

DEPARTMENT OF INDUSTRY, SCIENCE AND RESOURCES

MS23-000500

To: Minister for Industry and Science (For Information only)

LAUNCH OF THE RESPONSIBLE AI INDEX REPORT BY THE GRADIENT INSTITUTE

Timing: Routine

Recommendation:			
1. That you note the upcoming launch of the 2023 Responsible AI Index Report by the Gradient institute during the first week of April and the opportunity to make comments to the media.			
			Noted / Please Discuss
Minister:		Date:	
Comments:			
Clearing Officer:	s22	General Manager, Emerging Technologies Branch, Technology and Digital Division	Ph: s22 Mob: s22
Contact Officer:	s22	Manager, AI Policy Section, Emerging Technologies Branch, Technology and Digital Division	Ph: s22 Mob: s22
For Parliamentary Services' use only. Date Submitted to the Minister's office in PDMS:			17/3/2023

Key Points:

1. The Gradient institute will be launching the 2023 Responsible AI Index Report (the report) in the first week of April (embargoed copies of the report and infographic are at Attachment A). This is the second Responsible AI Index Report to be released, with the inaugural report released in 2022.
2. Looking at 2022, the 2023 report found:
 - a. At an overall level, there has been little change since 2021 in the overall performance of Australian organisations in developing and implementing Responsible AI systems. Performance is higher for those with the CEO leading the AI strategy.
 - b. Compared with 2021, more organisations are taking an enterprise-wide approach for the development of AI which is tied to the wider business strategy across all divisions.

- c. Encouragingly, awareness of Australia's AI Ethics principles has increased since 2021. There is also a high level of agreement with statements about how organisations have developed AI systems consistent with the intent of each principle.
 - d. Organisations that are more mature in their deployment of Responsible AI, are likely to see significant gains in terms of competitive advantage, with the benefits outweighing the costs.
3. This index is an opportunity to benchmark the success of our policies and initiatives, such as the recently launched Responsible AI Network. We will use this report moving forward on existing and new policy initiatives as a measure of success.
 4. The Gradient Institute has advised that they intend to pitch an exclusive interview to the Australian Financial Review. They have noted that they would like to have your comments for the interview as well. If you wish to be involved we can approach the Gradient Institute with your interest.

Background:

5. The Responsible AI Index report is a study of over 400 Australian-based organisations and their awareness of, intentions around, and use of Responsible AI practices. This report was launched in partnership with IAG and Transurban.
6. The report aims to provide a comprehensive assessment and to track the status, in Australian-based organisations, of: 1) AI maturity and, 2) the extent to which AI is being deployed responsibly to mitigate potential risks, and 3) to make recommendations as to how organisations may use AI responsibly.

Sensitivities and Handling: Nil.**Data referenced:**

s47E(d)

Consultation: Nil.**Attachment**

- A:** Responsible AI Index infographic and report

DRAFT



RESPONSIBLE AI INDEX 2022: REPORT

CONTENTS

1. BACKGROUND

2. THE RESPONSIBLE AI INDEX

3. THE STATE OF AI IN AUSTRALIA

4. AI ETHICS PRINCIPLES

5. PRINCIPLES IN PRACTICE

6. BRIDGING THE GAP

The 2022 Responsible AI Index is grounded in a robust quantitative methodology.

Objectives

Research Aim

To provide a comprehensive assessment and to track the status, in Australian-based organisations, of:

- 1) AI maturity and,
- 2) the extent to which AI is being deployed responsibly to mitigate potential risks, and
- 3) to make recommendations as to how organisations may use AI responsibly

Topics covered

MATURITY SEGMENTS

AI USAGE

AI ETHICS PRINCIPLES

ATTITUDES

ORGANISATIONAL STRATEGIES

ACTIONS TAKEN OR PLANNED

Audience

Sample

The sample for the study was made up of:

- Organisations based in Australia
- AI Strategy Decision Makers (e.g., CIOs, CTOs, CDOs, heads of data etc) working in organisations with 20 or more employees
- Covered a range of businesses by size, industry and location
- Organisations that had deployed AI in their business or were planning to do so in the next 12 months

Sample Size

N=439 RESPONDENTS

Methodology

Source

B2B online panel

Methodology


10-minute online survey

Timing

Fieldwork was conducted 24th October - 9th November 2022

Respondent organisations represent a range of industries which have been organised into seven different groups.

INDUSTRY GROUPS



Please note that due to an overrepresentation of respondents in the Technology & Telecommunications sector in the 2022 sample, the sample has been weighted back to the 2021 sample to enable comparability between the 2021 and 2022 Responsible AI Index.

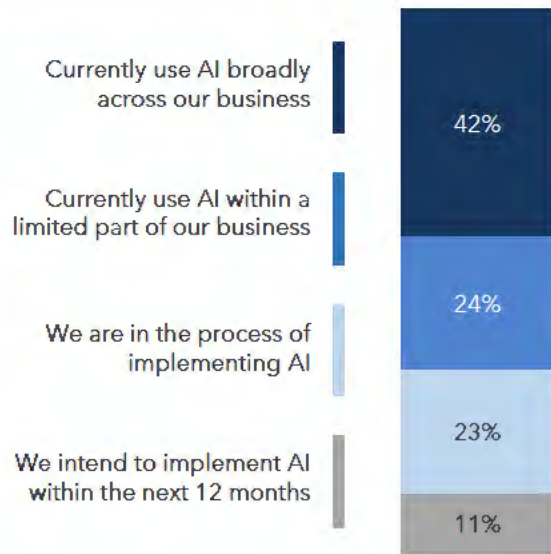
SAMPLE PROFILE

DISR - For release under the FOI Act

LEX 75414 - FOI - Document 1

The sample is based on roles with significant influence over the AI strategy of organisations with at least 20 employees. The sample covers a range of organisation sizes and locations, with a mix of AI usage. All organisations were either using AI or planning to implement AI in the next 12 months.


USE OF AI




JOB TITLE



BUSINESS SIZE




LEVEL OF INVOLVEMENT



GENDER




BUSINESS LOCATION



CONTENTS

1. BACKGROUND
- 2. THE RESPONSIBLE AI INDEX**
3. THE STATE OF AI IN AUSTRALIA
4. AI ETHICS PRINCIPLES
5. PRINCIPLES IN PRACTICE
6. BRIDGING THE GAP



INTRODUCING THE RESPONSIBLE AI INDEX

DISR - For release under the FOI Act

LEX 75414 - FOI - Document 1

To understand how organisations are using and developing AI in a responsible manner, a maturity model was created based on how respondents rated their organisation across a battery of statements about responsible AI and actions taken to implement AI in a responsible way, combined to provide a total score out of 100.

MATURITY MODEL


The maturity model is built on:

1. Self-assessed performance on 16 statements about Responsible AI across five categories:
 - i. Strategy & Leadership
 - ii. Governance
 - iii. People & Skills
 - iv. Data & Security
 - v. Monitoring & Review


2. A tally of the number of actions taken to implement AI responsibly from a total list of 13 actions

MODEL WEIGHTING

Weighted to account for ¾ of the model:




Weighted to account for ¼ of the model:



MODEL SCORE

Combined to provide a final score out of 100:



THE RESPONSIBLE AI MATURITY ASSESSMENT FRAMEWORK: PERFORMANCE ASSESSMENT

DISR - For release under the FOI Act

LEX 73414 - FOI - Document 1

The performance assessment component of the model is based on a self-assessed rating of performance (score 0-10) on the following categories:



STRATEGY & LEADERSHIP

- Having a leadership team that is clearly accountable for the impact of AI systems
- Having a leadership team that is demonstrably committed to the responsible use of AI
- Having a strategy in place for the responsible use of AI which stays up to date with emerging best practice and international frameworks, and is reviewed on an ongoing basis
- Having formal organisational routines (for example rewards, recognition, etc.) to incentivise responsible use of AI



GOVERNANCE

- Having appropriate mechanisms in place to allow individuals materially impacted by an AI-driven decision to understand and/or challenge that decision
- Scrutinising the systems and processes used by potential AI suppliers to ensure they are designed to not harm, deceive or cause unfair treatment of individuals, communities or groups
- Having robust processes to ensure all AI systems are compliant with relevant regulation and laws
- Having an ethical (or equivalent risks) framework in place to ensure AI-systems are formally assessed consistently against clear standards that account for its impacts on individuals, communities and groups



DATA & SECURITY

- Having robust systems and processes in place to ensure personal information used or created by AI systems is appropriately protected
- Reviewing underlying databases for potential bias to help ensure AI systems do not result in unfair treatment of or discrimination against individuals, communities or groups
- Having documented policies and processes in place to quickly respond to and resolve any adverse customer outcomes caused by the unauthorised use of AI systems



PEOPLE & SKILLS

- Including both technical and non-technical consultants or professionals (e.g. social scientists, psychologists, ethicists, legal experts) as well as customer representatives to review AI systems for the potential for harmful outcomes to customers
- Hiring/engaging a diverse (different cultures, genders, etc.) workforce to consider broader perspectives and consideration of risks into the development process
- Ensuring AI system designers and developers are appropriately skilled and knowledgeable about the ethical implications of their work, including risks of discrimination and bias and techniques to address these



MONITORING & REVIEW

- Routinely monitoring AI systems using clear metrics designed to trigger suitable corrective or remediation action when AI systems are not working as intended, for example monitoring of bias and the accuracy of decisions
- Where decisions have a material impact on individuals, communities or groups conducting a regular, independent peer review of all aspects of AI-systems and their impact

THE RESPONSIBLE AI MATURITY ASSESSMENT FRAMEWORK: PERFORMANCE ASSESSMENT

DISR - For release under the FOI Act

LEX 73414 - FOI - Document 1

The behavioural component of the model is based on the number of actions taken by the organisation out of the 13 possible options below. This component corrects any over-rating by respondents of their self-assessed performance:



STRATEGY & LEADERSHIP

- Reviewed global best practices and frameworks
- Engaged your leadership teams on issues around responsible AI



GOVERNANCE

- Reviewed the systems and processes used by AI vendors
- Evaluated the risks and opportunities for human rights
- Developed materials to aid decision making processes



DATA & SECURITY

- Reviewed underlying databases for potential bias



PEOPLE & SKILLS

- Hired technical consultants or professionals
- Consulted specialists in ethical AI
- Hired a more diverse workforce
- Hired non-technical consultants or professionals



MONITORING & REVIEW

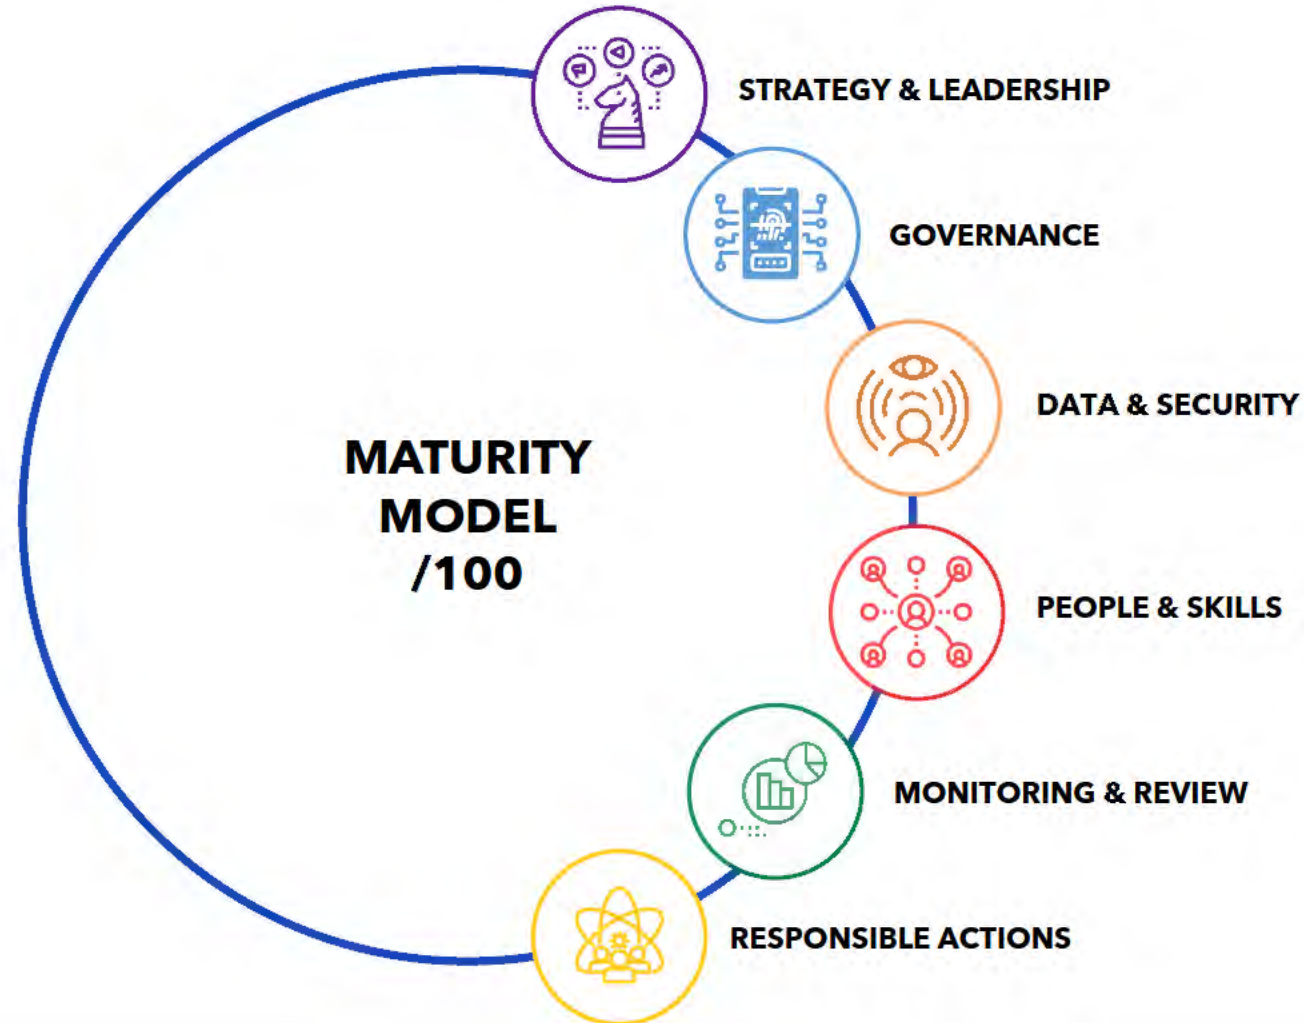
- Reviewed AI algorithms for potential bias
- Monitored outcomes for customers or employees
- Sourced legal advice around potential areas of liability

BUILDING THE MATURITY MODEL

DISR - For release under the FOI Act

LEX 73414 - FOI - Document 1

While the first five dimensions were given equal weight to each other in the model to represent the attitudinal component, the sixth behavioural dimension is weighted to be $\frac{1}{4}$ of the total score to reflect the importance of actions according with self-reported behaviours. After calculation, the total score was rebased to 100.



