


DIRECTOR ESSENTIALS: Digital Transformation

A small yellow arrow pointing to the right, positioned to the left of the main title.

AI AND BOARD GOVERNANCE

In partnership with
Data & Trust Alliance



About this Report

How Boards Can Use This Resource

About NACD

About the Data & Trust Alliance

Introduction

1| Dedicated AI Governance Is Necessary

2| What Is AI?

3| Evaluating AI Opportunity and Risk

4| AI Organizational Oversight Focus Areas

5| AI Oversight Checklist for Boards

About this Report

Artificial intelligence (AI) applications are already impacting business and society. It's time for boards to oversee AI's unique risks and help their firms to seize the opportunities that AI presents. This edition of Director Essentials outlines the boardroom's role in AI governance and helps directors become conversant in the basics of the underlying technologies. It then introduces the operational and regulatory matters that can shape a company's AI journey.

How Boards Can Use This Resource

- ▶ Become familiar with governance and high-level technical essentials related to artificial intelligence.
- ▶ Understand why boards need to pay attention to AI's associated risks and opportunities.
- ▶ Learn important terminology and applications of the technology.
- ▶ Explore the operational and regulatory challenges associated with AI.
- ▶ Begin AI governance in identified organizational oversight focus areas.

© Copyright 2023, National Association of Corporate Directors. All rights reserved.

Except as permitted under the US Copyright Act of 1976, no part of this publication may be reproduced, modified, or distributed in any form or by any means, including, but not limited to, scanning and digitization, without prior written permission from NACD.

This publication is designed to provide authoritative commentary in regard to the subject matter covered. It is provided with the understanding that neither the authors nor the publisher, the National Association of Corporate Directors, is engaged in rendering legal, accounting, or other professional services through this publication. If legal advice or expert assistance is required, the services of a qualified and competent professional should be sought. This material is for authorized users only and is subject to NACD's [Terms of Use](#).



About this Report

How Boards Can Use This Resource

About NACD

About the Data & Trust Alliance

Introduction

1| Dedicated AI Governance Is Necessary

2| What Is AI?

3| Evaluating AI Opportunity and Risk

4| AI Organizational Oversight Focus Areas

5| AI Oversight Checklist for Boards



About NACD



The National Association of Corporate Directors (NACD) is the premier membership organization for board directors who want to expand their knowledge, grow their network, and maximize their potential.

As the unmatched authority in corporate governance, NACD sets the standards of excellence through its research and community-driven director education, programming, and publications. Directors trust NACD to arm them with the relevant insights to make high-quality decisions on the most pressing and strategic issues facing their businesses today.

NACD also prepares leaders to meet tomorrow's biggest challenges. The NACD Directorship Certification® is the leading director credential in the United States. It sets a new standard for director education, positions directors to meet boardroom challenges, and includes an ongoing education requirement that prepares directors for what is next.

With an ever-expanding community of more than 23,000 members and a nationwide chapter network, our impact is both local and global. NACD members are driven by a common purpose: to be trusted catalysts of economic opportunity and positive change—in business and in the communities we serve.

▶ **To learn more about NACD, visit nacdonline.org.**

About the Data & Trust Alliance

Data & Trust Alliance

The Data & Trust Alliance is a nonprofit organization formed in 2020 by CEOs of leading businesses to develop and adopt responsible data and AI practices. Alliance members include American Express, CVS Health, GM, IBM, Johnson & Johnson, Mastercard, Meta, Nike, Pfizer, Starbucks, UPS, Walmart, and Warby Parker.

Among the tools and resources developed by the Alliance are algorithmic anti-bias safeguards for workforce decisions and a new set of criteria for due diligence processes to help M&A teams evaluate data- and AI-powered businesses. The Alliance is currently developing cross-industry 'data provenance' standards to bring transparency to AI systems.

The Alliance is co-chaired by Ken Chenault, General Catalyst chair and managing director, and former American Express chair and CEO; and Sam Palmisano, former IBM chair and CEO.

▶ **To learn more about Data & Trust Alliance, visit dataandtrustalliance.org.**

About this Report

How Boards Can Use This Resource

About NACD

About the Data & Trust Alliance

Introduction

1| Dedicated AI Governance Is Necessary

2| What Is AI?

