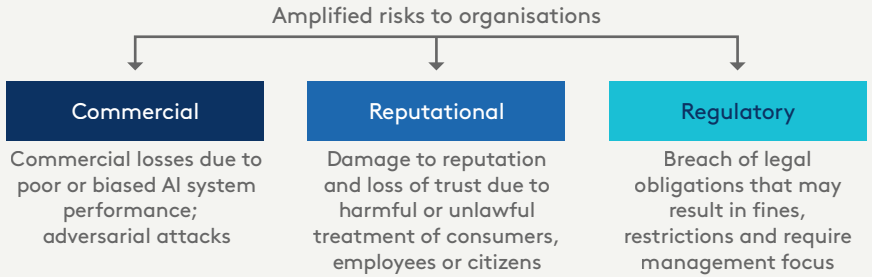

On the other hand, a lack of investment in AI capabilities also leaves organisations vulnerable to a range of other risks, such as a lack of competitiveness, higher costs, lack of new product and service delivery, poorer consumer service, as well as talent acquisition and retention challenges.

The risks of action and inaction must be carefully weighed by directors alongside the organisational strategy and the risk appetite of the organisation.

For more details on AI risks and opportunities, see **Chapter 2** of A Director's Introduction to AI.

1.3 AI AND GOVERNANCE IMPLICATIONS

Both the deliberate and 'shadow' AI use (see **Box 3**) throughout an organisation and its supply chains present directors with complex governance challenges.

KEY QUESTIONS FOR DIRECTORS

- How can we support experimentation and innovation with AI within the risk tolerance of the organisation?
- How is AI being used to support the delivery of the organisational strategy and related business goals?

BOX 3: What is 'shadow' AI use?

Shadow AI refers to employees' unauthorised use of AI applications for work-related purposes.

The recent availability – and relatively low cost – of capable cloud-based large language models such as ChatGPT means that a significant percentage of employees and contractors are leveraging Generative AI systems for their work without the explicit knowledge, permission or oversight of management.

A 2023 Information Audit and Control Association (ISACA) poll of IT governance professionals across Australia and New Zealand found widespread employee use of Generative AI (63 per cent of respondents), despite only 36 per cent of organisations expressly permitting its use. Just 11 per cent of respondents said that their organisation has a comprehensive policy for Generative AI use.

The phenomenon of shadow AI poses a range of amplified risks to organisations and their stakeholders, including breaches of privacy and confidentiality.

1.4 TRADITIONAL IT GOVERNANCE MAY NOT BE FIT-FOR-PURPOSE FOR AI

Some managers and directors may be tempted to place the oversight of AI systems within existing IT governance systems. However, **HTI's research strongly suggests that existing IT risk management frameworks and systems are largely unsuited for AI governance**.

This is because traditional IT governance focuses on point-in-time risk assessments of officially sanctioned systems, largely based on vendor assurances.

Such an approach has limitations in governing AI systems because:

SPEED AND RATE OF CHANGE

How organisations use AI is not a 'tomorrow' challenge – it is a 'now' challenge that involves rapidly advancing technologies.

OPACITY

Opacity in the sense of (1) the challenge of testing, validating, explaining and reproducing AI system outputs; and (2) difficulty identifying AI use within an organisation and its value chain.

DIVERSITY OF USE CASES

AI use crosses organisational barriers and reporting lines. Its use ranges from being used by frontline workers, being embedded within the core of the organisation's strategy and risk management approaches, and being embedded within existing systems (e.g. software updates) and supply chains. This decentralisation across the porous boundaries of the organisation makes AI use difficult to control.

AN UNCERTAIN POLICY, REGULATORY AND TECHNOLOGY ENVIRONMENT

These uncertainties are driven by local and international regulatory change, technology change, and a shifting threat environment.

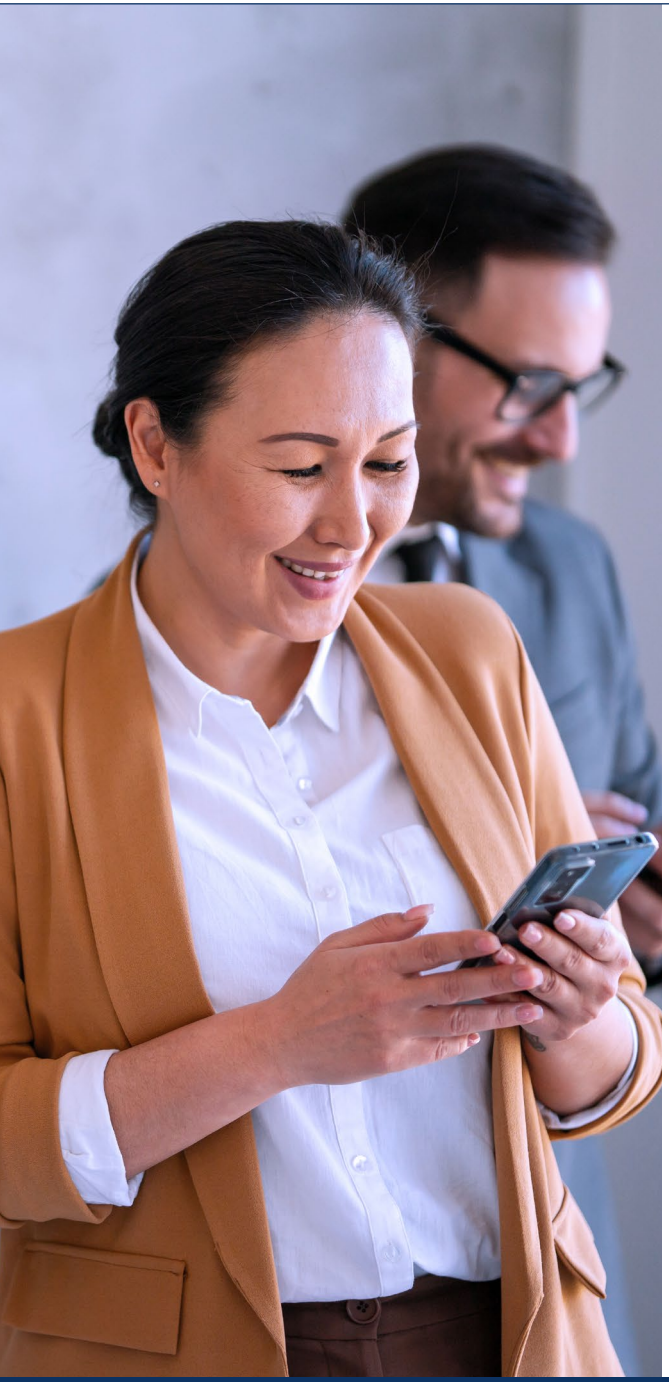

Organisations may be tempted to place the entirety of AI system oversight within a risk and compliance function. However, this can mean the significant opportunities of AI systems are not appropriately recognised.

