GLOBAL KNOWLEDGE BRIEF

The Artificial Intelligence Revolution

Part 3: Internal Audit's Role in AI Ethics

The Institute of Internal Auditors

Contents

About the Experts

Andrew Clark, Ph.D., CAP, GSTAT

Andrew is co-founder and chief technology officer at Monitaur. A trusted domain expert on the topic of ML auditing and assurance, he built and deployed ML auditing solutions at Capital One. He has contributed to ML auditing standards at organizations including ISACA and ICO in the UK. Before Monitaur, Andrew also served as an economist and modeling advisor for several very prominent crypto-economic projects while at Block Science.

Jim Enstrom, CIA, CRISC, CISA

Jim is senior vice president and chief audit executive, internal audit, at Cboe Global Markets, Inc. An accomplished business leader, he has extensive audit, compliance and risk management experience in areas such as financial reporting, business operations, and information technology. Prior to joining Cboe in 2009, Jim spent 13 years in public accounting, having worked at Arthur Andersen and Deloitte.

Tim Lipscomb

Tim is senior vice president, chief technology officer for Cboe Global Markets, Inc. He oversees software engineering and quality assurance for Cboe equities, options, and futures markets, as well as its Data and Access Solutions business. Previously, Tim was chief operating officer of Cboe Europe, where he oversaw the company's software engineering, infrastructure, and operational teams.

Ellen Taylor-Lubrano, Ph.D.

Ellen is machine learning team lead in the regulatory division of Cboe Global Markets, Inc. She joined Cboe in 2020 as the founder of the regulatory division's ML program, which applies ML/AI in the surveillance of financial markets. Prior to that, Ellen worked in fundamental scientific research and production software development.

INTRODUCTION

Amid rapid advancements in artificial intelligence (AI), concerns about ethics and related issues have prompted some to recommend a hiatus or slowdown in further development.[1] But despite calls for temporary halts, many organizations are ramping up AI use or planning to do so. Internal auditors will clearly have an important assurance and advisory role as organizations wrestle with AI choices and their implications.

Previous briefs in this series have focused on what internal auditors need to understand about AI and have revisited a landmark publication on the topic, The Institute of Internal Auditors' (IIA) *Artificial Intelligence – Considerations for the Profession of Internal Auditing*. Although it was published in 2017, this framework generally remains relevant and useful in most internal audit areas. "Internal audit can help an organization evaluate, understand, and communicate the degree to which artificial intelligence will have an effect (negative or positive) on the organization's ability to create value in the short, medium, or long term," according to the framework[2].

This third and final brief in the AI series addresses the ethical issues surrounding this multifaceted technology and what those issues mean to organizations and internal auditors. This brief also includes recommendations and insights from management and internal auditors already working on the frontlines of AI use.

[1] https://futureoflife.org/open-letter/pause-giant-ai-experiments/
[2] *Artificial Intelligence - Considerations for the Profession of Internal Auditing, Special Edition*, The Institute of Internal Auditors, 2017.

Risks and Opportunities

Internal Audit's Role as Adviser

Excitement Over AI Could Overshadow Ethical Considerations

The global artificial AI market size was valued at $136.55 billion last year and was expected to grow by a 37% compound annual growth rate from 2023 to 2030, according to Grand View Research, Inc.[3] This surge in interest, and the excitement and hype surrounding technologies such as generative AI, have spurred many software developers and organizations to rush ahead in their AI research or efforts. However, amid rapid advancements, many serious and distinct risks, including ethical and performance issues, may be overlooked. Internal auditors are well positioned to alert their organizations to these issues and to offer advice on the efficacy of current controls and the need for enhanced controls or guardrails. Indeed, the Partnership on AI, led by Google executives, published a paper that calls for internal audit to play a leading role in providing assurance over the processes involved in AI creation and deployment and ensuring they meet ethical expectations and standards.[4]