3| Evaluating AI Opportunity and Risk

4| AI Organizational Oversight Focus Areas

5| AI Oversight Checklist for Boards

Introduction

Artificial intelligence (AI) is a transformational technology with the potential to produce both significant value and harm for businesses and their stakeholders. With advances occurring rapidly amid widespread attention, it can be difficult for a company to cut through the hype to understand whether and how this technology can positively impact its business. However, most companies are likely already using some forms of AI in their operations through tools the organization has acquired or developed, or from services provided by suppliers.

AI poses unique opportunities and risks, both of which implicate boards' fiduciary responsibilities. Boards can continue to draw on the suite of governance principles,¹ IT governance frameworks, practices, and experiences honed through their oversight of cybersecurity, privacy, ethics and compliance, and emerging technology² to meet the challenges of AI governance.³

EMERGING TECHNOLOGY GOVERNANCE PRINCIPLES

1. Approach emerging technology as a strategic imperative, not just an operational issue.
2. Develop collective, continuous technology-specific learning and development goals.
3. (Re)align board structure and composition to reflect the growing significance of technology as a driver of both growth and risk.
4. Demand frequent and forward-looking reporting on technology related initiatives.
5. Periodically assess the organization's leadership, talent, and culture readiness for technological change.

On the other hand, artificial intelligence does require the board to shift its perspective, its engagement with management, and some of its practices. This publication outlines foundational AI knowledge, presents key board focus areas, and describes how the board can adapt to provide effective oversight. Boards will likely need to update committee responsibilities, increase AI-specific education, and engage more frequently and in greater depth with management in its evaluation of AI.



¹ NACD and Marsh McLennan, [Digital Transformation: How to Unlock Innovation and Growth with Emerging Technologies—A GNDI Quick Guide](#) (NACD and Marsh McLennan, 2022) p. 3.

² Steven De Haes, Laura Caluwe, Tim Huygh, and Anant Joshi, *Governing Digital Transformation: Guidance for Corporate Board Members* (Springer International Publishing, 2020).

³ NACD, [2023 Director's Handbook on Cyber-Risk Oversight](#) (Arlington, VA: NACD, 2023). Accessed August 14, 2023.

About this Report

How Boards Can Use This Resource

About NACD

About the Data & Trust Alliance

Introduction

1| Dedicated AI Governance Is Necessary

2| What Is AI?

3| Evaluating AI Opportunity and Risk

4| AI Organizational Oversight Focus Areas

5| AI Oversight Checklist for Boards

1 | Dedicated AI Governance Is Necessary



AI presents a challenge to traditional forms of governance due to its unique capabilities and the environment in which it is being developed. Once trained, AI systems can develop rules and propose choices without human input. As a result, traditional governance practices that are designed for oversight of technologies and decision-making processes that center on human judgment and use are not always fit for purpose.

Further complicating the challenge for boards as they oversee and govern AI use in their companies are three aspects of today’s AI systems:



1. AI technology advances exponentially and produces results at speed and scale.⁴



2. AI systems presently and generally lack transparency and “explainability” into the outputs, recommendations, and inferences they produce.⁵ This undermines trust and accountability, while also making it difficult to assure that model operations and outputs align with the company’s values, ethics, and purpose.⁶



3. There is currently no consistent regulatory regime to guide the technology’s development and use.⁷

⁴ Andrea Bonime-Blanc, “[Governing Exponential Tech: The Board’s Role in the Oversight of the Pandora’s Box of Generative AI](#),” *Directorship*, August 2023.

⁵ According to [c3.ai’s glossary](#), “explainability” is “...the concept that a machine learning model and its output can be explained in a way that ‘makes sense’ to a human being at an acceptable level.”

⁶ Mark Coeckelbergh, *AI Ethics* (Cambridge, MA: The MIT Press, 2020), pages 116–117.

⁷ Cecilia Kang, “[In US, Regulating AI Is in Its ‘Early Days.’](#)” *The New York Times*, July 21, 2023.



About this Report

How Boards Can Use This Resource

About NACD

About the Data & Trust Alliance

Introduction

1| Dedicated AI Governance Is Necessary

2| What Is AI?

3| Evaluating AI Opportunity and Risk


4| AI Organizational Oversight Focus Areas






5| AI Oversight Checklist for Boards



PREPARING FOR AI GOVERNANCE

NACD 2023 survey data reveals that 95 percent of directors believe AI will impact their business in some way in the next year.⁸ However, only 28 percent of respondents indicate that AI is a regular feature in board conversations.⁹ With boards anticipating a significant impact from AI on their businesses, boards must evaluate their own readiness and prepare themselves for this new oversight area in the five areas indicated below.¹⁰ After that, deeper learning is required, as described in this guide, concerning the technical concepts involved in sound management of AI.