THE RESPONSIBLE AI MATURITY ASSESSMENT FRAMEWORK: PERFORMANCE ASSESSMENT

DISR - For release under the FOI Act

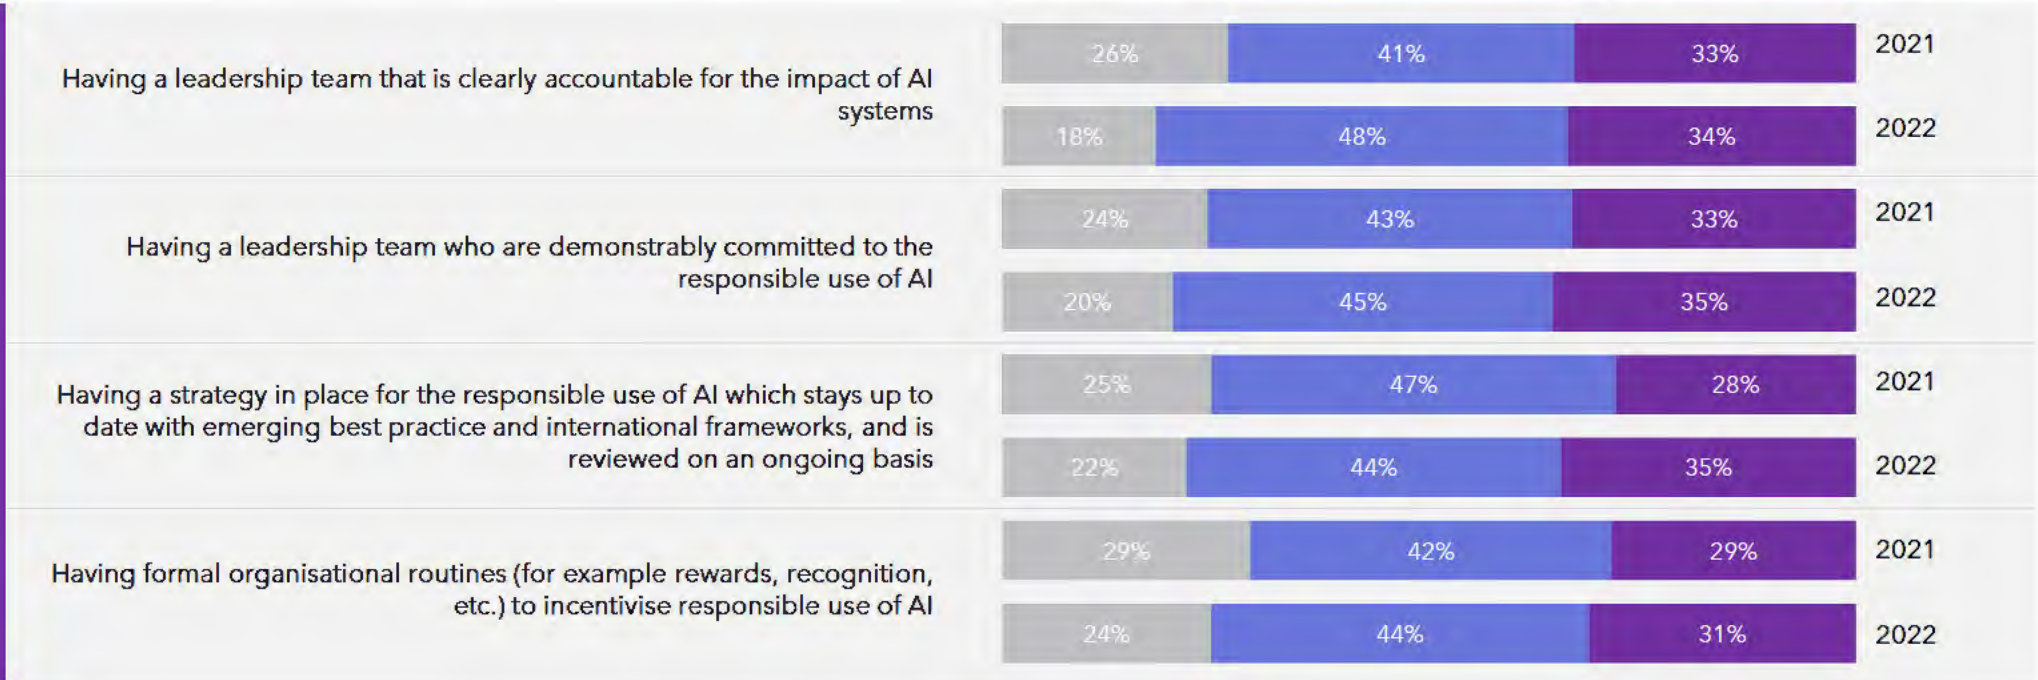
LEX 75414 - FOI - Document 1

Respondents rated their organisation's performance from 0-10 on the Strategy & Leadership dimensions through individual statements, with results in 2022 showing a similar distribution of scores across each attribute, with top 2 box scores typically higher than those in 2021.



STRATEGY & LEADERSHIP

% of respondents who rated themselves (0-6) / (7-8) / (9-10) in each area



The most notable uplift in performance is in having a regularly reviewed strategy in place for the responsible use of AI.

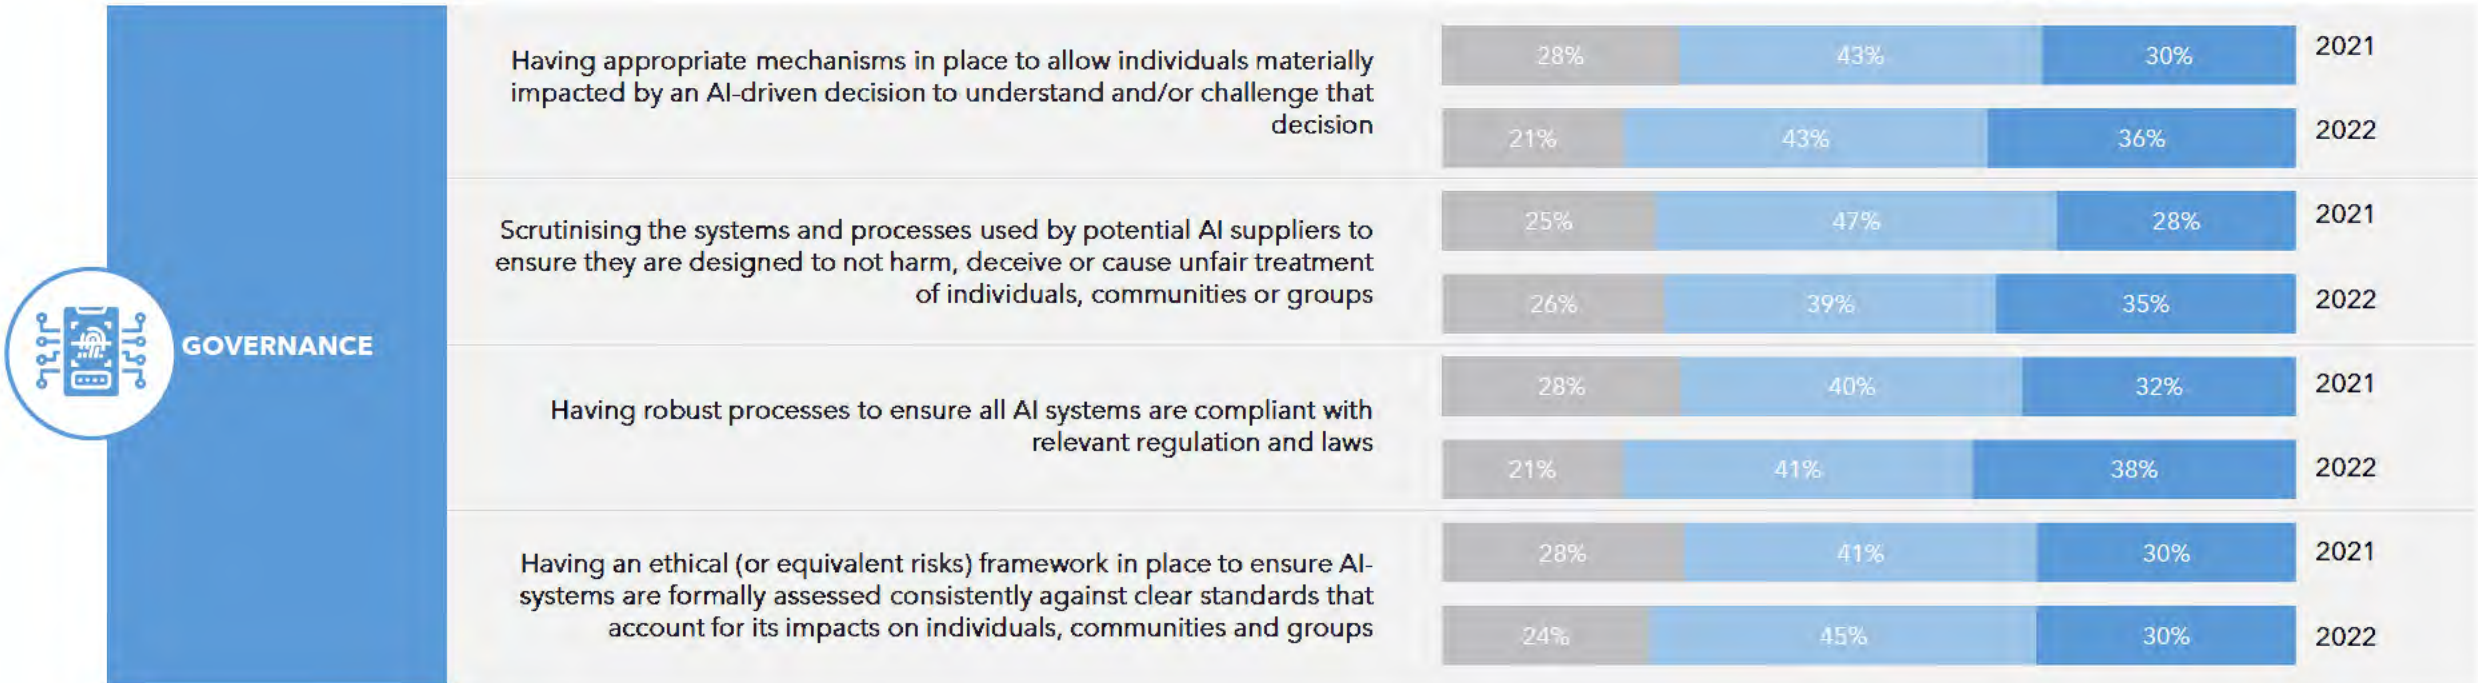
THE RESPONSIBLE AI MATURITY ASSESSMENT FRAMEWORK: PERFORMANCE ASSESSMENT

DISR - For release under the FOI Act

LEX 75414 - FOI - Document 1

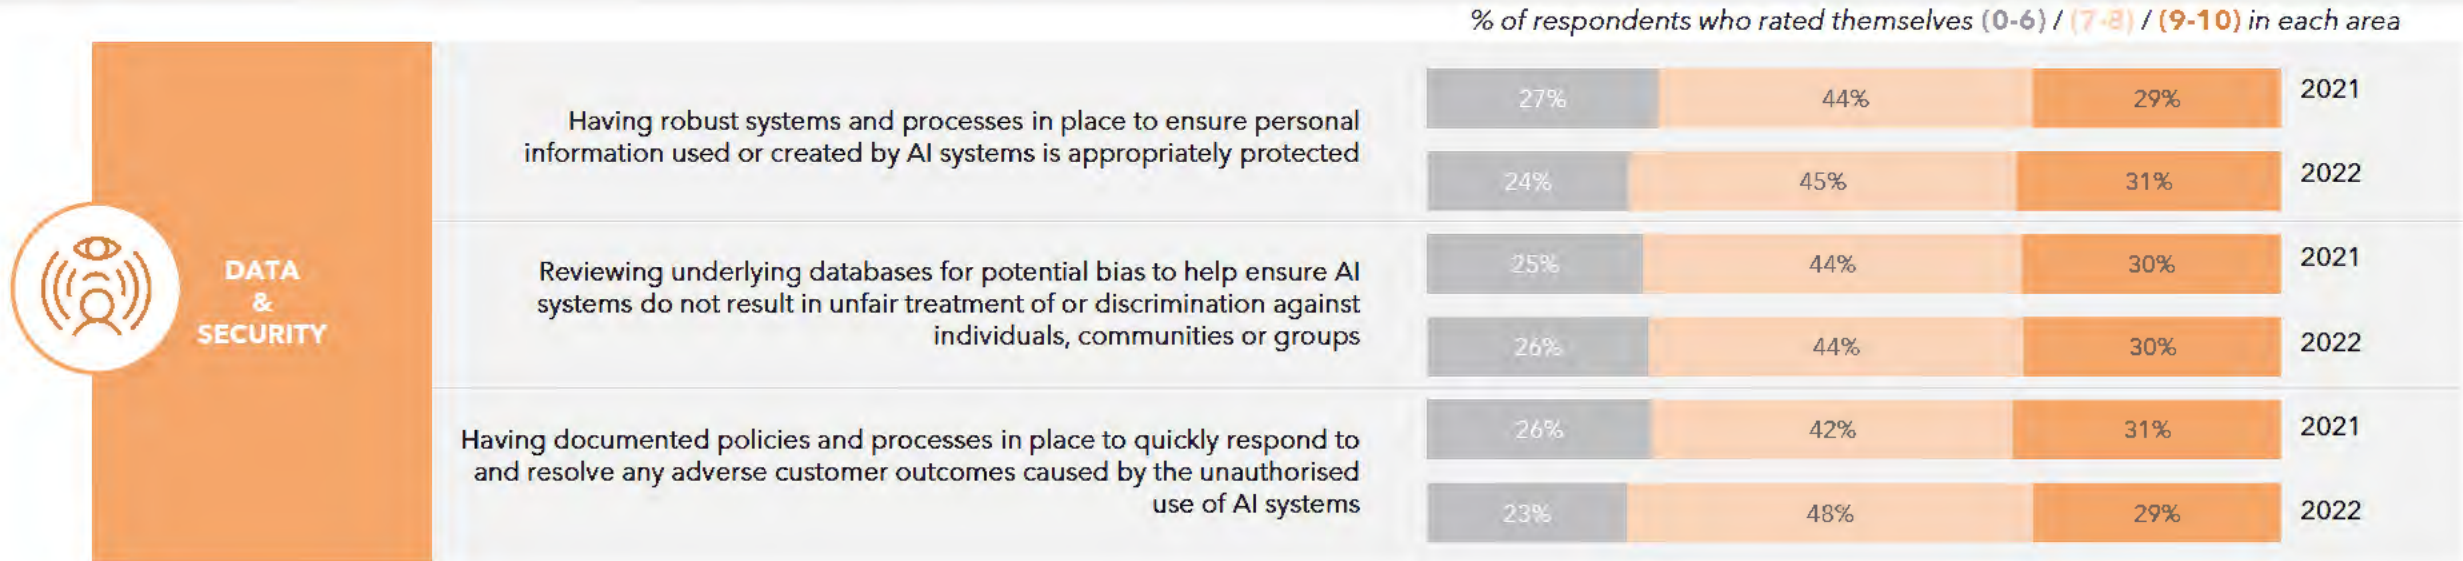
Respondents also rated their organisation's performance on Governance dimensions, showing similar distributions across statements in 2022, again typically higher than 2021 results, except having an ethical framework in place.

% of respondents who rated themselves (0-6) / (7-8) / (9-10) in each area



The most notable uplift in performance is in scrutinising the systems provided by AI vendors to ensure they do not cause harm or unfair treatment.

Respondents rated organisational performance on the Data & Security dimensions in 2022 similar to 2021 results.



Given recent high profile data breaches, it is concerning that around a quarter of organisations give themselves a relatively low score, below seven out of ten, on criteria relating to data and security.

