BOARD
PERSPECTIVES

ISSUE 167

The Director's Playbook for Generative AI

Generative AI (GenAI) is accelerating change faster than any predecessor technology. Yet, many questions remain about its trustworthiness. What is the board's role in balancing the opportunities and risks of GenAI and guiding organisations as the technology proliferates?

The big picture: So much has been written about GenAI lately it seems like a constant buzz inspiring both wonder and fear. According to McKinsey, the technology's value proposition is alluring:[1]

- GenAI is poised to boost performance and unlock up to US$4.4 trillion of value from sales and marketing, product research and development, customer operations, software engineering, and other business functions.

- With declining productivity growth over the last 20 years, GenAI and other technologies are expected to unleash a new wave of productive growth over the next two decades.

- Up to 70% of people in the global workforce have the potential to see up to 50% of their job functions automated, potentially freeing them up to perform more complex and interesting work.

The bottom line: As one recent publication noted, "AI won't replace people — but people who use AI will replace people who don't."[2]

---

[1] Sourced from McKinsey webinar, "The economic potential of generative AI: The next productivity frontier."
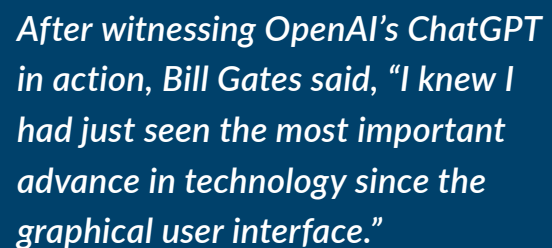
[2] "Augmented work for an automated, AI-driven world," IBM Institute for Business Value, 2023: www.ibm.com/thought-leadership/institute-business-value/en-us/report/augmented-workforce.

After witnessing OpenAI's ChatGPT in action, Bill Gates said, "I knew I had just seen the most important advance in technology since the graphical user interface." He added, "The development of AI is as fundamental as the creation of the microprocessor, the personal computer, the internet, and the mobile phone. It will change the way people work, learn, travel, get healthcare, and communicate with each other."[3]

Between the lines: The opportunities for using GenAI to enhance customer experiences, increase process efficiencies, innovate products and services, and improve productivity are immense.

Yes, but: There are risks and limitations:

- The risks of disinformation and misinformation span an ever-expanding list of potential, egregious abuses fuelling alarm and cries for regulation.

- The CEO of Google has warned of a "mismatch" between how fast AI is developing and how quickly society's institutions can adapt. With the speed at which all this is happening, small errors may be amplified.

- The CEO of OpenAI, which released ChatGPT late last year, recommends that Congress intervene to mitigate the risks of increasingly powerful models.

- Before he died, Steven Hawking noted that AI could become the best thing that ever happened to humanity. But he also warned that "thinking machines" surpassing human intelligence could present an existential threat.

> *After witnessing OpenAI's ChatGPT in action, Bill Gates said, "I knew I had just seen the most important advance in technology since the graphical user interface."*

Regardless of one's take, GenAI is here to stay. For directors interested in contributing value in the boardroom, understanding the opportunities, limitations and risks of GenAI models is now table stakes. Given the technology's accessibility, directors should immerse themselves in learning about GenAI and get hands-on with using it. They should also learn from experts inside and outside the organisation and from published articles that provide relevant content.

Our point of view: Armed with a baseline understanding, directors should consider the following questions when engaging CEOs and their teams in strategic conversations regarding GenAI.

---

[3] "The Age of AI Has Begun," Bill Gates. GatesNotes: The Blog of Bill Gates, March 21, 2023: www.gatesnotes.com/The-Age-of-AI-Has-Begun.

What is the business opportunity in deploying GenAI? This critical "Why should we care?" question should be considered strategically and tactically. Following are five high-level questions for advancing the conversation:

1. What are the implications of GenAI to our industry, and do we know what our competitors are doing with the technology?

