

GOVERNANCE INSIGHTS 2023

Get Smart on Artificial Intelligence (AI) and Corporate Governance: Key Considerations for Boards of Directors

On November 30, 2022, OpenAI released ChatGPT—a conversational, generative artificial intelligence (AI) chatbot trained on OpenAI’s foundational large language models—arguably catalyzing today’s new wave of interest in AI. This surge of interest is also driving an increase in competition and global investment in AI, currently projected to reach US\$200 billion by 2025, and an urgency to address the use and regulation of AI. Content-generating technologies such as OpenAI’s ChatGPT, Google’s Bard and Anthropic’s Claude, with their potential to drive innovation, increase efficiency and improve decision-making, have been hailed by proponents as transformative technologies, power enablers and skill levellers.

Conversely, critics have lamented the disruptive impact that generative AI could have on education, human intelligence and the labour market. Driving many of these debates is the fact that, at its core, AI is a powerful technology, unparalleled in its accessibility and potential for disruption. As Mustafa Suleyman, co-founder of Google DeepMind (an AI company) wrote, “AI is different...no technology this powerful has become so accessible, so widely, so quickly.” The scale and speed of adoption of generative AI systems and the growth of market interest in AI technologies are testament to this fact.

It is frequently suggested that companies that embrace innovation will be best positioned to seize the competitive advantages offered by generative AI. While business teams may be motivated to push for rapid adoption of new AI-related technologies, corporate directors mindful of their duty to manage risk may seek to tread more cautiously. Technological developments, including the rise of generative AI, do not alter the fundamental fiduciary responsibilities of corporate directors to make decisions honestly, prudently, in good faith and on an informed basis. For management teams, boards of directors and their advisers, effectively harnessing the opportunities while simultaneously managing the risks associated with the use of these technologies require (i) an intentional commitment to learning the current and potential uses of AI within the business, including a clear articulation of the organization’s goals in using AI; (ii) staying apprised of the risks associated with the organization’s use of AI; and (iii) developing a thorough AI governance strategy. In this chapter, we will discuss these three requirements and offer practical considerations for boards and their advisers.

Spotlight: AI and Generative AI: A Primer

Artificial intelligence is “the science and engineering of making machines intelligent.”

– John McCarthy (a founder of the discipline of artificial intelligence)

The term “artificial intelligence,” coined by John McCarthy in 1955, has a long history in computer science, cognitive science and philosophy, so caution is advised when considering what should be brought within its scope. That said, AI is generally thought of as any technology that allows computers to perform or mimic the cognitive functions that we usually associate with human intelligence. Research into AI predates the existence of the term. Alan Turing’s seminal paper, “[Computing Machinery and Intelligence](#),” was published in 1950, and the first proposal for artificial neural networks came in 1943 when [Warren McCulloch and Walter Pitts](#) developed the first models. The first running AI program

was created in 1955; it was called Logic Theorist and was able to prove complex theories in mathematical logic. Since then, AI has transformed the way we learn in school, edit photographs, buy groceries, use maps, book travel, conduct research, watch movies and listen to music. Many of our frequently used AI applications are built on technologies that are designed to perform a single task or a relatively limited set of related tasks. These are referred to as “narrow AI.” Narrow AI is one end of a spectrum; at the other end is “broad AI,” sometimes also called artificial general intelligence (AGI). Most of us use some form of narrow AI every day. In fact, virtually all AI that has been developed until very recently may be classified as narrow AI.

The past 15 years have witnessed a massive expansion in the development and use of a particular type of AI technology known as machine learning. Machine learning technology uses statistical algorithms to find patterns in data and to optimize performance in finding those patterns based on the data

consumed. This process is called model training and results in a trained model that can then be deployed to process new data.

Generative AI technologies, such as ChatGPT, are built on foundation models. These are large machine learning models trained on a vast quantity of data at scale. These models are typically trained on a variety of data points and on information covering a wide range of topics. Generative AI operates by establishing connections between the natural language input provided by the user and tokens within the relevant dataset. It creates associations between the words and then generates a natural language response, images or other content based on these associations.

Today's generative AI systems differ from predecessor AI technologies in a number of ways.

First, generative AI systems are designed for a wide range of use cases. Generative AI is not AGI, but its relative flexibility

is a significant shift along the narrow AI/broad AI spectrum. Although there are certainly drawbacks to this flexibility, these incredibly versatile foundational models are capable of performing multiple functions within a single organization. From a business perspective, this means an organization can implement a single AI system for multiple use cases within the business and across business units.

Second, these models are capable of efficiently creating new content in various forms (audio, images, code, text). This capability is incredibly valuable for many businesses because it can increase efficiencies across a business and can lower costs, if applied appropriately.