In meeting these challenges, directors need to engage with management to implement an iterative, integrated, flexible and adaptive governance approach which is:

HUMAN-CENTRED

This refers to governance mechanisms meaningfully and transparently tracking and reporting how AI systems are impacting key stakeholders (consumers, employees, suppliers, contracting parties, etc).

CROSS-FUNCTIONAL

AI governance cannot be achieved through the establishment of separate, disconnected roles or policies and procedures. AI governance needs to span various departments and roles, including those responsible for privacy, IT, legal, product design and development, procurement, HR, risk and strategy. It must also be led at a senior level within the organisation.

INTEGRATED

Effective frameworks will integrate all eight elements set out in **Section 2**, rather than cherry-pick one or two.

ITERATIVE AND ADAPTIVE

Given the speed of technological transformation, organisations should not rely on a 'set and forget' approach to AI governance. Governance systems and processes should be subject to regular review to monitor whether targets and outcomes are being achieved.

1.5 ALIGNING AI USE TO ORGANISATIONAL STRATEGY

The use of AI by organisations should be aligned to the broader organisational strategy. How AI is being used to achieve strategic objectives is core to the work of the board.

The organisational strategy should be regularly reviewed to clarify and adjust the role of AI and emerging technologies.

KEY QUESTIONS FOR DIRECTORS TO ASK

- How is AI currently being used to deliver business goals?
- What investments are we making in relation to the development and use of AI systems?
- How can we leverage AI in a responsible way to achieve our organisational strategy?
- What sorts of problems and challenges can or should AI systems be used to solve?
- Under what circumstances would we conclude that AI is not the right tool for the job?
- What is our overall assessment of the evolving balance between the risks and benefits of AI systems to drive business value?

1.6 AI-SPECIFIC RISK MANAGEMENT

As detailed in **Chapter 3** of A Director's Introduction to AI, directors have legal duties to effectively oversee the management and mitigation of organisational risks.

AI system use – or failure to make use of AI systems when appropriate – by organisations can pose a range of risks (see **Chapter 2** of A Director's Introduction to AI) that need to be carefully managed.

SUGGESTED DIRECTOR STEPS

- Understand current AI use, which can include the issue of an **AI inventory** (see **Box 4**).
- Review the organisational risk framework to test its application to AI use, noting increased scrutiny by stakeholders over how AI risks are being managed (see **Chapter 3** of A Director's Introduction to AI).
- Define and review the organisation's risk appetite and risk statement to cover AI use.
- Align risk management approaches with existing sectoral risk management obligations (such as that required for financial services organisations under section 912A *Corporations Act 2001* (Cth)).

SECTION 2: Practical steps for directors

This chapter focuses on the practical steps that directors can take in the boardroom and in conversations with management. It is structured around **eight key elements of effective AI governance frameworks.**[2]

ELEMENT 1:
Roles & responsibilities

ELEMENT 2:
Governance structures

ELEMENT 3:
People, skills & culture

ELEMENT 4:
Principles, policies & strategy

ELEMENT 5:
Practices, processes & controls

ELEMENT 6:
Supporting infrastructure

ELEMENT 7:
Stakeholder engagement & impact assessment

ELEMENT 8:
Monitoring, reporting & evaluation

2 See HTI's AI Governance Snapshot #1 Essential Components of AI Governance (HTI, 2024) for further information.

FOR PRACTICAL ASSISTANCE FOR DIRECTORS, THIS CHAPTER CONTAINS:

1. **Key questions for directors to ask** themselves and/or management; and
2. **A traffic light system** which assists directors process management's response to key questions:

   AMBER suggests there **may be some risk**, and advises that **directors should probe further** and assess management's position and response. An uplift in governance practice may be necessary.

   RED suggests there is **potential high risk**, and that directors should work with management to **address this risk** through implementing safe and responsible AI governance practices (as suggested in this guide).

The elements and questions featured in this section may also apply to the governance of some non-AI systems, or technology more broadly. This is deliberate – in ensuring that their organisations are adequately prepared to grasp the benefits and manage the risks of AI systems, directors can and should leverage existing governance knowledge and systems.

However, it is crucial that this knowledge and existing approaches are appropriately applied to the peculiar risks and concerns that the specific characteristics of AI systems create (see **section 1.4** and Chapter 2 of A Director's Introduction to AI).

Wherever possible, in each subsection we have highlighted where directors should look to the governance issues specific to AI.

This chapter is not intended to be a comprehensive guide. Regulatory requirements, guidance, and best practices in this area are rapidly evolving. Rather, directors should view these components as an ongoing conversation with management as AI governance continues to evolve.

2.1 ROLES & RESPONSIBILITIES

HTI's research suggests that there is little awareness amongst corporate leaders of where, how and why AI systems are being used across their businesses.[3] This lack of internal knowledge is a major barrier to AI governance efforts and amplifies AI risks.

Directors should be clear on which individual or body, at both the board and management level, has decision-making power and accountability for AI use.

While management will be responsible for AI implementation, the board has overall oversight over the organisation's AI governance.

In the absence of a structured approach to AI system accountability, most organisations adopt a form of 'guru-based governance', where responsibility sits with a single individual viewed as technically competent in AI. Such over-reliance on a single leader or a small set of technical personnel within the organisation is problematic, not least because it creates significant key person risk.

SUGGESTED DIRECTOR STEPS

1. Determine and document **which individual/ body at the board and management level has responsibility, and is ultimately accountable to the board, for decisions regarding AI use**. This includes a consideration of how to leverage existing governance structures (such as board and management committees) – see **section 2.2**.
2. Identify **who is currently involved in, and accountable to the board for, decisions relating to the procurement, development and use of AI systems.**
3. **Determine and record where in the organisation AI is already being used.** This could be in existing technology products. **An AI inventory (Box 4)** can provide a useful record of where AI is used within the organisation.
4. Consider whether decision-making processes applied by key accountable persons incorporate consideration of AI risk and opportunity.

BOX 4: What is an AI inventory?

An **AI inventory or register** is a structured, centralised, and up-to-date database of all AI systems that an organisation relies on, including those offered by third-party providers. The inventory should include details on the technical aspects of each system including:

- type of model and technology infrastructure being leveraged;
- the data used in both training and operation;
- the purpose and context of use;
- ongoing cost to the business; and
- the result of all recent risk and impact assessments.

An AI inventory is an essential asset for AI governance as it provides greater visibility into the mix of AI system types – their benefits, costs, and criticality, and the distribution of risk across systems and to stakeholders arising from AI use.

An AI inventory may also mitigate the risk of making incorrect or exaggerated claims about a product, service or company's AI use (known as 'AI-washing').

3 Lauren Solomon and Nicholas Davis, The State of AI Governance in Australia (HTI Report, 2023), 13.

KEY QUESTIONS FOR DIRECTORS TO ASK

- How are we tracking AI use within the organisation?
- Which individual or body at the board or management level is responsible for data governance?
- Which individual or body at the board or management level is responsible for decisions regarding the development and use of AI within the organisation?
- Which individual or body is responsible for making procurement decisions and identifying, assessing and reporting the risks associated with procurement? Are they tracking which procured products and services use AI?
- Is there an escalation protocol in place for proposed higher-risk AI uses?