 <p>Impacts on Strategy and Risk</p>	 <p>Committee Responsibility</p>	 <p>Board Composition</p>	 <p>Director Education</p>	 <p>Management Reporting</p>
Understand the company's current engagement with and use of AI and how to integrate AI into discussions of strategy and risk.	Leverage the strengths of existing committees to oversee the various components of AI where they would be most appropriately governed.	Evaluate board composition in context of changes stemming from AI. The technology may warrant new skills, experiences, and backgrounds on the board to provide effective oversight.	Invest in ongoing director-specific education, including management presentations, outside experts, personal engagement with AI research, or attendance at AI events.	Work with management to receive much more frequent, forward-looking updates about the company's AI initiatives.

⁸ NACD, [2023 Public Company Board Practices and Oversight Survey](#) (Arlington, Virginia: NACD, 2023), p. 3.

⁹ Ibid.

¹⁰ NACD, [Director FAQ: Board Oversight of Emerging Technologies](#) (Arlington, VA: NACD, 2018).

About this Report

How Boards Can Use This Resource

About NACD

About the Data & Trust Alliance

Introduction

1| Dedicated AI Governance Is Necessary

2| What Is AI?

3| Evaluating AI Opportunity and Risk

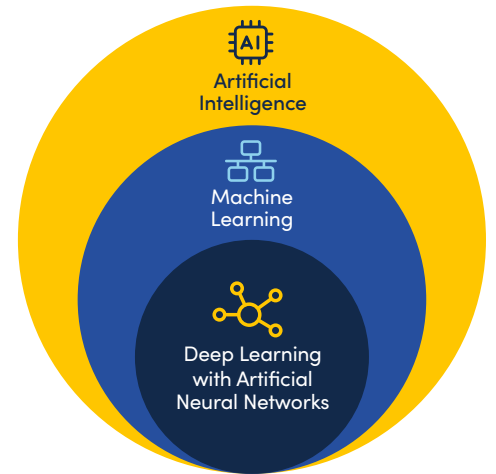
4| AI Organizational Oversight Focus Areas

5| AI Oversight Checklist for Boards



2| What Is AI?

Stanford professor John McCarthy coined the term artificial intelligence in 1955.¹¹ Since then, the field has evolved in multiple ways, propelled by advances in data collection and storage, chip design, and data analytics. The technology’s most significant advances have been driven by “machine learning” systems, which are not programmed but are trained on data sets and then continuously adapt. They “learn” by analyzing more data. This architecture, based on arrays of “neural networks” modeled on the neurons and synapses of the brain, is called “deep learning” (“deep” because it contains multiple layers of such neural networks).



WHAT ARE AI MODELS?

An AI model is “a learned representation of underlying patterns and relationships in data, created by applying an AI algorithm to a training data set. The model can then be used to make predictions or perform tasks on new, unseen data.”¹² Advances in machine learning systems have created models able to independently parse data, learn, and then make informed predictions based on what the model learned from the analyzed data. These models include the large language models (LLMs), which are also called “foundational” or “generative” models. Those models emerged at companies like Open AI, Deep Mind, Anthropic, Meta, IBM, and others.

HOW DOES AI WORK?

The foundation of AI performance is data, often from multimodal sources like text, speech, and images. The more and higher-quality data a system is fed, the better the AI. AI systems can only create as much value or be as trustworthy as the underlying data they are trained on and interact with while in operation.¹³

The use of deep learning with artificial neural networks is what currently gives foundational and generative AI models like ChatGPT their ability to produce content comparable to that created by humans, across a broad range of tasks.¹⁴ This technique also allows models to improve and scale at much greater speeds. The rapid improvement in AI capabilities and the continued accumulation of large amounts of data mean that performance can improve at a near-exponential scale. AI learns at the speed of data. This also means that models can quickly become obsolete that just months ago were best in class.¹⁵ This speed and scale of AI systems’ evolution both improves model performance across a range of applications and increases exposure to model risks and harms.

¹¹ Christopher Manning, Stanford University Human-Centered Artificial Intelligence, [Artificial Intelligence Definitions](#) (September 2020).

¹² See [“Key Terms for AI Governance,”](#) posted on iapp.org June 2023.

¹³ The Data & Trust Alliance and IBM, [Perspectives from the Front: Data Quality, AI Performance, and Trust](#) (New York City, New York: Data & Trust Alliance, September 2023).