Finally, these systems are built with user-friendly interfaces that are able to process and understand natural human language. For many companies, this distinction is the most important because the usability of generative AI tools allows companies to realize the benefits of AI faster and allows individual users to get up to speed quickly.

Understand the Use Cases for AI Within the Business

Some directors may experience a degree of discomfort taking on oversight responsibility for the use of technologies that are new, unfamiliar and constantly evolving. However, as companies explore use cases for generative AI and expand their daily use of AI, the board's responsibilities in this regard will be unavoidable. In fact, many of the evolving AI regulatory frameworks focus significant attention on the need for robust corporate governance structures to address risks associated with the use of AI. Although board members are not expected to become experts in the technology, directors' legal duty of care requires that board members make an effort to educate themselves and to analyze and consider AI technologies as carefully as any other person would in a similar situation. It is, therefore, important for directors to develop a sufficient understanding of AI technologies to be able to assess the costs and benefits of adopting new AI systems. Directors may wish to consider the following:

- **Define AI.** To frame internal discussions and guide the development of a governance framework that is both robust and directly applicable to the use of AI within the organization, consider seeking input from management to develop an internal working definition of AI and a clear statement of the goals the organization is trying to achieve with AI.
- **Increase the board's AI knowledge base.** Consider implementing board education initiatives on AI and generative AI, and leveraging internal and external resources to increase director understanding of AI and generative AI technologies and their use throughout the organization.

- **Assess current use cases.** Because AI has likely been used routinely for years, consider working with management to assess how the company currently uses AI, how recent technological developments may affect the company's use of AI, and how competitors and other industry participants use (or are expected to use) AI.
- **Assess strategic opportunities and key risks.** Consider working with management to determine how to mitigate the potential risks of using advanced AI technologies, without stifling innovation, and to assess the available opportunities for leveraging the potential of AI within the organization to achieve the organization's stated objectives.

Stay Apprised of Compliance Risks Associated with the Company's Use of AI

The use of AI, particularly generative AI, raises important compliance issues for corporate boards and their advisers. Regulators around the world are currently developing robust regulatory frameworks and initiatives for AI that seek to address the most pressing identified risks posed by AI while not discouraging innovation. Some of these regulatory frameworks specifically address generative AI. They all subject the development and use of AI technologies to stringent requirements designed to address privacy, security, confidentiality and bias concerns.

- To address recent developments and the rapid growth of AI, the Canadian government has introduced before Parliament the *Artificial Intelligence and Data Act (AIDA)*, a risk-based regulatory framework. Under AIDA, businesses will be held responsible for the AI activities under their control, no matter where they sit in the AI value chain (as a designer,

developer, provider or operator). They will be required to implement new governance mechanisms and policies that will consider and address the risks of their AI systems and give users enough information to make informed decisions (see also the [AIDA Companion Document](#), which provides insights into the Canadian government's approach to regulating AI systems). The Canadian government also launched the [Voluntary Code of Conduct on the Responsible Development and Management of Advanced Generative AI Systems](#) on September 27, 2023. The code outlines voluntary measures that organizations are encouraged to apply to their activities in the development and management of general-purpose generative AI systems. The measures are aligned with six core principles: accountability, safety, human oversight and monitoring, fairness and equity, transparency and validity and robustness. In addition, the government has noted that, although there is currently no comprehensive legal framework for AI, a number of existing laws apply to AI systems, such as the *Personal Information Protection and Electronic Documents Act* (PIPEDA). For more information on PIPEDA, please see our [Davies Governance Insights 2018](#).

- In addition, in the same bill that introduces AIDA, the Canadian government has proposed reforms to the federal private sector privacy law framework that would introduce transparency obligations for organizations that use AI systems to make predictions, recommendations or decisions about individuals that could have a significant impact on them.
- In Québec, recently enacted reforms of the province's privacy laws have brought into force both transparency obligations for organizations and new individual rights in relation to the use of AI systems that use personal information to make decisions without human intervention.

- The European Union's (EU's) [Artificial Intelligence Act \(AIA\)](#) sets out obligations for providers and users, depending on the level of risk an AI technology could pose to an individual's health and safety or fundamental rights. The AIA, expected to go into effect in 2025, will likely be the first comprehensive AI-focused legal framework worldwide. The AIA framework contains four levels of risk—unacceptable, high, limited and minimal—and the obligations apply to providers and users located outside the EU if the output produced by the AI system is intended to be used within the EU. This means that all parties involved in the development, use, import, distribution or manufacturing of AI systems will be held responsible under the AIA.

