

Directors & Boards



Thought Leadership in Governance Since 1976
directorsandboards.com

Second Quarter 2025

Is AI Inevitable?

TOM PETRO on BOARD READINESS FOR AI

SAMANTHA KAPPAGODA on AI GOVERNANCE

ALEC CRAWFORD on DEEP FAKES AND MISINFORMATION

MEGHAN ANZELC on AI AND THE WORKPLACE

Board Governance of AI and Emerging Technologies

Directors must be proficient in their knowledge of AI so the board can be vigilant and effective in its oversight role.

BY SAMANTHA KAPPAGODA

Five days after the November 2022 release of generative AI (Gen AI) tool ChatGPT, Sam Altman, the CEO of OpenAI, tweeted there were already more than one million users, with subsequent estimates of approximately one hundred million monthly users two months

later. This precipitated discussions in corporate boardrooms regarding Gen AI.

While emerging technologies, such as Gen AI, might provide potential innovation opportunities for competitive advantage with enhanced return on investment (ROI), they also present a potential risk that adverse consequences might undermine apparent benefits. Corporate board oversight should consider unintended consequences with applicable trade-offs of prospective benefits relative to risks of AI adoption and deployment, with respect to specific strategies and use-case applications, e.g., research and development (R&D), customer interaction, operational efficiencies and cost reduction, with corresponding financial, operational, compliance and reputational risks.

U.S. regulation of AI is currently decentralized across the states, without overarching federal regulation. One such example of state regulation is the 2024 Colorado AI Act, which imposes obligations on providers and deployers of AI systems to employ reasonable care to protect consumers from known or reasonably foreseeable risks of

While emerging technologies, such as Gen AI, might provide potential innovation opportunities for competitive advantage with enhanced return on investment, they also present a potential risk that adverse consequences might undermine apparent benefits.

algorithmic discrimination arising from use of “high-risk AI systems” as applied to the provision of education, employment, financial services, essential government services, health care, housing, insurance or legal services.

The 2024 EU Artificial Intelligence Act (EU AI Act), the most comprehensive AI regulation to date, imposed new requirements by classifying AI systems into the following four risk-based categories:

- “Prohibited” or “unacceptable” risk
- High risk
- Those triggering transparency obligations
- General purpose systems

This act is applicable to all participants in the AI value chain (e.g., product manufacturers, providers, deployers, importers, distributors and authorized representatives) and further applies to providers and deployers outside of the EU, if the AI systems or outputs in question are utilized within the EU’s jurisdiction. When the EU AI Act was passed, legal experts anticipated it would spur adoption of AI governance and ethics standards by other jurisdictions, similar to the impetus provided by the EU General Data Protection Regulation (GDPR) governing privacy.

AI AND CORPORATE STRATEGY

There is ongoing debate as to whether the mandate for oversight of AI adoption lies at the board level or at the board committee level. There is a strong argument that the risks associated with emerging technology adoption are within the purview of the entire board, including risks to comparative market share due to adoption timing relative to both incumbent competitors and potential new entrants. Some companies have taken the approach of quickly establishing multidisciplinary AI task forces. Others have not advanced beyond iterative proofs of concept. And some have adopted a wait-and-see approach, solely observing from the sidelines.

If this author may shift, for a moment, to an operating perspective as a data scientist, it is readily apparent that AI cannot be implemented off-the-shelf and there is no one-size-fits-all panacea across companies and industries. On the contrary, AI adoption is extremely context-specific, depending on a variety of company-related factors, such as industry or sector, size and maturity, public or private, regulated or unregulated, domestic or international (operations, domicile or participants in the value chain), risk tolerance and capacity, and what type of proprietary data may be contemplated for use, as well as development and maturity of the technology. In addition, use-case identification should incorporate

risk rankings and ROI estimates, and assess whether development is to be accomplished organically (i.e., in-house) or inorganically (i.e., by acquisition or third party/white label).

Returning to the governance perspective, while the oversight of AI adoption is a board-level responsibility, this article focuses on specific examples of oversight for board committees, in addition to identifying some operative issues, although it is by no means comprehensive.

Audit, risk and technology committees. Oversight of AI adoption, when delegated to a committee, has been commonly addressed by extending the responsibilities of existing audit or risk committees, rather than by establishing de novo technology committees. Some substantive initial questions posed by the audit committee are whether current and future use cases for AI might impact financial reporting or other areas of audit committee oversight.

In addition, the designated committee should inquire as to whether the vendor and subcontractor management



and cybersecurity programs of the company encompass use of Gen AI. Further inquiry should address whether there are adequate, appropriate and robust guardrails for employee utilization of Gen AI, although it is unclear if such protections alone are sufficient to safeguard intellectual property and confidential information.

By way of illustration, in 2023, a multinational electronics company prohibited employee use of Gen AI after a software engineer employee exposed proprietary code to a widely used Gen AI tool. According to media accounts, the exposed code would have been ingested and employed to train the GenAI model and could not retroactively be deleted.