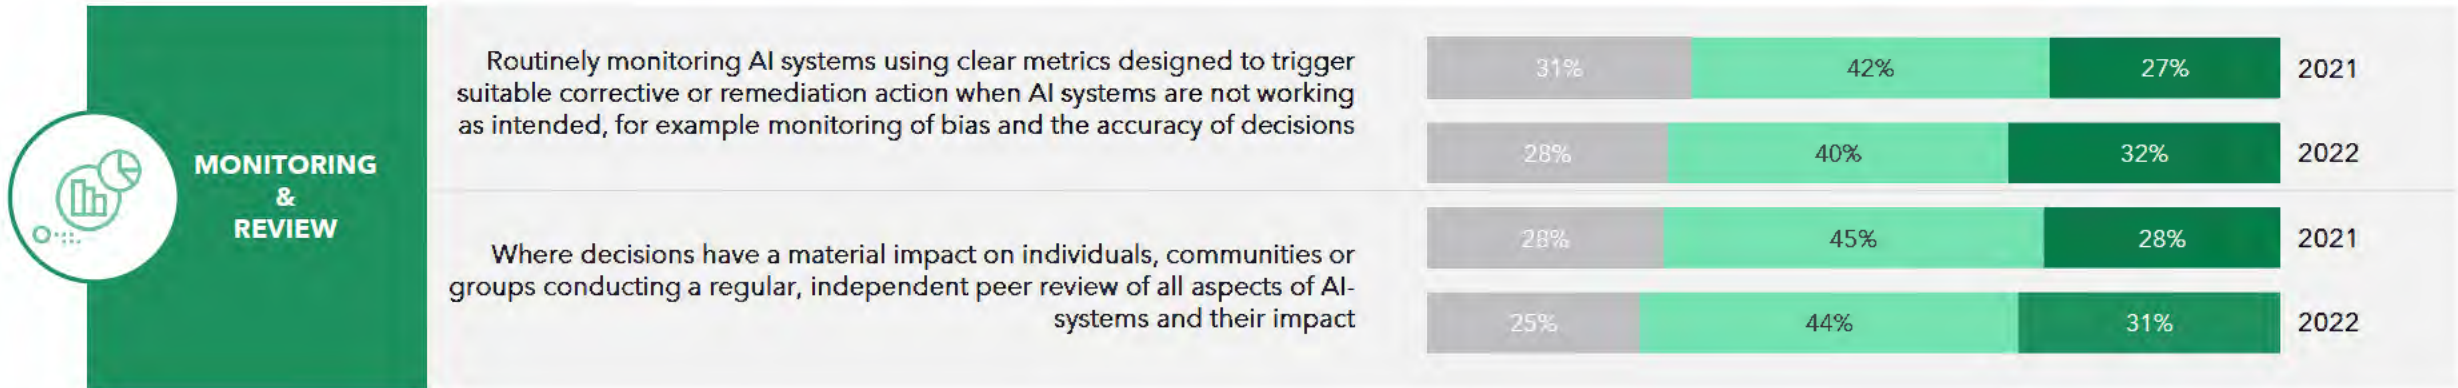
Respondents rated organisational performance on the Data & Security dimensions similar to 2021.

% of respondents who rated themselves (0-6) / (7-8) / (9-10) in each area



This relatively stagnant performance on people and skills may reflect the challenging labour market conditions during the COVID-19 pandemic with restrictions in place on the hire of international talent.

Monitoring & Review dimensions showing positive improvement in scores compared with 2021.



The most notable uplift in performance is in routinely monitoring AI systems using metrics to trigger remedial action when systems are not working as intended.

THE RESPONSIBLE AI MATURITY ASSESSMENT FRAMEWORK: ACTIONS TAKEN

DISR - For release under the FOI Act

LEX 73414 - FOI - Document 1

The behavioural component of the maturity model takes into account a range of practices that can be taken to support the responsible deployment of AI. Concerningly, fewer organisations are employing these practices in 2022, when compared to 2021 results.



This may indicate an appreciation of the challenges involved when developing AI responsibly and points to a need for resources to guide organisations towards frameworks and tools which can help them deploy AI systems responsibly.

The Index identifies four levels of maturity regarding an organisation's approach to Responsible AI.



Planning

- Early stage of AI deployment
- Focused on quickly reaping commercial benefits of AI automation without pausing to factor in ethical implications

Initiating

- Lack confidence to deploy AI
- Lack of knowledge about Responsible AI
- Lack leadership support



Maturing

- Implemented auditing processes for AI
- Strong focus on the moral and ethical implications of using AI technologies
- Uses external specialists and advisors

Developing

- Developed guidelines for responsible use of AI
- Strong culture of data protection and security




RESPONSIBLE AI MATURITY INDEX

DISR - For release under the FOI Act

LEX 73414 - FOI - Document 1

Most organisations sit within the Initiating and Developing groups. In 2022, the strongest shift is from Planning to Initiating, and there is a decline in the size of the Maturing group.

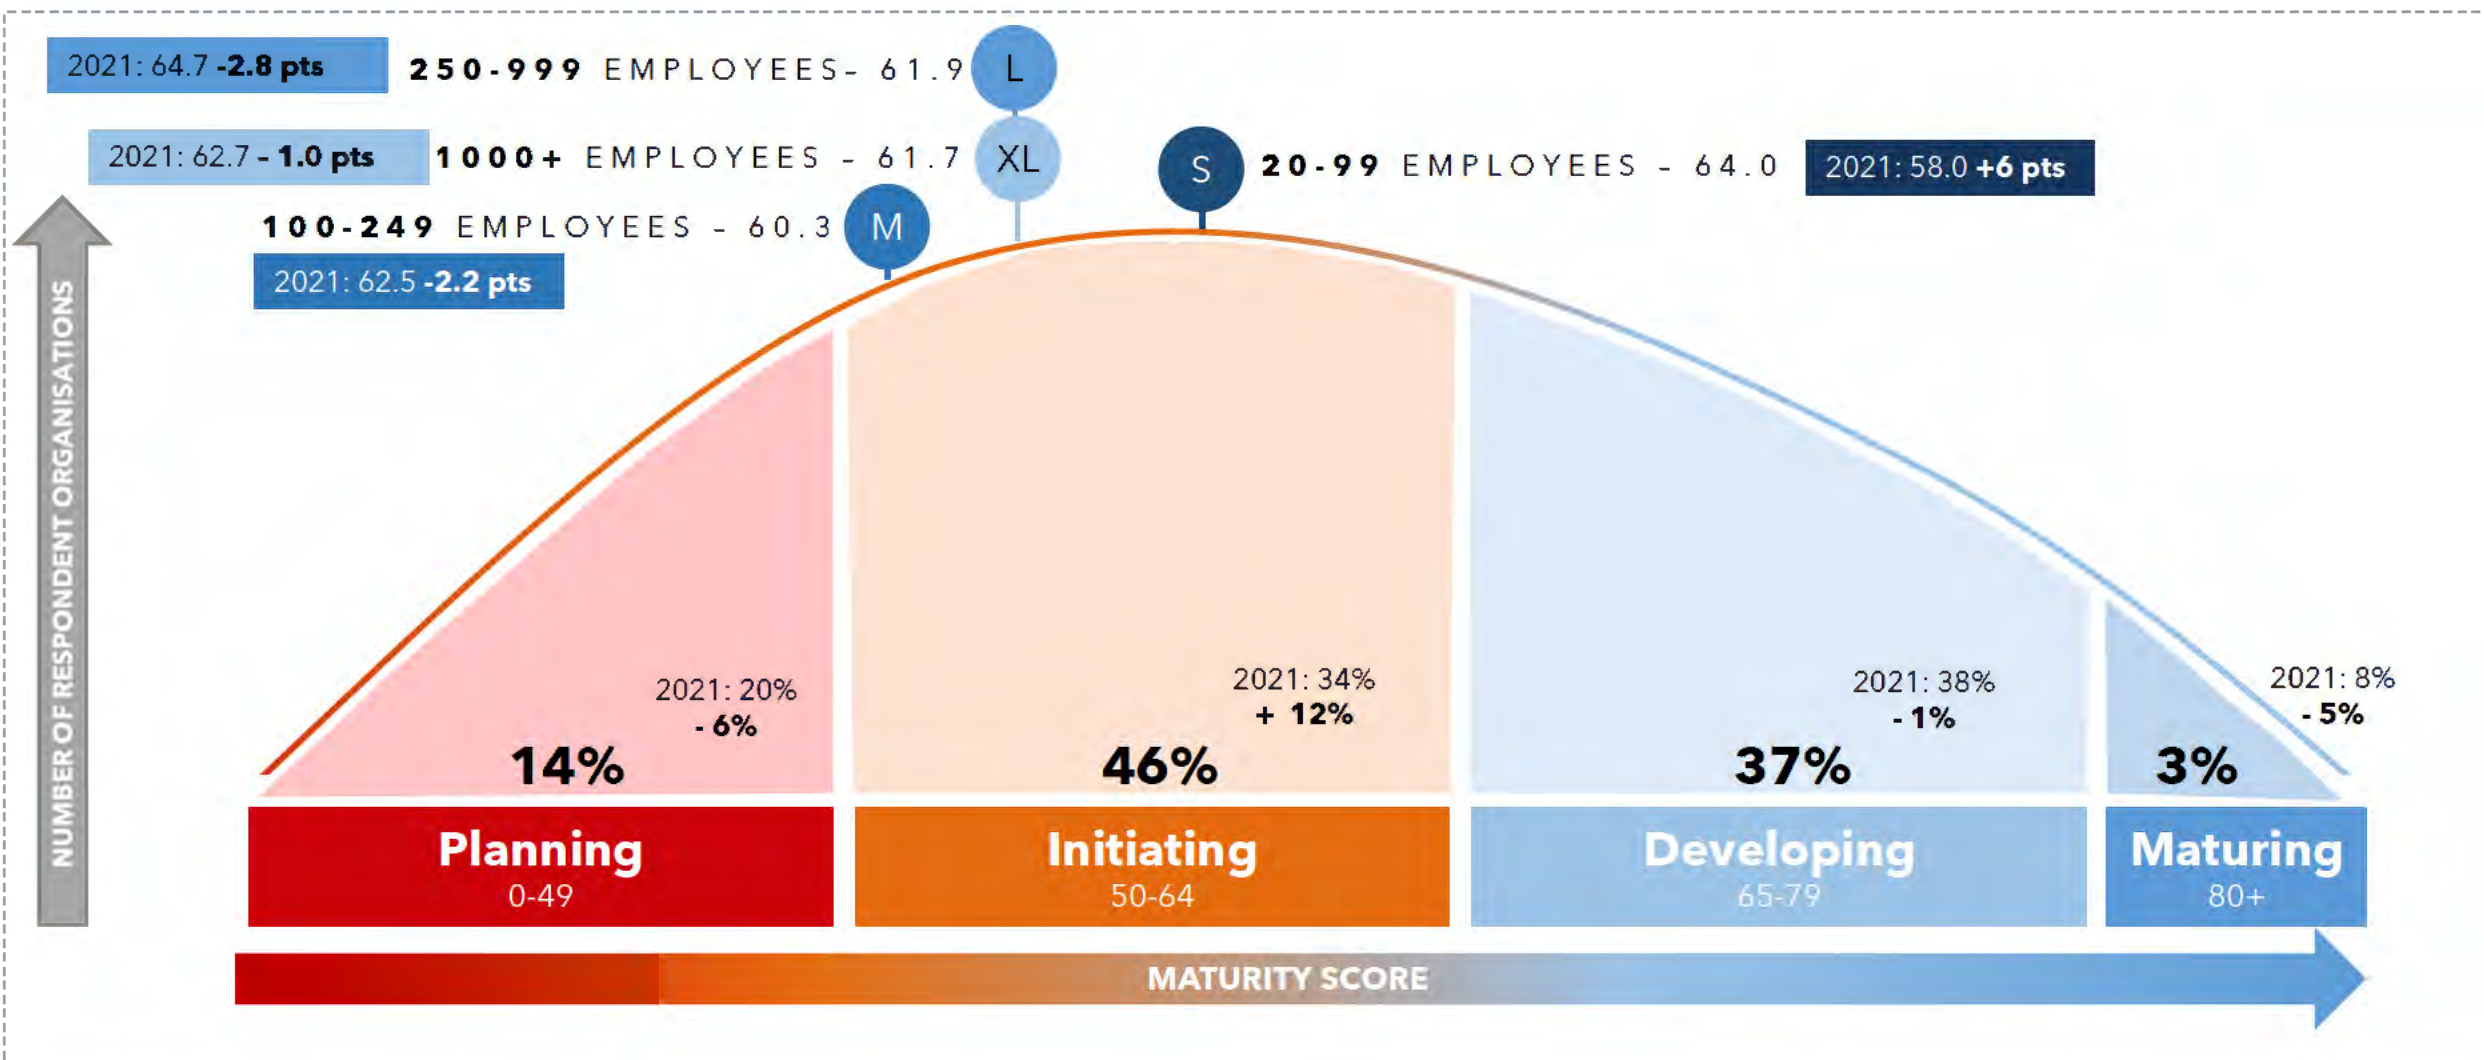


RESPONSIBLE AI MATURITY INDEX BY BUSINESS SIZE

DISR - For release under the FOI Act

LEX 73414 - FOI - Document 1

As with 2021, business size is not a strong indicator of maturity. Smaller companies are taking a more mature approach to Responsible AI, now with the highest maturity score overall.




RESPONSIBLE AI MATURITY INDEX BY INDUSTRY

DISR - For release under the FOI Act

LEX 73414 - FOI - Document 1

The Technology & Telecommunications industry shows a higher level of maturity compared to other industries, whilst Construction companies are the least mature, showing a higher likelihood to be planning their approach to Responsible AI, rather than initiating.