- Further, on May 31, 2023, the EU and the United States announced that they are working together to develop a voluntary code of conduct for AI to establish non-binding international standards for AI risk assessments, transparency and other requirements. There has also been considerable momentum in Europe to determine a regulatory framework for the foundation models that underlie generative AI systems such as ChatGPT. EU lawmakers have been working to include specific obligations in the AIA for foundation model developers, independent of the model's intended use, including compulsory model testing by independent third-party specialists. An alternative has been proposed by France, Germany and Italy, which circulated a joint paper on November 19, 2023, indicating that the three countries have agreed among themselves to require mandatory self-regulation for foundation model developers through codes of conduct.

- On October 30, 2023, U.S. President Joe Biden issued an [Executive Order on Safe, Secure, and Trustworthy Artificial Intelligence](#), which sets new standards for AI safety and security. The order outlines disclosure

obligations and industry-wide requirements for “AI systems,” which it broadly defines as “any data system, software, hardware, application, tool, or utility that operates in whole or in part using AI.” As in Canada, regulators in the United States have noted that many existing regulations apply to the use of AI, although there is currently no comprehensive regulatory framework.

Although none of these frameworks are expected to take effect until 2025, boards should understand and anticipate how these evolving regulatory frameworks may impact a company’s existing AI governance policies and practices. In addition, as we discussed in [Davies Governance Insights 2020](#), stakeholders increasingly expect boards to actively monitor enterprise risk as part of their oversight responsibilities. Companies and their boards may consider evaluating their existing AI risk management framework against the U.S. Department of Commerce’s National Institute of Standards and Technology’s (NIST) [AI Risk Management Framework](#). By outlining steps for identifying, prioritizing and managing AI risks, the NIST framework can help boards develop a baseline and prepare for guidance that may be issued by other U.S. agencies and regulatory bodies, as applicable. Additionally, boards can

- spearhead and oversee assessments of the key impacts and implications of evolving regulatory frameworks, both domestic and international, on the organization and its use of AI systems;
- consider hosting discussions with advisers and management to assess whether the company is developing or using AI systems in accordance with proposed regulatory requirements;
- consider what updates may be or may become necessary to existing AI-governance policies and enterprise risk management programs;

- consider establishing a committee that includes individuals with relevant expertise to oversee and report to the board on risks relating to the company’s use and management of AI and generative AI systems;
- consider whether all AI systems and tools used across the business and various business units, as applicable, will be governed by the same set of rules;
- charge advisers with providing regular updates on the evolving regulatory framework and confirm that the board is up to date on best practices related to AI oversight and compliance issues;
- consider AI regulatory developments in the context of the company’s current compliance obligations and oversee the development and implementation of appropriate systems, policies and controls; and
- consider the best approach to educating the board about the company’s exposure to risks related to the use of AI. This may include board briefings on significant AI-related incidents and investigations, as well as awareness briefings on the company’s critical uses of AI and any associated regulatory, financial, operational or reputational risks.

Develop a Thorough AI Governance Strategy


The latest generative and other machine learning-based AI technologies represent a significant technological advancement. The speed of their adoption across industries, across organizations, and by individuals in their daily lives is a testament to that reality. Today, many companies have begun to adopt more advanced AI systems and expand their use of AI in response to the growing popularity and transformative power of generative AI. However, AI applications can also

raise unique concerns related to data security, human rights, civil rights, information integrity, intellectual property rights and bias or discrimination. Perhaps now more than ever before, being a good corporate citizen requires the responsible use, development and management of AI.

As regulators race to keep up with the speed of adoption of generative AI and the speed of development of advanced (and potentially high-risk) AI systems, corporate boards and management teams are currently being forced to grapple with the impact of AI on people, processes and industries. Although this is more challenging without the clarity of comprehensive regulatory guidance, directors must make a good faith effort to implement oversight controls over AI. Boards may want to consider the following when developing a robust AI oversight strategy:

– **Safety and Security.** Consider whether existing AI systems are subject to frequent risk assessments and cybersecurity reviews. Ask management about the following cybersecurity protocols:

- > How frequently do we evaluate and revise our AI and cybersecurity policies?
- > How often are cybersecurity checks performed?
- > When was the last time the company evaluated its cybersecurity insurance policy?
- > Have cybersecurity issues increased with changes in the company's use of AI?
- > What steps are being taken to stay ahead of emerging risks, especially those related to cybersecurity and data privacy?
- > In the event of a major disruption or cybersecurity incident, what contingency plans are in place to ensure business continuity, and how regularly are these plans tested and updated to align with evolving risks and circumstances?




