THE IIA'S Artificial Intelligence Auditing Framework

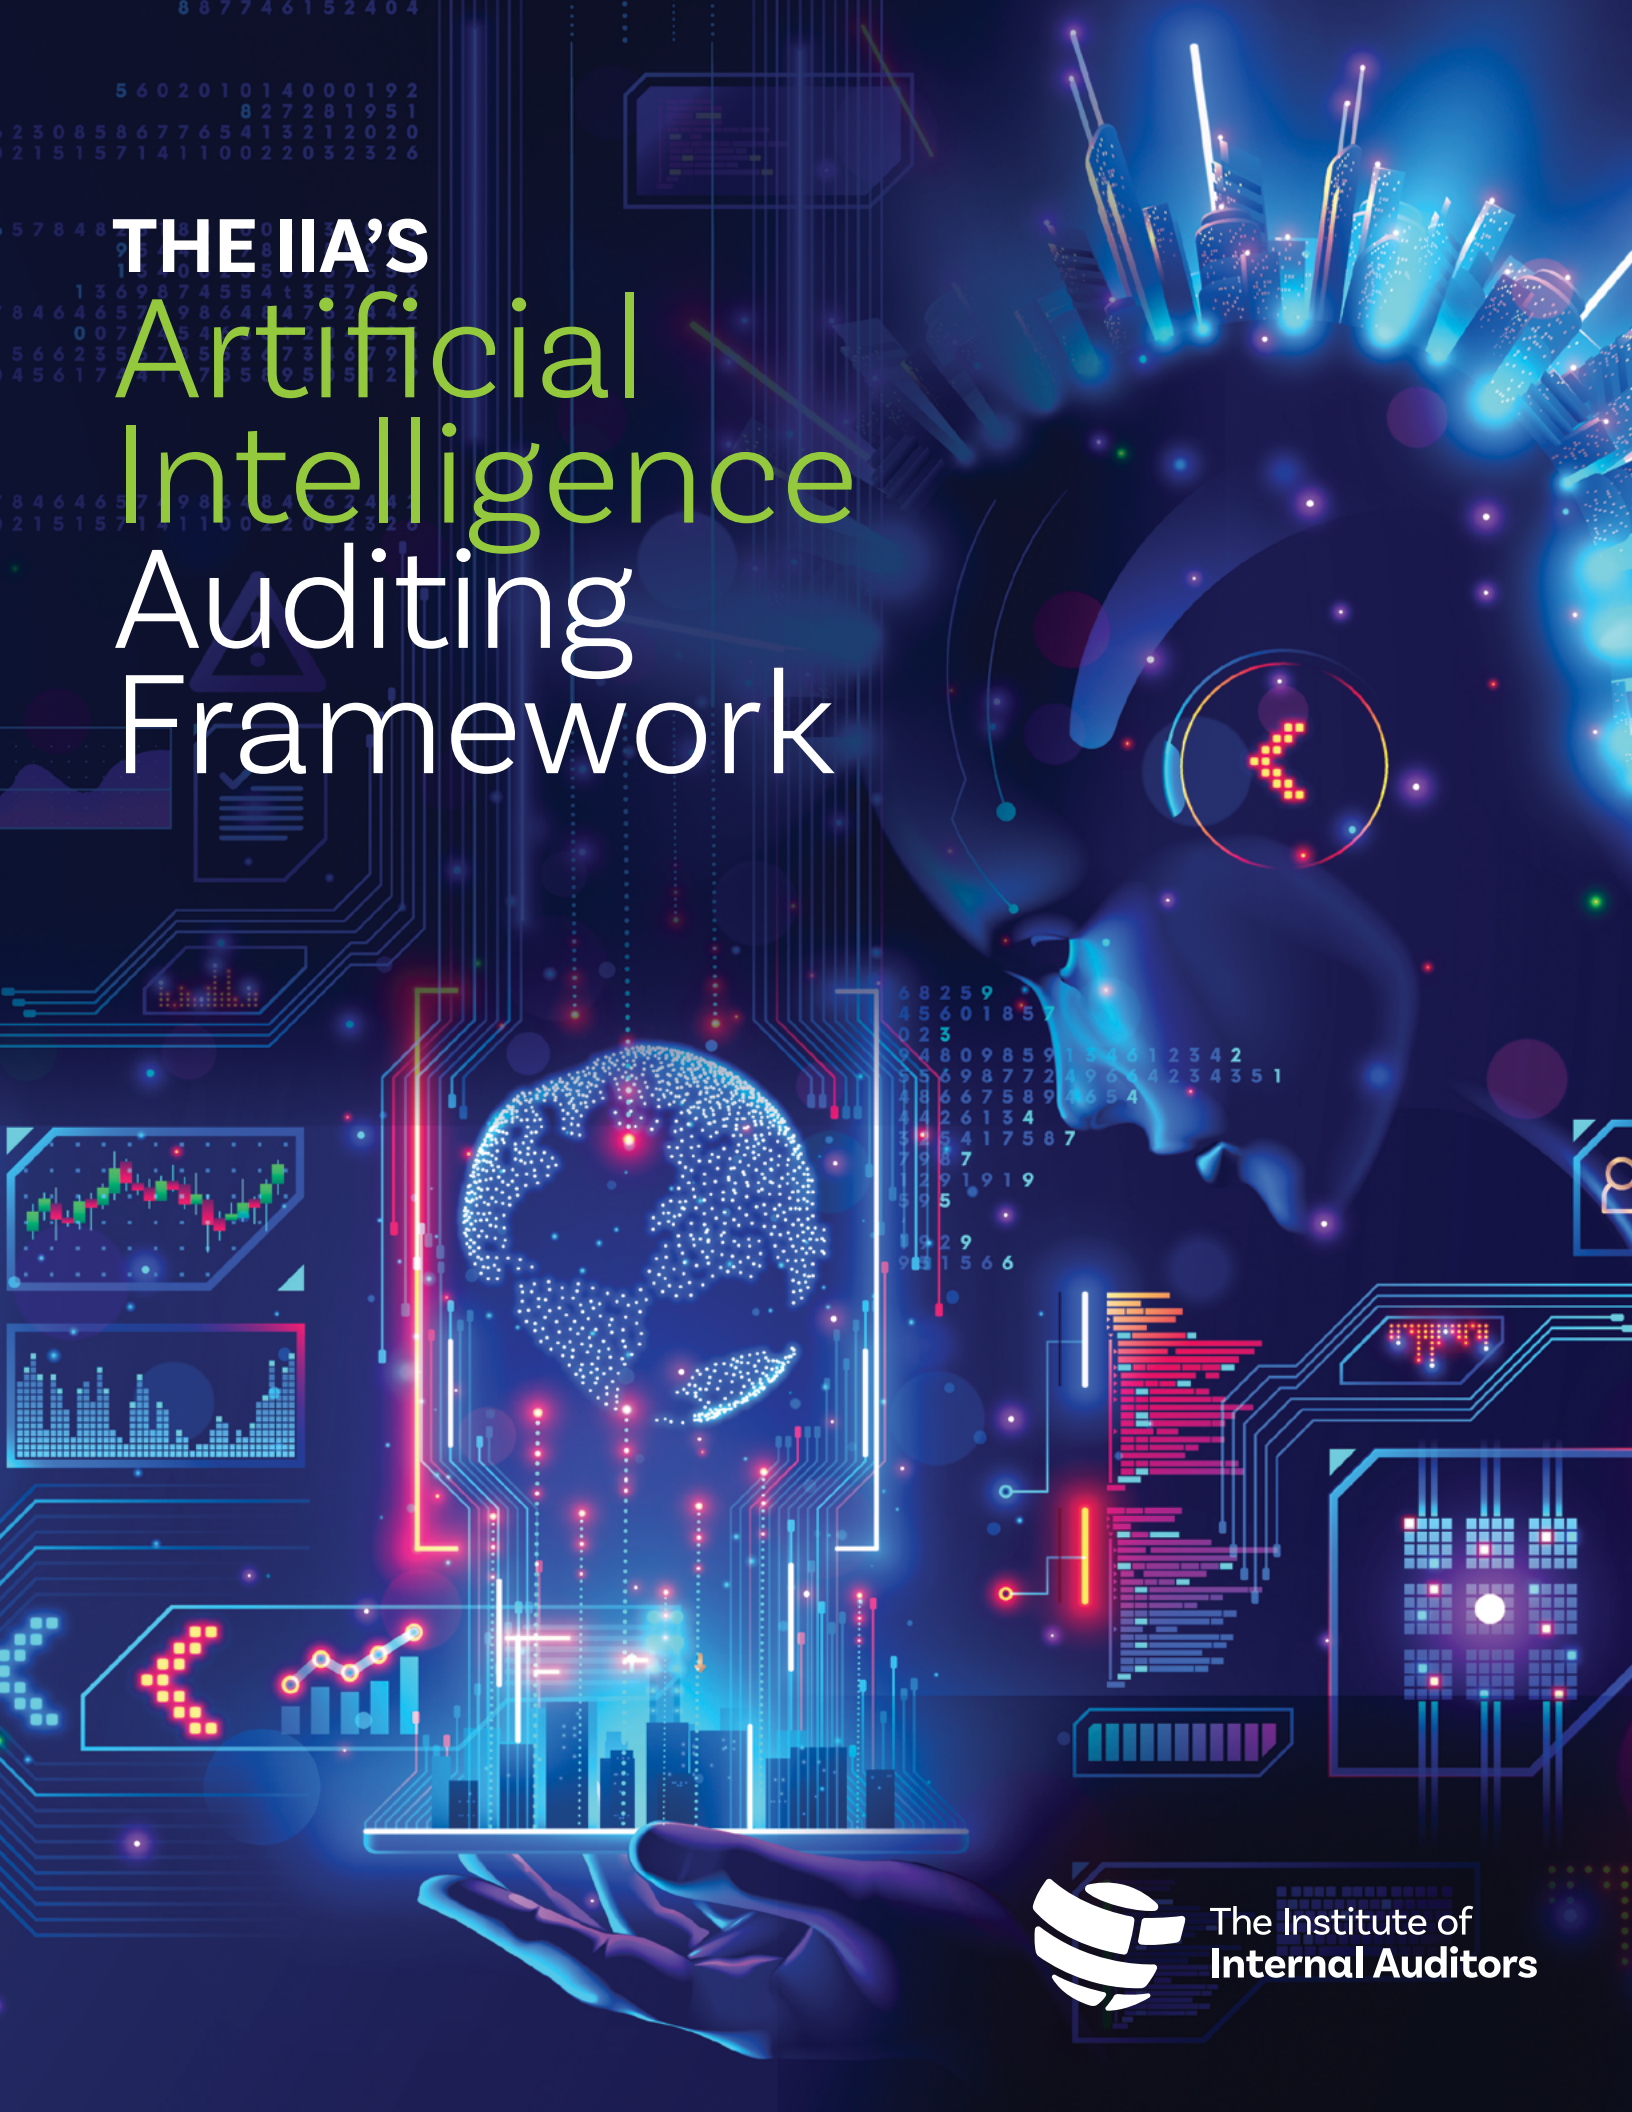

The Institute of Internal Auditors

Table of Contents

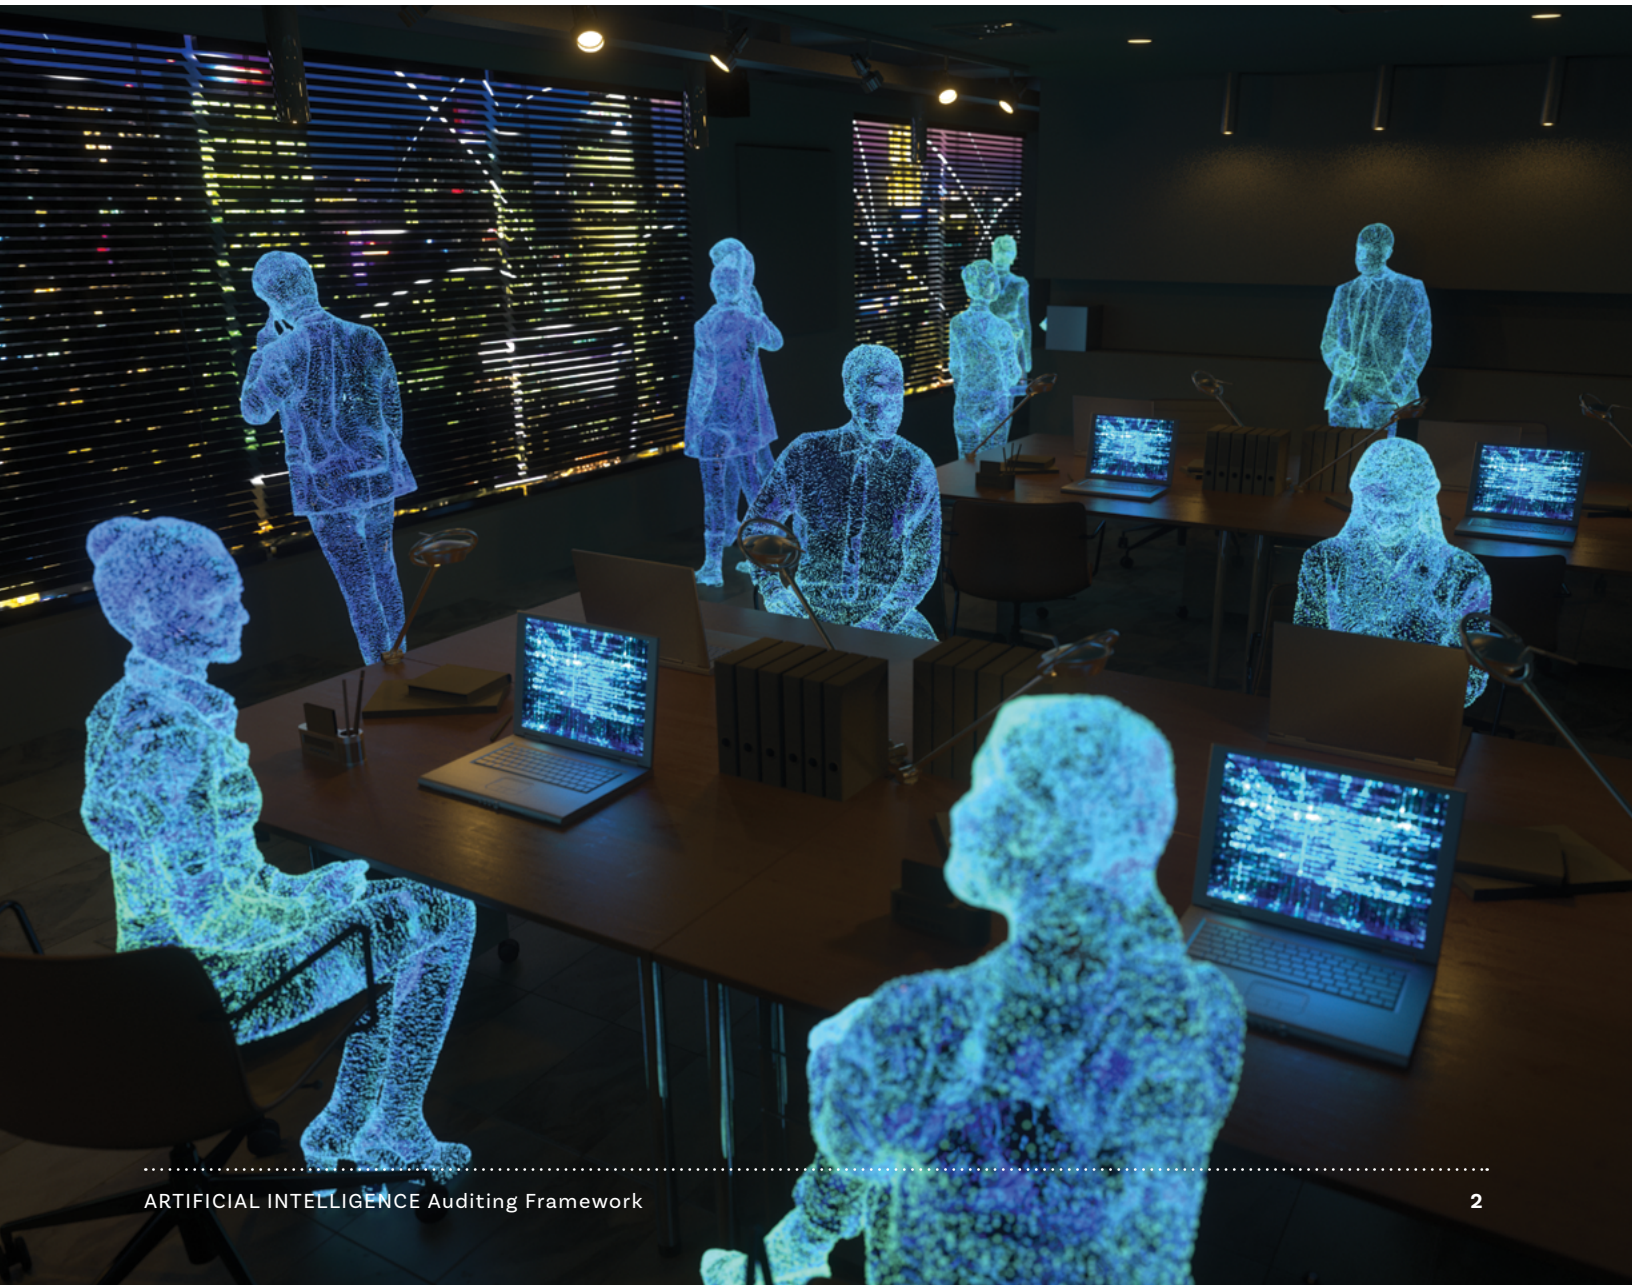

Introduction

Artificial intelligence (AI) is a broad term that has grown to encompass a wide range of existing and emerging technologies. While there is no single agreed-upon definition, AI generally refers to "systems endowed with the intellectual processes characteristic of humans, such as the ability to reason, discover meaning, generalize, or learn from past experience."[1] The current boom of AI applications demonstrates seemingly endless ways in which organizations can leverage AI-driven technologies to enhance the way we work, while posing both numerous and significant risks that are inherent, given the nature of the technology.

AI can be a daunting topic for an internal auditor, especially as organizations' AI adoption and use continue to grow. Now more than ever, organizations are looking to internal audit for increased guidance on AI. Whether as an advisor on risks and controls related to AI or in an assurance role on processes that use or rely on AI, it is vital that internal auditors increase their knowledge on the subject of AI.

Internal auditors are expected to provide assurance activities surrounding processes that can vary from simple business transactions to highly complex procedures that require deep understanding. The range and depth of AI literacy required to support assurance activities create continuous challenges for internal auditors who must continually develop their knowledge of AI to fully understand its risks and function to provide informed advisory and assurance.

AI as an auditing subject presents its own unique challenges, and its evolution means internal auditors must reassess risk and risk mitigation in the AI environment. That said, internal auditors already possess important foundational skills such as critical thinking, mapping processes, assessing risk, evaluating information technology controls, understanding organizational strategies, and providing independent assurance to the governance function.

The intention of The Institute of Internal Auditors (IIA) AI Auditing Framework is to help internal auditors in understanding the risk and identifying best practices and internal controls for AI. The framework will assist internal auditors in developing baseline knowledge. It is presented in four parts:

1. ***Overview – History and uses of AI.***
2. ***Getting Started – Understanding how an organization uses AI.***
3. ***AI Auditing Framework – Governance, Management, and Internal Audit.***
4. ***Practitioner's Guide and Glossary.***