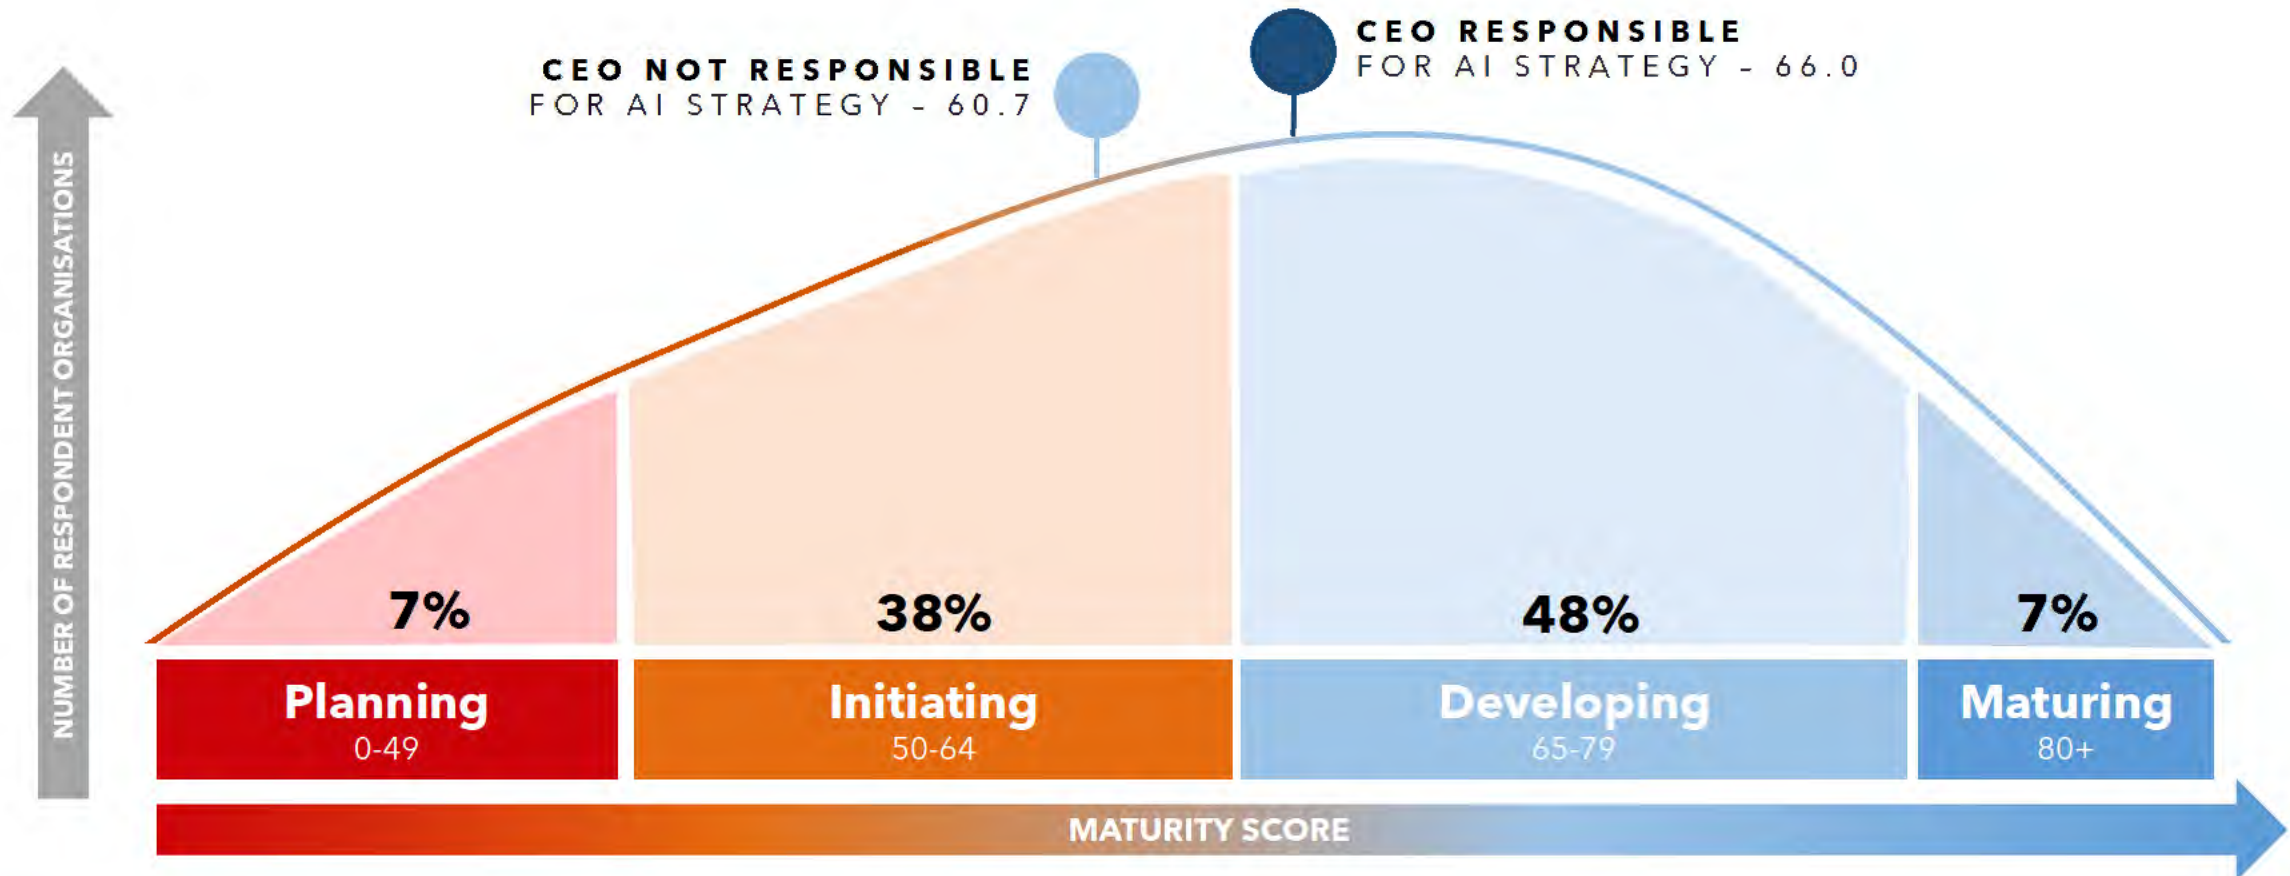


RESPONSIBLE AI MATURITY INDEX BY CEO INVOLVEMENT

DISR - For release under the FOI Act

LEX 73414 - FOI - Document 1

Organisations where the CEO is responsible for driving the AI strategy are more mature than those where the CEO is not taking the lead. More of these organisations are in the Developing and Maturing phase, compared to organisations overall, therefore showing a higher likelihood to be already deploying and using AI.

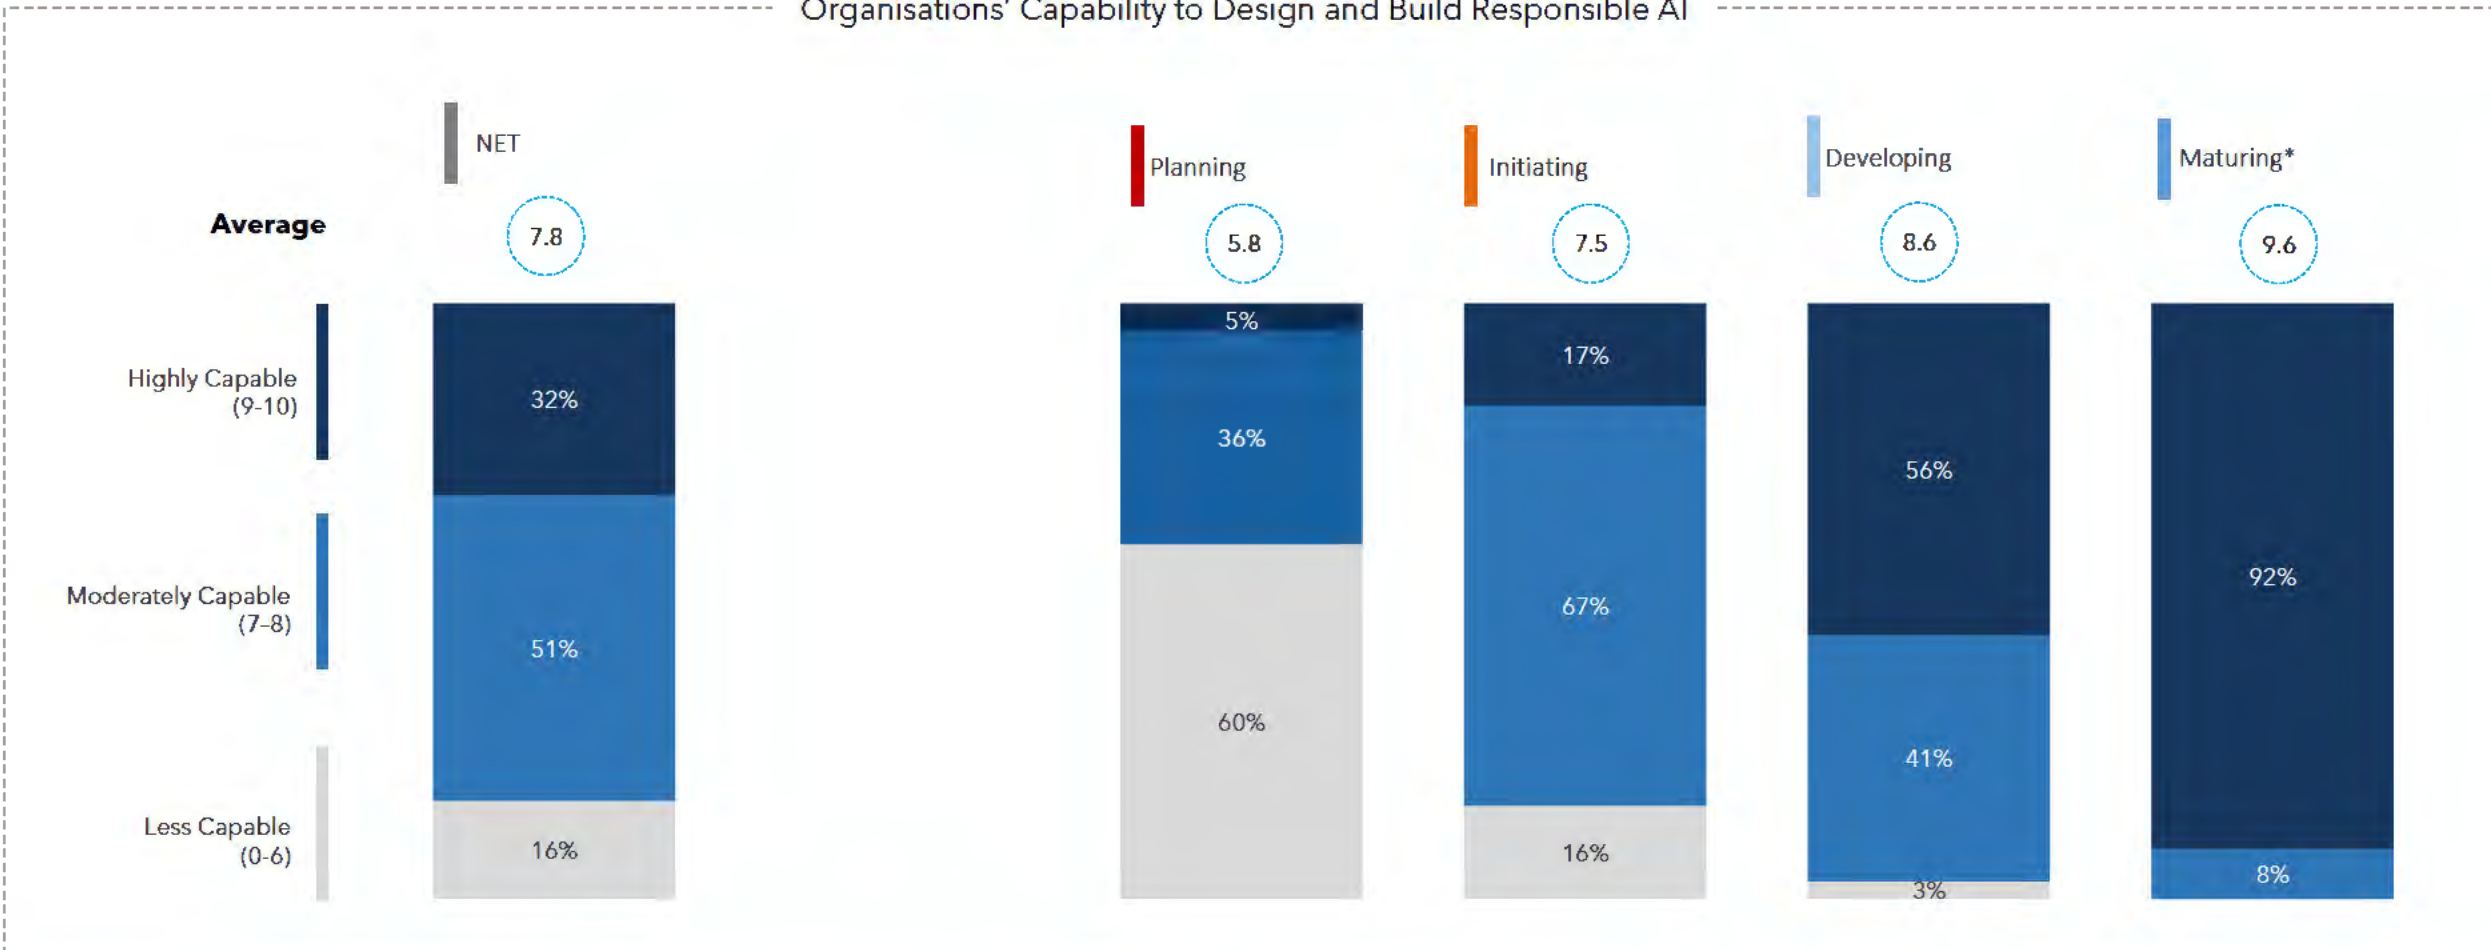


Organisations that have the CEO leading the AI strategy are more likely to invest in developing their culture and governance processes so as to elevate RAI practices to a level of standard routine.

CURRENT CAPABILITY TO BUILD RESPONSIBLE AI

The Developing and Mature segments highly rate their ability to design and build a responsible AI system.

Organisations' Capability to Design and Build Responsible AI



Australian organisations that are Planning to deploy AI recognise there are gaps in their capabilities to do this ethically and responsibly, whilst those Initiating are less unsure, but are still not completely confident. Those in the Developing and Maturing segments show more confidence in their capabilities.

CONTENTS

1. BACKGROUND
2. THE RESPONSIBLE AI INDEX
- 3. THE STATE OF AI IN AUSTRALIA**
4. AI ETHICS PRINCIPLES
5. PRINCIPLES IN PRACTICE
6. BRIDGING THE GAP




USE OF AI

DISR - For release under the FOI Act

LEX 75414 - FOI - Document 1

Overall use of AI in organisations is in line with 2021, but its now less limited to parts of the business, with more organisations using it broadly. This is especially true amongst the Developing cohort.

Organisation's use of AI



The Initiating cohort is becoming more confident in its ability to deploy AI responsibly before the technology becomes more widely deployed across the business.