2. Do we have a strategy for why, where, how and when we will deploy GenAI? What use cases are we considering, and how are we selecting and prioritising these opportunities? How are we measuring the value GenAI contributes?

3. Are we organised appropriately to roll out our strategy? How are we empowering our people to build, train and use GenAI?

4. Have we documented our organisation's guidelines and values for privacy, security, transparency, fairness, human vs. machine responsibilities, and other matters related to our GenAI deployments? Do our policies account for the need to govern and manage this technology differently than "classical AI"?

5. How do we know we are adhering to our guidelines and values? For example, do we have a cross-functional "ethics committee" that vets all plans and actions and monitors for unintended outcomes and consequences?

The bottom line: Gaining insights on the above questions enables the board to understand how and why management intends to position GenAI in the business. As for companies just beginning, they should start simple. Most important to the board is the recognition that all uses of GenAI are not created equal from a criticality standpoint, nor do they require the same degree of oversight.

What are the legal, regulatory and ethical issues we need to address? GenAI is on the radar of regulators and policymakers at the national, state and local levels as well as of other stakeholders due to the potential cyber, privacy and societal risks. The legal and regulatory environment varies by country and region.

- With legislative initiatives already underway and risk frameworks emerging around the world, the board should inquire how management keeps track of these market developments.

- Applicable requirements and guidelines are likely to include, but are not limited to, transparency, data security, fairness and bias recognition, accountability, ethical considerations, and continuous monitoring and improvement.

- These requirements and guidelines should be embedded into GenAI solutions and the company's internal policies supporting them.

Major players are also weighing in on responsible GenAI practices and managing risk.[4]

How are we sourcing and managing the data used by our GenAI model(s)? Directors should obtain an understanding from management regarding the following:

- Whether the organisation is using (a) publicly available models and domains, (b) foundation models that are fine-tuned with internal proprietary data or (c) fully customised models.

- The maturity and readiness of the existing data governance and data management infrastructure supporting the models deployed.

- The nature of the use cases and the interoperability and maturity of the supporting IT architecture and data ecosystem, which influence the selection or development of fit-for-purpose GenAI technology and the supporting data core.

A key point: The real power of GenAI will likely come from companies infusing it with internal proprietary data. Whether a company uses its own data, third-party data or data generally available in the marketplace will influence a model's risk profile.

Do we have the talent we need to implement GenAI? A key adoption challenge is finding and onboarding the right talent with the requisite expertise:

- Skilled technical practitioners in AI are scarce, especially those who understand how to incorporate the company's business requirements into a GenAI model.

- It takes a team effort to run an end-to-end, AI-embedded system or process, including business champions, data owners, senior program managers and developers, legal and compliance resources, and operational teams.

A reality check: Available skill sets heavily influence the mode of GenAI a company can deploy. While publicly available capabilities such as ChatGPT require no specialised expertise, they are far less secure, private and reliable. That is why most companies will likely choose the middle road — fine-tuning a foundation model, which requires lighter data science expertise through a low-code interface.

> *The real power of GenAI will likely come from companies infusing it with internal proprietary data. Whether a company uses its own data, third-party data or data generally available in the marketplace will influence a model's risk profile.*

---

[4] See "Microsoft AI/Responsible AI Principles and Approach" at www.microsoft.com/en-us/ai/our-approach?activetab=pivot1:primaryr5, Google's "Responsible AI Practices" at https://ai.google/responsibility/responsible-ai-practices/ and NIST's "AI Risk Management Framework" at www.nist.gov/itl/ai-risk-management-framework.

Do we have a governance framework that enables experimentation? Innovation involving any technology entails starting small with experimentation, learning by doing, keeping track of innovation in the marketplace and responding with agility to value-adding opportunities.

The board should inquire about the governance process — including the organisational structure — for overseeing GenAI advancements and experimentation across the company and monitoring industry developments.

- An adaptable GenAI governance framework functions through a small cross-functional, multidisciplinary team representing the data, engineering, security and operational aspects of GenAI models.