This framework, which leverages aspects of The IIA's Three Lines Model,[2] will include references to The IIA's International Professional Practices Framework (IPPF), which provide a basis of mandatory requirements and guiding principles for the profession of internal auditing. Applicable Standards should be reviewed for additional information. Related IIA guidance, such as Global Technology Audit Guides (GTAGs), are referenced to provide topic-specific content. Other relevant frameworks, such as the NIST Artificial Intelligence Risk Management Framework (AI RMF 1.0), are listed as additional resources for internal audit practitioners.

PART 1
Overview

History and Evolution

As an overview into the topic, it is important for internal auditors to understand the historical development of AI, the way AI is currently used across various industries, and what emerging AI trends internal auditors should consider.

The idea of AI dates to 1950, when British mathematician Alan Turing posed the question, "Can machines think?" in his paper, "Computing Machines and Intelligence."[3] He is considered one of the founders of AI in suggesting that machines eventually would be capable of human-like intelligence. Two years later, Arthur Lee Samuel, an American computer scientist at IBM, developed a program that could play the board game checkers by using programmed values to identify the best move.[4] The Dartmouth Summer Research Project on Artificial Intelligence in 1956 marked one of the earliest uses of the term AI, credited to John McCarthy, an American computer and cognitive scientist.[5]

The 1960s saw many advances in AI, including the use of robotics, problem-solving programs, and the first interactive computer program (also known as a natural-language understanding program or NLP) called ELIZA, developed by Joseph Weizenbaum, a German-American computer scientist and professor. ELIZA could be considered the first "chatbot," designed to simulate conversation with a human user.[6]

Development of AI in the 1970s included the first intelligent robot called WABOT, developed by the School of Science and Engineering at Waseda University in Tokyo, as well as continued work on NLPs by Indian-American computer scientist Raj Reddy.[7,8] Advances in the 1980s included development of a driverless Mercedes Benz van in 1986 under the supervision of Ernst Dickmanns, German leader of autonomous driving technology.[9]

The 1990s saw advances in AI-related technologies, including speech recognition software in Microsoft's Windows. IBM developed highly effective AI such as "Deep Blue," which made headlines in 1997 when it defeated chess grandmaster Garry Kasparov.[10]

By the 2000s, AI had become part of our daily lives, including applications such as Amazon's Alexa, Apple's Siri, and Google Assistant. The year 2023 marked the year of increased adoption of large language models (LLMs) such as ChatGPT, which have further elevated the capabilities of AI from simply forecasting outcomes to a variety of content creation.

