

What is AI, and what ethical risks does it present?

AI can be defined in several ways, but perhaps the most straightforward definition is the use of machines—that is, computers—to execute tasks that would otherwise require human intelligence. AI is also known by a variety of names, such as machine learning, deep learning, natural language processing, and computer vision. Many of us experience AI on a daily basis—for example, when we have to contend with a service provider's automated answering system to reach someone or obtain information.

As is typical of many aspects of technology, AI offers numerous benefits, including speed, consistency, and reduced labor costs. And as more and more enterprises use AI, and the use of AI expands into new areas, companies may find it increasingly important to use AI to avoid the competitive disadvantages of not using it.

However, AI poses risks as well as benefits. Some of these risks are unrelated to ethics, ranging from minor ones such as the failure of speech recognition to accommodate accents to major ones such as contamination of a water supply or the shutdown of a power grid. Real-world examples of AI gone awry include systems that discriminate against people based on their race, age, or gender, and social media systems that inadvertently spread disinformation.

Yet with AI now becoming a required business capability—not just a “nice to have”—companies no longer have the option to avoid AI's unique risks simply by avoiding AI altogether. Instead, they must learn how to identify and manage AI risks effectively. In order to achieve the potential of human and machine collaboration, organizations need to communicate a plan for AI that is adopted and spoken from the mailroom to the boardroom. By having an ethical

framework in place, organizations can create a common language by which to articulate trust and help ensure integrity of data among all of their internal and external stakeholders.

Of course, there are a number of risks that *are* ethical in nature. Like human intelligence, a critical component of AI is information, or data; without data, decisions would be made, and actions would be taken arbitrarily and without any logical basis. The following are just a few of the ethical and governance challenges associated with the use of AI:

- Multiple definitions of AI and related terms across the organization and its ecosystem
- Limited focus or alignment of AI purpose to the company's mission and values
- AI is deployed for narrow objectives without considering a broader aperture of how it can change the business for better
- Development of AI happens in an ad-hoc manner with limited standards or guardrails
- Data for AI design and development is acquired or used without any checks or testing for bias or authenticity
- AI is developed with a sole or primary focus on improving efficiency
- Outcomes of AI systems are not monitored for alignment with intended objectives

Having a common framework and lens to apply the governance and management of risks associated with AI consistently across the enterprise can allow for faster, and more consistent adoption of AI.

Some examples of how AI can be used, and the ethical and related factors arising from its use, are outlined below. 

Impacted stakeholder(s)	Employees	Customers	Regulators/Communities
Use case	Hiring	Credit decisioning	Autonomous driving
How AI can be used	<ul style="list-style-type: none"> • Screening/identifying right candidates for the job • Reduced human-introduced bias 	<ul style="list-style-type: none"> • Screening of customers for credit product applications • Identifying appropriate products for cross-/up-selling to customers 	<ul style="list-style-type: none"> • Self-driving vehicles—Cars or trucks that self-navigate and drive to a selected destination based on dynamic street conditions
Examples of potential risks	<ul style="list-style-type: none"> • Cyber-snooping and black-box concerns (AI explainability) • Bias based on past hiring trends, discriminatory data sets, etc. • If discrimination is proven, companies could be found liable under Title VII of the Civil Rights Act of 1964 or other laws prohibiting discrimination 	<ul style="list-style-type: none"> • Incorrect rejection of customer applications for credit products • Incorrect recommendations of products for cross-/up-selling purposes • Loss of business opportunity and customer trust • Improper usage of customer data, including scenarios where customers are not aware on how their data is used 	<ul style="list-style-type: none"> • Possible accidents and/or possible fatalities • Exploitation of vulnerabilities that lead to ineffective controls • Inability to adapt to new environments due to lack of training diversity
AI risk category	<ul style="list-style-type: none"> • Fairness 	<ul style="list-style-type: none"> • Trust and transparency 	<ul style="list-style-type: none"> • Safety



Why the board?

The ethical considerations outlined previously suggest the need for some type of oversight, but why is the board the right party to provide it?

The reasons seem straightforward. First, strategy and risk are among the key areas of board oversight. Second, AI is increasingly a critical tool in advancing and supporting strategy, but it carries risk. Thus, AI ethics oversight by the board is appropriate, even necessary. The board's oversight of AI can drive the right tone and set the guiding principles for an ethical framework for AI adoption that can be operationalized by the company to achieve responsible outcomes.