ORGANISATIONAL STRATEGY FOR AI

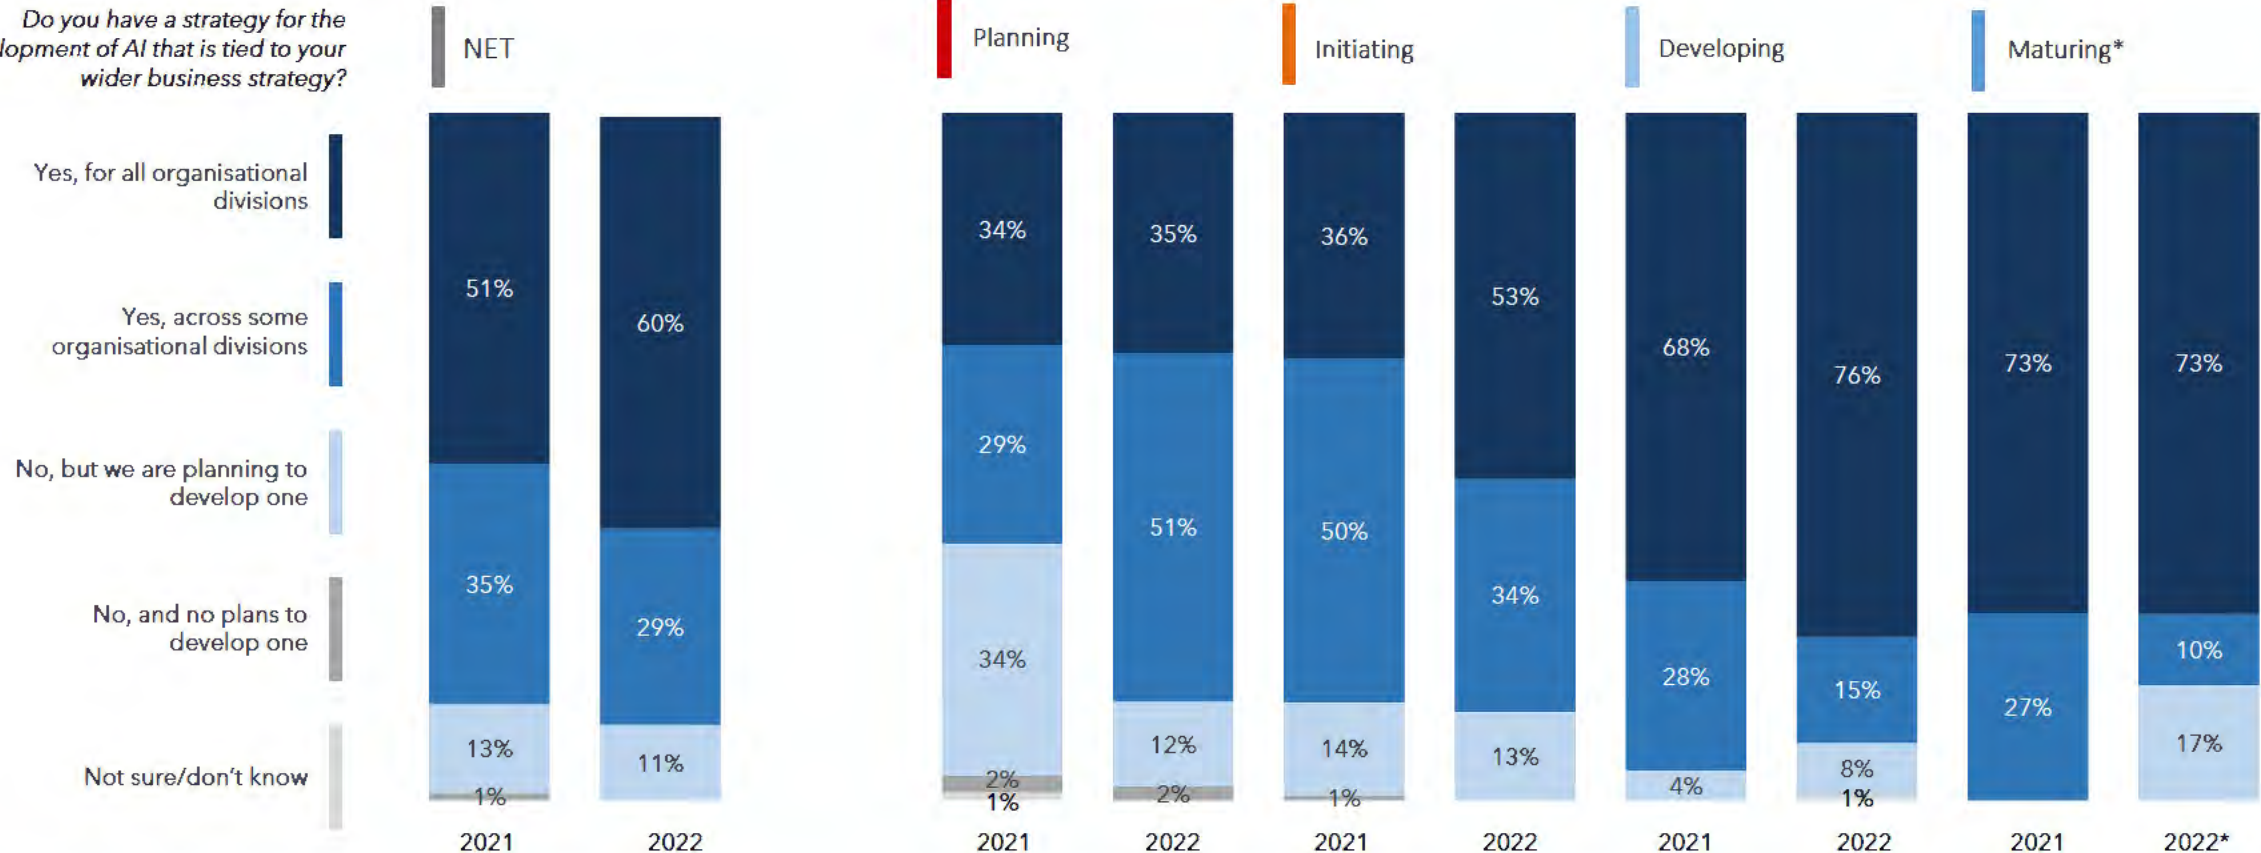
DISR - For release under the FOI Act

LEX 75414 - FOI - Document 1

6 out of 10 organisations now have an enterprise-wide AI strategy, with the Planning cohort shifting from recognition of need to developing AI strategies for specific business functions, whilst the Initiating cohort is now taking a more holistic business-wide view.

Strategy for AI development

Do you have a strategy for the development of AI that is tied to your wider business strategy?



The Initiating cohort need to continue the transition from opportunistic and tactical AI decision-making to a more strategic orientation.


KEY ROLES FOR DRIVING AI STRATEGY

Chief Information Officers remain the key figure in an organisation responsible for driving the AI strategy, but businesses with their CEO driving the AI strategy are much more likely to have an enterprise-wide AI strategy and as indicated previously, score higher on the RAI Index.


Who in your organisation is responsible for driving the organisation's AI strategy?



Organisations that have an enterprise-wide AI strategy tend to have a CEO personally invested in driving AI strategy



of organisations that have AI strategy tied to all divisions say their CEO drives AI strategy

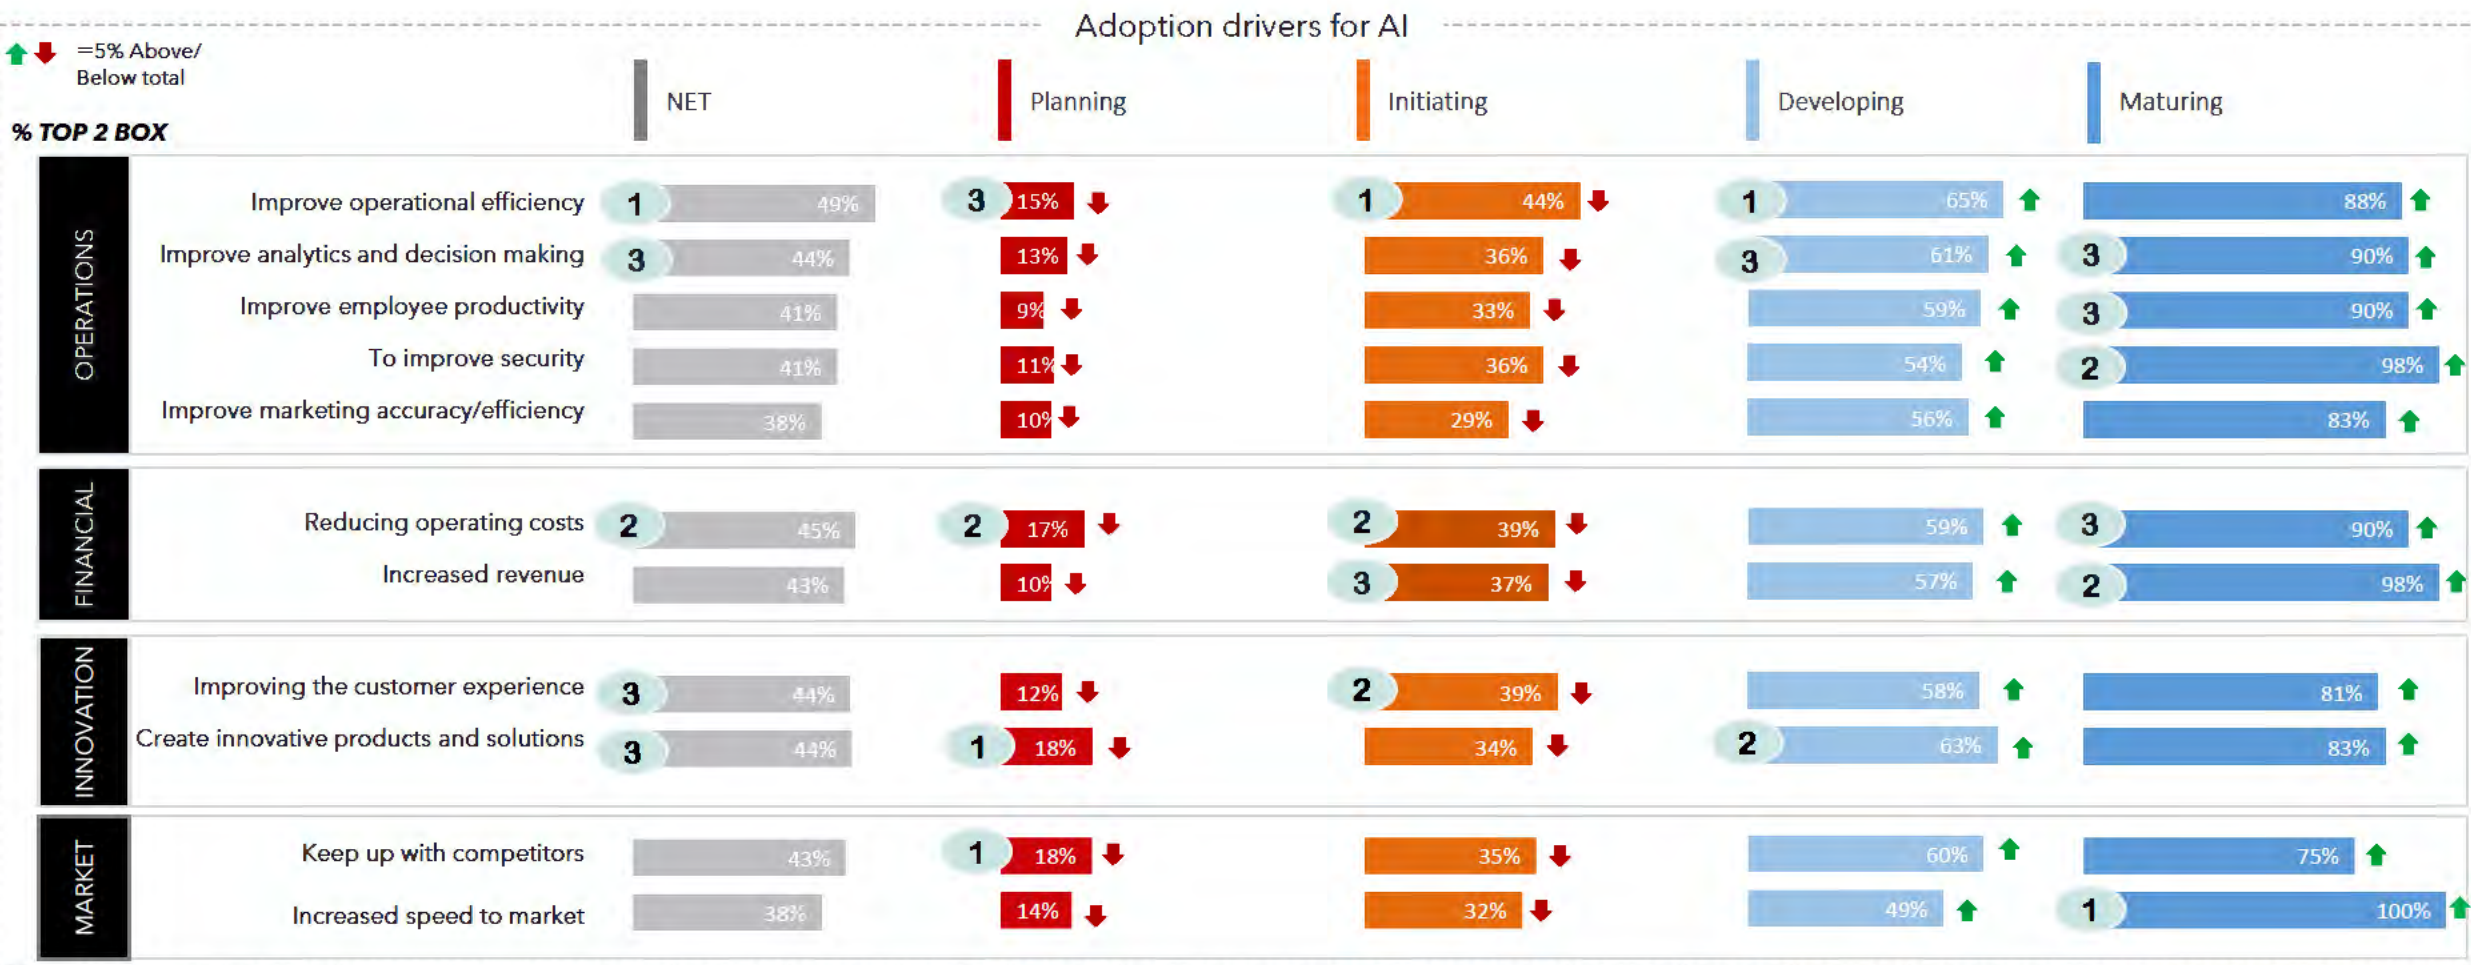


of organisations that have AI strategy tied to some divisions say their CEO drives AI strategy

Having a CEO driving AI strategy ensures accountability and a more strategic AI orientation across the business.

KEY ADOPTION DRIVERS FOR AI

Organisations are investing in AI in order to improve operational efficiency and reduce costs, with the more mature groups seeking to improve multiple facets of their whole business, including increased security



The AI use cases for the less mature segments may be more limited, though they may also not fully understand the full benefits of AI to their business.

OUTCOMES OF AI

DISR - For release under the FOI Act

LEX 73414 - FOI - Document 1


Organisations are discovering multiple benefits of AI, especially in improving security, enabling product and service innovation to help them keep pace with their competitors.



The achievement of security outcomes is important in the context of recent data breaches. However, those in the planning stage do not see this as an outcome for the business which reflects their inexperience. This group may benefit from case studies about how a responsible approach to AI can improve data security and privacy.

USE OF AI ACROSS BUSINESS AREAS

AI is most commonly used in IT, security and analytics with both Developing and Maturing segments significantly more likely to be using AI to support all areas of their organisation than less mature groups.



The main applications for in AI are in business functions which have processes that can be easily automated and scaled, such as analytics. AI is used less frequently in other areas which require more human interactions, such as CRM, HR and sales.

AI USE CASES FOR CUSTOMERS

DISR - For release under the FOI Act

LEX 73414 - FOI - Document 1


Organisations are looking to deploy AI to provide year-round, 24/7 service to customers; less mature groups limited AI use across customer touchpoints, focussing more on supporting customers interactions with digital channels.



The less mature groups see fewer opportunities to use AI to engage with customers.

AI PROJECTS MAKING IT TO PRODUCTION

Around half of AI projects succeed and make it into production, with success rates typically improving as an organisation matures.