Adoption Levels

According to IBM's Global AI Adoption Index 2023, 42 percent of companies surveyed reported using AI in their business and an additional 40 percent reported they are exploring AI.[11] The continued expansion of AI highlights why internal auditors must ensure they

are incorporating AI-related risks into the planning of their audits. Additionally, internal auditors should continually develop their knowledge of AI. Part 2, Getting Started, will delve deeper into considerations an internal auditor can use to identify AI usage within their organizations.

Just as the AI landscape continues to evolve, so do the ways in which AI is categorized. While there are different perspectives on how to group the various forms of AI, the following section provides a brief summary of common forms of AI either currently in use (Reactive Machine and Limited Memory) or strictly theoretical (Theory of Mind and Self-Aware).

From a functionality standpoint, IBM[12] categorizes AI into four types:

1. ***Reactive Machine AI***
2. ***Limited Memory AI***
3. ***Theory of Mind AI***
4. ***Self-Aware AI***

1. Reactive Machine AI

- Reactive AI is a type of AI with no memory that is designed to perform tasks based solely on human input or training. Sometimes referred to as Narrow or Weak AI, these systems rely on the "human in the loop" for human-generated programming that instructs the machine how to operate on its own. This programming is commonly referred to as the "algorithm," or a set of calculations that include aspects of both computer science and mathematics or statistics.

EXAMPLES:

- IBM's Deep Blue.
- Some machine learning applications are classified as Reactive Machine AI. Machine learning is often based upon statistical models that analyze data and produce predictive results from the input data. For example, an online retailer could use AI on their mobile application or website that suggests products based on a consumer's purchase history. The individual user's purchase history is the data set driving the output, which is tailored to that user.

AI Utilization

- **PwC's 2022 AI Business Survey (U.S.) — AI-supported decision-making is being utilized by 74% of technology leader survey respondents, 62% of operations and maintenance leaders, 61% of customer experience leaders, and 60% of strategy leaders.**
- **EY's 2023 Global CEO Outlook Pulse Survey — 99% of CEOs surveyed are making or planning significant investments in generative AI.**
- **McKinsey's 2023 State of AI — 79% percent of all global respondents report some exposure to generative AI and 22% said they are regularly using it.**

2. Limited Memory AI

- Limited Memory AI is less reliant on human interaction to produce results, giving it the unique ability to learn and improve based on training from larger datasets. Whereas Reactive Machine AI can only utilize currently available data, Limited Memory AI can incorporate both past and present data to improve performance in a range of new AI tools.
- Other examples of machine learning access much larger sets of data and perform more complex analysis. This is referred to as "Deep Learning," which is a subset of both machine learning and Limited Memory AI. It is differentiated by being less reliant on human interaction to produce results. Generative AI falls within this category and can be utilized to create content based on the programmed algorithms.

EXAMPLES OF GENERATIVE AI INCLUDE:

- ChatGPT, LLaMA, and Bard are examples of LLM chatbots that rely upon Deep Learning and can produce text content, while other forms of Generative AI can produce content such as music (MusicLM), art (DALL-E), and even computer coding (OpenAI Codex).

- Chatbots and virtual assistants are a commonly utilized form of Limited Memory AI that use natural language processing (NLP) and reinforcement learning to engage in human-like conversations with end-users. These tools are often utilized by organizations such as financial institutions that allow a user to troubleshoot their issues even outside of business hours.

Additionally, machine learning is often subdivided into four categories:[13]

1. ***Supervised Learning*** – past learning is applied to structured data with predetermined outcomes.
2. ***Unsupervised Learning*** – no predetermined "right" outcomes; rather looks for patterns in unstructured data.
3. ***Semi-Supervised Learning*** – contains elements of both supervised and unsupervised learning.
4. ***Reinforced Learning*** – dynamic programming where algorithms are trained through a system of rewards and punishments; it learns without human interaction.

Other types of AI:

Expert Systems simulate human judgment or behavior. They incorporate knowledge from multiple people into problem-solving, and in theory provide more effective solutions. Expert Systems are used in chemical research to analyze and predict molecular structure and in medicine to identify harmful bacteria.

Computer vision technology, combined with Deep Learning, enables machines to analyze images. It is currently being utilized in healthcare to detect and diagnose patient anomalies based on X-rays, MRIs, or CT scans. Facial recognition is another form of computer vision, which has a multitude of uses, including authentication when attempting to access a bank account or restricting physical access to buildings that house sensitive data.

While robotics and AI are distinct fields, the two are often paired to create emerging tools that can be used in the real world. For example, physical robots that use visual sensors and image processing use AI to learn to navigate their environment. Manufacturing, agriculture, and consumer packaged goods are all industries that also rely upon robotics paired with Limited Memory AI to create efficiency and boost productivity within their operations. The use of robotic and AI-assisted surgeries within the healthcare industry promotes more precision that can result in faster patient recovery.

3. Theory of Mind AI

While Theory of Mind AI does not exist today, current research aims to develop AI systems that understand and interact with nuanced factors like emotions and motivations in a human-like manner. Ongoing work in this area includes efforts to develop systems that can analyze and engage with humans based on input from their voices, facial expressions, and sentiments in real-time and respond in a human-like manner.

4. Self-Aware AI

Self-Aware AI, like Theory of Mind AI, is currently theoretical and does not exist in practice. Most closely related to recent conversations on the possibility of "AGI" or artificial general intelligence, this hypothetical version of AI would be uniquely self-aware with what many envision as a rich, inner consciousness that matches or exceeds what humans are capable of. While a popular talking point in recent months, there is continued debate on the viability of this level of AI functionality.

PART 2
Getting Started

As organizations continue to deploy AI in various ways, internal auditors must be proactive and collaborate closely with management to understand their organization's overall strategy for AI, how AI is currently being utilized, and what future utilization is planned. Especially during the planning process, internal auditors should begin by researching and gathering relevant information regarding the potential use of AI under review from multiple internal and external sources.

Important internal information can include:

- Policies and procedures that reference AI, which can be gathered and reviewed to better understand organizational processes.
- An organization's documented strategic initiatives or the strategic plan, including aspects of AI.
- Recent board reports containing the vision and information on how leadership and the board are discussing topics like AI usage and risk-related concerns.
- Information obtained during ongoing risk assessment meetings with stakeholders.

External resources can provide an additional frame of reference as internal auditors begin reviewing their organization's AI strategy. Valuable external resources may include:

- The IIA's three-part Global Perspectives & Insights: The Artificial Intelligence Revolution.[14]
- The IIA's Artificial Intelligence 101 Series.
- The IIA's Analytics, Automation, and AI Virtual Conference.
- Foundational IT and Cybersecurity Audit resources such as The IIA's Certificates on Auditing the Cybersecurity Program and IT General Controls.
- The IIA's Practice Guides and Global Technology Audit Guides (GTAGs).
- NIST's AI Risk Management Framework (AI RMF 1.0).
- National Cybersecurity Centre's Guidelines for Secure AI System Development.[15]
- White House's AI executive order of October 2023.[16]
- IBM's AI governance eBook.[17]

Entity-Level Control Environment: Execution and Strategy

Once auditors have equipped themselves with these resources, asking the question, "How is AI being used?" is a simple, effective starting inquiry when gathering information. The answer to this question will most likely entail asking multiple individuals or departments because many organizations do not have centralized management of AI, nor do they have established policies (including defining what is AI), procedures, or a strategy regarding acceptable use of AI.

For organizations where AI has been developed and deployed, an internal auditor should have a discussion with the AI/data science team. That discussion should include asking them to explain which AI/algorithms have been deployed, including their function, sources of data used, use, limitations, risks, and ethical implications. Internal auditors should also begin to understand what existing controls are in place to help manage the risks posed by AI – or, if management has implemented new controls related to its use and deployment of AI systems. Gaining a preliminary understanding of the design of the controls used to manage AI-related risk is an important step that can be performed in concert with these initial discussions.

For organizations where it is unclear if or how AI is being utilized (formally or informally), the organization's IT function is a good starting point because, as noted in the Adoption Levels section in Part 1, technology leaders appear to have a higher tendency to experiment with and utilize AI in their department. If IT confirms that AI is being used, or if the initial inquiries determine AI is being used in the organization, the next logical inquiry is to determine to what extent AI is utilized.

While an initial conversation with the AI/data science team or IT management is a good first step, the discussion should not be limited to those groups. From those initial discussions, internal auditors may learn that other departments or individual users are using AI for their specific function, which would necessitate additional conversations. Working with management to review or collaborate on creating an inventory of which departments are currently using AI and updating that list frequently is advised. The inventory should include other key aspects such as the goal or objective of the AI, who uses it, who manages it, the specific AI tools in use, risk considerations, and who oversees it. The process of reviewing or collaborating with management to develop an AI inventory could also be accomplished during the annual risk assessment process.

Most internal auditors work closely with their chief financial officer (CFO) related to testing internal controls over financial reporting, or other executives such as the chief information security officer (CISO), chief information officer (CIO), etc., so having that professional relationship with members of the C-suite should provide another opportunity for initial conversations on AI. Important questions internal auditors can pose to their executives include:

- "Has an AI strategy been established, and if so, what are the details of that strategy (including aspects such as using AI to maximize efficiency of operations or using AI to reduce costs)?"
- "Has the C-suite determined who is accountable for managing AI-related risks?"
- "What role does the C-suite play in engaging the Board of Directors (or equivalent) for AI governance considerations?"

At this point, internal auditors will have:

- Researched AI within their organization and reviewed external resources.
- Conducted initial AI conversations with management, including their AI/data science team or IT management (or both) and the executive leadership team (CFO, CISO, CIO, etc.).
- Collaborated with management in reviewing or developing an inventory to capture how AI is being utilized (or planned for future use).
- Started the process of understanding what AI governance is in place.

Accomplishing these four tasks would indicate that internal audit has taken the first steps in establishing a baseline knowledge of AI in the organization. It would also provide an opportunity for internal audit to highlight any immediate observations that should be communicated to management in a timely manner.

Data

After internal auditors have a fundamental understanding of how AI is being used, they should develop more robust knowledge of AI usage within the organization. Because the algorithms utilized to power AI are dependent upon large volumes of data (also called "big data"), determining what organizational data is being used within any given AI application and how that data is managed is critical. An algorithm is a set of rules for the AI to follow and is what enables a machine to quickly process vast amounts of data that a human cannot reasonably process with the same ease or speed. Given AI's ability to rapidly ingest and respond to large amounts of diverse data sets, the architecture, performance, and accuracy of the algorithms involved is very important.

Algorithms are initially developed by humans, so human error and biases (both intentional and unintentional) could impact the performance of the algorithm. Part 3 of this framework will provide more details on risks related to algorithm errors and biases.

Outside of AI, many organizations already have developed a strategy for collecting, storing, using, managing, and protecting data. AI is like other data-driven applications in that the same important aspects about data are relevant and should be considered, which include integrity, privacy, confidentiality, validity, accuracy, and completeness.