Furthermore, the applicable committee should inquire whether the company has a comprehensive, secure technology adoption policy that governs hybrid workplace environments, spanning employer-owned electronic devices deployed at work, personal devices at work, employer-owned devices at home (with or without the deployment of a VPN) and personal devices at home.

Compensation and human capital committees. Governance of human capital management is complex, particularly given shifts in the demographic composition of the labor force, as well as changes in employee preferences and the workplace after the COVID-19 pandemic. Human capital strategy must be integrated with AI strategy, particularly if the company intends to develop technology in-house, given the relative scarcity of appropriately skilled employees in domestic markets. In this context, this committee can pose questions about hiring priorities and strategies, as well as appropriate metrics and key performance indicators (KPIs) for hiring and managing scientific and technical employees (on-shore versus offshore), in addition to upskilling and reskilling the workforce as a whole.

Prior to the widespread use of algorithms, human resource professionals conducted or had input into interview candidate selection, hiring, salaries and advancement. Currently, many of these processes have become entirely or partially automated. Although this automa-

tion might increase processing speed and scale relative to human decision-making, algorithms may risk replicating or perhaps even amplifying human biases, disproportionately affecting protected classes. There is little regulatory oversight of automated employment tools, with the exception of New York Law 114 on Automated Employment Decision Tools (AEDTs), which prohibits employers and employment agencies from deploying AEDTs that have not undergone annual bias audits with public disclosure regarding such audits, and further requires providing certain notices to prospective or current employees.

During a recent panel entitled “Inhuman Capital: AI, the Workplace and the Board” at The Character of the Corporation in November 2024, I highlighted an illustrative example for this committee. Some years ago, a global technology company developed and deployed a machine-learning model to screen employment candidates for a software engineering position, by ranking those candidates. However, journalists subsequently reported gender bias in this model, which had been trained on a corpus comprised of resumes of primarily male employees in technical positions at the company. The model had apparently penalized candidates whose resumes listed degrees granted by women’s colleges and extracurricular activities such as “women’s chess club.” Deployment of this model was cancelled shortly thereafter.

Another concern is with respect to the bias implications of AEDTs incorporating the widely held viewpoint that gaps in employment history are sub-optimal for prospective employees. It is worth considering whether, without a “human in the loop,” an AEDT might provide lower rankings to those individuals or groups that have experienced unavoidable gaps in employment (due to raising children and absence of daycare or recurring health-related issues), therefore disproportionately affecting women or individuals with disabilities.

ESG committees. There has been much discussion regarding Gen AI use cases in ESG management, including accelerating data collection and curation, as well as providing advanced data analytics to assist in compliance and reporting. However,

there has been less emphasis on the issue of associated power consumption in the discourse.

Investment in data centers has increased substantially in recent years in the United States and other major economies. Training large-scale AI models requires extensive computational resources, with substantial corresponding increases in power and water consumption by data centers, as well as by data transmission networks. According to industry estimates, data centers utilize approximately 2% of global electricity usage, which will likely increase alongside accelerating Gen AI adoption. A recent report by the U.S. Energy Information Administration forecasts electricity consumption will continue to grow at a 2% year-on-year rate in 2025 and 2026, matching 2024 growth, following relatively little change in the prior two decades.