¹⁴ Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin, [“Attention Is All You Need,”](#) revised on arxiv.org on Aug. 2, 2023.

¹⁵ Ryan McManus, [“The Atomic Unit of Strategy,”](#) posted on dialoguereview.com on May 27, 2023.

About this Report

How Boards Can Use This Resource

About NACD

About the Data & Trust Alliance

Introduction

1| Dedicated AI Governance Is Necessary

2| What Is AI?

3| Evaluating AI Opportunity and Risk

4| AI Organizational Oversight Focus Areas

5| AI Oversight Checklist for Boards



3 | Evaluating AI Opportunity and Risk

Many companies have already begun to adopt artificial intelligence and integrate it into their workflows and strategies. Whether your company is at the beginning of its journey or further down the path, board members should understand the key areas of AI opportunity and risk, along with the fundamentals for overseeing the technology.

STRATEGIC BENEFITS AND OPPORTUNITIES

As with most new and evolving technologies, AI presents significant opportunities to explore.¹⁶ It is imperative that directors understand these benefits, begin a dialogue about them with their management teams, and specifically ask questions about how, or if, AI can be a strategic asset to the company.

Directors who were in management during the rise of the Internet are likely to draw parallels between today's AI and Web 1.0. Boards must approach AI oversight as an organization-wide strategic imperative and not simply as an IT operational issue or focus exclusively on risks. It is a common pitfall in technology governance to lose sight of potential value to the customer or other stakeholders.¹⁷ This means that while directors do not need a deep understanding of AI operations, boards should always seek to understand how AI can generate value for the business and the customer and orient discussions on the topic to that end. This will spark a much more productive conversation about AI between management and the full board and can help the company proactively adopt beneficial AI tools as strategic enablers before they become strategic threats.

COMMON APPLICATIONS FOR AI

- ▶ **Process Automation:** Like prior forms of computing, AI systems can be used to drive efficiency and productivity through automation. AI's strengths in data processing and prediction make it a great candidate for automating many tasks across a range of business functions including HR, sales, marketing, legal, finance, supply chain, and IT.¹⁸
- ▶ **Improved Research and Development:** AI systems, particularly deep learning systems, have driven R&D and innovation across multiple industries over the past decade, and generative models promise to supercharge these advances. An example is "generative design" where AI is used in R&D to discover and develop new drugs and materials.¹⁹
- ▶ **Human-AI Pairing and Task Augmentation:** Research showed that human-AI pairing and task augmentation drove the most significant performance improvements across more than 1,500 companies. While many headlines draw on the fear of job loss to AI, augmentation may unlock opportunities for enhanced human creativity and productivity.²⁰

¹⁶ Michael Chui, Eric Hazan, Roger Roberts, Alex Singla, Kate Smaje, Alex Sukharrevsky, Lareina Yee, and Rodney Zempel, "[The Economic Potential of Generative AI](#)," posted on mckinsey.com on June 14, 2023.

¹⁷ NACD and Marsh McLennan, [Digital Transformation: How to Unlock Innovation and Growth with Emerging Technologies—a GNDI Quick Guide](#) (2022), p. 6.

¹⁸ Michael Chui, Eric Hazan, Roger Roberts, Alex Singla, Kate Smaje, Alex Sukharrevsky, Lareina Yee, and Rodney Zempel, "[The Economic Potential of Generative AI](#)" posted on mckinsey.com on June 14, 2023.

¹⁹ Adam Zewe, "[A Smarter Way to Develop New Drugs](#)," posted on news.mit.edu on April 26, 2022.

²⁰ H. James Wilson and Paul R. Daugherty, "[Collaborative Intelligence: Humans and AI Are Joining Forces](#)," *Harvard Business Review*, July-August 2018.

About this Report

How Boards Can Use This Resource

About NACD

About the Data & Trust Alliance

Introduction

- 1| **Dedicated AI Governance Is Necessary**
- 2| **What Is AI?**
- 3| **Evaluating AI Opportunity and Risk**
- 4| **AI Organizational Oversight Focus Areas**
- 5| **AI Oversight Checklist for Boards**




RISK & LIABILITY MITIGATION

It is the board's role to ensure that management is appropriately monitoring and mitigating the risks and harms in AI models. Algorithmic harm can be introduced via poor training-data quality, technical flaws in the algorithm, human error in development, or interacting with the algorithm outside intended use cases.²¹ Failure to address AI risks can result in further regulatory, reputational, or legal risks down the road, as well as failing to produce the expected value for the business.