Big data means more than just large amounts of data — big data refers to data that reaches such high volume, variety, velocity, and variability that organizations invest in system architectures, tools, and practices specifically designed to manage the data. Much of this data may be generated by the organization itself, while other data may be publicly available or purchased from external sources. For comprehensive guidance on understanding and auditing big data, including a discussion of opportunities and risks, and a sample work program, see The IIA's "GTAG: Understanding and Auditing Big Data."

Another critical aspect of both data usage and related AI applications is whether data is hosted or processed by a party outside of the organization. Internal auditors must always consider the risks related to third (and fourth) party transactions because the vendors' internal control environments may not be as comprehensive as the organization's environment (or the desired vendor

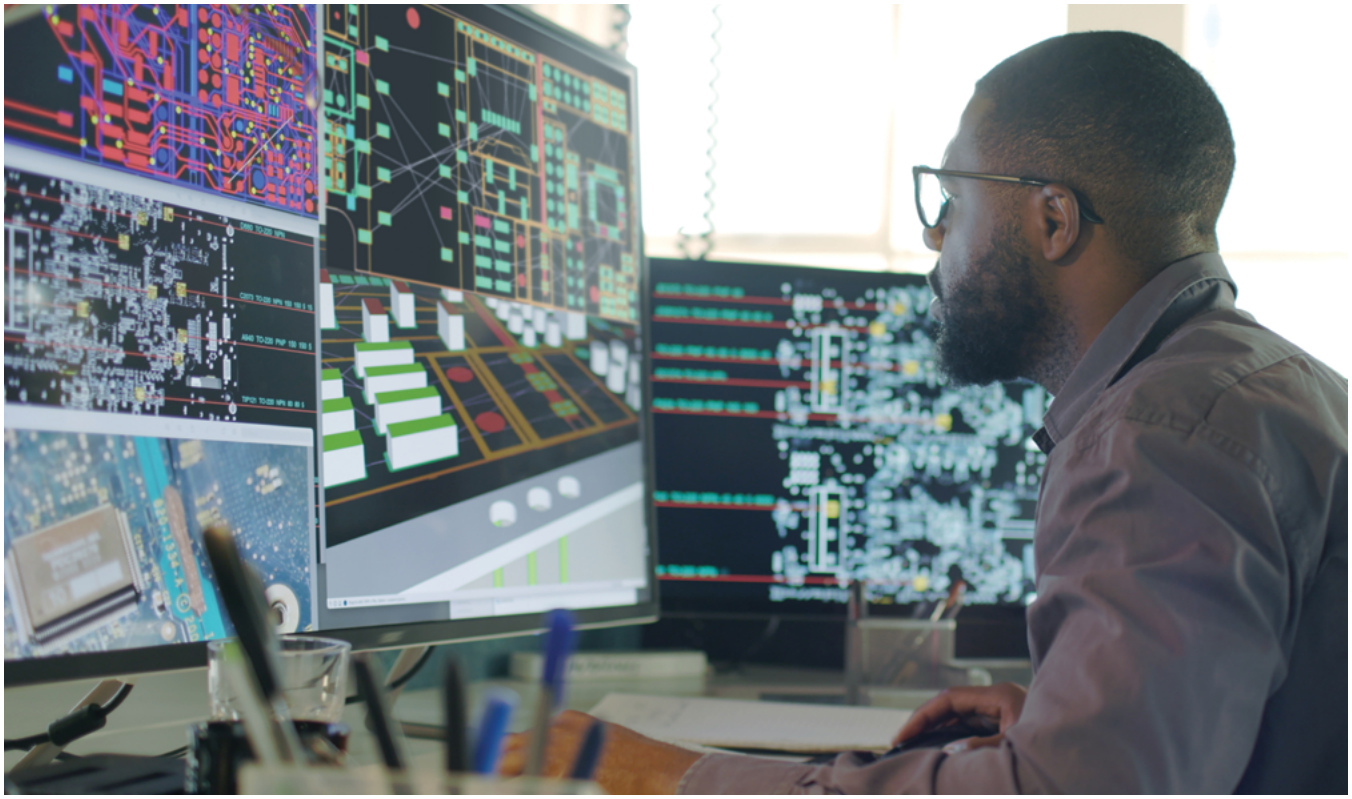

control environment). The IIA's Practice Guide "Auditing Third-party Risk Management" provides internal auditors with a more detailed approach regarding the risks related to utilizing external vendors.

Another important aspect of data is user access. Understanding who can edit or make changes to data is critical in that manipulating data sets from an input standpoint can certainly impact the downstream output of AI. Understanding and documenting administrator user access to AI-reliant data is also imperative. The IIA's "GTAG: Auditing Identity and Access Management" provides a closer look at internal audit considerations related to how the organization ensures users have appropriate access to IT resources.

Cybersecurity

Cybersecurity must also be considered as it relates to restricting unauthorized users from accessing data and ensuring privacy, confidentiality, and protection of data. The adoption and evolution of AI is forcing organizations to reemphasize their cyber resilience capabilities. As AI becomes more powerful and more decisions are handed off to new, complicated, and opaque algorithms using huge data sets, protecting these systems from outside, malevolent forces is critical to organizational success. Cyber resiliency is vital for any organization that utilizes AI.

Internal auditors are typically involved with testing the effectiveness of IT internal controls. This familiarity of how the organization has implemented cybersecurity-related internal controls can assist internal auditors in validating that those same controls are being used to protect AI-related data. Examples of cybersecurity controls include:

- Use of encryption.
- Presence of anti-virus software.
- Utilization of intrusion prevention/detection systems.
- Security event logging of both prompts and responses.
- Ensuring a penetration test is performed periodically to proactively look for vulnerabilities.
- Employee training in best practices to detect and avoid phishing, smishing, or other social engineering schemes.

For additional details, see The IIA's "GTAG: Auditing Cybersecurity Operations: Prevention and Detection."

Internal auditors need to determine where AI-reliant data is stored (internally, externally, or both) and consider what cybersecurity controls are in place. For externally stored data, a Service Organization Company (SOC) report should be obtained to learn about the vendor's control environment. Management should be aware of any control deficiencies found on the SOC report and ensure that those deficiencies do not put the AI-reliant data at risk. Service-level agreements (SLAs) with vendors should include a "right to audit" clause.

PART 3

AI Auditing Framework

The first version of The IIA's AI Auditing Framework was issued in 2017. It provided internal audit professionals with an approach for performing AI advisory and assurance services in a systematic and disciplined manner. This updated version of the framework modernizes the content with examples from the current AI environment, while providing additional details to assist internal auditors as both advisors and assurance providers. The framework has three domains:

The IIA's AI Auditing Framework

- ***Governance***
- ***Management***
- ***Internal Audit***

The framework links to The IIA's Three Lines Model: the Governing Body (Governance) oversees management (First and Second Lines), while the role of internal audit is covered in the third domain, which includes both independent assurance (Third Line) and advisory activities.

The IIA's AI Auditing Framework is intended for use by internal auditors. However, the framework's Governance and Management domains outline activities and functions outside of internal audit required to manage AI within an organization. The main objective of the framework is to equip internal auditors with essential baseline knowledge of AI to serve their organization as 1) an advisor to management to consult on the overall approach on how AI is managed, executed, and monitored and/or as 2) an assurance provider to audit the processes and controls management has established to manage, execute, and monitor AI.

Organizational maturity of AI usage contributes to how internal audit will be utilized. For example, less AI mature organizations may need internal audit to assume an advisor role in the initial exploration of AI, while a more AI mature organization would likely engage internal audit to provide assurance activities, such as evaluating established processes and internal controls for operating effectiveness. To perform both roles successfully, internal audit needs a solid understanding of how AI should be managed and how the organization is currently managing it.

Governance — the framework's first domain — is based on an organization's approach to strategic planning of AI and in providing oversight and monitoring over how AI is planned, managed, and executed by management. The governing body relies upon information that is provided to them by the internal audit function. Internal auditors should strive to develop a trusted

advisor relationship with governing bodies such as the audit committee, board, or equivalent governing body, and this relationship should include emerging topics such as AI that presents new oversight challenges.

The framework's **Management** domain outlines an approach that the organization would use when planning and executing AI within the organization. The internal control environment surrounding AI is established by management in the "First Line." It also includes strategic aspects such as setting goals and objectives related to the overall AI strategic plan. Internal audit must ensure it understands the strategic direction of AI for the organization, and management's approach for managing AI.

The framework's Management domain also contains "Second Line" monitoring aspects of AI, such as what aspects enterprise risk management should consider when monitoring the "First Line." Internal audit is often expected to participate in an organization's risk assessment process. This domain will be relevant to internal audit as it maintains knowledge of AI-related risks.

The framework's third domain, **Internal Audit,** includes aspects of both advisory activities to management and in providing assurance services in an audit capacity ("Third Line"). Internal audit can utilize the framework as a starting point in both roles when tasked with participating in AI assignments.

Because AI is evolving rapidly, the framework will require periodic updates. This evolution, combined with the complex nature of AI, means internal audit likely will be able to provide only limited assurance. The framework by itself may not cover all aspects of AI, but it does provide a solid foundation for internal auditors as they develop fundamental knowledge of AI as an audit topic.

Governance

AI governance refers to the structures, processes, and procedures implemented to direct, manage, and monitor the AI activities of the organization. Governance includes helping to ensure that AI activities, decisions, and actions are consistent with the organization's values, as well as its ethical, social, and legal responsibilities. It also includes providing oversight to ensure that those employees with AI responsibilities have the necessary skills and expertise.

As reflected in the Three Lines Model, internal audit functions as the "Third Line," providing independent and objective assurance on the validation of internal controls used by the organization to manage risks, including all aspects of AI. Internal audit can provide AI-related advisory services to the organization, but from a governance standpoint, the governing body relies heavily on the assurance activities provided by internal audit to better understand organizational operating effectiveness.

Governance of AI is vital. Two of the most important roles governance plays are evaluating how well the organization is managing AI operations and whether the organization's AI strategic goals and objectives are being achieved in a manner that is consistent with established values. As presented in previous sections, there are a number of AI-specific risks; however, one of the main considerations is providing oversight that AI is being utilized in a way that will not cause harm.

Strategy

A strategic plan allows an organization to clarify and communicate the direction and vision required to achieve its goals; the same is true with an AI strategy. Each organization's AI strategy should be unique, based on its approach to capitalizing on the opportunities AI provides while being mindful of an organization's specific circumstances such as the details of current technology services or ongoing data governance initiatives. A thoughtful and methodical AI strategic approach will support an organization's ability to focus its resources and promote alignment across all employees while mitigating potential risks.

Two important points to keep in mind:

1. Planning an AI strategy is not a one-time event; it is an iterative process that should be performed periodically. Internal audit should work with management to determine a timetable for reviews of AI strategy.

2. An AI strategy should not be planned in isolation; given the range of potential data sources and use cases, organizational AI strategies should be cross-functional. Given the critical importance of AI, board-level involvement and oversight is likely to occur, because AI has the potential to dramatically alter or modify business strategies.