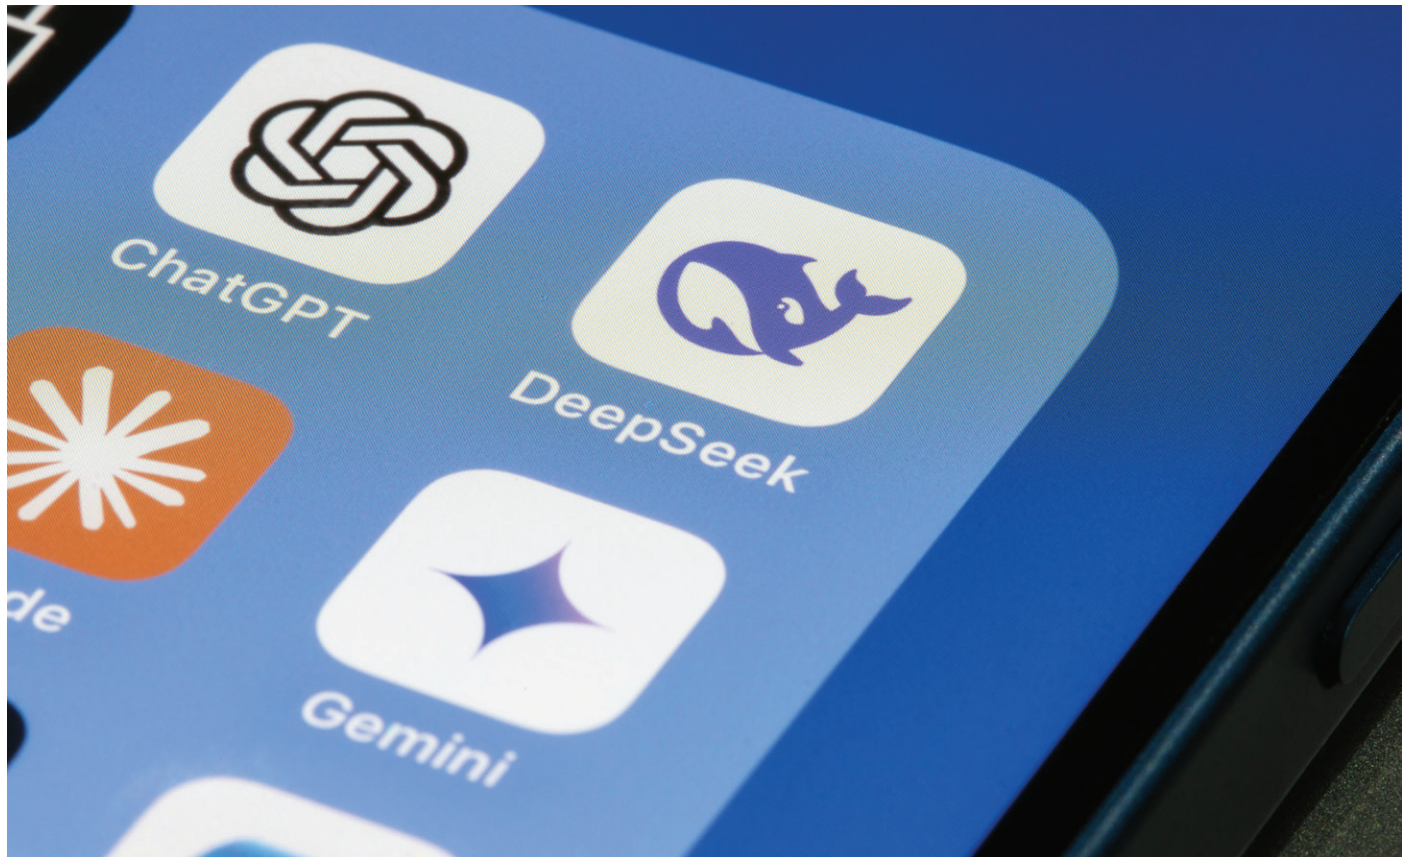
Governance and nominating committees. Optimal board composition is a perennial discussion for governance and nominating committees, as more recently has been whether the board has sufficient technology expertise to fulfill its obligations. If a seat in the boardroom were to be designated for subject matter expertise in AI or other emerging technologies, this raises the question of whether such a designated director should also be employed as a science and technology executive concurrent with board service, since expertise can quickly become stale, particularly in a rapidly evolving environment.

It also seems prudent for the entire board to strive for technology fluency — or at the very least, technology conversancy — rather than unduly rely on one director to lead all such discussions. In addition, when seeking independent external experts and consultants to provide periodic technology briefings, boards should consider whether such advice actually encompasses board-level governance or merely focuses on operational or research considerations.

2025 DEVELOPMENTS

There have already been several changes this year in the technology landscape, including policy, regulation and escalation in capital expenditure plans.

The prior administration’s executive order on AI, which had prioritized governance of AI devel-



opment, was rescinded. This was followed shortly by a new executive order on AI, signaling a shift to deregulation and competitive advancement, accompanied by the announcement of the Stargate Initiative, a \$500 billion public-private sector initiative for AI infrastructure development.

At the end of January, the Chinese AI firm DeepSeek released an open-source model (R1), with claims of matching or even outperforming competitors at a fraction of the development cost of the Open AI model. This event generated great interest, triggering a selloff in U.S. equity markets, while also raising concerns regarding intellectual property, as well as privacy and security associated with collection and transmission of personal data to China.

Subsequent to the release of DeepSeek R1, *The Wall Street Journal* reported the four largest technology firms would continue to increase capital expenditures (CapEx) after record outlays in 2024. Microsoft, Google and Meta were projecting combined CapEx of at least \$215 billion for their current fiscal years, corresponding to an annual increase of more than 45%. Meanwhile, Amazon was projecting total CapEx across its businesses growing to more than \$100 billion, with most of that increase in AI.

In February, DeepSeek stated plans to release a more advanced model (R2) by May 2025; Anthropic released Claude 3.7 Sonnet, which they described as their most “intelligent” model thus far; and OpenAI released GPT 4.5. In addition, the initial requirements of the 2024 EU AI Act came into effect, starting with prohibition of “unacceptable risk” AI systems.

The technology landscape has evolved rapidly in 2025, with prospective acceleration ahead. This demonstrates the need for boards to exercise continued vigilance in their oversight role and for directors to be proficient in their knowledge of AI and other emerging technologies, particularly in the context of oversight of implementation and deployment of high-risk and safety-critical use case applications. ■

Samantha Kappagoda is an independent director and nominating committee chair of Credit Suisse Funds, chief data scientist of Numerati Partners, visiting scholar at New York University Courant Institute of Mathematical Sciences and a member of the business board of the governing council at The University of Toronto.

The author thanks Michael Kwak, David Sharrow and David K.A. Mordecai for their help in the development of this article.