AMBER

- Accountability for AI systems rests entirely with technical teams and/or relatively junior levels of management.
- Management is aware of internal AI system use, but has not assessed or documented employee or contractor use of third-party systems.
- Limited guidance/policy on use of AI and appropriate guardrails.
- Risk management frameworks are applied, but are not tailored to the amplified and new risks associated with AI.

RED

- AI understanding is highly concentrated in a few personnel.
- Management cannot confirm where and why AI systems are being used across the organisation.
- It is unclear who is responsible for the procurement, management, and outcomes of mission-critical AI systems.
- Existing risk management frameworks are not applied to AI use and procurement decisions.

2.2 GOVERNANCE STRUCTURES

Due to the complexity, rapid increase in use, and constant evolution of AI systems, it is critical that boards take a structured governance approach that appropriately leverages both diverse perspectives and expert insight.

At an early stage of AI adoption, organisations can use committee structures to individually review AI systems. However, as AI use proliferates across an organisation, other approaches, including risk triage and self-assessment of low-risk uses, may be necessary for effective governance.

SUGGESTED DIRECTOR STEPS

1. Determine **which existing or new board and management governance structure** (such as board and/or management committees) would most appropriately support AI oversight – see suggestions in **Box 5**.
2. Review **board and management committee charters** to determine whether and how they incorporate AI issues.
3. Consider how **external experts can be leveraged within existing governance structures.** For instance, the board should consider whether the relevant board or management committee should schedule briefings from an external AI expert (and if so, whether this will be on an ad hoc basis or on a regular/rolling basis). Having a more formal external advisory panel of experts may assist some organisations.
4. Consider **the nature and frequency of management reporting** to the board/relevant board committee.

BOX 5: What board committee is appropriate for AI governance issues?

In many cases, the board Risk Committee, which has a broad remit over organisational risk, will be best placed to have overall oversight over AI governance issues, with the more granular, operational-level AI issues left to management. Of course, such committees do not absolve the board from retaining overall responsibility for effective oversight.

At the management level, some larger organisations, such as Telstra (see **Case study 1**) and Microsoft (see **Case study 4**), have created AI Committees or Offices to assess and review current and potential AI use.

It is important to recognise that there is no 'one size fits all' and that each organisation's approach to AI governance structures is unique and dependent on the nature of the organisation, taking into account factors including size, sector and the role and significance of AI to the organisation.

CASE STUDY 1: Governance structures, Telstra[4]

Telstra has introduced specific governance structures to respond to the unique characteristics and challenges of AI systems. These structures provide advice, approvals, and create clear lines of oversight for Telstra's implementation of AI systems.

In 2019, Telstra introduced new operational procedures to give effect to its Responsible AI Policy, which included the creation of an AI Model Register for all AI use cases in Telstra and a review of all high-impact AI use cases by the Risk Council for AI & Data.

- **Risk Council for AI & Data (RCAID)** – RCAID is a cross-functional body with experts from across Telstra's business, including its legal, data, cyber security, privacy, risk, digital inclusion and communications teams. It provides a single, dedicated body to provide advice and approval. Any AI systems that are assessed as potentially having more than a low impact on stakeholders using an impact assessment process, including third-party systems, must be reviewed and either approved by RCAID, or escalated.

  Employees submit their AI use case proposals to RCAID, which meets fortnightly or otherwise as needed. RCAID assesses any potential risks, including any significant impacts on stakeholders. RCAID either approves the use case, makes recommendations to mitigate any risks, or escalates it to the Executive Data & AI Council if a decision cannot be reached, or if the use case is considered to represent higher levels of risk. Feedback from employees is that the process is 'absolutely essential' and 'great for getting advice'. By asking the right questions early, the RCAID process aims to avoid subsequent issues.

- **Executive Data & AI Council (Council)** – The Council, which is comprised of executives from each business function, has oversight and responsibility for the use of data and AI in Telstra. It provides oversight over RCAID and its operations and receives escalations from RCAID for individual use cases representing higher levels of risk. RCAID reports monthly to the Council on approved use cases.

- **Audit and Risk Committee (Committee)** – All significant risks are reported to the Committee which briefs the board twice a year on key issues (including in relation to data and AI). This enables effective board oversight of any significant AI matters.

As AI is increasingly adopted throughout Telstra, it is considering how to scale its AI governance processes, including ensuring responsible AI by design in its development of new systems, and exploring options for streamlining its review processes, such as self-assessment for low-risk uses of AI.

4 See HTI's AI Governance Lighthouse Case Study: Telstra (HTI, 2024) for further information.

KEY QUESTIONS FOR DIRECTORS TO ASK

- Which existing board and management committees are most appropriate for supporting oversight of AI?
- Do the relevant board and management committee charters/Terms of Reference need to explicitly stipulate board oversight of AI?
- Should the relevant board and management committee leverage external expertise? If yes, how?
- How, and how often, does management report on AI to the board/relevant board committee?

AMBER

- Committees related to AI have poorly defined responsibilities, decision-making authority and/or reporting requirements.
- There is a lack of cross-functional representation of business units on management committees or councils.
- Limited use of external experts.
- Limited reporting to the board on AI or reporting as an isolated one-off exercise.

RED

- The risks and benefits of AI do not receive board and/or board committee oversight.
- There are no processes for key stakeholders (such as employees or consumers) to have their views represented.
- No use of external experts.
- No board reporting (or only at board request).