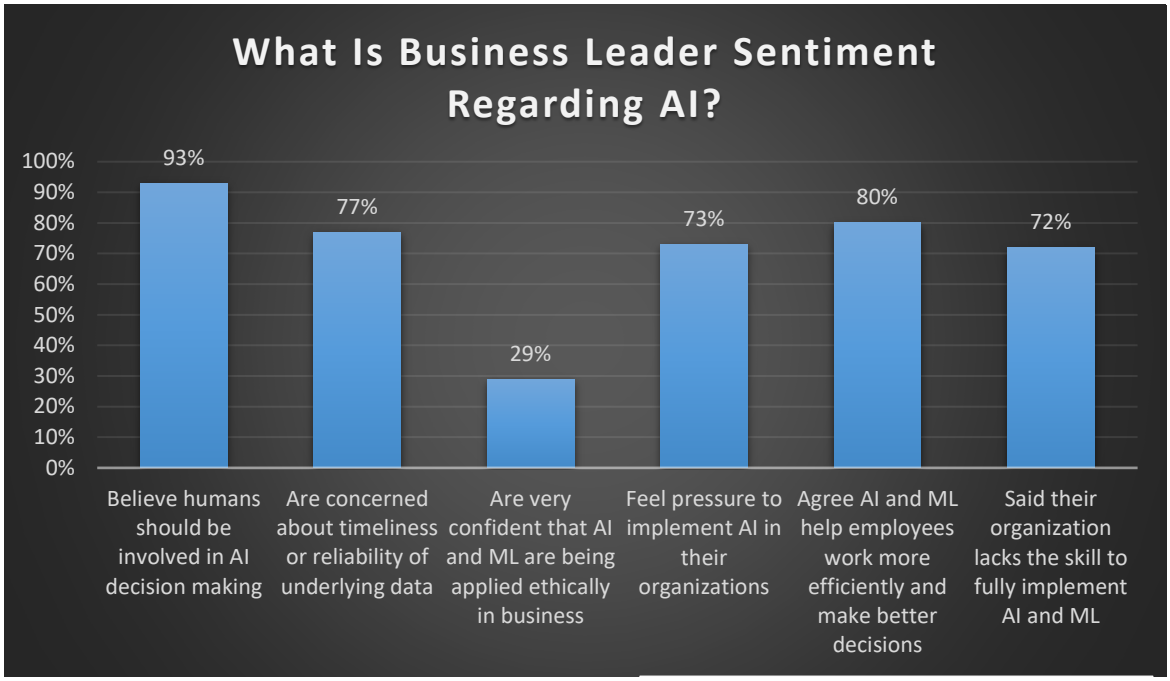

Source: Workday Survey, June 2023.

With that in mind, it's important for organizations and internal auditors to understand AI's risks and limitations, and what impact they might have on a business's use of AI. "There's a misconception that AI is really smart," said Andrew Clark, co-founder and CTO, Monitaur, an AI governance software company. Unfortunately, generative AI, which is receiving much of the current focus among the media and organizations, is only as smart as the data that it has been trained on and, at least

[3] *Artificial Intelligence Market Size, Share & Trends Analysis Report By Solution, By Technology (Deep Learning, Machine Learning), By End-use, By Region, And Segment Forecasts, 2023 – 2030*, Grand View Research, Inc., June 2023.
[4] "Closing the AI Accountability Gap: Defining an End-to-End Framework for Internal Algorithmic Auditing," The Partnership on AI, January 2020.

in the technology's early stages, that training may include random social media posts, web content, and other material that has not been authenticated.

Use of public generative AI programs can expose private or confidential company, customer, or business partner data. And because generative AI is so easy to use, these capabilities are accessible to everyone from veteran cybercriminals to amateur hackers. While cybersecurity efforts can mitigate some of the potential damage these efforts can cause, awareness of the elevated risk, from organizational level to the individual employee, is essential for proper cybersecurity.

Generative AI can also incorporate intentional or unconscious biases. When a regulatory organization works to identify problematic activity, for example, there are ethical and legal considerations about whether the data approaches being used might be biased against certain members or types of trading activity, noted Ellen Taylor-Lubrano, machine learning team lead--regulatory, Cboe Global Markets. On another front, researchers have found high error rates in using AI facial recognition systems to identify people of color, women, and young people, making misidentification more likely and increasing the chances for people to be wrongly accused of crimes. AI may also be subject to knowledge gaps and inaccuracies. For example, while AI systems can be trained to detect illnesses, they may not recognize a disease such as melanoma in someone with skin characteristics that were not included in its original data set.[5]

Current generative AI models also are not transparent about their sources, so without knowing the origins of the information it generates, users may expose themselves to legal, copyright, and intellectual property risks. Equally alarming, it may produce "facts" that the system has made up (called hallucinations) when trying to respond to a prompt. Generative AI, "is meant to mimic a human, not to be correct," Clark said. Internal auditors can advise organizations on the best ways to address such errors or omissions or their unintended consequences.

The user friendliness of generative AI can be another risk for organizations. In the past, models were typically built by people with advanced degrees or knowledge of systems who had expertise in automating those models, Clark said. Today, it's possible for people with little or no understanding of models, systems, or the data they are using to leverage a tool such as generative AI and ask it to make a prediction or a decision using information that may be incomplete or lacks proper context.