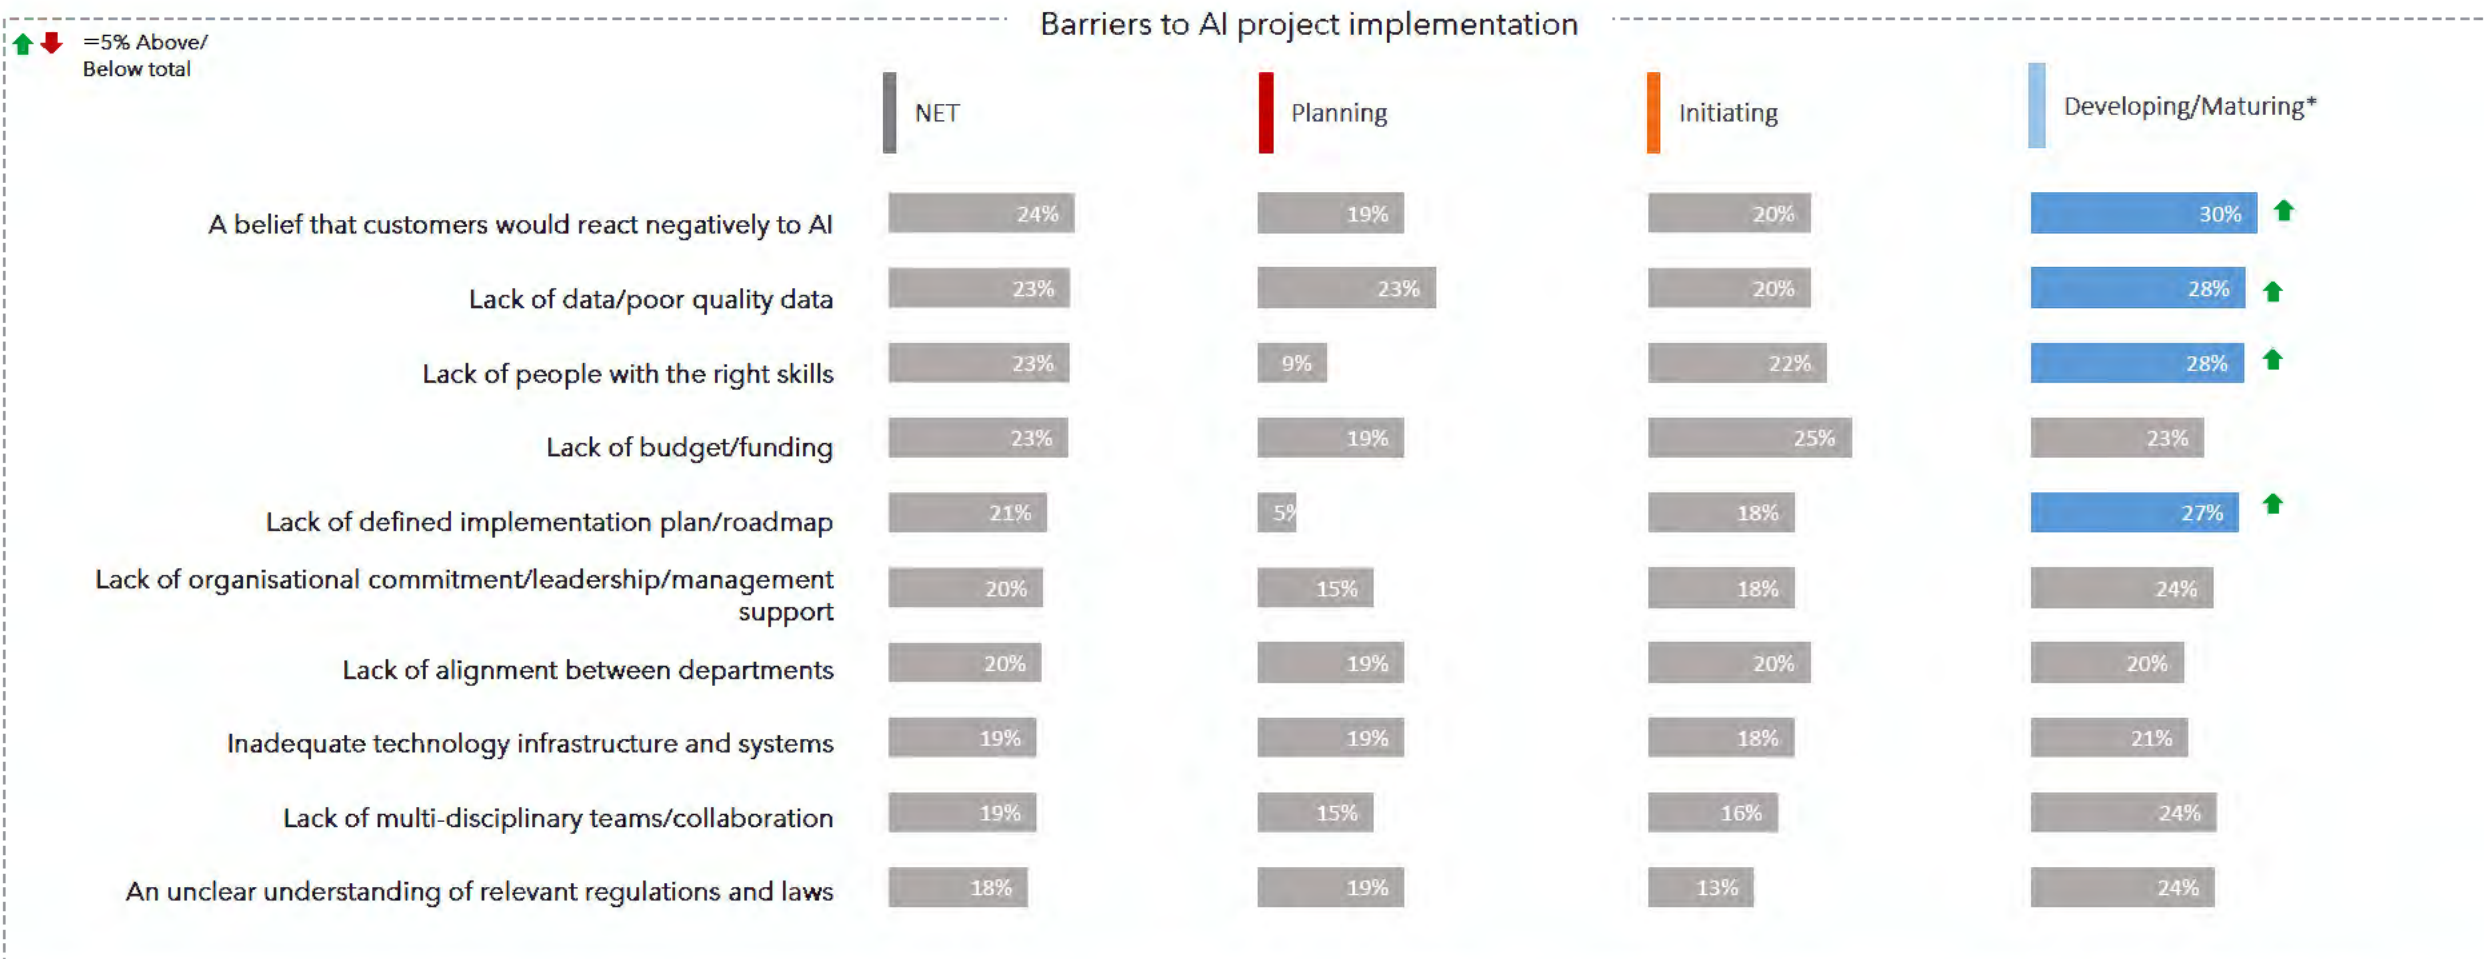
Practices to test AI systems for reliability and safety in a controlled environment using accurate and representative data should help to mitigate any unintended, negative impacts of projects before they make it into production.

MAIN REASONS WHY AI PROJECTS DO NOT MAKE IT TO PRODUCTION

DISR - For release under the FOI Act

LEX 75414 - FOI - Document 1

AI projects do not make it into production due to a combination factors including data availability and quality, funding, skill gaps and technology infrastructure requirements.




Organisations should consider factors such as data quality and quantity, and infrastructure requirements when piloting and testing AI systems, using resources that provide guidance on how to transition from pilot studies to a production scale deployment.

OTHER REASONS WHY AI PROJECTS DO NOT MAKE IT TO PRODUCTION

DISR - For release under the FOI Act

LEX 75414 - FOI - Document 1


As an organisation's approach to responsible AI matures, barriers around the co-ordination of strategy and initiatives emerge.



Those in the planning phase find transparency and explainability to be challenging and would benefit from using tools and guidelines to overcome these barriers.

CONTENTS

1. BACKGROUND
2. THE RESPONSIBLE AI INDEX
3. THE STATE OF AI IN AUSTRALIA
- 4. AI ETHICS PRINCIPLES**
5. PRINCIPLES IN PRACTICE
6. BRIDGING THE GAP



The elements of the Department of Industry's AI Ethics Principles were incorporated into the questionnaire, and are examined in more detail in this section of the report to identify the gap between attitudes towards responsible AI and the steps that organisations are taking to implement AI responsibly.



HUMAN PRINCIPLES

HUMAN, SOCIAL & ENVIRONMENTAL WELLBEING

Throughout their lifecycle, AI systems should benefit individuals, society and the environment

- AI system objectives should be clearly identified and justified
- AI systems should be used to benefit all human beings, including future generations
- Positive and negative impacts should be accounted for throughout the lifecycle of all legitimate internal business process AI systems

HUMAN-CENTERED VALUES

Throughout their lifecycle, AI systems should be inclusive and accessible, and should not involve or result in unfair discrimination against individuals, communities or groups

- AI systems need to be aligned with human values and enable an equitable and democratic society
- Must respect, protect and promote human rights
- Should be designed to augment, complement and empower human cognitive, social and cultural skills

FAIRNESS

Throughout their lifecycle, AI systems should respect human rights, diversity, and the autonomy of individuals

- AI systems need to be fair and enable inclusion throughout their lifecycle
- Should be user-centric, designed to allow all people to interact with it
- Measures should be taken to ensure AI produced decisions are compliant with anti-discrimination laws



PRIVACY AND RELIABILITY

PRIVACY PROTECTION & SECURITY

Throughout their lifecycle, AI systems should respect and uphold privacy rights and data protection, and ensure the security of data

- Ensuring respect for privacy and data protection, including proper data governance and management
- Also ensures appropriate data and AI system security measures are in place, including the identification of potential security vulnerabilities and assurance of resilience to adversarial attacks

RELIABILITY & SAFETY

Throughout their lifecycle, AI systems should reliably operate in accordance with their intended purpose

- Ensures AI systems are reliable, accurate and reproducible
- AI systems should adopt safety measures that are proportionate to the magnitude of potential risks
- Responsibility should be clearly and appropriately identified, for ensuring that an AI system is robust and safe



INTEGRITY

TRANSPARENCY & EXPLAINABILITY

There should be transparency and responsible disclosure to ensure people know when they are being significantly impacted by an AI system, and can find out when an AI system is engaging with them

- Transparency through responsible disclosure when an AI system is significantly impacting on a person's life
- Information provided in a timely manner, with reasonable justifications for the AI systems outcomes
- Aims to ensure people have the ability to find out when an AI system is engaging with them

CONTESTABILITY

When an AI system significantly impacts a person, community, group or environment, there should be a timely process to allow people to challenge the use or output of the AI system

- Knowing that redress for harm is possible, when things go wrong, is key to ensuring public trust in AI
- Needs to be sufficient access to the information available to the algorithm and inferences drawn, to make contestability effective


ACCOUNTABILITY

Those responsible for the different phases of the AI system lifecycle should be identifiable and accountable for the outcomes of the AI systems, and human oversight of AI systems should be enabled

- Organisations/ individuals should be identifiable and ensure responsibility for AI systems and their outcomes both before and after their design, development, deployment and operation
- They must consider the appropriate level of human control or oversight for the particular AI system or use case


Awareness of the Department of Industry AI Ethics Principles is higher in 2022 across all maturity segments.

Awareness of Australian Federal Department of Industry's AI principles



There is a continued opportunity to increase awareness of the AI Ethics Principles amongst the less mature groups to further increase industry engagement on the benefits of a responsible approach to AI. This should be supported by guidance on how to practically implement the eight Principles.

Even though most organisations claim to have formal AI standards in place, these may not be across all functions where AI is used.



As organisations become more mature in their development and use of AI, they learn how to develop and apply standards across all business functions where AI is used.

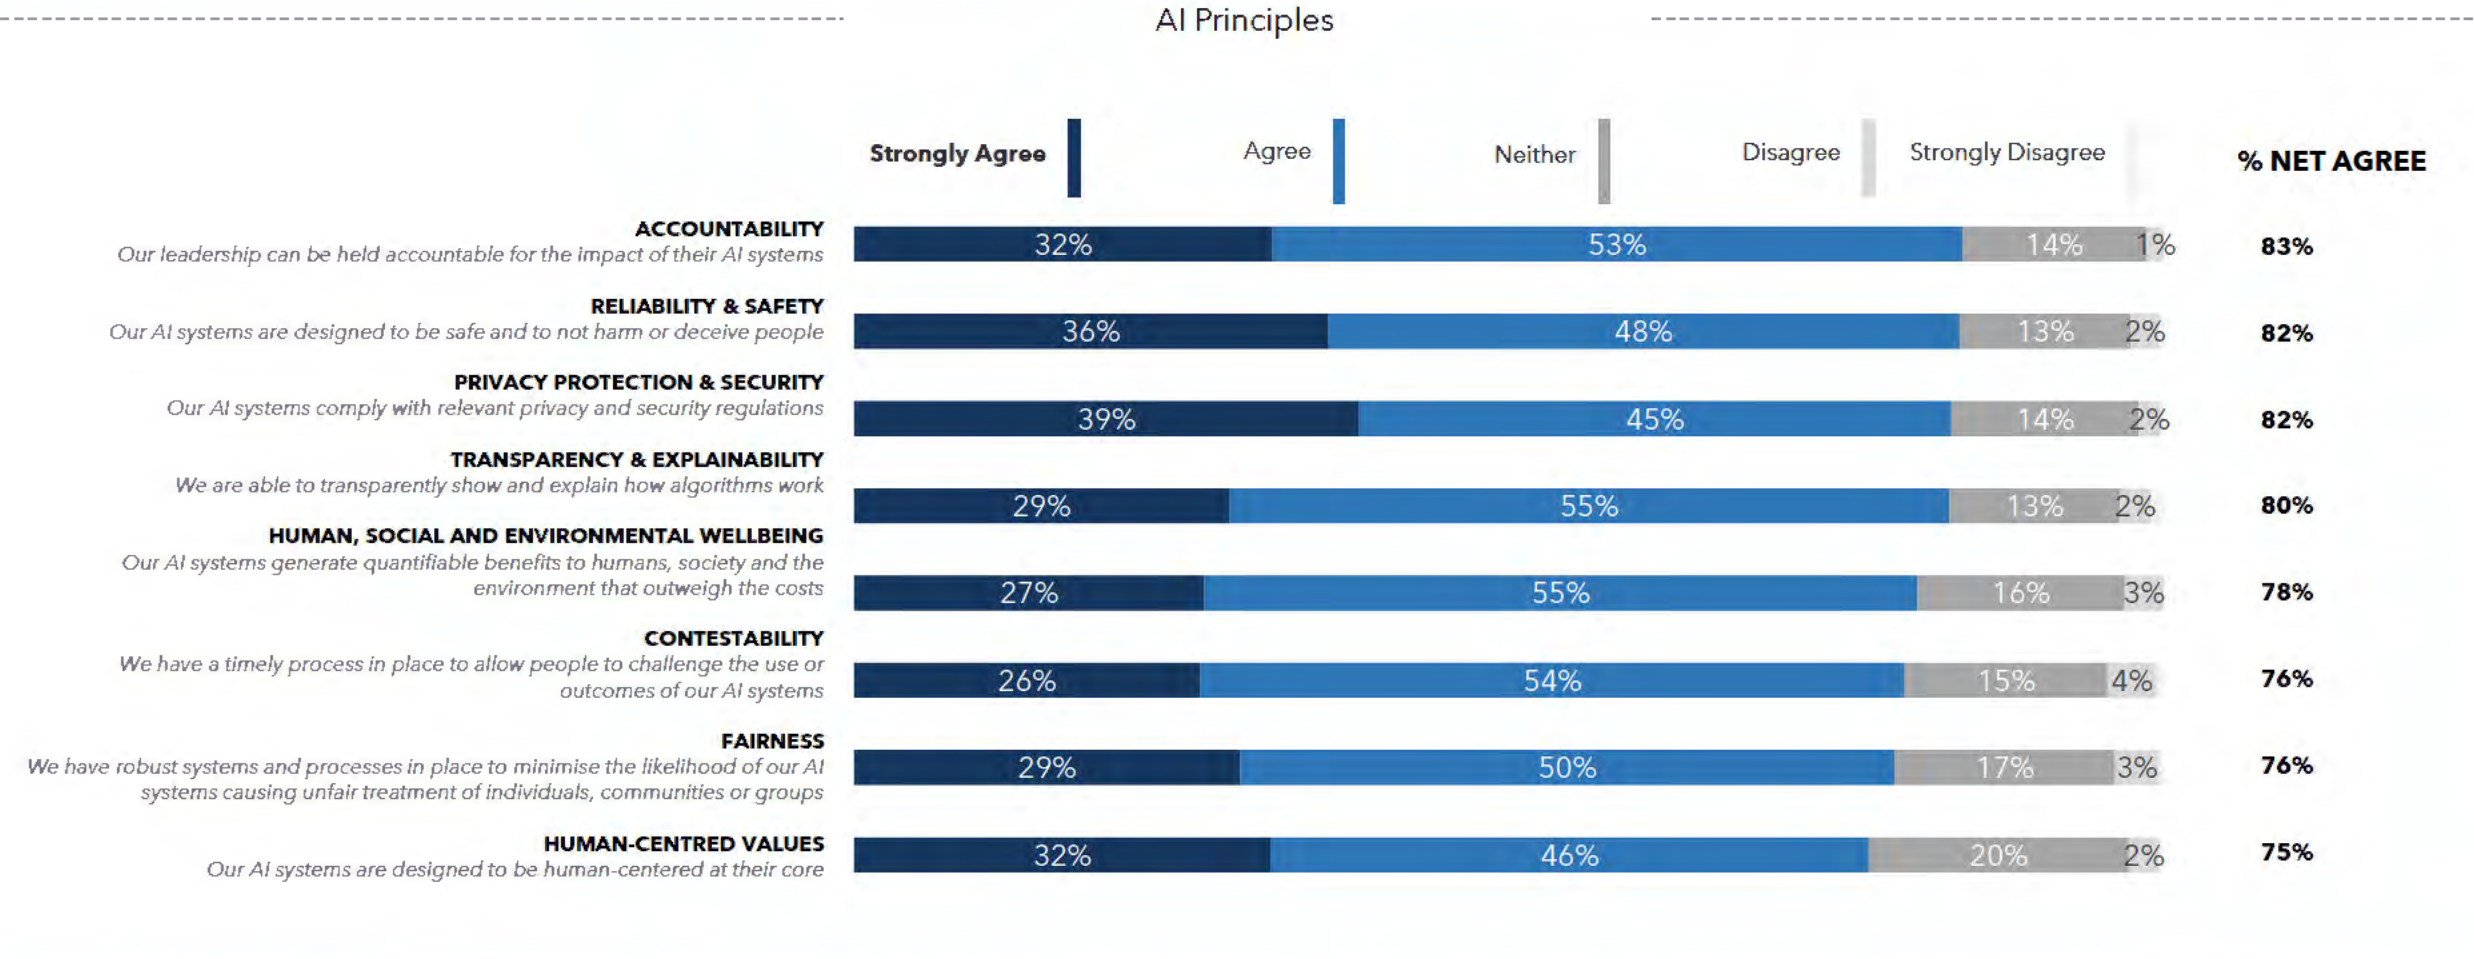
CONTENTS

1. BACKGROUND
2. THE RESPONSIBLE AI INDEX
3. THE STATE OF AI IN AUSTRALIA
4. AI ETHICS PRINCIPLES
- 5. PRINCIPLES IN PRACTICE**
6. BRIDGING THE GAP



AI PRINCIPLES


Most respondents agree that their organisation is broadly following the stated intent of the Australian AI Ethics Principles, most notably for Privacy and Protection, Reliability and Safety, and Accountability. Agreement is lowest for Human-Centred Values, Fairness and Contestability.