Overseeing AI ethical risks

Areas of risk

From the board perspective, some primary areas of ethical risk include:

- **Fairness:** Will the use of AI result in discriminatory outcomes? Do AI systems use datasets that contain real-world bias and are they susceptible to learning, amplifying, and propagating that bias at digital speed and scale?
- **Transparency and Explainability:** Are AI systems and algorithms open to inspection and are the resulting decisions fully explainable? How is the use of AI that leverages individual data communicated and explained to impacted individuals before, during, and after business interactions?
- **Responsibility and Accountability:** Who is accountable for unintended outcomes of AI? Does the organization have a structured mechanism to recognize and acknowledge unintended outcomes, identify who is accountable for the problem, and who is responsible for making things right?
- **Robustness and Reliability:** What measures for reliability and consistency do AI systems need to meet before being put into use? What are the processes to handle inconsistencies and unintended outcomes?
- **Privacy and Trust:** Do AI systems retain and/or generate trust with customers, employees, and other external stakeholders? Does the AI system generate insights and actions for individuals that they do not expect, leading to concerns and questions of trust and propriety?
- **Safety and Security:** Do AI results help to maintain or increase safety, and have they been tested for errors in controlled environments? Have risks to human life, social, and economic systems been identified and mitigated?

As suggested previously, the risks flowing from AI can harm a wide range of companies in a variety of ways. Customer-focused businesses can be faced with customer frustration or alienation if customer experiences incorporating AI do not perform as expected, such as experiencing session disruptions or nonmeaningful responses to unique and/or sensitive customer questions. Such experiences can directly result in a loss of trust from customers, leading to reputational damage. Regulators can find companies to be in violation of laws or simply adopting predatory or exploitative practices when they use data or algorithmic approaches that do not incorporate ethical checks. Suppliers can experience lost time and product spoilage if “just-in-time” inventory systems using AI do not perform as expected. And so on. In other words, faulty AI systems and processes can negatively—and seriously—impact companies.

It is also noteworthy that the areas of risk outlined previously may involve perception as much as, if not more than, reality. Even if an AI “system” is operating well, the perception that it is, for example, generating unfair or unreliable outcomes can be as damaging as actual unfair or unreliable outcomes. Thus, it is important that the board address AI ethical risks from both perspectives—1) principles- and values-driven and 2) operating standards.

Managing risk

While the ethical risks associated with AI may differ from other risks in some respects, they have some common elements that many boards are accustomed to handle. These elements include the following:

- Determining where and how AI is used throughout the organization
- Evaluating whether the use of AI is appropriate in the circumstances and whether AI is yielding desired benefits, without increasing the organization's risk exposure
- Assessing the ethical and other risks associated with the company's use of AI and setting, or overseeing the setting of, appropriate “guardrails” regarding the use of AI
- Becoming familiar with the leadership of each part of the organization that uses AI, as well as the individuals with direct responsibility for AI
- Evaluating whether the leadership and resources allocated to AI are adequate and, if not, where any inadequacies exist
- Working with management to provide any needed resources and to address any inadequacies
- Considering the engagement of independent advisers to assist the company and/or the board to define and maintain a robust approach to AI ethics and/or to supplement the board's skill sets
- Periodically revisiting the above elements to determine whether “tweaks” or more substantial adjustments are called for ➔

These and other actions taken to manage AI ethical risks can generally be undertaken as part of an ongoing, integrated risk assessment process conducted by the board, rather than on a separate, stand-alone process. For example, periodic reviews of AI and associated risks might be added to the items routinely addressed in the enterprise risk management process, and the people responsible for overseeing the use of AI might be included in periodic talent reviews.

Other matters associated with AI ethical risk oversight include where responsibility for such oversight resides. Does the full board exercise this oversight responsibility, or does it more properly reside within a committee of the board—and, if so, which committee? The answer may differ widely among companies, depending upon their respective industries, how they use AI, and the nature of the skill sets possessed by the board and its committees. Consequently, each board needs to make these decisions based upon the individual characteristics of the company, its business model, and its board members.