It is important for boards to understand that managing risk in AI tools entails difficult decisions that will require tradeoffs. For example, optimizing a tool with privacy-enhancing features may sacrifice transparency about how the algorithm's explainability and transparency work, which could impact compliance with certain trustworthy or responsible AI standards. These are important decisions because the AI model does not have any context or prior understanding of the world in which it exists. Individuals who are impacted by the decisions generated by AI may then be vulnerable to outputs generated in unjust, unexplainable, and harmful ways. For example, a model may not have been trained to understand the concept of human or civil rights, and thus, the model may inadvertently discriminate against certain classes of people.²²

Boards as AI stakeholders must be prepared to ensure that the development and use of AI remains safe and aligned to the company's core values, ethics, purpose, and mission, all while taking the appropriate measures to mitigate harm. The organization's values and ethics should be documented and regularly referenced by the board when engaging with management in oversight of AI.



²¹ The Data & Trust Alliance and IBM, *Perspectives from the Front: Data Quality, AI, and Trust* (New York City, New York: Data & Trust Alliance, September 2023).

²² E. Tabassi, *Artificial Intelligence Risk Management Framework (AI RMF 1.0)* (Gaithersburg, MD: NIST, 2023) https://tsapps.nist.gov/publication/get_pdf.cfm?pub_id=936225 (Accessed August 14, 2023)

About this Report

How Boards Can Use This Resource

About NACD

About the Data & Trust Alliance

Introduction

1| Dedicated AI Governance Is Necessary

2| What Is AI?

3| Evaluating AI Opportunity and Risk

4| AI Organizational Oversight Focus Areas

5| AI Oversight Checklist for Boards



COMMON AI DEVELOPMENT AND USE RISKS

- i. **Bias:** During the process of developing and testing an AI model, the assumptions, preferences, and historical biases that exist in the training data or among developers can become embedded within the algorithm. This could lead to the production of discriminatory outputs by the AI.
- ii. **Explainability & Interpretability:** Because machine learning systems learn from their own experience, rather than being preprogrammed, models can be “black boxes,” meaning that it is not always possible to understand and explain how the inputs and operations within a model produced an output. Lack of explainability prevents users from challenging AI-influenced outcomes and could expose the company to litigation and reputational risk.²³
- iii. **Privacy:** AI models can jeopardize individuals’ privacy and introduce legal risk if insufficient care is taken to implement legally required privacy protections for the data used to train or operate the model.
- iv. **Security and Resilience:** AI that is entrusted to support or control critical systems—transaction processing, supply chain, financial operations, etc.—introduce new vulnerabilities to those systems if the AI fails, is hacked, or has insufficient human oversight. Harm can also be caused to individuals via exposure of their data from a vulnerability or breach in the AI system. AI models also suffer from their own vulnerabilities, such as prompt injection attacks, and may require new controls to protect the model's integrity and performance.²⁴
- v. **Misuse:** Once AI models have been released for public use there is the potential for misuse of the tool for malicious activities. Models have been used for activities like malware development, phishing campaigns, and the creation of disinformation and fake news. Models can also be used to produce untrue, synthetic data that can be fed back into the algorithm, which can then poison the model and render it ineffective and inoperable.²⁵
- vi. **Model Drift:** AI models continue to evolve and change as they are used. If not monitored effectively, models can begin performing outside their intended use which heightens the risk for stakeholder, reputational, and legal harm.

²³ Mark Coeckelbergh, *AI Ethics* (Cambridge, MA: The MIT Press, 2020), pages 116–117.

²⁴ OWASP prompt injection definition and page: https://owasp.org/www-project-top-10-for-large-language-model-applications/Archive/0_1_vulns/Prompt_Injection.html.

²⁵ Iliia Shumailov, Zakhar Shumaylov, Yiren Zhao, Yarin Gal, Nicolas Papernot, and Ross Anderson, “The Curse of Recursion: Training on Generated Data Makes Models Forget,” posted on ArXiv.org on May 31, 2023. <https://doi.org/10.48550/arXiv.2305.17493>.

About this Report

How Boards Can Use This Resource

About NACD

About the Data & Trust Alliance

Introduction

1| Dedicated AI Governance Is Necessary

2| What Is AI?