This level of agreement is encouraging but does not align with the overall Responsible AI Index scores, which may indicate a gap between strategic intent and the actions taken by organisations to put the Australian AI Ethics Principles into practice.

HUMAN PRINCIPLES

Degree of importance placed on human principles and fairness, and the tangible actions taken to address these, increases with maturity. Developing and Maturing segments have undertaken the most substantive steps to help reduce bias and risk.




Most organisations, especially the less mature, are not taking the necessary actions to elicit and assess potential impacts of AI systems, incorporate diversity, and measure and improve system fairness.

PRIVACY AND RELIABILITY

DISR - For release under the FOI Act


LEX 75414 - FOI - Document 1

More mature cohorts are placing higher importance on system reliability, privacy and safety than less mature organisations, showing a higher likelihood to have taken tangible actions.



There is a gap between most organisations' strategic intent and the actions undertaken to protect systems against attacks, and monitoring systems to ensure they operate safely and reliably.

Most organisations have not taken any actions to ensure transparency and explainability, contestability and accountability, even though these are deemed to be important considerations. Encouragingly, some those in the Planning phase are taking practical steps to hold their leadership to account.



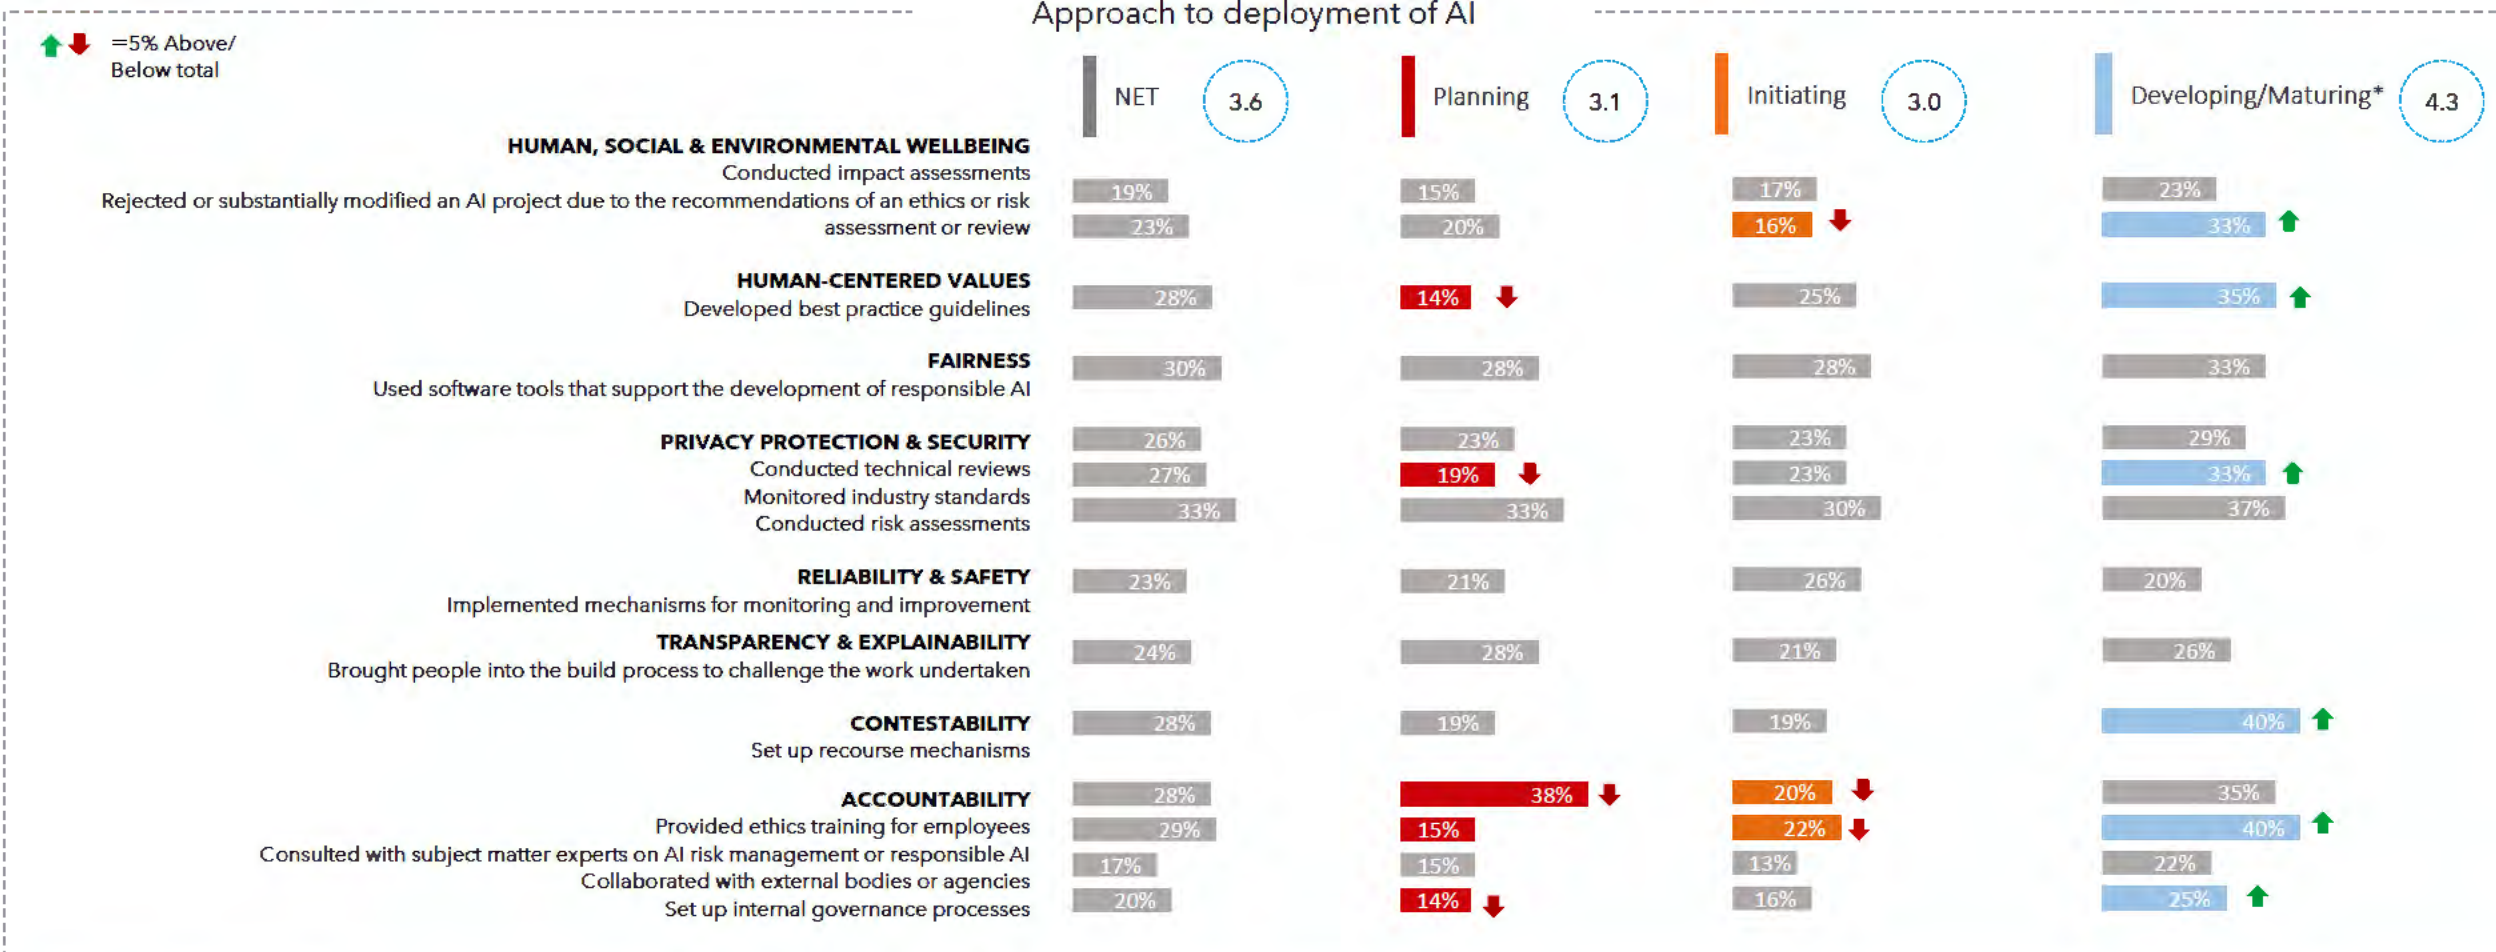
This indicates that organisations need support to understand how to document design decisions, explain how models operate and make decisions, establish recourse mechanisms and implement accountability practices.

RESPONSIBLE AI PRACTICES UNDERTAKEN BY AI USERS

DISR - For release under the FOI Act

LEX 75414 - FOI - Document 1


Among organisations who have deployed AI, most have only implemented a limited number of responsible AI practices. Even the more mature segments have only implemented an average of 4.3 practices out of a total of 15 that were shown to respondents.



This may indicate lack of understanding about the potential risks and benefits of AI, a lack of resources or expertise to implement responsible AI practices, and a lack of incentives to do so.

At an organisational level, there is a level of concern about a range of impacts, including the reputational risk of negative customer feedback and damage to the brand.


Concerns surrounding the organisational impacts of AI



Concern is lowest about potential bias in decision making from an organisation's AI systems, which, consistent with other data, suggests that principles and practices relating to human-centred values are less of a priority.

CONCERNS ABOUT ORGANISATIONAL IMPACTS


The Maturing and Developing cohorts remain more concerned than the Planning and Initiating about a range of potential negative organisational impacts of AI systems.



This suggests that as an organisation's approach to responsible AI matures, the likelihood of adverse events crystallising becomes more apparent, necessitating the adoption of risk mitigation practices.

There are also concerns at a societal level about the potential impacts of AI systems, including negative outcomes for individuals, groups and communities.


Concerns surrounding the organisational impacts of AI



There is less concern about the potential impact of reduced business competition and unequal access for different segments of Australian society.

CONCERNS ABOUT SOCIETAL IMPACTS

The more mature cohorts are significantly more concerned than the Planning and Initiating about a range of potential negative impacts of AI systems on society.



Again, this points to the need for strategies and practices to be implemented to reduce the impact and probability of these risks materialising.

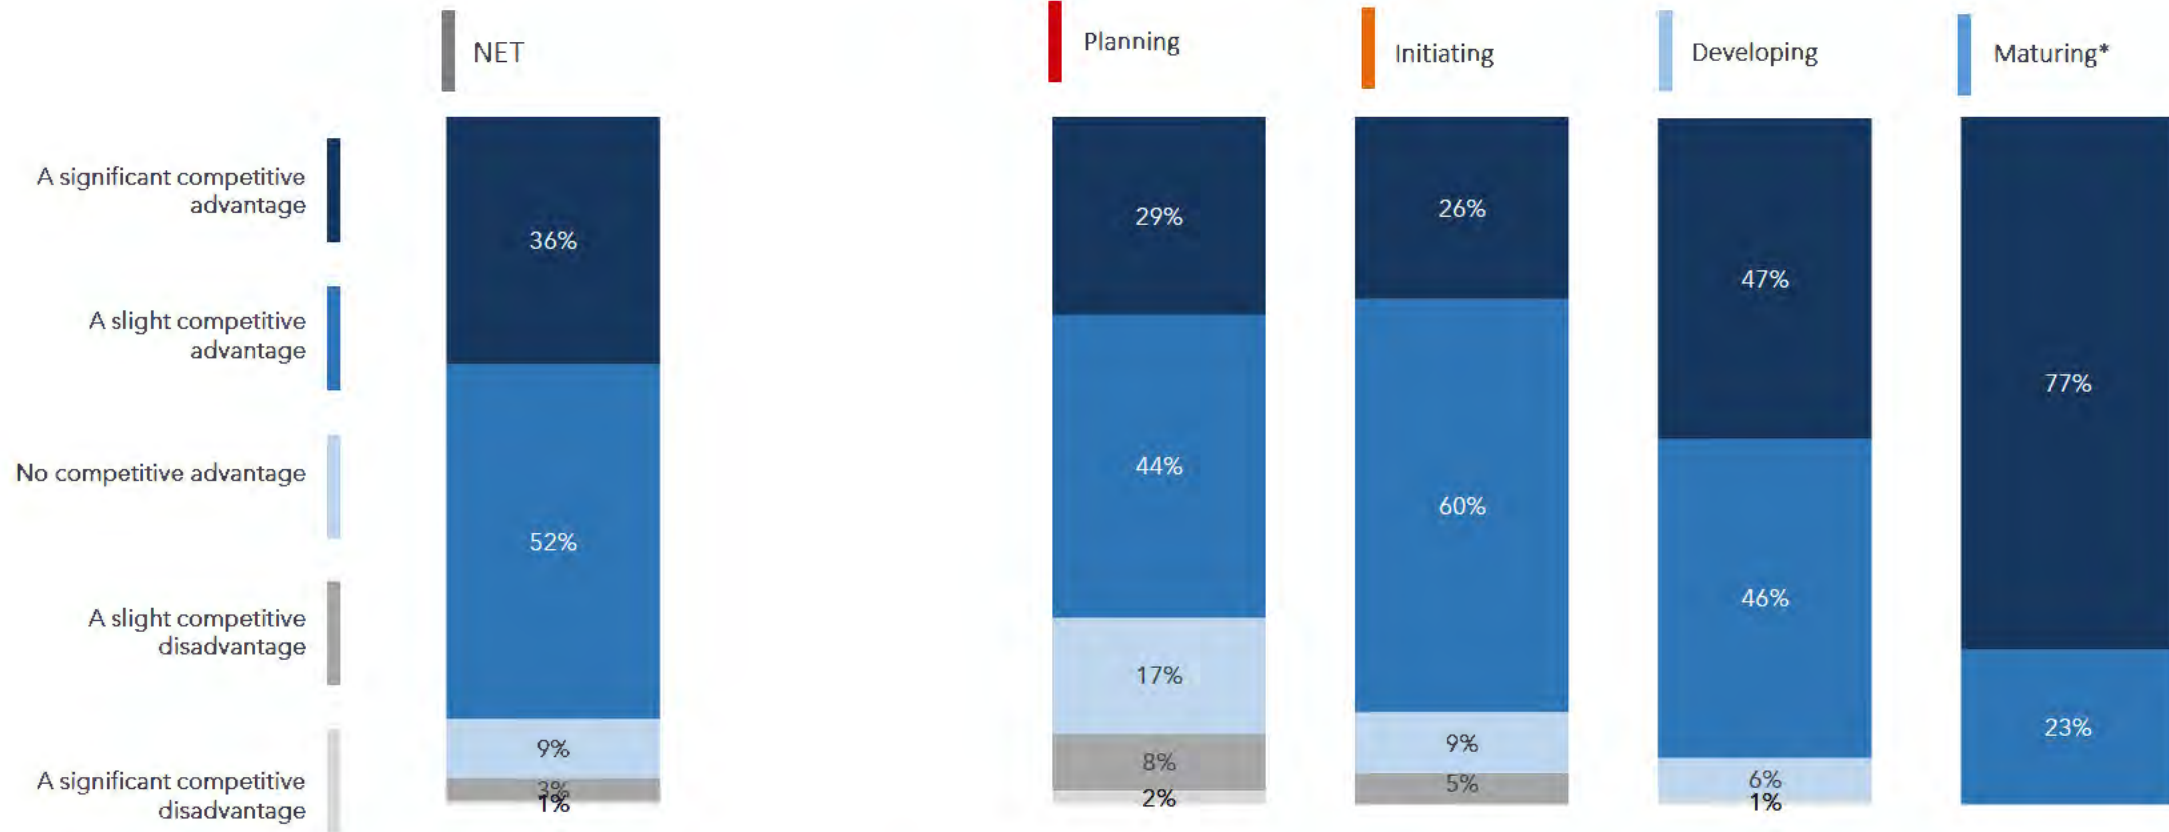
COMPETITIVE ADVANTAGE FROM RESPONSIBLE AI

DISR - For release under the FOI Act

LEX 75414 - FOI - Document 1

As organisations move from the Initiating to Developing phase of maturity, they are more likely to gain a significant competitive advantage through taking a responsible approach to AI.