Addressing these points will help ensure that AI initiatives support the organization's overall objectives and align with stated organizational values. Formulating goals for AI allows organizations to frame important strategic considerations, including the answer to baseline questions such as, "Why are we using AI?" and "What are we attempting to achieve?" AI goals should be developed like other "SMART" organizational goals — specific, measurable, achievable, relevant, and time-based — to avoid adopting AI tools and services without a clear scope of the organization's reason for doing so.[18]

Desired attributes of AI should be included when setting goals, objectives, and expectations. The organizational expectations or objectives may include the following desirable attributes for artificial intelligence:

Desirable Attributes for Artificial Intelligence

- **Effective**
- **Valid**
- **Reliable**
- **Safe/Secure**
- **Unbiased**
- **Transparent**
- **Ethical**
- **Explainable**
- **Private**
- **Compliant with laws**
- **Fair**
- **Confidential**
- **Responsible**
- **Accurate**
- **Efficient**
- **Accountable**

The organization's overall attitude and approach toward risk and risk management should be a primary consideration when developing or updating the AI strategic plan and goals. Having a higher risk appetite in pursuit of AI goals may not be appropriate for an organization that is risk averse in other aspects, whereas organizations with historically high-risk tolerance may be more willing to accept AI-related risks. Regardless of an organization's risk tolerance, it is essential to recognize and map AI risks during AI strategic planning.[19]

Management – First and Second Lines

In developing the AI strategy, management is responsible for ensuring that internal controls have been designed properly and are functioning effectively to mitigate risk. As described in previous sections, effective internal controls are a critical requirement of AI. Many organizations test and report the results of IT controls on a quarterly and/or annual basis. Management should be aware of any internal control issues that also can have an impact on the use of AI, especially related to areas of the internal control environment that are already being evaluated such as:

- Data integrity and data governance.
- User access.
- Cybersecurity.
- System development life cycle.
- Change management.
- Back-up/recovery controls.

COBIT and COSO are examples of internal control frameworks that can be utilized by organizations to assist with their approach and evaluation of the internal control environment.[20,21]

First Line Management

Leadership

Defining roles and responsibilities related to AI-based initiatives will support the organization in determining what resources are required to operate effectively. Identifying executive ownership, while incorporating input from the other members of the C-suite, will help ensure accountability.

An AI Leadership Team of cross-functional members is another way organizations can monitor and communicate AI initiatives and support accountability. Such a team should include:

- AI and/or data science managers.
- The organization's CISO.

- Key IT personnel.
- Legal (to provide direction on regulatory considerations).
- Finance/accounting to track the costs and ROI of AI projects.
- Risk management.
- Compliance.

Internal audit, with its breadth of knowledge about the organization, is uniquely positioned to serve as an advisor to support AI initiatives and should be considered as an AI Leadership Team member. Internal audit's participation should be structured to ensure that its independence as an assurance provider is not compromised.

A well-thought-out planning process will support the organization when executing AI projects. Employees involved in executing the projects need to be aware of the most critical risks, including unwanted results. Highlighting and ensuring that the daily execution of projects includes awareness around social, ethical, environmental, and economics aspects is important. Additionally, fostering an environment that encourages employees to openly discuss ideas and concerns related to AI initiatives can help to create a culture of transparency, awareness, and mutual responsibility to support ambitious AI projects.

Policies and Procedures – Internal Use and Business Applications

Defining, adopting, and disseminating robust organizational policies and procedures around the use of AI within the organization is another important aspect of an organization's AI strategy. Clear policies and procedures provide direction to the employees directly involved in AI initiatives and to employees who may use AI as part of their daily work responsibilities. Developing an AI acceptable use policy should be a top organizational priority. It should include aspects of cybersecurity best practices, intellectual property/legal considerations, and the risks associated with various AI tools. The policy should be supplemented by a documented process that users must follow when requesting the use of AI. Using a formal approval process for AI use also will support the organization's efforts to maintain an inventory of users or departments who utilize AI.

Policies and procedures that clarify the guidelines and expectations used to develop, deploy, and monitor AI initiatives formalize the process. They provide a baseline to validate if projects are being performed in a manner that is consistent with the organization's approved policies, ethics, and overall organizational risk culture. Internal auditors are in a unique position to provide immediate feedback on this topic, given their knowledge and experience providing assurance over key policies and procedures. In many cases, as a starting point, existing policies and procedures may provide reasonably effective measures to mitigate risks posed by AI development. For example, the AI systems that are being developed may be subject to existing System Development Life Cycle (SDLC) or change management control processes. Over time, as organizations evolve and elevate AI use cases, more mature or new controls will certainly need to be considered.

Accordingly, policies and procedures that clarify expectations and guidelines for third parties involved in AI initiatives are also important. Coordination between the teams managing AI and the organization's group that manages third-party relationships (such as legal) will promote consistent AI vendor relationships. Because third parties are an extension of the organization's processes, maintaining a good understanding of vendors' control environments will be critical. Where available, management should obtain AI vendors' SOC reports to understand their control processes and to be aware of any concerns such as audit findings. Use of any third parties as it relates to developing AI capabilities or ongoing support of AI initiatives should be clearly defined and monitored, including SLAs that contain the right to audit.

Once these policies and procedures have been outlined, organizations can promote cross-functional buy-in of AI policies and procedures by sharing drafted organizational policy documentation such as the acceptable usage policy across all staff and inviting feedback during an open comment period. Organizations also should plan for resources needed to train staff on these new policies to ensure that employees are ready to adopt and adhere to newly defined roles, controls, and responsibilities related to AI usage.[22]

IT Resources to Support AI

Effective IT resource optimization is required to support AI initiatives and should be budgeted by management accordingly. The use of AI requires intensive computer asset performance to sustain

reliable processing. Examples of IT resources capabilities used to support an organization's AI initiatives include:

- ***Central processing units (CPU)*** – the "brains" of the computer; processors that execute commands or instructions.[23]
- ***Graphics processing units (GPU)*** – more capable brains that can process many pieces of data simultaneously with additional mathematical capabilities; able to produce graphics, imagery, and are more prevalent in creative production AI.[24]
- ***Storage*** – location of the data that is required by AI for processing. Storage is commonly measured in terabytes (1,000 gigabytes) or petabytes (1,000 terabytes); for context, a 5-10 minute, high-definition video measures approximately one gigabyte (1 billion bytes); on-site hosted servers or cloud-based solutions are examples of where data may be stored.
- ***Memory*** – also called RAM (random access memory); location where shorter-term data is stored that is available more quickly than storage data; measured in gigabytes, where individual computer workstations have 8 to 48 gigabytes of RAM; the more complex the AI that is running, the more RAM is required.
- ***Supercomputers*** – fastest processing computers that are used for high-performance computing and contain multiple CPUs.
- ***Workstations*** – includes desktops and laptops with technical specifications that support the requirements of the AI that is being utilized.
- ***Software*** – platforms, programs, and applications that are used to develop, deploy, and manage AI; development software. Examples include Microsoft Azure AI, IBM Watsonx.ai, and Google Cloud AI Platform; deployment software, which is used to integrate AI into existing applications; examples include Docker and MLflow.
- ***Networking connectivity*** – this is a broad category that includes the hardware, software, and services that allow users to share digital resources and exchange information; examples include file servers and routers.

While internal auditors are not expected to know all the technical specifications and details of AI requirements, they should have a basic knowledge of IT resources.

Staffing and Training

Proper staffing is an important element of an organization's AI strategy. Human resources should collaborate with management to ensure that employees with the required AI experience are recruited throughout the organization. AI experience should be prioritized not only for the employees who are tasked with managing the day-to-day aspects of AI, but also for leadership who will direct AI initiatives.

Because AI is developing so rapidly, it is important that the organization's employees are aware of advances and the corresponding risks. Organizations should ensure that general AI awareness training is provided to all employees and that more technical training opportunities, such as seminars, online training, or educational courses, are available to employees who focus on AI initiatives.

As mentioned above in the Policies and Procedures section, implementing training on the formal acceptable AI use policy and including AI in the employee handbook and in new-hire orientation are good ways to increase organizational awareness of AI along with possible risks. By integrating training initiatives focused on AI and digital literacy, organizational policies and procedures, and upskilling opportunities, organizations can support AI initiatives via direct investment in current and incoming staff members. The implementation and outcomes of these initiatives should be monitored by internal audit as part of an organization's AI controls.

Execution

Risk Management by First and Second Lines

Part 2 discussed the importance of identifying AI risks related to security, integrity, privacy, and confidentiality of data, and addressing these concerns should be a focus as the organization executes AI projects. AI algorithms rely upon accurate and reliable data and project teams should closely monitor input data. Organizations have multiple ways to validate the completeness of data being used in AI projects,

including ensuring record totals match and analyzing error reporting when data is transferred between systems. Management should design and monitor internal controls that detect anomalies with data quality or completeness.

Other important data considerations include restricting user access only to those employees who are working on an AI project, which includes administrator access. Determining user roles and ensuring proper segregation of duties is critical as well. For example, database administrators oversee underlying input data and should not have access to modify the algorithms that process that data; a task that is traditionally a developer's responsibility.

When an AI project is being implemented, it is important that the organization ensure the project is transparent, explainable, responsible, and auditable:

- ***Transparency*** – able to easily understand in simple terms the purpose of the AI or algorithm.
- ***Explainability*** – able to explain the mechanics, calculations, or results processed by the AI or algorithm.
- ***Responsibility*** – using the AI or algorithms in an ethical, safe, fair, and trustworthy manner.
- ***Auditability*** – as AI applications may begin to replace or augment certain key compliance or other important business processes, maintaining traceability through effective audit logs or related information will be an important component to AI development, as assurance over these processes will likely be needed for many of the potential use cases.

AI project management should define the following aspects for each initiative:

- ***Objectives, roles, and timing*** – what the initiative intends to achieve, who participates, and when it occurs.
- ***Resource requirements*** – what technology and/or staffing resources are needed to achieve success.
- ***Data requirements*** – what input data is required by the AI or algorithm(s).
- ***Privacy, legal, and regulatory requirements*** – what are the related compliance requirements.
- ***Risk assessment*** – what are the relevant risks that threaten achievement of project objectives or unwanted results, such as biases, unethical treatment, or misuse.

- ***Success metrics or key performance indicators (KPIs)*** – how project success is monitored and quantified.

- ***Testing requirements*** – at a point in time, how to validate the AI or algorithm is performing as designed and what changes are required; this will include both end-users (who will ultimately utilize the AI) and AI/data science professionals; identifying and communicating issues will be critical in this stage. Reviewing how third-party developers test and confirm the effectiveness of their algorithms is an important consideration.

- ***Testing requirements*** – from an ongoing standpoint, as AI output, by its very nature changes due to data input, consideration should be given to ongoing testing or quality assurance over the model – depending upon the use case, this concept may need to be built into the business requirements, or AI design.

Ongoing monitoring of AI projects should be performed by management to ensure the initiative is proceeding as planned and to identify any issues or concerns that have occurred. As indicated in the Three Lines Model, management plays a vital role in the internal control environment by providing the first level of actions to mitigate risk. Monitoring at the project level is important because that is where issues are initially detected. Management reporting to both executive leadership (C-suite) and the board should be part of this process. It is important to not just monitor the project's overall progress, but also to identify and report on any negative results, such as any ethical concerns or breach of sensitive information. Including an assessment of any third parties to ensure they are fulfilling their responsibilities in the AI project is also important.

