What should be included in my organization's AI policy?: A data governance checklist

AUTHORS

 Molly Reynolds

 Joel Ramsey

Further reading:

- The board says we need an AI strategy. How do we start?
- Can HR use AI to recruit, manage and evaluate employees?
- What's new with artificial intelligence regulation in Canada and abroad?

Organizations developing, procuring or permitting personnel to use AI systems for business purposes should have a responsible AI policy to guide consistent, compliant decision-making. Not only will this enable an organization to maintain compliance with developing legal and regulatory standards, but it will also serve an increasingly important role as the organization plans strategies and transactions involving the use or procurement of AI.

Industry and legal frameworks have coalesced around several core principles for responsible AI, many of which can intersect with existing corporate policies, departments and workstreams. A key intersection is the alignment of the organization's data governance regime with its approach to responsible AI. In this article, we identify a checklist of data governance considerations that arise in procuring, developing and using AI systems throughout the responsible AI lifecycle.

The essential elements of data governance for AI

Among the core principles for the responsible and ethical use of AI, the following tenets should inform an organization's data governance considerations:

1. Accountability

The AI Policy should set out roles and responsibilities of various employees, teams, managers, officers, and the board with respect to the adoption, use and oversight of AI systems. This stems from the guiding principles of accountability and human oversight of AI operations and outputs. If AI systems are provided and managed by a third party, the AI policy should address how the organization will ensure these roles and responsibilities are carried out by the third party through effective oversight and contractual rights.

The AI Policy should articulate who is responsible for:

- determining which records should be created to document the use case for an AI system, the due diligence performed during procurement, development and/or oversight of the tool, how the system is functioning and how it is adapted over time (including how its functions may need to be limited to ensure continued compliance with the policy);
- deciding how long such diligence and monitoring records should be retained and where; and
- updating the organization's record retention policy and schedule accordingly.

2. Transparency

The AI Policy should require that appropriate information regarding AI systems used or developed by the organization can be provided to relevant stakeholders. This may include external or internal disclosures of whether AI is used for particular decisions, responding to individual or regulatory requests for explanations of how the system works, or justifying the fairness and accuracy of automated decisions or recommendations. Contracts with third parties providing AI systems should have mechanisms to obtain this information through appropriate governance, oversight and audit rights.

The data governance function should have input into:

- which records are created to map the source of the data used to train the AI system, and the type and source of data inputted into the tool as part of its use;
- the required logical, architectural and technical documentation of how the AI system processes inputs to make the recommendations or decisions envisioned by the use case;
- the frequency of internal or external reviews of the AI system's operation in the ordinary course or in response to questions or complaints raised by users or other stakeholders and how the results of such reviews and any remediation steps are recorded; and
- where the above records will be stored, how long they will be retained, who has access to them, and which teams or personnel should be engaged if the records are requested by regulators or other external actors.

3. Security

The AI Policy should identify all existing information on security-related policies, procedures, and training that apply to the use of AI systems, as well as any additional security protocols required for such tools.

> Internal data governance specialists and existing policies and frameworks can be leveraged to support responsible AI without duplicating compliance, reporting and risk mitigation regimes.

Organizations should set expectations for how data inputted into and generated from AI systems will be protected throughout its lifecycle, whether it can be used for further training of the AI systems, and how it will be securely destroyed at the end of that lifecycle.

Data security incident response plans should be adapted if necessary to provide guidance to responding to breaches involving AI systems. This may include 1) updating applicable regulators, stakeholders, or individuals to be notified of an incident; 2) modifying the types or content of breach records kept; and 3) ensuring that post-incident debrief and remediation efforts address considerations specific to the affected AI system (for more on cybersecurity, read "A sword and a shield: AI's dual-natured role in cybersecurity").

4. Risk mitigation

The AI Policy should identify existing or custom risk management frameworks that must be used to assess the nature and risk profile of any AI system. The frameworks should include risk assessment, acceptance and mitigation processes from initial proposal through the use and ultimate decommissioning of the AI system. These processes should 1) consider risks associated with bias, accuracy, security, performance and reputation; and 2) be designed, developed and used in a way that prioritizes the mitigation of risks in accordance with defined mitigation practices applicable to the evaluated risk level of such AI system. These risk assessments and processes should be integrated into the organization's third-party risk management program and apply to all phases of the procurement cycle and ongoing oversight of third parties.

The documentation of the above risk management processes (including records of diligence, testing, results, remediation and corporate decision-making) is a key component of AI governance. The organization may need to update internal policies and provide additional training to individuals across functions to ensure that the processes are followed and documented, and that risk management outcomes are appropriately reported (for more on AI governance, read "The board says we need an AI strategy. How do we start?").

Conclusion

It is clear from this brief canvass of just four responsible AI principles that data governance and records management are core components of an effective AI strategy. While the implementation of such a strategy requires thought, resources and a cross-functional approach, it is possible to integrate it with existing corporate risk management frameworks. Internal data governance specialists and existing policies and frameworks can be leveraged to support responsible AI without duplicating compliance, reporting and risk mitigation regimes. The effort and coordination will set up an organization proactively for success as it adapts to an ever-changing AI landscape, including for future transactions, cyber-threat planning and response, and successful litigation.

This article was published as part of the Q4 2024 Torys Quarterly, "Machine capital: mapping AI risk".

To discuss these issues, please contact the author(s).