In addition, companies may need to consider the composition of their boards and/or committees based in part upon the nature and extent of their usage of AI. A company that makes extensive, impactful use of AI might consider seeking a board member who has some degree of familiarity with AI and/or data science. As noted above, a potential alternative to having board expertise in the area is the engagement of advisers or consultants to assist the board in overseeing AI ethics and/or supplementing the board's existing skill sets.

See "Questions for the board to consider asking" for suggestions as to areas of inquiry for boards concerning the use of AI and associated ethical risks.

Conclusion

As with many other aspects of technology, AI is becoming indispensable to companies that are focused on long-term growth and value generation. AI is also increasing the impact of data risks and generating new risks, such as those from unintended consequences. Appropriate oversight by and guidance from the board can help to identify, assess, and manage these risks. ➤



Questions for the board to consider asking:

1. How can AI impact our business now and in the future? What is our AI strategy?
2. What is our approach to AI governance? How are we driving trust in our company's use of AI?
3. What are our principles and/or framework to deploy and use AI in a responsible way? Have we communicated these internally and externally? How are they being embedded in AI initiatives across the business?
4. Who oversees the use of AI? Does that person or group have adequate and appropriately skilled resources?
5. What guardrails have we established to address the challenges associated with ethics and governance of AI?
6. Are our uses of AI appropriate? Are they achieving the desired results? Has the use of AI created unanticipated risks, including ethical challenges?
7. Do we have adequate skill sets on the board or management team to properly oversee the use of AI? Do we need to seek out director candidates with relevant skill sets?
8. How are we collaborating with our ecosystem of business partners, suppliers, customers, regulators, and other constituents to align on approaches to trustworthy AI?



Authors



Vic Katyal
Principal and Chief Operating Officer
Risk & Financial Advisory
Deloitte & Touche LLP
vkatyal@deloitte.com



Cory Liepold
Principal
Risk & Financial Advisory
Deloitte & Touche LLP
cliepold@deloitte.com



Satish Iyengar
Senior Manager
Risk & Financial Advisory
Deloitte & Touche LLP
siyengar@deloitte.com

Deloitte AI Institute



Nitin Mittal
Principal, AI Strategic Growth Offering Leader
Analytics & Cognitive Leader
Deloitte Consulting LLP
nmittal@deloitte.com



Irfan Saif
Principal, AI Strategic Growth Offering Leader
Risk & Financial Advisory
Deloitte & Touche LLP
isaif@deloitte.com

Contact us



Carey Owen
National Managing Partner
Center for Board Effectiveness
Chief Talent Officer, Risk & Financial Advisory
Deloitte & Touche LLP
coven@deloitte.com



Maureen Bujno
Managing Director and Audit & Assurance Governance Leader
Center for Board Effectiveness
Deloitte & Touche LLP
mbunjo@deloitte.com



Debbie McCormack
Managing Director
Center for Board Effectiveness
Deloitte LLP
dmccormack@deloitte.com



Krista Parsons
Managing Director
Center for Board Effectiveness
Deloitte & Touche LLP
kparsons@deloitte.com



Bob Lamm
Independent Senior Advisor
Center for Board Effectiveness
Deloitte LLP
rlamm@deloitte.com



About this publication

This publication contains general information only and Deloitte is not, by means of this publication, rendering accounting, business, financial, investment, legal, tax, or other professional advice or services. This publication is not a substitute for such professional advice or services, nor should it be used as a basis for any decision or action that may affect your business. Before making any decision or taking any action that may affect your business, you should consult a qualified professional adviser. Deloitte shall not be responsible for any loss sustained by any person who relies on this publication.

About the Center for Board Effectiveness

Deloitte's Center for Board Effectiveness helps directors deliver value to the organizations they serve through a portfolio of high quality, innovative experiences throughout their tenure as board members. Whether an individual is aspiring to board participation or has extensive board experience, the Center's programs enable them to contribute effectively and provide focus in the areas of governance and audit, strategy, risk, innovation, compensation, and succession.

About Deloitte

Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited, a UK private company limited by guarantee ("DTTL"), its network of member firms, and their related entities. DTTL and each of its member firms are legally separate and independent entities. DTTL (also referred to as "Deloitte Global") does not provide services to clients. In the United States, Deloitte refers to one or more of the US member firms of DTTL, their related entities that operate using the "Deloitte" name in the United States and their respective affiliates. Certain services may not be available to attest clients under the rules and regulations of public accounting. Please see www.deloitte.com/about to learn more about our global network of member firms.