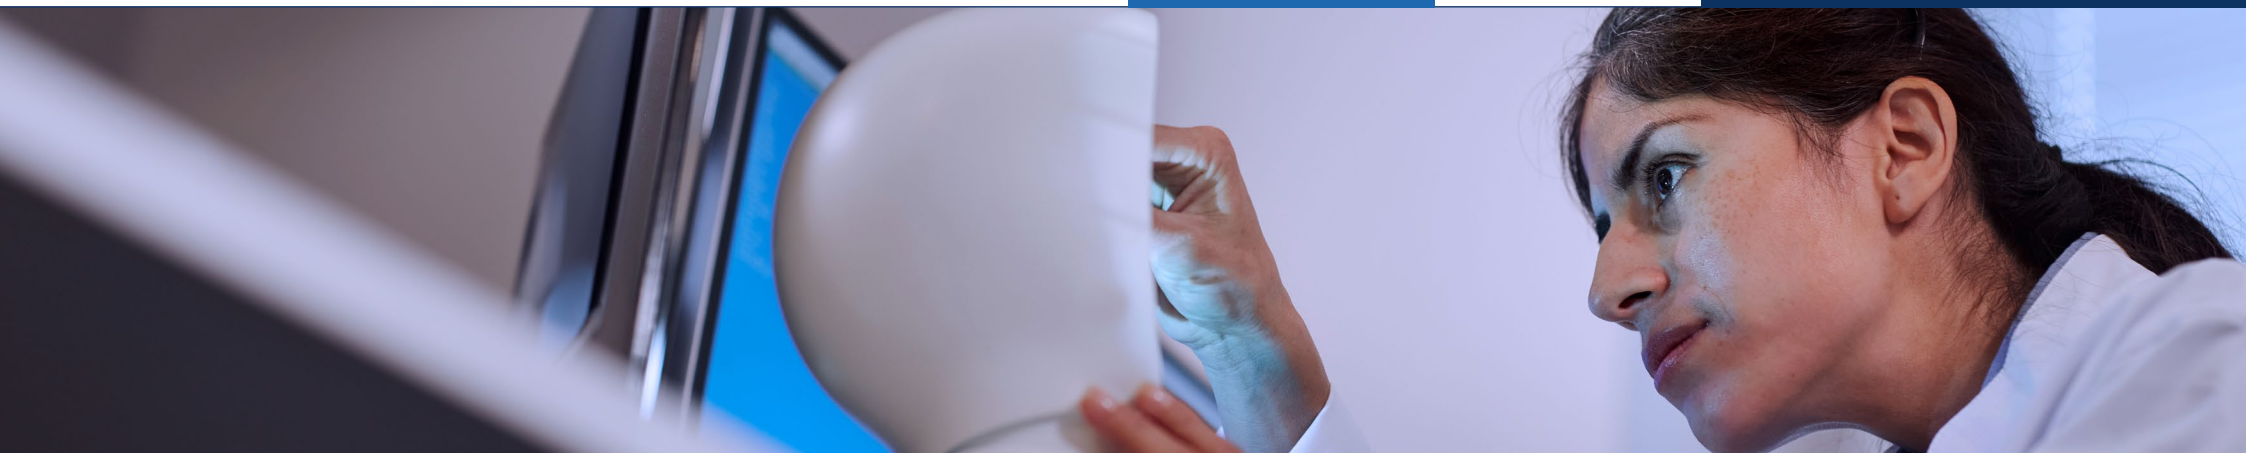

2.3 PEOPLE, SKILLS & CULTURE

AI has the potential to increase productivity across a broad range of functions throughout organisations. However, it also has the ability to transform roles, and therefore has a significant impact on the workforce more broadly (see **Box 6** in **section 2.2.2** of A Director's Introduction to AI). It is critical that organisations have the right talent and culture to promote safe and responsible AI use, and to navigate the related workforce impacts.

SUGGESTED DIRECTOR STEPS

1. Verify that management has **assessed the skills, capabilities and training required** across the organisation to benefit from AI systems and manage risks.

2. **Invest in appropriate management and director training** on the strategic opportunities, risks, and appropriate governance approaches related to AI systems.

3. Discuss with management the **impact the workforce and workforce planning**, such as the impact on hiring, promotion and skills development.

BOX 6: How much should directors know about AI?

While directors are not expected to be AI experts, a base understanding of AI, its risks and potential liability that may arise from these risks, is important. We set these out in A Director's Introduction to AI.

Expectations of board and management AI capability and competency will depend on how prevalent AI is within your organisation and sector – a higher baseline level of knowledge will likely be required for those operating in the technology industry and/or those significantly impacted by AI. Given the dynamic pace of change and innovation within AI, directors should consider how their base AI understanding accommodates recent developments.

In this highly technical area which is constantly evolving, directors should leverage external expertise, such as the establishment of Advisory Boards, or the inclusion of external stakeholder presentations and perspectives within board meetings.

It is not recommended that an AI expert be appointed to the board in lieu of hiring and/or developing appropriate management or director expertise.

KEY QUESTIONS FOR DIRECTORS TO ASK

- What baseline level of AI knowledge (i.e. minimum viable understanding) is required across the organisation?
- What AI capabilities are required by key accountable people?
- What AI-related training do staff receive at different levels and across functions?
- What training can directors receive to increase knowledge of AI risks and opportunities?
- How will AI impact the skills required of our workforce? Are there opportunities for training and redeployment?
- Have we communicated AI impacts to our workforce?
- What consultation or communication is taking place with our workforce on potential AI impacts?

AMBER

- Management views AI systems as a purely technical concern for the IT or data and analytics teams.
- Limited understanding of AI capabilities across the organisation.
- One-size-fits-all staff training.
- Limited engagement with staff when developing or deploying AI systems.

RED

- Lack of clarity about the required AI skills and capabilities of key accountable people.
- Evidence that staff members find it hard to 'speak up' when AI systems fail to operate as intended, particularly if staff are unaware of the system's intended outcomes.
- No understanding of workforce AI capability needs and investment required.
- No communication to employees about AI impacts.

BOX 7: AI training – Telstra and KPMG examples

The increasing adoption of AI systems, such as Generative AI, increases the need for training to enable employees to take advantage of the benefits of these systems whilst avoiding their risks. It is important to design and provide a training program appropriate for the uses of AI within an organisation.

Organisational training can provide employees with the minimum viable understanding of the organisation's use of AI systems. For example, Telstra requires all employees to undertake training on data and AI risks and governance as part of their annual 'Business Essentials Training'. Additional training is also available for interested employees through its Data & AI Academy. For more detail on Telstra's approach to AI governance, see Case study 1.

Training for employees is also necessary when new AI tools are introduced. For example, KPMG Australia is providing all its staff with training on KymChat, its internal Generative AI agent, regarding what it can do and how to engineer successful prompts. For more detail on KPMG's approach to AI governance, see Case study 2.

2.4 PRINCIPLES, POLICIES & STRATEGY

Guiding principles set the foundation for AI application, whilst **policies** provide practical and operational guidance on AI use. Both are required for effective AI governance.

Directors should also consider how AI fits within broader organisational strategy and how AI can be leveraged to meet business objectives (see **section 1.5**). This is particularly important as AI systems increasingly penetrate core business functions, with the most rapid growth in strategy, corporate finance and risk.

SUGGESTED DIRECTOR STEPS

1. **Require that AI is considered and, where appropriate, embedded within the organisation's strategy.** Organisations should set clear strategic objectives as to how AI will be used to deliver organisational goals. 'AI for AI's sake' should be avoided.
2. Engage with management to discuss **how high-level safe and responsible AI principles, such as Australia's AI Ethics Principles, have been made actionable** via specific policies.
3. **Introduce an organisational AI use policy** to facilitate safe and responsible AI use and reduce shadow AI use.
4. Integrate AI into **relevant policies (such as privacy, data governance, cyber and procurement)** for a holistic strategic and risk management approach. These policies should be reviewed periodically for currency.