Monitoring and reporting should also include disclosing any project-specific internal control issues or analysis of internal control issues that were from other areas of the organization that might affect the AI project. Enterprise risk management and/or compliance should also be part of the control issue monitoring process from a "Second Line" perspective.

Support from the Second Line in Risk Management

The main goals of an organization's enterprise risk management process are to understand how risks may threaten the achievement of objectives and then take actions to mitigate those risks. Risk categories include strategic, financial, environmental, market, social, ethical, technological, economic, political, legal, and regulatory. AI is a topic that is generally considered a technological risk; however, it is important to recognize that AI risk can fall within any of the aforementioned categories, necessitating a robust risk management process for AI projects that considers both technological and non-technological concerns.

The IIA's AI Auditing Framework provides risk management considerations to support organizational AI projects in following best practices around AI risk management. Where appropriate, other existing frameworks should be considered, in particular, NIST's Artificial Intelligence Risk Management Framework. Internal auditors often collaborate with risk management professionals in such activities as the organization's annual risk assessment process; therefore, it is vital that internal auditors understand AI-related risks and continue to increase their knowledge base. Additionally, internal auditors should consider AI-related risks at the engagement level, that is, when auditing processes that include some aspect of AI.

Identification

Identifying AI-related risks may be a new task for many organizations. Ideally, enterprise risk management, (along with internal audit, compliance, and legal), will participate in the initial discussions of all AI initiatives to help frame risks surrounding the AI project. As mentioned in the Strategy section, a cross-functional AI Leadership Team is an effective way to proactively identify potential risks or threats prior to them being realized while ensuring that controls and risk mitigation techniques are in place across the organization.

Organizations that have already established an effective enterprise wide risk assessment process should consider performing an initial AI-focused risk assessment. If a separate AI risk assessment is not feasible, organizations at a minimum should ensure that AI is included during the overall risk assessment process.

For example, organizations engage the executive leadership team on a periodic basis to identify risks in various areas. In many instances, the discussions or surveys include specific questions, such as, "What are the organization's biggest strategic risks?" To bring

AI to the forefront of the risk assessment process, organizations should highlight AI as an emerging risk area, share ongoing feedback from the AI Leadership Team and/or internal staff, and gather input from executives accordingly. The risk identification stage in the risk management process is important because it may highlight risks that were not previously identified.

Organizations that have established a clear AI strategy, with defined objectives and goals, are providing the context enterprise risk management needs to assist in AI risk identification. This context allows enterprise risk management to develop an inventory of risks that threaten the achievement of those objectives and goals, allowing organizations to embed in their strategic plans safeguards against potential harm from the use of AI. It is important for organizations to be mindful that the risk landscape around AI continues to evolve rapidly, causing unwanted, negative consequences from unaccounted-for risks that may include:[25]

- Biased or discriminatory outcomes that may unfairly affect specific segments of the population.
- Compromised privacy or confidentiality.
- Lack of accountability.
- Lack of transparency.
- Lack of explainability.
- Financial harm or economic inequality.
- Environmental harm.
- Misinformation or manipulation.
- Copyright infringement.

The "Black Box"

While much of the risk identification, assessment, and mitigation processes for AI projects follow existing best practices, it is important to note that AI's "black box" poses a distinct risk. The term refers to the lack of transparency into AI systems and the ways they make decisions. Deep learning models in particular can be difficult to understand given the complex processing performed by algorithms paired with potentially limited to no visibility or understanding of how an output was produced. This can pose specific challenges as enterprise risk management (and internal auditors) try to capture documentation needed to support the risk management cycle defined above. Risk professionals and internal auditors can address the "black box" directly by:

Identifying and clearly communicating where they may have missing or incomplete information within an AI project.

- ***Example:*** If an organization is using a third-party AI vendor that does not provide detailed information on an algorithm's training data, this should be documented and disclosed as a potential risk.

Continually assessing and updating the board on potential impacts related to the identified information gaps.

- ***Example:*** Once the lack of vendor documentation on the training set has been documented, internal auditors should update the board in the face of a risk-related consequence (such as multiple examples of biased AI output that suggest problems with training data).

Presenting direction on how to mitigate risks associated with "black box" knowledge gaps previously documented and assessed.

- ***Example:*** Based on the presented assessment and consequences, the organization makes the decision to migrate to a new AI vendor with more transparent data documentation.

Assessment

Assessing and analyzing identified AI-related risks should follow a similar process that an organization uses for reviewing other risks – impact and likelihood should be considered first. Impact of AI-related risks may be difficult to quantify due to the numerous considerations such as legal, regulatory, social, financial, environmental, and ethical ramifications. Damage to brand reputation is another consideration for impact.

The combination of impact and likelihood results in inherent risk, which is a measure of risk that exists without the consideration of internal controls. After assessing inherent risk, residual risk should be the next determination, which includes consideration of how well risks are mitigated.

As an example of assessing the security of AI as an objective, cyber threats may be identified as a significant risk. To address cyber risks, organizations deploy cybersecurity controls with the goal of reducing the inherent risk down to an acceptable level of residual risk. If enterprise risk management assesses that residual risk has not been reduced to an acceptable level, the organization must determine how to proceed.

Risk prioritization is the process an organization uses to rank risks in order of importance, that is, the most impactful risks are addressed first. Organizations have limited resources but face unlimited risks, so ensuring that AI-related risks are prioritized within the broader entity-wide analysis of risk is important. How AI-related risks rank within individual organizations will vary based on their risk assessment process, how much they utilize AI, and the maturity level of their internal control environment. Simply put, there is no "one size fits all" approach to assessing AI-related risks.

Mitigation

Risk mitigation is an action (or actions) that management takes to reduce risk to a more acceptable level. In many instances, organizations choose to treat AI-related risks through mitigation actions such as the addition of internal controls; however, there are other possible risk responses as described in the table below.[26]

Numerous factors influence how an organization determines how to respond to AI-related risks. Therefore, having a defined, repeatable risk response process in place is vital. AI-related risks may change during the course of a project, so an organization should continuously revisit how it responds to and mitigates risks.

Internal Audit – Advisory and Assurance Activities

After describing how an organization should approach AI in the previous two framework domains, there is one remaining domain – Internal Audit.

The first two domains give a baseline for how internal audit can provide both advisory and audit services to

Basic Risk Responses

RESPONSE	CHARACTERISTICS	DEFINITION
Treat	Reduce, Mitigate, Enhance, Exploit, Leverage, Optimize	Apply controls to reduce inherent risk to an acceptable residual level or apply other measures to maximize and take advantage of potential possible variances in outcomes.
Tolerate	Accept, Pursue	Determine whether potential benefits warrant taking the risk, having established measures considered necessary to mitigate or leverage likelihood and/or impact.
Transfer	Share, Spread	Spread risk either by transferring some or all of it to a third party (such as through insurance or outsourcing), or applying the resources of multiple teams to hedge against possible losses.
Terminate	Avoid	Terminate or avoid risk by abandoning the planned action or eliminating the goal altogether, prioritizing other goals in preferences.

the organization. The Governance and Management domains contain the details an internal auditor should use to advise the organization on moving towards those practices or to form a basis to evaluate how the organization is approaching, utilizing, managing, and monitoring AI.

"Reasonable assurance" is a term that is often referenced within the internal audit profession. From an internal control standpoint, reasonable assurance means there is a high likelihood the controls mitigate risk, but it is not absolute. The same rationale should be considered for internal auditors who are tasked with providing assurance surrounding AI.

Challenges

Several aspects of AI make assurance activities difficult for internal auditors, including:

- AI (or more specifically, the algorithms) is inherently highly complex – a more difficult "black box" problem.
- The capabilities and risks of AI are multiplying at a rapid pace.
- AI as an audit topic is evolving with limited tools or widely adopted approaches.
- There are limited training opportunities available to enhance auditing AI skill sets.

AI as an audit topic can seem overwhelming, but focusing on the following considerations will help internal auditors develop a positive, confident mindset:

- Internal auditors are not expected to be experts on every audit topic; rather, having a disciplined, methodical approach with a focus on critical thinking and identifying risk should be the objective for all audits, not just AI. Familiarity and a working knowledge of AI is vital; however, knowing all technical aspects of AI is not likely. It may be necessary to engage outside technical resources to assist with more technical aspects such as deciphering algorithms.

- Because AI is so highly complex and changing, it is unlikely that internal auditors will ever have a mastery of knowledge on the topic; think of auditing AI as a progression, not as a destination – increase the understanding of AI over time.

- Be willing to ask relevant questions about AI within the organization:

    - How does AI help us reach our strategic goals?
    - What risks are involved and how are we mitigating them?
    - Are there adequate internal controls surrounding AI-related processes?
    - Is the data that will be utilized for AI complete, accurate, and reliable?
    - How is AI tested prior to deployment to ensure biases do not exist?
    - How is AI tested after deployment to ensure biases do not exist?
    - How is AI governed?
    - How does the organization ensure adequate AI training and awareness exist?

As described in Part 2, understanding organizational use of AI starts with research and discussion. It is vital that internal auditors leverage the professional relationships they have developed. Being transparent with both management and the governing body is important. Explain in simple terms how you are thinking about AI as a topic and how you plan to engage the organization to learn more about it.

Internal audits of AI are a relatively new responsibility for many organizations. While as assurance providers, internal auditors are not expected to be experts on the topic of AI, they must identify opportunities to increase their knowledge and awareness of the subject. Gaining a better understanding of the more technical aspects of AI, such as algorithms, will be important for future professional education.

While AI certainly has complex elements, it is important to remember that it produces some form of output from the input it receives. From an assurance standpoint, internal auditors may never have absolute knowledge of all the inner workings of AI; however, helping an organization 1) evaluate what they are doing to ensure the input data is as accurate as possible, then 2) understanding how that output is scrutinized should be the practitioners' main objectives. Internal auditors apply these concepts today when performing IT audits of business applications. The common thread is the notion of traceability – ensuring the data and output is aligned with the business objectives and requirements of the AI use case.

Endnotes

1. Britannica, "Artificial Intelligence."
2. The IIA's Three Lines Model.
3. A.M. Turing, "Computing Machinery."
4. A.L. Samuel, "Checkers."
5. McCarthy, Dartmouth.
6. Weizenbaum, "ELIZA."
7. Humanoid Robot, "Waseda University."
8. Hsu, "Raj Reddy."
9. "Prometheus Project."
10. Britannica, "Deep Blue."
11. IBM AI Adoption Index.
12. IBM "Different types of AI."
13. Certes, "Types of AI."
14. IIA Global Perspectives & Insights.
15. National Cybersecurity Centre.
16. White House Fact Sheet.
17. IBM AI Governance e-book.
18. Doran, "Smart Way."
19. Moyer, ISACA, "Quantitative Approach."
20. COBIT, "Framework."
21. COSO, "Guidance."
22. Deloitte, "3 Lines."
23. "Gartner Glossary."
24. Intel, "GPU."
25. Forbes, 15 Biggest Risks.
26. IIA, CRMA.