Competitive Advantage from Responsible AI



While organisations may be tempted to make some ethical sacrifices in order to expedite their AI projects and keep pace with competitors, the evidence shows that those that invest in a responsible AI approach believe this has provided a significant competitive advantage.


OVERALL COST-BENEFIT OF RESPONSIBLE AI

DISR - For release under the FOI Act

LEX 75414 - FOI - Document 1

Overall, more organisations now believe the benefits of taking a responsible approach to AI outweigh the costs, with this change occurring across all maturity cohorts, especially those in Initiating.


Costs and Benefits of Responsible AI



Organisations may need resources and support to quantify the benefits of designing and building responsible AI systems in order to build the business case and obtain leadership support.

CONTENTS

1. BACKGROUND
2. THE RESPONSIBLE AI INDEX
3. THE STATE OF AI IN AUSTRALIA
4. AI ETHICS PRINCIPLES
5. PRINCIPLES IN PRACTICE
- 6. BRIDGING THE GAP**



SUMMARY OF STATE OF RESPONSIBLE AI MATURITY

TOPIC	OBSERVATION	IMPLICATION
Responsible AI Maturity	At an overall level, there has been little change since 2021 in the overall performance of Australian organisations in developing and implementing Responsible AI systems. Performance is higher for those with the CEO leading the AI strategy.	The development of Responsible AI systems represents a significant organisational challenge and requires leadership commitment to develop appropriate culture and governance processes.
AI Strategy	Compared with 2021, more organisations are taking an enterprise-wide approach for the development of AI which is tied to the wider business strategy across all divisions.	This helps with the setting of strategic goals and coordinating individuals within the organisation to promote Responsible AI in terms of governance, policy and incentives.
Principles	Encouragingly, awareness of Australia's AI Ethics principles has increased since 2021. There is also a high level of agreement with statements about how organisations have developed AI systems consistent with the intent of each principle	There is a continued opportunity to increase awareness of the AI Ethics Principles amongst the less mature groups to further increase industry engagement on the benefits of a responsible approach to AI.
Responsible AI Practices	There is a significant gap between perceptions of how AI systems have been designed and how they perform, and the actions that have been taken to ensure AI systems are developed responsibly.	Organisations need practical help and support to implement AI responsibly, including clear signposting to the types of resources which are available, so that an approach can be chosen which works best for the organisation.
Benefits of RAI	Organisations that are more mature in their deployment of Responsible AI, are likely to see significant gains in terms of competitive advantage, with the benefits outweighing the costs.	Organisations may be tempted to make some ethical sacrifices in order to expedite their AI projects and keep pace with competitors. However, the evidence indicates that there are significant returns to be gained from investing in a responsible approach to AI development, including increased competitiveness.

REVIEW OF RESPONSIBLE AI TOOLS AND GUIDELINES

DISR - For release under the FOI Act

LEX 75414 - FOI - Document 1

The National Artificial Intelligence Centre has worked with The Gradient Institute, with support from Fifth Quadrant, to conduct a review of responsible AI tools and guidelines. The purpose of the review is to help businesses put the [Australian AI Ethics Principles](#) into practice in their organisations. The full report and summary of this review can be downloaded from the NAIC's website. See links below.



Image of full report with link to
NAIC



Image of summary with link to
NAIC

IAG is using artificial intelligence to predict whether a motor vehicle is a total loss after a car accident, improving customer experience by reducing insurance claims processing times from over three weeks to just a few days.

How Predictive Total Loss Works

Predictive Total Loss automates business processes to deliver proactive and transparent customer communications that keep customers informed at each stage of the motor total loss experience. It removes manual processing steps to settle customers' claims sooner. In addition to the positive impact to customer advocacy, Predictive Total Loss has:

- Put customers at the centre of the business with tailored specific customer problems and pain points when a customer experiences a claim. This reduces claim times and improves customer experience after a car accident.
- Reduced claim cycle time and improved customer activity.
- Automated aspects of the total loss processes each month. This freed up tens of thousands of manual processes from the claims teams, improved overall efficiencies of the claims teams.

@Chris Dolman - does this need updating/ replacing?
 @Artak Amirbekyan - is there a publicly available Transurban case study to include?

Application of AI Ethics Framework

Prior to deployment, Predictive Total Loss was evaluated using IAG's established AI ethics framework and the Australian Government's voluntary AI Ethics Principles to identify potential issues or risks prior to go-live, including:

- Human, social and environmental wellbeing: Making sure the objective of the project was to benefit IAG's customers, with no other conflicting objectives, and clearly documenting this to assist with ongoing monitoring.
- Reliability and safety: Experimentation to verify that customers had a positive experience and setting conservative thresholds for modelling to help reduce the likelihood of wrongly predicted total losses.
- Fairness: Careful consideration of the potential benefits and harms of the system, including the distribution of benefits and harms across the population.

THANK YOU

For more information please contact:

Steve Nuttall

Director, Fifth Quadrant

E: snuttall@fifthquadrant.com.au

Ph: (+61 2) 9927 3306



THE RESPONSIBLE AI INDEX




Bridging The Responsible AI Gap: 2022

The 2022 Responsible AI Index, sponsored by IAG and Transurban, measures and tracks how well organisations are designing and implementing Responsible AI systems. It is based on a survey of over 400 executive decision makers responsible for AI development.

The mean Responsible AI Index score for Australian-based organisations is 62 out of 100 (unchanged since 2021). Organisations where the CEO is responsible for driving the AI strategy are more mature than those where the CEO is not taking the lead.

With only 3% in the maturing segment, this demonstrates there is significant room for improvement, and raises a concern that many organisations are not yet including responsible practices and ethics into the design and deployment of their AI systems.



THE RESPONSIBLE AI MATURITY SEGMENTS

Less Mature (0-64)

More Mature (65+)

- Early stages of AI deployment
- Focussed on quickly reaping commercial benefits of AI automation
- May have lack of confidence, leadership support and/or knowledge of responsible AI
- Implemented auditing processes for AI and/or developed guidelines for responsible use of AI
- Strong focus on the moral and ethical implications of using AI technologies
- Strong culture of data protection and security
- Use external specialists and advisors

RESPONSIBLE AI STRATEGY

ORGANISATIONS ARE TAKING A MORE STRATEGIC APPROACH TO AI

Compared with 2021, more organisations are taking an enterprise-wide approach for the development of AI which is tied to the wider business strategy across all divisions.



PUTTING PRINCIPLES INTO PRACTICE

GROWING AWARENESS OF AUSTRALIA'S AI ETHICS PRINCIPLES

The Australian Department of Industry, Science and Resources has developed eight AI Ethics Principles to ensure that AI is safe, secure and reliable. Compared with 2021, more organisations are aware of the Principles



THE RESPONSIBLE AI GAP


There is a significant gap between perceptions of how AI systems have been designed and perform, and the actions that have been taken to ensure AI systems are developed responsibly.

AGREEMENT WITH STATEMENTS ABOUT AI PERFORMANCE	PRINCIPLE	ACTIONS TAKEN
81% Our AI systems generate quantifiable benefits to humans, society and the environment that outweigh the costs."	HUMAN, SOCIAL AND ENVIRONMENTAL WELLBEING <i>Throughout their lifecycle, AI systems should benefit individuals, society and the environment.</i>	22% Have identified and assessed the risks and opportunities for human rights.
77% Our AI systems are designed to be human-centered at their core.	HUMAN-CENTRED VALUES <i>Throughout their lifecycle, AI systems should be inclusive and accessible and should not involve or result in unfair discrimination.</i>	23% Have reviewed underlying databases for potential bias.
79% We have robust systems and processes in place to minimise the likelihood of our AI systems causing unfair treatment of individuals, communities or groups.	FAIRNESS <i>Throughout their lifecycle, AI systems should respect human rights, diversity and the autonomy of individuals.</i>	26% Have reviewed AI algorithms for potential bias.
84% Our AI systems comply with relevant privacy and security regulations.	PRIVACY PROTECTION AND SECURITY <i>Throughout their lifecycle, AI systems should respect and uphold privacy rights and data protection and ensure the security of data.</i>	24% Have hired technical consultants or professionals.
84% Our AI systems are designed to be safe and to not harm or deceive people.	RELIABILITY AND SAFETY <i>Throughout their lifecycle, AI systems should reliably operate in accordance with their intended purpose.</i>	25% Have monitored outcomes for customers or employees.
84% We are able to transparently show and explain how algorithms work.	TRANSPARENCY AND EXPLAINABILITY <i>There should be transparency and responsible disclosure to ensure people know when they are being significantly impacted by an AI system, and can find out when an AI system is engaging with them.</i>	26% Have developed supporting materials to explain the AI inputs and decision-making processes.
80% We have a timely process in place to allow people to challenge the use or outcomes of our AI systems.	CONTESTABILITY <i>When an AI system significantly impacts a person, community, group or environment, there should be a timely process to allow people to challenge the use or output of the AI system.</i>	23% Have set up recourse mechanisms (i.e. if an AI negatively affects a member of the public).
85% Our leadership can be held accountable for the impact of their AI systems.	ACCOUNTABILITY <i>Those responsible for the different phases of the AI system lifecycle should be identifiable and accountable for the outcomes of the AI systems, and human oversight of AI systems should be enabled.</i>	23% Have engaged the business leadership on the issues around responsible AI.

BENEFITS OF RESPONSIBLE AI


RESPONSIBLE AI PROVIDES A COMPETITIVE ADVANTAGE

Those that invest in a responsible AI approach believe this has provided a significant competitive advantage.



THE BENEFITS OF RAI OUTWEIGH THE COSTS

Most organisations believe the benefits of taking a responsible approach to AI outweigh the costs



BRIDGING THE AI RESPONSIBILITY GAP

REVIEW OF RESPONSIBLE AI TOOLS AND GUIDELINES

The National Artificial Intelligence Centre has worked with The Gradient Institute, with support from Fifth Quadrant, to conduct a review of responsible AI tools and guidelines. The purpose of the review is to help businesses put the Australian AI Ethics Principles into practice in their organisations.

- Downloads
 - 2022 Responsible AI Index Report
 - 2022 Responsible AI Review of Tools and Guidelines Report
 - 2022 Responsible AI Review of Tools and Guidelines Infographic

Get your score today with the free Responsible AI Self-Assessment Tool

DEPARTMENT OF INDUSTRY, SCIENCE AND RESOURCES

MS23-000284

To: Minister for Industry and Science (For Decision)

REQUEST FOR NATIONAL SCIENCE AND TECHNOLOGY COUNCIL (NSTC) TO COMMISSION RAPID REPORT ON ARTIFICIAL INTELLIGENCE

Timing: Urgent.

Recommendation:			
1. That you sign the letter to the Chief Scientist at <u>Attachment A</u> requesting the NSTC to commission a rapid response report on the opportunities and impacts of Artificial Intelligence (AI).			
			Signed / Not signed
Minister:		Date:	
Comments:			
Clearing Officer:	s22	Head of Division, Technology and Digital	s22
Contact Officer:	s22	General Manager, Emerging Technologies Branch	s22
For Parliamentary Services' use only. Date Submitted to the Minister's office in PDMS:			16/2/2023

Key Points:

1. OpenAI's launch of ChatGPT in November 2022 has seen growing community awareness of generative Artificial Intelligence (AI), and the opportunities and risks associated with widespread use of the technology. In particular, media coverage has focussed on job security and concerns about the integrity of academic credentialing.
2. s34
3. s34 you have the opportunity to ask the National Science and Technology Council (NSTC) to commission a rapid response report on the opportunities and impacts of AI.

4. Following discussions with your office, we recommend you direct the NSTC to provide expert advice to the following two questions:
 - a. What are the opportunities and risks of applying large language models and multimodal machine learning technologies over the next two, five and ten years?
 - b. What are some examples of strategies that have been put in place internationally by other advanced economies since the launch of models like ChatGPT to address the potential opportunities and impacts of AI?
5. Based on these questions, Professor Genevieve Bell has provided the expert insight as to three things the response is likely to focus on:
 - a. "A short description of what generative AI actually is (ie: what is ChatGPT and why it is part of a broader set of work)
 - b. What is the emergent opportunity space, and current moves globally. (I think it is far harder to predict what this space will look in the 2-10 year area, as my best guess is that we are in the middle of the Gartner hype curve ascent right now, and everyone is a bit exuberant).
 - c. And what are the emergent responses (again bearing in mind the above)."
6. If you agree to this request, we have prepared a letter for your signature at Attachment A to Australia's Chief Scientist, Dr Cathy Foley, who is Executive Officer for the NSTC.
 - a. The next meeting of the NSTC is currently scheduled for 30 March 2023 and the agenda will include discussion of its forward workplan.

Rapid Report:

7. A rapid report typically takes four to six weeks to develop and deliver to the department upon receipt of your request. They provide a short (no more than 1500 words) review into specific areas, commissioned independently and peer reviewed by experts.
8. NSTC reports are typically commissioned to only provide advice, rather than policy recommendations. Upon delivery of the report to the department, the decision to release publicly will be made in conjunction with your office.
9. Professor Genevieve Bell is best placed in the NSTC to sponsor this report. These questions have been sent to her for her expert opinion.

Sensitivities and Handling:

10. ^{s34}

Data referenced: Nil.

Consultation: YES

11. Office of the Chief Scientist, Science Division.

ATTACHMENT

A: Proposed letter to Dr Foley to commission the AI rapid report



**THE HON ED HUSIC MP
MINISTER FOR INDUSTRY AND SCIENCE**

MS23-000284

**Dr Cathy Foley AO PSM
Australia's Chief Scientist
GPO Box 2013
CANBERRA ACT 2601**

Dear Dr Foley

I am writing to ask the National Science and Technology Council (the Council) to commission a rapid response report on recent developments in Artificial Intelligence (AI).

In particular, I welcome the Council's views on large language models and multi modal models, such as ChatGPT, Bard and Stable Diffusion. I am seeking the Council's expert advice and opinion in response to the following questions:

1. What are the opportunities and risks of applying large language models and multimodal machine learning technologies over the next two, five and ten years?
2. What are some examples of strategies that have been put in place internationally by other advanced economies since the launch of models like ChatGPT to address the potential opportunities and impacts of AI?

I would appreciate the delivery of this advice to my department by the end of April 2023.

Yours sincerely

Ed Husic MP