In addition to monitoring potential concerns with internal use of AI, organizations should also consider external threats. The same models behind technologies such as ChatGPT can be used to create tools that can produce malicious software and code, scam pages, and phishing emails. They can also be used to identify organizational vulnerabilities, as well as train new types of cybercrime tools, among other functions.[6] What's more, AI could make it easier for hackers to develop malware that can steal data or exert control over it.

While these threats may sound daunting, there are also risks in failing to embrace AI. If others surge ahead in AI use, an organization may be perceived as less tech-oriented or future-focused by current or potential customers or talent, giving competitors an advantage. AI also offers tangible benefits that can enable companies to streamline and enhance processes, thereby boosting productivity, improving customer service, minimizing costs, and potentially opening new service, market, or product opportunities. In addition, in many situations AI can help organizations identify risks or threats or spot new opportunities. AI may offer organizations access to a huge internal knowledge base faster and more efficiently than a straight search would do, according to Tim Lipscomb, senior vice president and chief technology officer at Cboe Global Markets. If an organization is using a manual information-gathering process, it may not be able to make the best decisions or respond to threats or opportunities as would be the case if it were using AI.

[5] "Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification," Joy Buolamwini and Timnit Gebru, *Conference on Fairness, Accountability and Transparency*, 2018.
[6] "Surge in Generative AI Tools for Cybercrime Sparks Concerns," GRC Report, August 10, 2023.

Turn to Fundamental Auditing Concepts

Adapting Three Lines and Other Existing Models

Using Fundamental Assurance Approaches for New Technology

While a technology may be new, many of the details of putting it to work may not be. For example, decision models and machine learning have long been used in the financial sector, noted Jim Enstrom, senior vice president and chief audit executive, Cboe Global Markets. (See the sidebar on "Minding Your Model Risk Management" on page 7.) IT auditors have had to address a myriad of risks, including ethical uses in the past, and AI is no different in this regard. It is critical, therefore, to ensure internal audit has a seat at the table to understand the strategic use of AI within the organization.

If we view AI systems through the lens of software development processes, internal auditors can go back to fundamental concepts, Enstrom said. Traceability, for example, should be a consideration if the AI system will be making decisions or working autonomously, while auditability will also be key. Just as internal auditors work with teams across their organizations to understand their work, the internal audit team will also have to work with engineers, data scientists, and programmers to understand what the systems are doing, the sources of data used as inputs, what requirements were used to build the model, and what artifacts can be used to defend decisions the model makes. "We have to think about new ideas for auditing AI, fueled by agile and iterative approaches, and working collaboratively with the first and second lines. Yet we also have a clear opportunity to leverage existing tools, methodologies, and approaches as a starting point," he said.

Taylor-Lubrano notes that because organizations have long used statistical models, they can regard machine learning and other examples of AI as a new version of those models. If approaches to ethical issues or other risks that were used in the past are no longer adequate, organizations will have to rethink their approaches. "We have a nice opportunity to add ethics to the debate now that AI has put a spotlight on it," Enstrom added.

That includes applying existing review criteria to AI systems. Because his organization is currently using AI as an assistive technology, with humans reviewing output, "we're treating AI essentially as a vendor," said Lipscomb. "We go through the appropriate vender onboarding processes and control structures around that, then we would expect a third-line review of the process."

The Three Lines Model

Under the IIA's Three Lines Model[7], effective risk management starts at the top, with management, as the first line, as risk owners and further clarifies roles, including those of the board. This governance framework can serve as a tool to help a company consider how to navigate opportunities and risks presented by AI. "We use it as part of our governance framework," Enstrom said. Among other things, it can aid in understanding roles and responsibilities for AI, including board oversight. As the independent, objective third line, internal audit reports to the audit committee, but can also offer perspective to the full board on ethics and other concerns. It can also advise on how changes driven by AI might alter the organization's risk profile.

The Three Lines Model can also help organizations recognize the necessity for each line to assess and monitor risk within its own purview, Enstrom noted. If AI is being used autonomously without rigorous human review of its output or decision making, the risk might be high, which may mean that management should implement enhanced quality assurance procedures or other controls within the first line. For the second line, chief risk or compliance officers may need to determine how best to establish adequate assurance and control, which would also be a consideration for internal audit's assurance

[7] *The IIA's Three Lines Model: An Update of the Three Lines of Defense*, The Institute of Internal Auditors, 2020.

role, as the third line. In light of any new changes, internal audit could also raise questions on how autonomous technology is being rolled out, if it is a priority on the board's agenda, and how it may be managed going forward.

The bottom line is that no matter how many changes AI may drive, "we have an opportunity to add value by positioning internal audit as a key element of the AI governance framework, leveraging our knowledge and experience around controls, and what we know as a profession; this all carries forward," Enstrom said.