PART 4

Practitioner's Guide and Glossary

The practitioner's guide is a simple checklist that internal auditors can utilize to begin their assessment of how the organization approaches, uses, manages, and reports on AI. Internal auditors can use the key points outlined in the Governance, Management, and Internal Audit sections of Part 3 to develop their audit plan or as considerations in an advisory role. Many of the aspects or considerations in the assurance section below are closely linked to items previously listed under the other domains.

This checklist is intended to provide a quick-start guide, but it should be customized based on organizational considerations, such as the extent to which AI already is being used and whether formal AI strategic planning, policies, procedures, processes, and reporting have been established.

Aspects or Considerations	Status/ Results
Create a vision, strategy, and prioritization for AI and update frequently.	
Link AI initiative to organizational strategic objectives. (This may include revenue enhancing use cases, or internal applications to reduce cost or improve efficiencies.)	
Ensure that ethics, bias, social, and legal aspects are included in the strategy.	
Determine how to measure success of AI initiatives, including goals and ROI.	
Ensure that the AI strategic plan is consistent with the organization's risk culture.	
Ensure that the AI strategic plan is consistent with the organization's values.	
Ensure that the AI strategic plan is formally communicated to the board.	
Ensure that the strategic plan includes AI resource optimization.	

Aspects or Considerations	Status/ Results
Ensure that the internal control environment is conducive for supporting AI. Consider what immediate policy changes are needed to support AI growth – adding a question about AI use in the third-party vendor management policy, for example.	
Define executive management responsible for overseeing AI initiatives.	
Establish a cross-functional AI Leadership Team to monitor all AI initiatives.	
Ensure that legal and compliance teams monitor all current and emerging regulatory requirements.	
Define the role of internal audit as an advisor and/or assurance provider.	
Ensure that Three Lines Model is in place and includes AI.	
Ensure that CISO (or equivalent) is involved in all AI initiatives.	
Ensure that third-party roles in AI initiatives are clearly defined and monitored.	
Ensure that finance/accounting tracks ROI on AI initiatives.	
Develop an AI acceptable use policy that is required for all employees.	
Develop policies and procedures for executing and maintaining AI initiatives.	
Develop policies and procedures for AI initiatives that utilize third parties.	
Ensure IT resources are sufficient to support AI initiatives and controls.	
Ensure staffing levels are sufficient to support AI initiatives and controls.	
Ensure HR recruiting has a focus on hiring practices for professionals with AI experience.	
AI leadership maintains required AI management knowledge.	
AI operational employees maintain required AI technical knowledge.	
All employees complete training regarding acceptable use and risks of AI.	
Include subject of AI in employee handbook and in new-hire orientation.	
Ensure that fair social, environmental, and economic aspects are considered in all AI-related projects.	
Ensure that AI-related data is secure, private, and confidential.	
Ensure that AI-related data is transparent, explainable, and responsible.	
Define objectives, goals, timing, and resource requirements for AI projects.	
Define operating responsibilities for all relevant employees in AI projects.	
Ensure that user access to AI is commensurate with job duties.	
Define data requirements and privacy considerations for AI projects.	

Aspects or Considerations	Status/ Results
Define applicable legal and regulatory requirements for AI projects.	
Perform AI project risk assessment to identify possible threats to success.	
Define possible biases, including ethical and social considerations for AI projects.	
Define success metrics or project key performance indicators for AI projects.	
Establish reporting parameters such as frequency, content, and milestones for AI projects.	
Establish testing approach to validate AI is working as intended prior to and after going live.	
Report on achievement of metrics/KPIs to executive leadership and board.	
Ensure reporting includes disclosure of bias, ethical, or social concerns.	
Ensure reporting includes compliance with legal and regulatory requirements.	
Ensure reporting includes disclosure of any unintended or negative results.	
Ensure reporting includes disclosure of possible data loss or privacy breaches.	
Ensure that related internal controls are evaluated and reported periodically.	
Include AI as part of the enterprise risk management (ERM) process.	
Identify risks that threaten AI strategic goals and objectives.	
Identify risks that may have ethical, social, environmental, or financial implications.	
Identify risks that are related to the use of third parties for AI.	
Ensure that a process is in place to capture new or emerging risks.	
Ensure that employees with AI risk management responsibilities are properly trained.	
Perform an AI-based risk assessment and update periodically.	
Prioritize AI-related risks based on severity score (impact and likelihood).	
Ensure there is a process in place to select appropriate risk responses, including monitoring progress of responses.	
Ensure the organization is engaged with the board regarding AI strategy, goals, and objectives.	
Ensure the organization provides periodic updates to the board regarding AI in a manner that is clear and easily understandable.	
Ensure the organization engages the board regarding risk management approach for AI.	
Perform initial internal and external research of AI.	
Determine if formal AI strategy has been developed.	
Conduct initial discussions with established organizational relationships (such as IT and CFO) to understand how AI is currently being utilized and managed.	

Aspects or Considerations	Status/ Results
Conduct initial discussions with AI/data science team (if applicable) and/or IT management.	
Create an inventory of current and planned AI uses.	
For current uses of AI, develop understanding of how it is being used, goals, and objectives.	
For planned uses of AI, develop understanding of approach, how risks are assessed, and plan for testing prior to deployment.	
Develop understanding of the following aspects of AI-related input data: • Governance. • Architecture. • User Access. • Cybersecurity Controls. • Processing Controls (integrity, accuracy, completeness). • Third-Party Considerations (SOC reports).	
Verify how AI is tested and reviewed to ensure it achieves its objectives and is free from biases, both pre-deployment and post-deployment.	
Verify that AI initiatives have clear objectives, and goals, and that projects are managed by an appropriate level of leadership.	
Verify that periodic reporting to the governing body is performed by management.	
Verify that AI is considered as a part of the enterprise risk management process, and includes risks related to: • Ethics. • Social and Economic Considerations. • Environmental Aspects. • Financial Implications. • Legal and Regulatory Violations.	
Verify that policies and procedures have been developed that outline how AI should be used and managed by the organization, including an AI acceptable use policy.	
Develop an understanding how an organization supports learning and training of AI to raise knowledge and awareness for all employees.	

Related IIA Standards

6.1 Internal Audit Mandate

6.2 Internal Audit Charter

6.3 Board and Senior Management Support

7.1 Organizational Independence

7.2 Chief Audit Executive Qualifications

8.1 Board Interaction

8.2 Resources

8.3 Quality

8.4 External Quality Assessment

Glossary

Definitions of terms marked with an asterisk are taken from the Glossary of The IIA's International Professional Practices Framework®, 2017 edition. Other definitions are either defined for the purposes of this document or derived from the following sources:

IBM. "Explainers." *IBM*. https://www.ibm.com/topics.

Institute of Risk Management. "Risk Culture." *Institute of Risk Management*, 2023. https://www.theirm.org/what-we-say/thought-leadership/risk-culture/.

Anderson, Urton; Michael J. Head; Steve Mar; Sridhar Ramamoorti; Chris Riddle; Mark Salamasick; Paul J. Sobel. *Internal Auditing: Assurance & Advisory Services, 5th Edition.* (Lake Mary, FL: The Internal Audit Foundation, 2022.) https://www.theiia.org/en/products/bookstore/internal-auditing-assurance-and-advisory-services-5th-edition/.

ISACA. "Glossary," *ISACA*. 2022. https://www.isaca.org/resources/glossary.

NIST Computer Security Resource Center. "Glossary." Gaithersburg, MD.: *NIST*. https://csrc.nist.gov/glossary.

Joint Task Force. *NIST SP 800-53: Security and Privacy Controls for Information Systems and Organizations, Revision 5, Appendix A: Glossary.* Gaithersburg, MD: NIST, September 2020. https://doi.org/10.6028/NIST.SP.800-53r5.

Grassi, Paul; Michael E. Garcia; James L. Fenton. *NIST SP 800-63-3: Digital Identity Guidelines, Appendix A: Definitions and Abbreviations.* Gaithersburg, MD: *NIST*, June 2017. https://doi.org/10.6028/NIST.SP.800-63-3.

Sawyer's Internal Auditing: Enhancing and Protecting Organizational Value, 7th Edition. (Lake Mary, FL: The Internal Audit Foundation, 2019.) https://www.theiia.org/en/products/bookstore/sawyers-internal-auditing-enhancing-and-protecting-organizational-value-7th-edition/.

Techopedia.com. "TechDictionary." https://www.techopedia.com/dictionary.

Algorithm – A clearly specified mathematical process for computation; a set of rules that, if followed, will give a prescribed result. (NIST Glossary).

Artificial intelligence – An advanced computer system that can simulate human capabilities, such as analysis, based on a predetermined set of rules. (ISACA).

Assessment – The process of identifying risks to organizational operations (including mission, functions, image, reputation), organizational assets, individuals, other organizations, and the Nation, resulting from the operation of an information system. Part of risk management incorporates threat and vulnerability analyses, and considers mitigations provided by security controls planned or in place. Synonymous with risk analysis. (NIST Glossary).

Audit Committee – A committee of the board charged with recommending to the board the approval of auditors and financial reports. (Sawyer's).

Backup – Files, equipment, data, and procedures available for use in the event of failure or loss, if the originals are destroyed or out of service. (ISACA).

Backup and recovery – Refers to the process of backing up data in case of a loss and setting up systems that allow that data recovery due to data loss. Backing

up data requires copying and archiving computer data, so that it is accessible in case of data deletion or corruption. Data from an earlier time may only be recovered if it has been backed up. (Techopedia).

Big data – A term used to refer to the large amount of constantly streaming digital information, massive increase in the capacity to store large amounts of data, and the amount of data processing power required to manage, interpret, and analyze the large volumes of digital information. (Internal Auditing, 5th Edition).

Black box testing – A testing approach that focuses on the functionality of the application or product and does not require knowledge of the code intervals. (ISACA).

Board* – The highest-level governing body (e.g., a board of directors, a supervisory board, or a board of governors or trustees) charged with the responsibility to direct and/or oversee the organization's activities and hold senior management accountable. Although governance arrangements vary among jurisdictions and sectors, typically the board includes members who are not part of management. If a board does not exist, the word "board" in the Standards refers to a group or person charged with governance of the organization. Furthermore, "board" in the Standards may refer to a committee or another body to which the governing body has delegated certain functions (e.g., an audit committee).)

Chatbot – A chatbot is an artificial intelligence program that simulates interactive human conversation by using key pre-calculated user phrases and auditory or text-based signals. Chatbots are frequently used by organizations to provide 24-hour customer relationship management (CRM) services. This type of software bot can also be used as an intelligent virtual assistant. (Technopedia).

Computer science – Computer science is the study of both computer hardware and software design. It encompasses both the study of theoretical algorithms and the practical problems involved in implementing them through computer hardware and software. The study of computer science has many branches, including artificial intelligence, software engineering, programming and computer graphics. (Technopedia).