KEY QUESTIONS FOR DIRECTORS TO ASK

- How does our current and intended use of AI support our overall strategy?
- Are the AI principles, policies and strategy adaptable, scalable and broad enough to capture a wide range of current and potential AI use cases within the organisation?
- How clearly documented is the organisation's approach to AI use?
- What AI-specific policies are in place to guide AI use across the organisation and its supply chain?
- Do we have a clear policy on the use of Generative AI and the risks posed by shadow AI use?
- Do our existing privacy, data governance, cyber and procurement policies address AI? Are these fully aligned with how we intend to leverage AI systems in our strategy?

AMBER

- The organisation has set AI principles without detailed guidance for employees about expected practices, responsibilities or frameworks when making decisions about the use of AI.
- Privacy, data governance, cyber and procurement polices are inconsistent with, or not integrated with, the AI policy.
- The organisation has set AI principles without considering how it fits with broader organisational strategy.

RED

- No policy or guidance on AI use.
- AI principles and policies do not align with broadly accepted principles of safe and responsible and ethical AI use (such as Australia's AI Ethics Principles).
- Presentation of a single AI policy without comprehensive review of other highly relevant organisational policies (e.g. privacy and data governance policies).
- Outright bans on all AI as a medium/long term strategy – it is likely some employees will still use AI on personal devices.

2.5 PRACTICES, PROCESSES & CONTROLS

While many organisations have adopted principles around ethical or responsible AI, this alone is insufficient. Clear practices, processes, and controls aligned to the specific characteristics of AI systems are necessary to implement and enforce the organisation's approach to safe and responsible AI across the value chain.

SUGGESTED DIRECTOR STEPS

1. Require relevant **controls** for AI use, and that these controls are regularly reviewed and updated for alignment with best practice.

   Controls include:
   - **Risk appetite statement and risk management framework**: The risk appetite statement and risk management framework should be reviewed and updated so that they incorporate and address AI risks. A fit-for-purpose risk management framework should include processes to determine high-risk and low-risk AI applications.
   - **AI impact assessments** should be undertaken to identify, assess and respond to opportunities and potential risks/harms arising from AI use (see Box 8).
   - **Compliance frameworks and policies:** Risk and other relevant policies such as privacy, and data. cyber and procurement, should be updated to account for regulatory, commercial, and reputational risks arising from AI use.
   - **Other relevant policies and templates/precedents and contracts** should be reviewed and updated to incorporate safe and responsible AI practices.

2. Confirm with management that there are processes in place to assess supplier and vendor risk.

BOX 8: AI Impact Assessments[5]

AI Impact Assessments identify, assess, and respond to opportunities and potential risks and harms arising from AI use. They are generally undertaken when a potential AI use is being proposed, and are often scrutinised by the relevant board or management body tasked with reviewing AI use cases within the organisation.

ISO Standard 42005 (currently under development) aims to provide guidance for organisations performing AI system impact assessments, including consideration of key documentation and appropriate stages of the AI system lifecycle.

Best practice is to undertake stakeholder engagement to consider the full gamut of AI opportunities as well as risks and harms, and to address blind spots and bias. Enacting the principle of *'nothing about us without us'* is crucial to stakeholder engagement as part of robust AI governance systems.

Examples of AI Impact Assessments include:

- **Public sector AI governance framework** including the Canadian Government Algorithmic Impact Assessment tool and associated transparency requirements; UK Algorithmic Transparency Recording Standard; and NSW Government Artificial Intelligence Assurance Framework.
- **Voluntary AI risk management frameworks** such as the NIST AI Risk Management Framework and ISO Standard 42001:2023
- **Corporate policies** such as Microsoft's Responsible AI Impact Assessment Template, and Atlassian's Responsible Technology Review Template.

5 See HTI's AI Governance Snapshot #2 *Putting people at the centre of AI – impacted communities and missing voices* (HTI, 2024) for further information.

KEY QUESTIONS FOR DIRECTORS TO ASK

- What is our risk appetite for AI use? Have we updated our risk appetite statement?
- What AI Impact Assessment and risk management tools or frameworks are we currently using?
- Does our risk management framework incorporate risks arising from AI? Does it differentiate between high-risk and low-risk AI applications?
- What steps are we taking to be confident that we are meeting our legal and regulatory obligations for the use of AI and associated data collection, storage, and use?
- Do we have robust testing and piloting approaches for AI systems under real-world conditions?
- What process are we using to assess supplier and vendor risk?
- What notification requirements are there for suppliers to advise of AI use or introductions to products?
- What capacity do we have to reject updates (such as software products) if deemed not to be in line with organisational policy on AI use?

AMBER

- Management are unaware of international standards around AI risk management.
- Controls are only considered and implemented at a single point in time.
- There is no interrogation or independent verification of vendor claims regarding AI performance or risks.

RED

- Lack of specificity as to the AI impact and risk assessment tools being used by the organisation.
- The organisational risk appetite statement does not include AI.
- The organisation's risk management framework does not include AI risks.
- Absence of a process to assess supplier and vendor AI risks and respond to them.

2.6 SUPPORTING INFRASTRUCTURE

Fundamental to any AI governance framework is the supporting infrastructure. Supporting infrastructure includes the systems required to deliver the required governance practices. Given AI systems are dependent on data, having effective data governance in place is crucial.

SUGGESTED DIRECTOR STEPS

1. Verify that management has an appropriate **AI system and data inventory in place** (see Box 4).
2. Confirm that **data governance policies have been reviewed and updated** to account for AI systems' specific characteristics.
3. Confirm that **cyber security policies have been reviewed and adapted to address AI systems** and mitigate novel attacks and misuse.

KEY QUESTIONS FOR DIRECTORS TO ASK

- Where, how, and why is AI being used across our organisation? Have we created an AI inventory?
- What internally- and externally-sourced data is being – or could be – used as an input or for training to AI systems?
- Have we reviewed the legality of the collection, storage, and use of the data used within our organisation and as input for AI systems?
- How do our data governance and cyber security policies and practices support the responsible use of AI?
- Does the system architecture enable transparency or explanation of decisions made by AI?

AMBER

- Lack of investment in systems and processes which can provide transparency and traceability of AI system use and performance.
- Incomplete or irregularly updated data inventory.
- Limited management understanding of risks associated with key AI vendors / products.

RED

- Absence of a data inventory or data governance policy.
- No periodic purging of data.
- Cyber security policy and practices which do not account for AI use.

CASE STUDY 2: Data governance and Generative AI, KPMG Australia[6]

In March 2023, KPMG launched KymChat, an internal Generative AI system for its employees. KymChat was originally designed to assist staff in locating the right expert across the business. However, in light of its success and flexibility, its functionality was expanded. KymChat is now used for a variety of uses including answering questions about internal policies, compiling thought leadership, and (in approved cases), preparing draft advice to clients.

KPMG partnered with Microsoft to build KymChat on Microsoft Azure's OpenAI Service. This meant that client data did not leave the KPMG environment such that KPMG standards for privacy, confidentiality, and data protection were maintained. Internal data access was also strictly regulated according to employee work relevance.