- Overall governance involves trust, ethical use, risk management, third-party ecosystem, legal and regulatory compliance and policies, and standards and controls.

- Each team tasked with a model should be responsible for its efficacy.

The need: A GenAI review and approval process is imperative. As applications proliferate through the organisation, the CEO should ask every direct report how this capability is being applied, where it is being used, what decisions and processes are affected most by it, and who bears responsibility for the efficacy of its outcomes.

What monitoring mechanisms do we have in place? The board should ensure that management has implemented a process that assures GenAI model outcomes align with intended results and comply with relevant regulatory requirements. When this is not possible due to the complexity of the correlations in the model, embedding self-check mechanisms and conducting human reviews of AI-generated output are alternatives. Internal audit can also serve as a check and balance.

- Human oversight supported by automated alerts and controls is an integral part of any GenAI solution, particularly when the model is connected to hardware or software, or there is a significant impact on sensitive decisions (e.g., employment, healthcare services access or protection of vulnerable parties).

- As models inevitably change over time, they should be evaluated periodically for unreliable or biased content. Effective monitoring provides decision-makers with real-time alerts of trends indicating the emergence of anomalies or errors and enables continuous model improvement.

How do we set accountability in a GenAI model? Model owners, including those responsible for their design, development and operation, should be held accountable for their proper functioning. As noted earlier, the context for accountability means that, throughout their life cycle, GenAI models perform in accordance with their prescribed, intended role that conforms with applicable regulatory frameworks. The question for the board regarding these models is, "How do we know they are working as intended?"

Important takeaway: A GenAI model's performance should be supervised like an employee's performance. The aforementioned monitoring mechanisms and emphasis on sufficient transparency facilitate the supervisory process along with appropriate enforcement policies.

How do we manage the risks? GenAI takes issues resident in social media to a whole new level as its early implementations have exposed its shortcomings. While the summaries provided by the models are easy to read, the source and provenance of content are not always evident.

- Ownership rights are a major concern, as the legal and regulatory landscape shifts at lightning speed to address how original content is differentiated from proprietary content. How can one be sure that the output received from GenAI is not infringing on the intellectual property (IP) of others? Who owns the output from GenAI models — and can that output be copyrighted?

- Potential bias and prejudice in text and images is an issue.

- Images or videos created by GenAI can appear realistic but are deceptively false ("deepfakes"), making it almost impossible to discern truth from fiction. These models hallucinate or drift, meaning they can deliver results not backed by the data to which they have access.

These issues can lead to misinformation — inaccurate and misleading content — and blatant plagiarism. They can also lead to disinformation (e.g., fake news and mimicking people or specific individuals through falsified photographs, videos and voice scams). These issues also open the door to more sophisticated cyber threats and deceptive social engineering strategies. Fake product reviews, mimicked voices, and fake images and videos will likely become commonplace.

Unless the data set is controlled, the misinformation, falsehoods, opinions and conjectures so prevalent in social media and elsewhere could become part of the written record from which GenAI draws.

Yes, but: If a company uses its own data in conjunction with an appropriate governance framework, these issues become less of a risk. Unless the data set is controlled, the misinformation, falsehoods, opinions and conjectures so prevalent in social media and elsewhere could become part of the written record from which GenAI draws. Lack of data control also triggers the ownership and copyright issues mentioned above.

Important takeaway: Boards should ask management how these issues will be addressed.
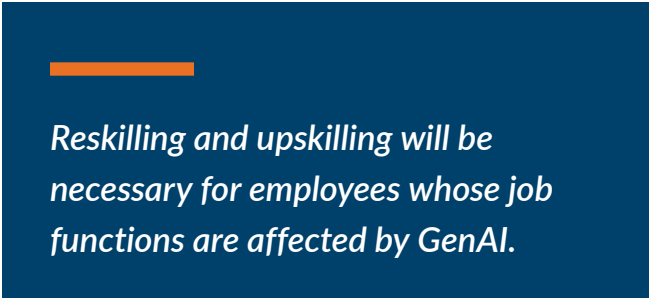
- Controlled data sets in a closed source model and requiring attribution of GenAI content can engender transparency and confidence. But it also introduces the risk of unintentional bias on the part of those who define the data sets and program the algorithms for creating content from the data.