- > In the event of a cybersecurity incident, how will the incident response team communicate with the rest of the organization?
- **Risk Management.** Consider whether management has established appropriate risk management systems. Ask management about internal processes for identifying and mitigating any new or emerging AI or cyber-related risks posed by the company's use of new AI systems or generative AI:
 - > What are the key risks that management has identified and what risk mitigation strategies or resources are currently available or under development?
 - > Who is primarily responsible for monitoring AI compliance and risk mitigation?
 - > How frequently are internal risk assessments conducted?
 - > How does the internal risk management system incorporate feedback and insights from employees at various levels within the organization, and what mechanisms are in place to foster a culture of risk awareness and reporting?
 - > In the current context of rapid technological advancements and changes in the business environment, how adaptable is the current risk management framework?
- **Effective Oversight.** Consider what systems have been developed to monitor the outcomes and impact of generative AI technologies and new AI systems that the company implements. Ask management about the structures that are currently in place or being developed:
 - > What is the organization seeking to accomplish with new AI systems and generative AI technologies?
 - > Has the company developed policies and procedures for responding to whistleblower complaints regarding AI issues, material AI-related incidents and issues arising from the company's AI vendors?
- > Are procedures in place to ensure that the board has sufficient information to perform AI oversight activities?
- > Are there any specific risks to the business that are created or exacerbated by the company's use of new AI systems or generative AI technologies?
- > How are we measuring success in terms of the use, integration and monitoring of generative AI systems?
- > Who is primarily responsible for each of the organization's major AI systems?
- **Stakeholder Engagement.** Maintain regular communication with key shareholders, customers, suppliers and communities in which the company operates. Ask management about the processes currently in place to promote dialogue with key shareholders and other stakeholders:
 - > Is the company attuned to customer expectations regarding the use of customer data to train and operate the company's AI systems?
 - > Does the company listen to and, where appropriate, respond to community members' concerns about potential misuse and negative consequences of the company's use of AI that affect the company's employees, key stakeholders, customers or the environment?
 - > How does the organization communicate its AI story, including current use cases and strategic opportunities?
 - > Does stakeholder engagement occur only in response to a crisis or does the board receive regular reports on ongoing processes that enable the company to engage with stakeholders?

Concluding Thoughts

In prior editions of *Davies Governance Insights*, we have discussed how boards and senior management might respond to the ever-changing environment in which their organizations operate. We posited that evolving a business into a “next generation governance organization” is one way to build organizations that are more resilient, agile and innovative in the face of a rapidly changing environment. We defined a next generation governance organization as one that is focused on its business strategy, is people-centred and proactively engages with shareholders and other stakeholders.

Today’s transformative AI technologies and the creation of next generation governance organizations seem to go hand in hand. The context may be different but the

principles remain the same. Generative AI is increasingly considered necessary for organizations that are focused on forward-looking strategic value creation. Boards of directors serving these organizations will continue to be called upon to respond to the environment that is evolving around and in response to new AI technologies. In this chapter of *Davies Governance Insights 2023*, we add the following element to the concept of a next generation governance organization: next generation governance organizations are led by management teams and boards of directors that remain intentionally committed to learning. For directors in today’s environment, applying this intentionality requires learning how AI is currently being used by their organization, staying apprised of the risks associated with the company’s use of AI and developing a thorough AI governance strategy.



Key Contacts

Toronto



Patricia L. Olasker

416.863.5551
polasker@dwpv.com



Aaron J. Atkinson

416.367.6907
aatkinson@dwpv.com



Brett Seifred

416.863.5531
bseifred@dwpv.com

Montréal



Franziska Ruf

514.841.6480
fruf@dwpv.com

New York



Jeffrey Nadler

212.588.5505
jnadler@dwpv.com

Contributors

Researching and writing this report is a project undertaken by Davies Ward Phillips & Vineberg LLP and not on behalf of any client or other person. The information contained in this report should not be relied upon as legal advice.

Aaron Atkinson (Activism)

Jesany Michel (AI)

Ivana Gotzeva (AI)

Alexander Max Jarvie (AI)

Patricia Olasker (Activism)

Brandon Orr (Activism)

Alexandria Pike (ESG)

Sarah Powell (ESG)

Brett Seifred (Diversity)

Matthew Sherman (Diversity)

Ghaith Sibai (Activism)

Zachary Silver (ESG)

Mathieu Taschereau (Activism)

About Davies

Davies is a law firm focused on high-stakes matters. Committed to achieving superior outcomes for our clients, we are consistently at the heart of their largest and most complex deals and cases. With offices in Toronto, Montréal and New York, our capabilities extend seamlessly to every continent. Contact any of our lawyers to talk with us about your situation.

Visit us at dwpv.com

DAVIES

TORONTO

155 Wellington Street West
Toronto ON Canada
M5V 3J7

416.863.0900

MONTRÉAL

1501 McGill College Avenue
Montréal QC Canada
H3A 3N9

514.841.6400

NEW YORK

900 Third Avenue
New York NY U.S.A. 10022

212.588.5500