KPMG follows strict data governance processes to assess data for ownership rights, lineage, provenance and bias so that they are not inappropriately or unlawfully using anyone's data. The following steps are taken:

- **Permission:** KPMG does not include data within KymChat without investigating and confirming that it has the right and permission to use that data.
- **Anonymisation:** If KPMG material was originally prepared for a client, and where consent is provided, the material is sanitised so that any identifying or confidential information is removed.
- **Legal review:** KPMG's legal team provides sign off on the use of data by KymChat.

KPMG's experience shows how important data governance is for Generative AI. Often, these systems are developed by giving them as much data as possible. However, an organisation's existing internal knowledge, such as policies, precedents and other work products are often unstructured and have never formally had data governance applied to them.

Without quality data governance, AI is unlikely to deliver effective and safe outcomes. Organisations need to ask themselves about the source of the data and whether it is trusted, the quality of the data and how to assess that, and whether the organisation has the right to use the data in its AI systems.

6 See HTI's AI Governance Lighthouse Case Study: KPMG Australia (HTI, 2024) for further information.

2.7 STAKEHOLDER ENGAGEMENT & IMPACT ASSESSMENT

Because AI systems can transform aspects of an organisation and its relationships with stakeholders, it is critical that organisations engage with these stakeholders to explain and manage impact.

Further, in contrast to traditional IT projects, the impact of AI on stakeholders (in terms of the outputs or outcomes of the AI system) may change or evolve during different parts of the AI lifecycle. For instance, after deployment, the accuracy or predictive ability of AI models decreases (known as 'model drift').

Finally, given the decreased explainability of AI systems, effective engagement may help stakeholders better understand how these systems operate, and enable organisations to better respond to their concerns.

It is important to recognise the disproportionate and negative impact of AI bias on vulnerable and marginalised populations (see **Chapter 2** of A Director's Introduction to AI). Engagement with these groups should be prioritised.

SUGGESTED DIRECTOR STEPS

1. **Identify and engage with stakeholders to understand AI's impact and stakeholder expectations** of AI use and governance.
2. Request that management **review AI system design and assessment processes and policies to confirm they incorporate accessibility and inclusion practices** (so as to reduce the risk of bias).
3. Consider whether AI-generated results/outcomes are explained to stakeholders and whether an appeal process is available.

BOX 9: Director guidance on stakeholder engagement

The AICD's Elevating stakeholder voices to the board: A guide to effective governance assists directors in identifying and elevating key stakeholder voices to the board.

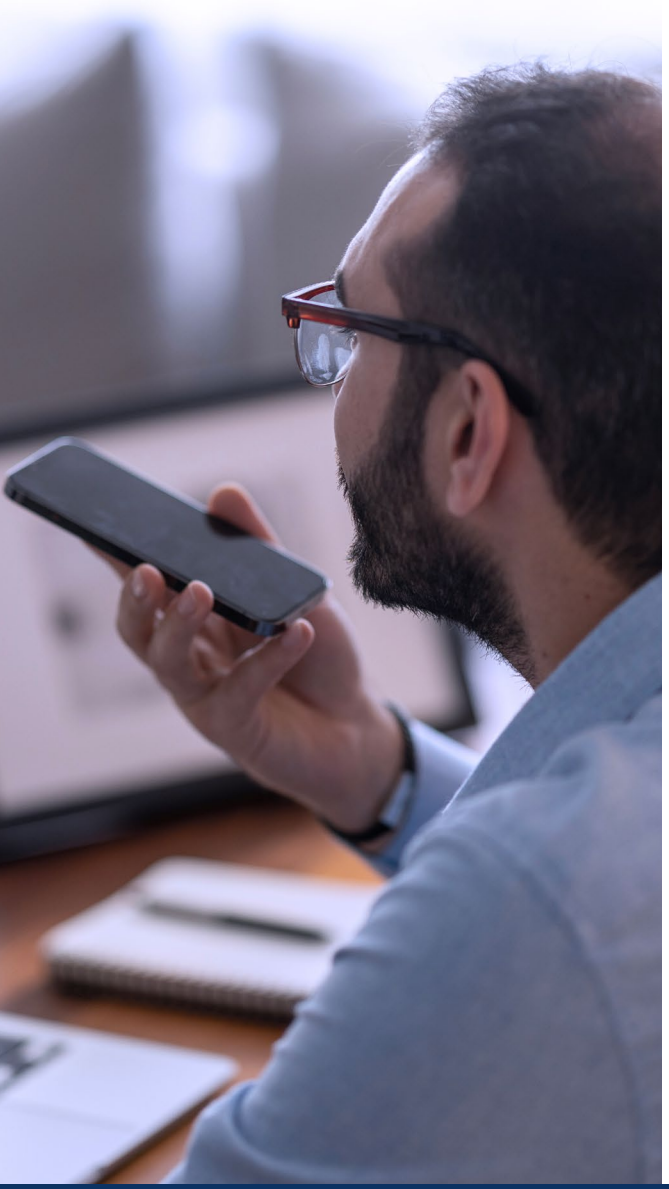

KEY QUESTIONS FOR DIRECTORS TO ASK

- How does our AI Impact Assessment incorporate stakeholder views? (See **Box 9**)
- What processes do we have in place to understand the potential AI harms arising to impacted stakeholders?
- How are we ensuring the voices of potentially vulnerable stakeholders are represented in engagement mechanisms?
- How do we include the participation of stakeholders in the development of safe and responsible AI principles and policies and governance frameworks?
- What processes are in place for impacted stakeholders to request reasons, contest, or provide redress for decisions made by AI systems?

AMBER

- Stakeholder consultation tends to involve repeated engagement with a narrow set of stakeholders, which are not representative of potentially impacted groups.
- Lack of compensation for the participation of civil society stakeholders.
- Stakeholder engagement is rushed or sought late in the process.

RED

- Management suggests that stakeholders do not understand or have ill-informed views.
- The organisation has no stakeholder engagement process.
- Management assesses AI system impact and risks internally without engaging potentially impacted parties, even for high-risk applications.

CASE STUDY 3: Stakeholder engagement, University of Technology Sydney[7]

UTS undertook a novel consultation process with UTS students, tutors and academics, using the principles of Deliberative Democracy (DD) to collectively determine the principles that should govern the use of analytics and AI at UTS. This structured process involves the creation of Deliberative Mini-Public (DMP), which has the ability to influence policy and decision-making, includes representative and diverse viewpoints, and provides for open dialogue and deliberation. Importantly, the DMP must be sanctioned by senior leadership with a commitment that its recommendations matter.

The 20 members of the DMP were selected from 131 applicants using stratified sampling to ensure a representative and balanced mixture of gender, faculty, and students. Across five workshops run over seven weeks, the DMP identified the principles that should govern UTS' AI use: accountability/ transparency, bias/ fairness, equality and access, safety and security, human authority, justifications/ evidence, and consent. This in-depth process provided non-tokenistic engagement that gave participants responsibility for the outcomes.