- There is also the issue of managing data from external sources, such as strategic suppliers, third-party providers and channel partners.

What are the change management issues? With resistance to change a formidable risk[5] for many organisations and the deployment of GenAI proliferating, employees need to understand the ground rules for responsible and ethical use of the technology. Management should communicate GenAI's strengths and limitations, the intention to deploy the technology thoughtfully and with purpose, the initial use cases planned, and how those use cases align with broader strategic efforts, such as environmental, social and governance (ESG) and diversity, equity and inclusion (DEI) initiatives.

Reskilling and upskilling will be necessary for employees whose job functions are affected by GenAI. The policies and ground rules for the technology's use should align with applicable laws and regulations and the need to protect the company's IP (e.g., trade secrets and other confidential information). These rules and policies should also address the related impact on cybersecurity and privacy risks and reinforce monitoring protocols and accountabilities.

Why it matters: GenAI adds yet another disruptive force for business along with the COVID-19 pandemic, supply chain disruptions, workplace upheavals, talent shortages, inflation and rising interest rates. The good news is that off-the-shelf software interfaces of GenAI are so readily available it is almost equivalent to software as a service (SaaS).

> *Reskilling and upskilling will be necessary for employees whose job functions are affected by GenAI.*

The bottom line: Smart directors should prepare themselves for the journey.

GenAI is yet another wake-up call for board members. It is not the end game. Private equity firms are looking for the next OpenAI. Data scientists are experimenting with graphic databases that can lead to the next AI leap. The market can expect a plethora of AI startups. With the metaverse, digital transformation initiatives, quantum computing, ever-changing customer experiences, evolving regulatory landscape, and a strong focus on modernising legacy infrastructure to foster agility and innovation, the need to be technology-engaged in the boardroom is clear.

In this digital world framed by the internet, digital devices, smart devices, the cloud and ever-increasing connectivity, mobility, and computing power, directors rooted in the analogue age and unable or unwilling to make this transition need not apply.

---

[5] *Executive Perspectives on Top Risks 2023 and 2032,* Protiviti and NC State University's ERM Initiative, December 2022: www.protiviti.com/us-en/survey/executive-perspectives-top-risks-2023-and-2032.

How Protiviti Can Help

GenAI is a game changer. Across all industries, from technology to healthcare to financial services to consumer products, organisations are adopting it and other forms of AI, intelligent automation, and advanced analytics to improve processes, drive new business opportunities, and achieve and sustain competitive advantage. We partner with organisations to solve business challenges with cutting-edge AI services that improve speed, enhance precision and optimise customer experiences.

As companies progress on their digital transformation journey, AI is a key element that promises to deliver critical insights leading to accelerated innovation and success while protecting vital business assets. For example, our services improve efficiency, speed and reliability while enabling employees to focus on higher-value activities. We use machine learning (ML) to match candidates to upskill on high-demand skills, enabling them to excel in more rewarding roles. Our AI services team works with business domain experts to create custom solutions that offer data-driven insights to solve emerging business problems.

About the Author

Christine Livingston
Managing Director, Protiviti

Christine is responsible for Protiviti's technology AI/ML and innovation solutions. With over a decade of experience in AI/ML deployment, she has delivered hundreds of successful solutions, including many first-in-class AI-enabled applications. She has helped several *Fortune* 500 clients develop practical strategies for enterprise adoption of new and emerging technology, including the creation of capability-driven AI-enabled technology road maps. She focuses on identifying emerging technology opportunities, developing innovation strategies and incorporating AI/ML capabilities into enterprise solutions.

Contact Christine at christine.livingston@protiviti.com.

protiviti®