OSLER

AI governance: navigating the path ahead

DEC 11, 2023 7 MIN READ

Related Expertise

- Artificial Intelligence
- Technology

Authors: Simon Hodgett, Sam Ip

Without question, during the past year, Canada has witnessed remarkable and radical change in the artificial intelligence (AI) sector, with the emergence and unprecedented speed of adoption of generative AI technologies. The rapid adoption has been accompanied by a wave of legislative and regulatory activity that has the potential to dramatically affect the use and adoption of AI and related technologies.

Revolutionary technology and regulatory changes present challenges for businesses. AI offers benefits to businesses, but also brings new risks. Leadership teams need to ensure accountability regarding the use of AI, including by implementing appropriate policies and procedures. Understanding intellectual property risk and protection is key. At the same time, organizations need to take practical steps to address risks associated with AI, including through robust compliance programs, transparency around the use of AI within the organization and the adoption of ethical standards and codes of conduct.

In 2024 and beyond, businesses will need to proactively consider the benefits of AI, take steps to mitigate risks and navigate an increasing complex regulatory and ethical environment.

AI regulatory changes are proliferating

Over just a short period, a number of legislative changes and new guidance have been proposed that have the potential to have a profound impact on the adoption and use of AI technologies.

At the federal level, the Canadian government has proposed the *Artificial Intelligence and Data Act* (AIDA), a key aspect of the proposed Digital Charter for Canadians. AIDA has been through first and second reading and is currently being considered by the Standing Committee on Industry and Technology.

The Digital Governance Council has proposed standards for Canadian fintech innovators that are aimed at addressing risks associated with AI and machine learning (ML) solutions for financial institutions. The Office of the Superintendent of Financial Institutions and the Global Risk Institute have issued a joint report providing guidance on AI for financial services institutions.

Finally, the Canadian federal government has also worked with a variety of corporate

organizations to prepare a Voluntary Code of Conduct on the Responsible Development and Management of Advanced Generative AI Systems. Undoubtedly more changes are to come.

Key considerations for navigating AI governance

The adoption and use of AI technologies require a comprehensive strategy that flows through the organization. As a starting point, organizations should establish a risk management framework that is proportionate to the scale and impact of their AI activities and that is practical to implement.

The first step involves assessing how the organization intends to use AI and whether the organization is a developer, deployer or user of AI solutions. A user of AI solutions will not have the same concerns or be exposed to the same risks as the developer of those same technologies. In establishing a risk management framework, organizations should determine their overarching AI strategy and objectives, informed by the organization's risk appetite. From there, implementation of the framework and overall responsibility for safe, secure and trustworthy AI adoption is best delegated to one or more lead persons.

Finally, establishing policies and protocols for the organization's use of AI is critical. This process should determine how safety, fairness and equity, as well as transparency, will be taken into account. Policies and protocols should dictate permitted use cases. Human oversight is also essential, as are guidelines for how the performance of the AI system will be assessed. It is also important to determine how issues will be reported, escalated and resolved within the organization.

There are a variety of available resources that provide guidance for the adoption of AI risk management frameworks. For example, the AI Risk Management Framework (NIST) sets out principles that organizations can incorporate to adapt existing governance strategies to manage specific AI risk.

In establishing a new framework, organizations should assess their existing policies and consider whether there is a need to adopt new AI-specific policies or whether existing policies can be adapted. For example, existing policies relating to confidentiality may need to be clarified in the context of the use of AI. This should include an emphasis on the need for caution in submitting sensitive, confidential or proprietary information into generative AI services. Such caution is essential not only in connection with the use of generative AI, but also in providing training data to build and enhance the AI models that underlie such services.

Specific policies may be required for technical uses of AI, including basic rules regarding the use of generative AI to assist with coding. Such policies should identify low-risk development activities, such as the development of code for uses internal to the organization. They should also ensure that appropriate safeguards are in place for higher risk uses, such as a requirement for human review of code that is generated for higher risk settings.

In addition to policies and procedures, organizations should develop communications and training strategies relating to AI. Developing internal communication and training regarding the proper use of AI tools is important in establishing compliance protocols in line with the risk management framework. Organizations also need to be prepared to explain to external parties the steps taken to mitigate potential biases and fairness issues in the AI systems that the organization builds or uses.

Given the complex and rapidly evolving environment, both legal and ethical considerations need to be taken into account at the appropriate stages of the development lifecycle for an

AI solution. This is particularly valuable earlier in the development process where measures to address risk may be more cost-effective and practical. Foundational responsible AI principles, such as ethics by design, can be integrated from the onset. Given the importance of appropriate privacy considerations in the use of AI, privacy experts, particularly those familiar with privacy by design frameworks, should be consulted to provide guidance to development teams working with AI.

Most technology organizations have implemented intellectual property protections that apply to software and associated source code. However, AI models, their attributes and architecture may not be adequately protected by existing frameworks. Organizations need to carefully consider how to protect these models and how to appropriately protect ownership of enhancements and improvements as solutions evolve. When dealing with third parties and commercial arrangements, ensuring the terms of contracts governing these arrangements define AI-specific elements – such as models, weights and biases within them and the model architectures – and specifically allocate ownership is an important element of intellectual property protection. Common intellectual property and data definitions within contracts do not always precisely allocate ownership for such items.

In addition, commercial contracts for the supply of AI solutions are an effective avenue to ensure the particular arrangement reflects the organization's AI policies and legislated requirements. The risk allocation in commercial contracts can also reflect the organization's risk appetite with respect to AI applications, as well as the organization's expectations relating to specific use cases. Contracts can be used to ensure that the AI vendor commits to evolving standards, that their solution is explainable and that they commit to testing their solution for bias.

Data powers AI. A comprehensive risk management framework needs to address the collection, use and disclosure of data in compliance with applicable regulations, including privacy laws. Policies addressing the transfer and disclosure of data should be aimed at ensuring that such transfers and disclosures only occur with necessary consents and contractual permissions. Protocols must also ensure the prior assessment of the availability, quantity and suitability of data sets to support solutions that provide useable and legally compliant results. For example, data inputs should ensure a sufficiently large sample size. There should also be measures in place to minimize and address unwanted biases to provide for outcomes that are fair and equitable.

With so many evolving aspects to AI, including the plethora of standards, guidelines and emerging regulations, organizations are undoubtedly overwhelmed. Fortunately, there are solutions being developed to track AI requirements. Organizations should be on the lookout for these solutions as they begin to emerge into the market and are incorporated into standard use. These tools may assist with translating high level policy requirements into actionable controls, assessments of system impact, system guardrails and model inventories.

Looking Forward

The speed of innovation in AI applications is breathtaking and the complexity of the associated policy, ethical and legal considerations promises to grow exponentially in the coming weeks, months and years. Organizations need to be prepared for this rapidly shifting technological frontier and take proactive steps to ensure that appropriate policies and protocols are in place, consistent with regulatory requirements and the business's risk appetite.