Students and staff felt empowered by this process, building engagement and trust. Staff commented that they had never been involved in such a meaningful consultation at the university before. Meanwhile, students reported that they felt privileged to be part of it and developed a sense of ownership of the process and outcome.

To build on its partnership with students, UTS initiated a series of 'Student Partnership in AI' workshops, using deliberative democracy principles and processes:

- **Generative AI:** This workshop explored how such technology could be used responsibly to assist learning outcomes for students, and issues surrounding the use of automated software to detect AI writing in assessments.
- **Predictive AI:** This workshop discussed a pilot machine learning model to identify and support students who may withdraw from UTS before census date.

Each workshop had 20 participants recruited to maximise the diversity of voices, such that all faculties were represented at both undergraduate and postgraduate level, providing UTS with the opportunity to hear a wide range of student voices and consider their feedback and concerns about these technologies.

Having built both trust and credibility with staff and students, UTS leadership proceeded to develop its AI policies and a dedicated governance committee in the forum of the Artificial Intelligence Operations Board. This body is tasked with developing institutional knowledge and insights about the use, management, and control of AI for the purposes of teaching, learning and operations at UTS, and responsible for endorsing the use of AI systems across the university.

7 See HTI's AI Governance Lighthouse Case Study: University of Technology Sydney (HTI, 2024) for further information.

2.8 MONITORING, REPORTING & EVALUATION

Both the value and risk of AI systems arise from their ability to learn and adapt. After deployment, AI models can experience 'model drift' or degradation in performance. Both AI systems themselves and overarching AI governance frameworks are, therefore, not 'set and forget' – they require regular re-assessment against key performance indicators (KPIs) and metrics, new or potential regulation, and broader market and technological developments.

SUGGESTED DIRECTOR STEPS

1. Verify that management has implemented **a risk-based monitoring and reporting system** for AI systems that are mission-critical and/or could cause significant harm, including AI systems and vendor systems.
2. Establish **clear metrics and outcomes to track and measure the performance** of the AI governance framework.
3. Develop and implement a **monitoring and reporting framework and frequency**.
4. Consider **seeking internal and external assurance.**

KEY QUESTIONS FOR DIRECTORS TO ASK

- What KPIs are we using to assess whether the AI governance framework is performing as intended?
- What is the appropriate performance framework and reporting frequency to enable the organisation to capitalise on opportunities and address risks?
- How are we identifying and responding to errors in our AI systems?
- How are we using internal and external audit as a check and balance?
- What are the limitations of our internal and external audit processes? Are these clearly disclosed in our reporting?

AMBER

- Absence of key indicators of responsible AI use and performance at an organisational and system level.
- No clear process for consideration of internal or external audit recommendations.
- Monitoring system implemented without clear KPIs.

RED

- No consideration or investment in ongoing monitoring, reporting and evaluation of AI systems.
- Management are unaware of where AI is present in mission-critical systems.
- No clear line of reporting of risks to the board.

CASE STUDY 4: A vendor's view: Microsoft's holistic approach to AI governance[8]

Microsoft's Responsible AI Governance Framework has six responsible AI principles at its core, being: (1) accountability; (2) inclusiveness; (3) reliability and safety; (4) fairness, (5) transparency; and (6) privacy and security.

These principles are supported by policies and standards, and practical implementation practices in the form of training, tools and testing. Microsoft then institutes checks and balances in the form of monitoring and auditing to ensure compliance. Finally, practices are reported for transparency and tracking.

However, principles and policies are not enough. Effective governance structures are critical to effective oversight of the implementation of responsible AI. Microsoft uses a three-tiered system which comprises:

- **Aether**, whose research-led working groups provide subject-matter expertise on emerging trends with respect to Microsoft's AI principles.
- **The Office of Responsible AI (ORA)**, which sets company-wide policies and practices for responsible AI and ensures internal roles and responsibilities are clearly defined. ORA also ensures readiness to adopt responsible AI practices within Microsoft and supports customers and partners to do the same. It operates the intake and triage function for sensitive use cases and also formulates and advocates for responsible AI public policy externally.
- **The Responsible AI Strategy in Engineering (RAISE) group**, which enables Microsoft's engineering teams to implement responsible AI processes through systems and tools.
- **The Environmental, Social and Public Policy Board Committee** provides oversight of its responsible AI program at the board level.

Through adopting a multi-disciplinary holistic governance approach, Microsoft seeks to embody responsible AI principles and practices within its company and across their value chain.

8 See Microsoft's Responsible AI website for more information and resources.

Conclusion

AI systems offer a wide range of potential benefits, but directors need to take care when overseeing its deployment by confirming that the organisation has implemented effective, safe and responsible AI governance practices.

Applying the lessons from this resource is an important starting point, but it is not intended to be a comprehensive guide.

Directors need to work with management to carefully consider the regulatory and governance implications of AI systems specific to their organisations and their industry. They must also stay up-to-date on key regulatory and policy developments to maintain a 'minimum viable understanding' of AI governance.

Whatever the future of AI regulation brings, there is already a broad range of existing legal obligations that apply to an organisation's use of AI systems which must be complied with.

By understanding the specific challenges and characteristics of AI and its impact on governance, directors will guide their organisations to deploy AI systems safely and responsibly for maximum strategic and competitive advantage.