3| Evaluating AI Opportunity and Risk

4| AI Organizational Oversight Focus Areas

5| AI Oversight Checklist for Boards



4 | AI Organizational Oversight Focus Areas

Like digital transformation, AI impacts the entire organization and therefore requires a holistic approach to governance and oversight. Board governance of AI operations involves directors asking the right questions to ensure that the use and development of AI tools and systems remains aligned with the company’s overall mission, purpose, values, and ethics. Effective AI governance within the organization will blend strategy, corporate ethics, and values together with policies, procedures, and controls that guide responsible model development and use in the organization.

DATA GOVERNANCE & QUALITY

Data governance and quality assurance have long been established practices in many companies. However, their foundational importance takes on new significance with AI because models can operate and create results without human intervention. Thus, the performance, trustworthiness, and safety of the AI model is directly correlated with the quality of data on which it is trained. An organization is only able to develop trustworthy AI if the data used in training and operation of the model is of a sufficiently high quality.²⁶ In fact, tasks related to data preparation and selection such as cleaning, transforming, and splitting will likely require the majority of time in model development and training. Data quality is a complex topic. Boards should focus their questions to management on accuracy, completeness, consistency, uniqueness, provenance, and fitness for purpose.²⁷

The board should request that management define data quality standards that meet the organization’s needs, along with metrics and information to validate alignment of the data used in AI systems with these standards. The board should also ensure that structures and processes are in place to prepare and validate the quality of both internal and external data sources before they are used in AI models.

CAPITAL ALLOCATION

AI model development or tool procurement can represent a significant capital expenditure. Licensing or acquiring finalized, purpose-specific AI tools can be economical but may be less purpose specific, while developing, training, and deploying original AI models is expensive but could lead to tailored outcomes.

Boards should also recognize that capital allocation with AI is not limited to technology investments but has enterprise-wide scope. For instance, boards must ask management how much the organization may need to spend on workforce retraining or increased compliance costs—two nontechnical, significant investments required to make the application of AI successful. One area where the board can update its processes is M&A transactions. Where AI represents a significant component of a target company’s business model, boards should insist that the due diligence process incorporate evaluations of AI value drivers and risk factors—from algorithmic diligence and data diligence to cultural diligence.

²⁶ The Data & Trust Alliance and IBM, *Perspectives from the Front: Data Quality, AI, and Trust* (New York City, New York: Data & Trust Alliance, September 2023).

²⁷ *Ibid.*, pages 2–3.

About this Report

How Boards Can Use This Resource

About NACD

About the Data & Trust Alliance

Introduction

1| Dedicated AI Governance Is Necessary

2| What Is AI?

3| Evaluating AI Opportunity and Risk

4| AI Organizational Oversight Focus Areas

5| AI Oversight Checklist for Boards



Scrutinizing capital allocations is a core governance function, and boards should subject the expenses tied to AI to the same rigor as other capital investments. And, given the high cost of model training, scrutiny of AI expenses can be an opportunity to ask questions about the broader governance of AI development and deployment within the company. A lack of good governance and development practices will lead to wasted capital on an unsafe AI model that must be retrained or abandoned, leading to loss on investment or a prolonged timeline to return on investment.

AI GOVERNANCE AND MANAGEMENT FRAMEWORK

AI oversight can be adapted from existing frameworks that have worked across other functions. Common models that can be adapted for AI governance include the Three Lines Model,²⁸ which is frequently used in audit and cybersecurity, or a Hub and Spoke model.²⁹ Whatever framework is chosen, it should be coupled with new organizational structures, technical processes, and controls that direct the use and development of AI within the organization. Typical practices include the use of an AI impact assessment prior to and during model development or AI tool implementation, as well as periodic algorithmic audits to certify that the model remains safe.³⁰

The board of directors should validate whether management has a governance framework with accountability and controls implemented within the organization. By asking questions and validating the framework, the board will be directing the organization toward responsible engagement with AI with appropriate management and accountability.

ETHICAL AND TRUSTWORTHY AI USE AND DEVELOPMENT

Oversight of corporate ethics and compliance is a core fiduciary duty of the board.³¹ Thus, the board has an obligation to oversee and validate whether AI is used and developed in a way that aligns with the characteristics of trustworthy AI. Trustworthy AI systems include those that are “valid and reliable, safe, secure and resilient, accountable and transparent, explainable and interpretable, privacy-enhanced, and fair with harmful bias managed.”³² Several responsible and trustworthy AI frameworks have been developed, with many consistent principles across frameworks.³³

The board should direct management to develop, acquire, or implement only those AI systems that include the principles and characteristics of trustworthy AI and that also align with the company’s ethics and values. The goal of governance in the responsible development of safe systems is to instill trust, deliver benefits to the business, reduce harm to the greatest extent possible to all stakeholders, and align to the company’s core values, strategy, mission, and purpose.