Minding Model Risk Management

Model risk management addresses the risks that may result when decisions are made using models that are incorrect or improperly used. The goal of model risk management is to identify, measure, and mitigate or prevent the use of inaccurate data, assumptions, methodologies, processes, or interpretations. The banking sector has well established model risk management paradigms that are used to monitor models for credit, finance, and marketing activities, Clark noted. (See OCC 2011-12, Supervisory Guidance on Model Risk Management, from the Office of the Comptroller of the Currency.) As an Office of the Comptroller handbook on the topic notes, "sound model governance includes board and management oversight, policies and procedures, a system of internal controls, internal audit, a model inventory, and documentation."[8] Organizations can leverage these recommendations aimed at the banking industry, Clark advised, and avoid having to build their own risk model management systems from scratch. Effective model risk management is one factor in speeding adoption of AI and machine learning, "by creating stakeholder trust and accountability through proper governance and risk management," according to EY.[9]

[8] *Safety and Soundness: Model Risk Management*, Version 1.0, Comptroller's Handbook, Office of the Comptroller of the Currency, August 2021.
[9] "Understand Model Risk Management for AI and Machine Learning," Gagan Agarwala, et al., May 13, 2020, EY.

Using AI Within Internal Audit

Improving Effective Assurance with New Technology

Understanding AI Privacy and Accountability Considerations

In addition to understanding the AI implications for their organizations, internal auditors will also have to consider how best to use generative AI and other tools in their own audits, and what kinds of privacy risks to consider. For example, in working with generative AI, "it is essential to ensure that the data entered into ChatGPT is anonymized and that sensitive information is not shared or stored on the platform," according to an Internal Auditor article[10]. "Additionally, internal auditors need to ensure they have the appropriate consent and authorization to use the data in ChatGPT." The article details how internal auditors can use AI in planning, testing, reporting, and monitoring, and underscores the importance of leveraging the capabilities of tools such as ChatGPT while protecting the confidentiality and privacy of sensitive data.

Key Questions to Consider

Clark recommends that organizations develop a strategic understanding of what AI does or can mean to them. Internal audit can recommend that organizations address issues such as:

- Where and how is AI being used?
- What is the company trying to model? What is the purpose of that model?
- Are there solutions other than machine learning tools that can help us reach our goals?
- What risks are involved?
- How is or should the organization be automating decision making with models?
- Are there adequate monitors and risk management controls around AI?
- Is there a second line function dedicated to model risk management? If so, are there existing model-risk-management systems that can be used with AI tools?
- How does AI affect the audit scope and process?

Organizations should be certain to address ethical issues if an algorithm is being used in a process that makes consequential decisions about people. When that is the case, they should ask:

- Are there protections or laws in place? If so, how can the organization ensure that processes using AI are complying.
- If there are no external compliance considerations, are there still steps that should be followed to ensure the company is doing the right thing, according to its own values?

Internal audit can treat these considerations with the same care as external mandates, making sure there is a process to monitor and validate compliance and reporting on related compliance concerns.

[10] "*On the Frontlines: AI in 'IA,*'" Alex Rusate, *Internal Auditor*, May 17, 2023.

Conclusion

Because of weighty ethical issues related to AI, Clark advises that organizations that are not confident in the outcomes that the systems may produce should take a step back before implementing them. Instead, he recommends tackling AI initially as a research and development (R&D) project, giving the company a chance to explore how the technology fits its needs and identify potential risks.

Digital transformation is exciting, but internal auditors should keep a clear-eyed view of any technology's risks and limitations and focus on providing relevant advice and assurance. Amid the hype surrounding any new technology, "we need to be the ones asking which business problems it will actually solve and which data privacy issues and other risks may be involved," Clark said.

About The IIA

The Institute of Internal Auditors (IIA) is a nonprofit international professional association that serves more than 235,000 global members and has awarded more than 190,000 Certified Internal Auditor (CIA) certifications worldwide. Established in 1941, The IIA is recognized throughout the world as the internal audit profession's leader in standards, certifications, education, research, and technical guidance. For more information, visit theiia.org.

Disclaimer

Global Knowledge Briefs are intended to address topics that are timely and relevant to a global internal audit audience, and each topic covered is vetted by members of The IIA's volunteer North American Content Advisory Committee. Subject-matter experts are primarily identified and selected from The IIA's list of Global Guidance Contributors.

To apply to be added to the Global Guidance Contributors list, email Standards@theiia.org. To suggest topics for future Global Knowledge Briefs, email Content@theiia.org.

Copyright

October 2023

The Institute of Internal Auditors

Global Headquarters
The Institute of Internal Auditors
1035 Greenwood Blvd., Suite 401
Lake Mary, FL 32746, USA
Phone: +1-407-937-1111
Fax: +1-407-937-1101