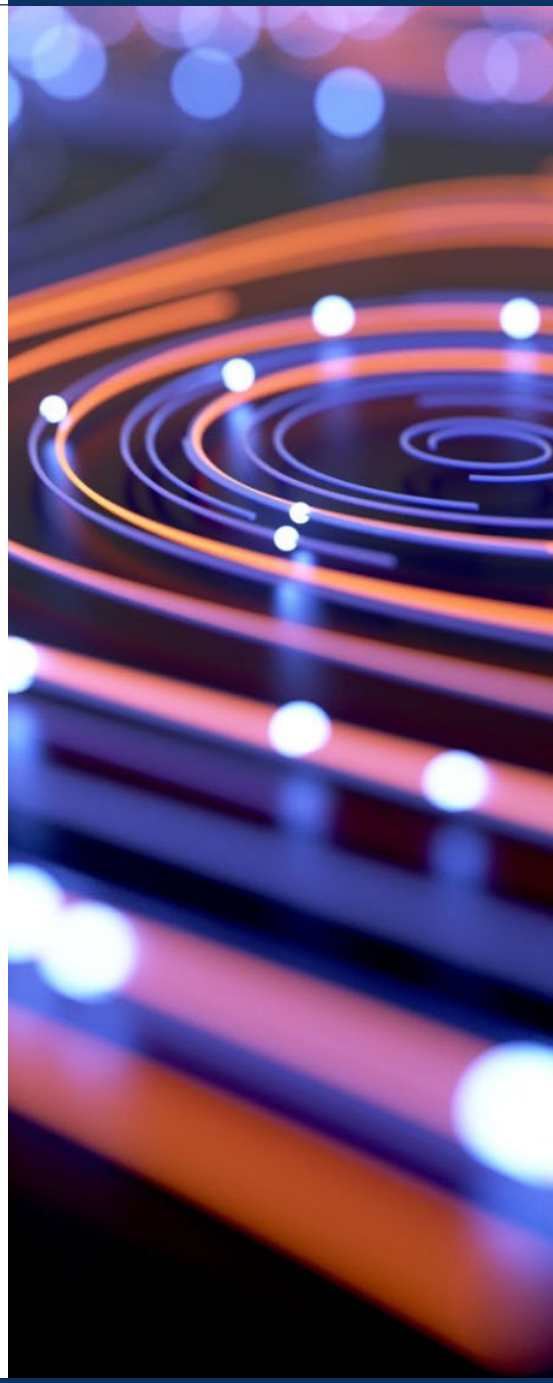

Appendix - Additional resources

- AICD (2019), Ethics in the Boardroom - a guide to decision making
- Australian Human Rights Commission, Technical Paper: Addressing Algorithmic Bias (2020); Guidance Resource: Artificial intelligence and discrimination in insurance pricing and underwriting (2022); and HRIA Tool: AI in Banking (2024).
- Atlassian (2023), Responsible Technology Review Template and No BS Guide to Responsible Tech Review.
- Gradient Institute and CSIRO (2023), Implementing Australia's AI Ethics Principles: A selection of Responsible AI practices and resources.
- Human Technology Institute (HTI) - Lauren Solomon and Nicholas Davis (2023), The State of AI Governance in Australia and Insight Summary.
- HTI (2024) AI Governance Snapshot Series: Essential Components of AI Governance; Putting people at the centre of AI – impacted communities and missing voices.
- HTI (2024) AI Governance Lighthouse Case Study Series: Telstra; KPMG Australia; University of Technology Sydney.
- ISO/IEC 42001 (2023): Artificial Intelligence Management System.
- KPMG & The University of Queensland (2023), Trust in Artificial Intelligence A global study.
- Microsoft (2022), Responsible AI Impact Assessment Template and Responsible AI Impact Assessment Guide.
- NIST (2023), Artificial Intelligence Risk Management Framework (AI RMF 1.0).

Acknowledgements

The AICD would like to acknowledge our Guide co-authors:

Professor Nicholas Davis MAICD
Co-Director, Human Technology Institute
University of Technology Sydney

Lauren Solomon
Lead, AI Governance (until April 2024)
Human Technology Institute,
University of Technology Sydney

Llewellyn Spink
AI Corporate Governance Specialist
Human Technology Institute,
University of Technology Sydney

The AICD would also like to acknowledge and thank the following people who were involved in the review of the Resource:

- Alison Kitchen MAICD
- Kee Wong FAICD
- Phil Coffey GAICD
- Wendy Stops GAICD

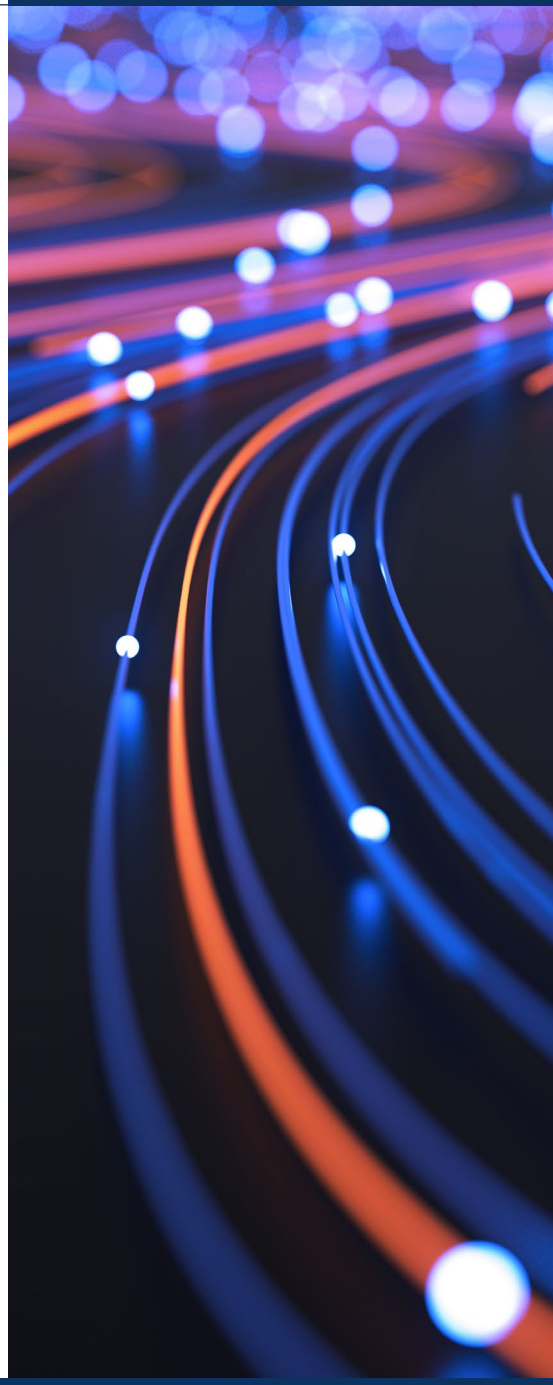

ABOUT AICD

The Australian Institute of Company Directors is committed to strengthening society through world-class governance. We aim to be the independent and trusted voice of governance, building the capability of a community of leaders for the benefit of society. Our membership includes directors and senior leaders from business, government and the not-for-profit sectors.

ABOUT HTI

The UTS Human Technology Institute (HTI) is an impact-oriented institute building human values into new technologies. Bringing together policy, legal and technical experts, HTI provides independent expert advice, policy development, capability building, and data science solutions to support government, industry and civil society.

DISCLAIMER

The utmost care has been taken to ensure this document accurately reflects the legislative and regulatory landscape as at the date of publication. However, this is an area subject to constant regulatory and legal change. The material in this publication does not constitute legal, accounting or other professional advice. While reasonable care has been taken in its preparation,the AICD and HTI do not make any express or implied representations or warranties as to the completeness, reliability or accuracy of the material in this publication. This publication should not be used or relied upon as a substitute for professional advice or as a basis for formulating business decisions. To the extent permitted by law, the AICD and HTI exclude all liability for any loss or damage arising out of the use of the material in the publication. Any links to third party websites are provided for convenience only and do not represent endorsement, sponsorship or approval of those third parties, any products and services offered by third parties, or as to the accuracy or currency of the information included in third party websites. The opinions of those quoted do not necessarily represent the view of the AICD or HTI. The AICD and HTI reserve the right to make changes without notice where necessary.

For more information on A Director's Guide to AI Governance:

T: 1300 739 119

E: **policy@aicd.com.au**

JOIN OUR SOCIAL COMMUNITY

aicd.com.au

CCT1403_24