²⁸ Jonas Schuett, “[Three Lines of Defense against Risks from AI](#),” (2022).

²⁹ The Data & Trust Alliance and IBM, [Perspectives from the Front: Data Quality, AI Performance, and Trust](#) (New York City, New York: Data & Trust Alliance, September 2023).

³⁰ Microsoft, [Microsoft Responsible AI Standard, v2 GENERAL REQUIREMENTS for EXTERNAL RELEASE](#) (Microsoft Corporation, 2022).

³¹ NACD, [Director Essentials: Strengthening Compliance and Ethics Oversight](#) (Washington, DC: NACD, 2015).

³² E. Tabassi, [Artificial Intelligence Risk Management Framework \(AI RMF 1.0\)](#) (Gaithersburg, MD: NIST, 2023) https://tsapps.nist.gov/publication/get_pdf.cfm?pub_id=936225 (Accessed August 14, 2023)

³³ Examples of frameworks include the NIST AI Risk Management Framework, OECD AI Principles, IAPP Trustworthy AI, and EU Ethics Guidelines for Trustworthy AI. Corporate examples can be found in Microsoft’s [Responsible AI Standard](#) and Google’s [Responsible AI Practices](#).

About this Report

How Boards Can Use This Resource

About NACD

About the Data & Trust Alliance

Introduction

1| Dedicated AI Governance Is Necessary

2| What Is AI?

3| Evaluating AI Opportunity and Risk

4| AI Organizational Oversight Focus Areas

5| AI Oversight Checklist for Boards



IMPACTS ON WORKFORCE

AI has the potential to change business models, automate entire functions within the organization, and augment the tasks performed by a majority of a company's employees. Research in this area has revealed that AI can significantly impact work tasks of up to 80 percent of US employees.³⁴ Likewise, IBM research estimates that "more than 120 million workers in the world's 12 largest economies may need to be retrained or reskilled."³⁵ Boards should request information from management on AI's impact on the company's employees while also asking questions about how AI could affect the company's workforce size; recruiting; efforts around diversity, equity, and inclusion; and costs for retraining and skill development.

Research has shown that many HR functions are already using AI for screening and recommending job candidates, recommending compensation and employee learning plans, and predicting attrition rates. Consequently, boards should inquire about the degree to which the company's HR function uses these systems and then update their human capital oversight practices to account for how AI is impacting the workforce. The board should also evaluate the organization's current workforce and leadership to assess their readiness for AI-driven organizational change and consider how AI may impact the company's future workforce needs for long-term success.

AI REGULATORY COMPLIANCE

AI regulation at present is nascent and highly fragmented. Significant regulations include New York City's recently passed law that will require companies to independently audit any AI tools that are used in hiring decisions.³⁶ Similar laws are proposed in New Jersey, California, Vermont, and Washington, DC.³⁷

The European Union's proposed AI Act, a law similar in scope and potential impact to the EU General Data Protection Regulation, proposes a classification system for AI tools with a variety of requirements and obligations tailored to a risk-based approach.³⁸ Tools would be authorized in the EU market only if they meet the requirements outlined in the proposed law. Parallel to these regulations, some companies are adopting voluntary commitments for the safe and responsible development and use of AI where no clear, national regulatory framework exists, such as the United States.³⁹

³⁴ Tyna Eloundou, Sam Manning, Pamela Mishkin, and Daniel Rock, [GPTs Are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models](#) (Aug. 22, 2023).

³⁵ Annette LaPrade, Janet Mertens, Tanya Moore, and Amy Wright, [The Enterprise Guide to Closing the Skills Gap Strategies for Building and Maintaining a Skilled Workforce](#) (Armonk, NY: IBM Corp., 2019).

³⁶ The New York City Council - [File #: Int 1894-2020](#). Posted on Legistar.council.nyc.gov. and accessed August 14, 2023.

³⁷ Steve Lohr, ["A Hiring Law Blazes a Path for A.I. Regulation,"](#) the *New York Times*, May 25, 2023, sec. Technology.

³⁸ European Parliament, [BRIEFING EU Legislation in Progress: Artificial Intelligence Act](#) (European Union, 2023).

³⁹ The White House, ["FACT SHEET: Biden-Harris Administration Secures Voluntary Commitments from Leading Artificial Intelligence Companies to Manage the Risks Posed by AI,"](#) posted on whitehouse.gov on July 21, 2023.; and the National Institute of Standards and Technology, [Artificial Intelligence Risk Management Framework](#) (NIST, January 2023).

About this Report

How Boards Can Use This Resource

About NACD

About the Data & Trust Alliance

Introduction


1| Dedicated AI Governance Is Necessary

2| What Is AI?

3| Evaluating AI Opportunity and Risk

4| AI Organizational Oversight Focus Areas

5| AI Oversight Checklist for Boards



Even though the regulatory environment for AI remains complex and emergent, boards should take a proactive approach and strive to remain ahead of AI regulatory requirements. The board should work with management to anticipate the impact of upcoming AI regulations on the business, with the understanding that this space is constantly evolving and tends to eventually align with what emerges from the European Union. Compliance failures represent a significant risk, and boards have a responsibility to oversee the company's compliance function. Fulfilling this responsibility for AI regulation will be challenging, but regulations may become a lever to ensure that companies are engaging with AI systems safely and responsibly.

AI REGULATORY APPROACHES TO MONITOR

As regulations develop, it can be helpful for the board to assign responsibility for tracking AI regulatory matters to a chief legal officer or general counsel.

- ▶ **Horizontal regulation of AI reaches across sectors and use cases.** The leader in horizontal regulation is the European Union, whose AI Act classifies several AI use cases as “high-risk” and subjects them to a set of stringent compliance requirements.
- ▶ **Vertical regulation addresses specific sectors or use cases.** Examples include the July 2023 New York City law that will require employers to conduct bias audits of AI and other algorithms for recruiting, hiring, or promotion.
- ▶ **Voluntary approaches involve nonbinding guidelines and best practices, combined with enforcement of existing laws.** Examples of this approach include the US federal government's development principles in the “[Blueprint for an AI Bill of Rights](#),” the National Institute of Standards and Technology's (NIST) [AI Risk Management Framework](#); and warnings from [several federal regulators](#) that they will use their preexisting enforcement authority.



About this Report

How Boards Can Use
This Resource

About NACD

About the Data & Trust
Alliance

Introduction

1| Dedicated AI
Governance Is
Necessary

2| What Is AI?

3| Evaluating AI
Opportunity
and Risk

4| AI Organizational
Oversight Focus
Areas

5| AI Oversight
Checklist for Boards



5 | AI Oversight Checklist for Boards

1. Determine to what degree the company currently engages with AI throughout the business.
 - Which business units and functions use AI tools?
 - In what way are the tools used?
 - What type of tools are being used? Are we building them, or were they procured?
 - Is our management team aware of how and when employees are using publicly available AI tools and models to complete tasks?
2. Discuss AI with management to understand how they are thinking about the technology.
 - What use cases are we exploring?
 - How does AI impact the company's strategy and approach to risk?
 - What organizational gaps hinder the development of AI that may give us a strategic advantage?
3. Integrate AI into board strategy and risk discussions.
 - What are AI's potential impacts on company strategy, business model, and industry?
 - Who within the organization should report to the board on AI matters? Do we have a consensus on metrics and information the board would like to see regarding AI initiatives?
 - How often does the full board need to be reported to on these matters?
 - Which committee will perform primary oversight of AI, if any?
4. Discuss potential changes to oversight structures, processes, or practices related to oversight focus areas:
 - Data quality
 - Capital allocation
 - AI governance and management framework
 - Ethics oversight and responsible development
 - Workforce impacts
 - Regulatory compliance
5. If we are using AI within the organization, is the company compliant with all laws currently governing the use of AI and monitoring legal developments to ensure compliance with upcoming regulations and rules?
6. Through which compliance framework is the company testing its models? (See: NIST AI Framework, etc.)
7. Evaluate the board's structures, practices, and composition to determine if new AI expertise may be needed to oversee the organization's strategy with regard to AI.
 - Are investments in AI-specific education for the board necessary?
 - Are board structures and processes appropriate to ensure responsibility for AI?
 - Have we delegated responsibility to a committee to oversee this issue?
 - Which components of AI oversight are best handled by different committees? (e.g., risk, algorithmic auditing, workforce impacts, regulatory and compliance oversight, etc.).
 - Do committee charters require updates to reflect these new responsibilities?



1515 N. Courthouse Road

Suite 1200

Arlington VA 22201

Phone 571-367-3700

▶ nacdonline.org