Cybersecurity – Cybersecurity refers to any technology, measure, or practice for preventing cyberattacks or mitigating their impact. Cybersecurity aims to protect individuals' and organizations' systems, applications, computing devices, sensitive data, and financial assets against simple and annoying computer viruses, sophisticated and costly ransomware attacks, and everything in between. (IBM).

Data governance – Data governance refers to the process of managing the quality of data within an organization in order to ensure that, at all times during its life-cycle, data is accurate, available, consistent, secure, and usable. Business analysts and data scientists search for information across the enterprise to gain insight and understanding of that information, supporting business needs. (IBM).

Data integrity – The property that data has not been altered in an unauthorized manner. Data integrity covers data in storage, during processing, and while in transit. (NIST Glossary).

Deep learning – Deep learning is an iterative approach to artificial intelligence (AI) that stacks machine learning algorithms in a hierarchy of increasing complexity and abstraction. Each deep learning level is created with knowledge gained from the preceding layer of the hierarchy. (Technopedia).

Facial recognition – Facial recognition is a type of biometric technology that uses data to verify the presence of a human being's face in a digital capture. There are two main uses for facial recognition software: recognition and authentication. (Technopedia).

Governance* – The combination of processes and structures implemented by the board to inform, direct, manage, and monitor the activities of the organization toward the achievement of its objectives.

Internal control – An overarching mechanism that an enterprise uses to achieve and monitor enterprise objectives. (NIST Glossary).

Large language Model – A large language model (LLM) is a type of machine learning model that can perform a variety of natural language processing (NLP) tasks such as generating and classifying text, answering questions in a conversational manner, and translating text from one language to another. The label "large" refers to the number of values (parameters) the language model can change autonomously as it learns. Some of the

most successful LLMs have hundreds of billions of parameters. (Technopedia).

Machine learning – Machine learning (ML) is the sub-category of artificial intelligence (AI) that builds algorithmic models to identify patterns and relationships in data. In this context, the word machine is a synonym for computer program and the word learning describes how ML algorithms become more accurate as they receive additional data. (Technopedia).

NIST – National Institute of Standards and Technology.

NIST Artificial Intelligence Risk Management Framework – As directed by the National Artificial Intelligence Initiative Act of 2020 (P.L. 116-283), the goal of the AI RMF is to offer a resource to the organizations designing, developing, deploying, or using AI systems to help manage the many risks of AI and promote trustworthy and responsible development and use of AI systems. The Framework is intended to be voluntary, rights-preserving, non-sector-specific, and use-case agnostic, providing flexibility to organizations of all sizes and in all sectors and throughout society to implement the approaches in the Framework. The Framework is designed to equip organizations and individuals – referred to here as AI actors – with approaches that increase the trustworthiness of AI systems, and to help foster the responsible design, development, deployment, and use of AI systems over time.

Risk – Something that threatens achievement of an objective.

Risk appetite* – The level of risk an organization is willing to accept.

Risk culture – Risk culture is a term describing the values, beliefs, knowledge, attitudes, and understanding about risk shared by a group of people with a common purpose. This applies to all organizations - including private companies, public bodies, governments, and not-for-profits. (Institute of Risk Management).

Robotics – Robotics is the engineering and operation of machines that can autonomously or semi-autonomously perform physical tasks on behalf of a human. Typically robots perform tasks that are either highly repetitive or too dangerous for a human to carry out safely. (Technopedia).

Service Organization Company (SOC) Report – Audit report, completed by independent assessor, which evaluates an organization's internal control environment; can be provided by vendors to customers for assurance purposes that their internal controls are operating effectively.

Speech recognition – Speech recognition, also known as automatic speech recognition (ASR), computer speech recognition, or speech-to-text, is a capability that enables a program to process human speech into a written format. While it is commonly confused with voice recognition, speech recognition focuses on the translation of speech from a verbal format to a text one, whereas voice recognition just seeks to identify an individual user's voice. (IBM).

References

Ambrozi, Austin. "11 Challenges Of Adopting AI In Business (And How To Address Them Head-On)," *Forbes*, October 24, 2023. https://www.forbes.com/sites/forbesbusinesscouncil/2023/10/24/11-challenges-of-adopting-ai-in-business-and-how-to-address-them-head-on/?sh=6710c8474bfe.

Ankers, Damon. "Types of Artificial Intelligence: A Detailed Guide." *Certes IT Service Solutions*. https://certes.co.uk/types-of-artificial-intelligence-a-detailed-guide.

Appel, Gil; Juliana Neelbauer; David A. Schweidel. "Generative AI Has An Intellectual Property Problem." *Harvard Business Review*, April 7, 2023. https://hbr.org/2023/04/generative-ai-has-an-intellectual-property-problem.

Billington, James. "The Prometheus Project: The Story Behind One of AV's Greatest Developments." *ADAS & Autonomous Vehicle International*, August 22, 2018. https://www.autonomousvehicleinternational.com/features/the-prometheus-project.html.

Britannica. "Artificial Intelligence." *Encyclopedia Britannica*. 2023. https://www.britannica.com/technology/artificial-intelligence.

Britannica. "Deep Blue Computer Chess-Playing System." *Encyclopedia Britannica*. 2023. https://www.britannica.com/topic/Deep-Blue.

COSO. "Guidance, Internal Control Integrated Framework." *COSO*. 2023. https://www.coso.org/guidance-on-ic.

Deloitte. "Modernizing the three lines of defense model: An internal audit perspective." Deloitte, 2023. https://www2.deloitte.com/us/en/pages/advisory/articles/modernizing-the-three-lines-of-defense-model.html.

Doran, George T. "There's a SMART Way to Write Management's Goals and Objectives." *Management Review*, 70, November 1981, 35-36. https://community.mis.temple.edu/mis0855002fall2015/files/2015/10/S.M.A.R.T-Way-Management-Review.pdf.

EY. "The CEO Outlook Pulse – October 2023," EY, 2023. https://www.ey.com/en_us/ceo/ceo-outlook-global-report#:~:text=The%20CEO%20Outlook%20Pulse%20

Gartner. "Gartner Glossary." *Gartner*. 2023. https://www.gartner.com/en/information-technology/glossary/cpu-central-processing-unit.

Hsu, Hansen. "Meet 2021 CHM Fellow Honoree Raj Reddy." *Computer History Museum*. https://computerhistory.org/blog/meet-2021-chm-fellow-honoree-raj-reddy/.

Humanoid Robotics Institute. "History of Humanoid Robot in Waseda University." Waseda University. https://www.humanoid.waseda.ac.jp/history.html.

IBM. "eBook: Build responsible AI workflows with AI governance." *IBM*. https://www.ibm.com/account/reg/us-en/signup?formid=urx-51898.

IBM. "IBM Global AI Adoption Index." *IBM*, 2023. https://www.ibm.com/watson/resources/ai-adoption.

IBM. "Understanding the Different Types of Artificial Intelligence." *IBM*, October 2023. https://www.ibm.com/blog/understanding-the-different-types-of-artificial-intelligence/.

The Institute of Internal Auditors. *CRMA Study Guide and Practice Questions, 3rd Edition*. The IIA. 2023. https://www.theiia.org/en/products/bookstore/crma-study-guide-and-practice-questions-3rd-edition/.

The Institute of Internal Auditors. *Global Perspectives & Insights: the Artificial Intelligence Revolution*. The IIA. 2023. https://www.theiia.org/en/content/articles/global-perspectives-and-insights/2023/global-perspectives-insights-the-artificial-intelligence-revolution/.

The Institute of Internal Auditors. *The IIA's Three Lines Model: An update of the Three Lines of Defense*. The IIA. 2020. https://www.theiia.org/en/content/position-papers/2020/the-iias-three-lines-model-an-update-of-the-three-lines-of-defense/?msclkid=f2923355c01e11ecb401fe1dc46cbc38.

The Institute of Internal Auditors. *International Professional Practices Framework. 2017 ed.* (Lake Mary, FL: The Institute of Internal Auditors, 2017). https://www.theiia.org/en/products/bookstore/international-professional -practices-framework---ippf---2017-edition/.

Intel. "What is a GPU?" Intel. https://www.intel.com/content/www/us/en/products/docs/processors/what-is-a-gpu.html.

ISACA. "COBIT, An ISACA Framework." *ISACA*. 2023. https://www.isaca.org/resources/cobit.

ISACA. "Glossary." *ISACA*. 2022. https://www.isaca.org/resources/glossary.

Marr, Bernard. "The 15 Biggest Risks of Artificial Intelligence." *Forbes*. June 2, 2023. https://www.forbes.com/sites/bernardmarr/2023/06/02/the-15-biggest-risks-of-artificial-intelligence/?sh=16756d8b2706.

McCarthy, J; M.L. Minsky; N. Rochester; C.E. Shannon. "A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence." *BibSonomy*. https://www.bibsonomy.org/bibtex/24550126962fd8014daa80db1ffae4df2/mhwombat.

Moyer, Steven; Gunter Brunhart; Richard Dubs, Thomas Erickson, Robert Skalamera, Rob Kepner, Marty Meyer, "A (Kind of) Quantitative Approach to Organizational Risk Tolerance." *ISACA*, July 8, 2021. https://www.isaca.org/resources/news-and-trends/isaca-now-blog/2021/a-kind-of-quantitative-approach-to-organizational-risk-tolerance.

National Cyber Security Centre. "Guidelines for secure AI system development." *National Cyber Security Centre*. 2023. https://www.ncsc.gov.uk/collection/guidelines-secure-ai-system-development.

NIST. *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*, Gaithersburg, Md.: NIST, 2023. https://nvlpubs.nist.gov/nistpubs/ai/nist.ai.100-1.pdf.

NIST Computer Security Resource Center. "Glossary." Gaithersburg, Md.: NIST. https://csrc.nist.gov/glossary.

PwC. "PwC 2022 AI Business Survey (U.S.)." *PwC*. https://www.pwc.com/us/en/tech-effect/ai-analytics/ai-business-survey.html.

QuantumBlack AI by McKinsey. "The State of AI in 2023: Generative AI's Breakout Year." *McKinsey*. 2023. https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-in-2023-generative-AIs-breakout-year.

Samuel, A.L. "Some Studies in Machine Learning Using the Game of Checkers." *IBM Journal of Research and Development*. July 1959. https://ieeexplore.ieee.org/document/5392560.

Techopedia.com. "TechDictionary." https://www.techopedia.com/dictionary.

Turing, A.M. "Computing Machinery and Intelligence." *Mind*, Volume LIX, Issue 236, October 1950, Pages 433–460. https://academic.oup.com/mind/article/LIX/236/433/986238.

Weizenbaum, Joseph. "ELIZA—A Computer Program For the Study of Natural Language Communication Between Man and Machine." *Communications of the Association for Computing Machinery*. January 1966. https://dl.acm.org/doi/10.1145/365153.365168.

White House. "Fact Sheet: President Biden Issues Executive Order on Safe, Secure, and Trustworthy Artificial Intelligence." *The White House*. October 30, 2023. https://www.whitehouse.gov/briefing-room/statements-releases/2023/10/30/fact-sheet-president-biden-issues-executive-order-on-safe-secure-and-trustworthy-artificial-intelligence/.

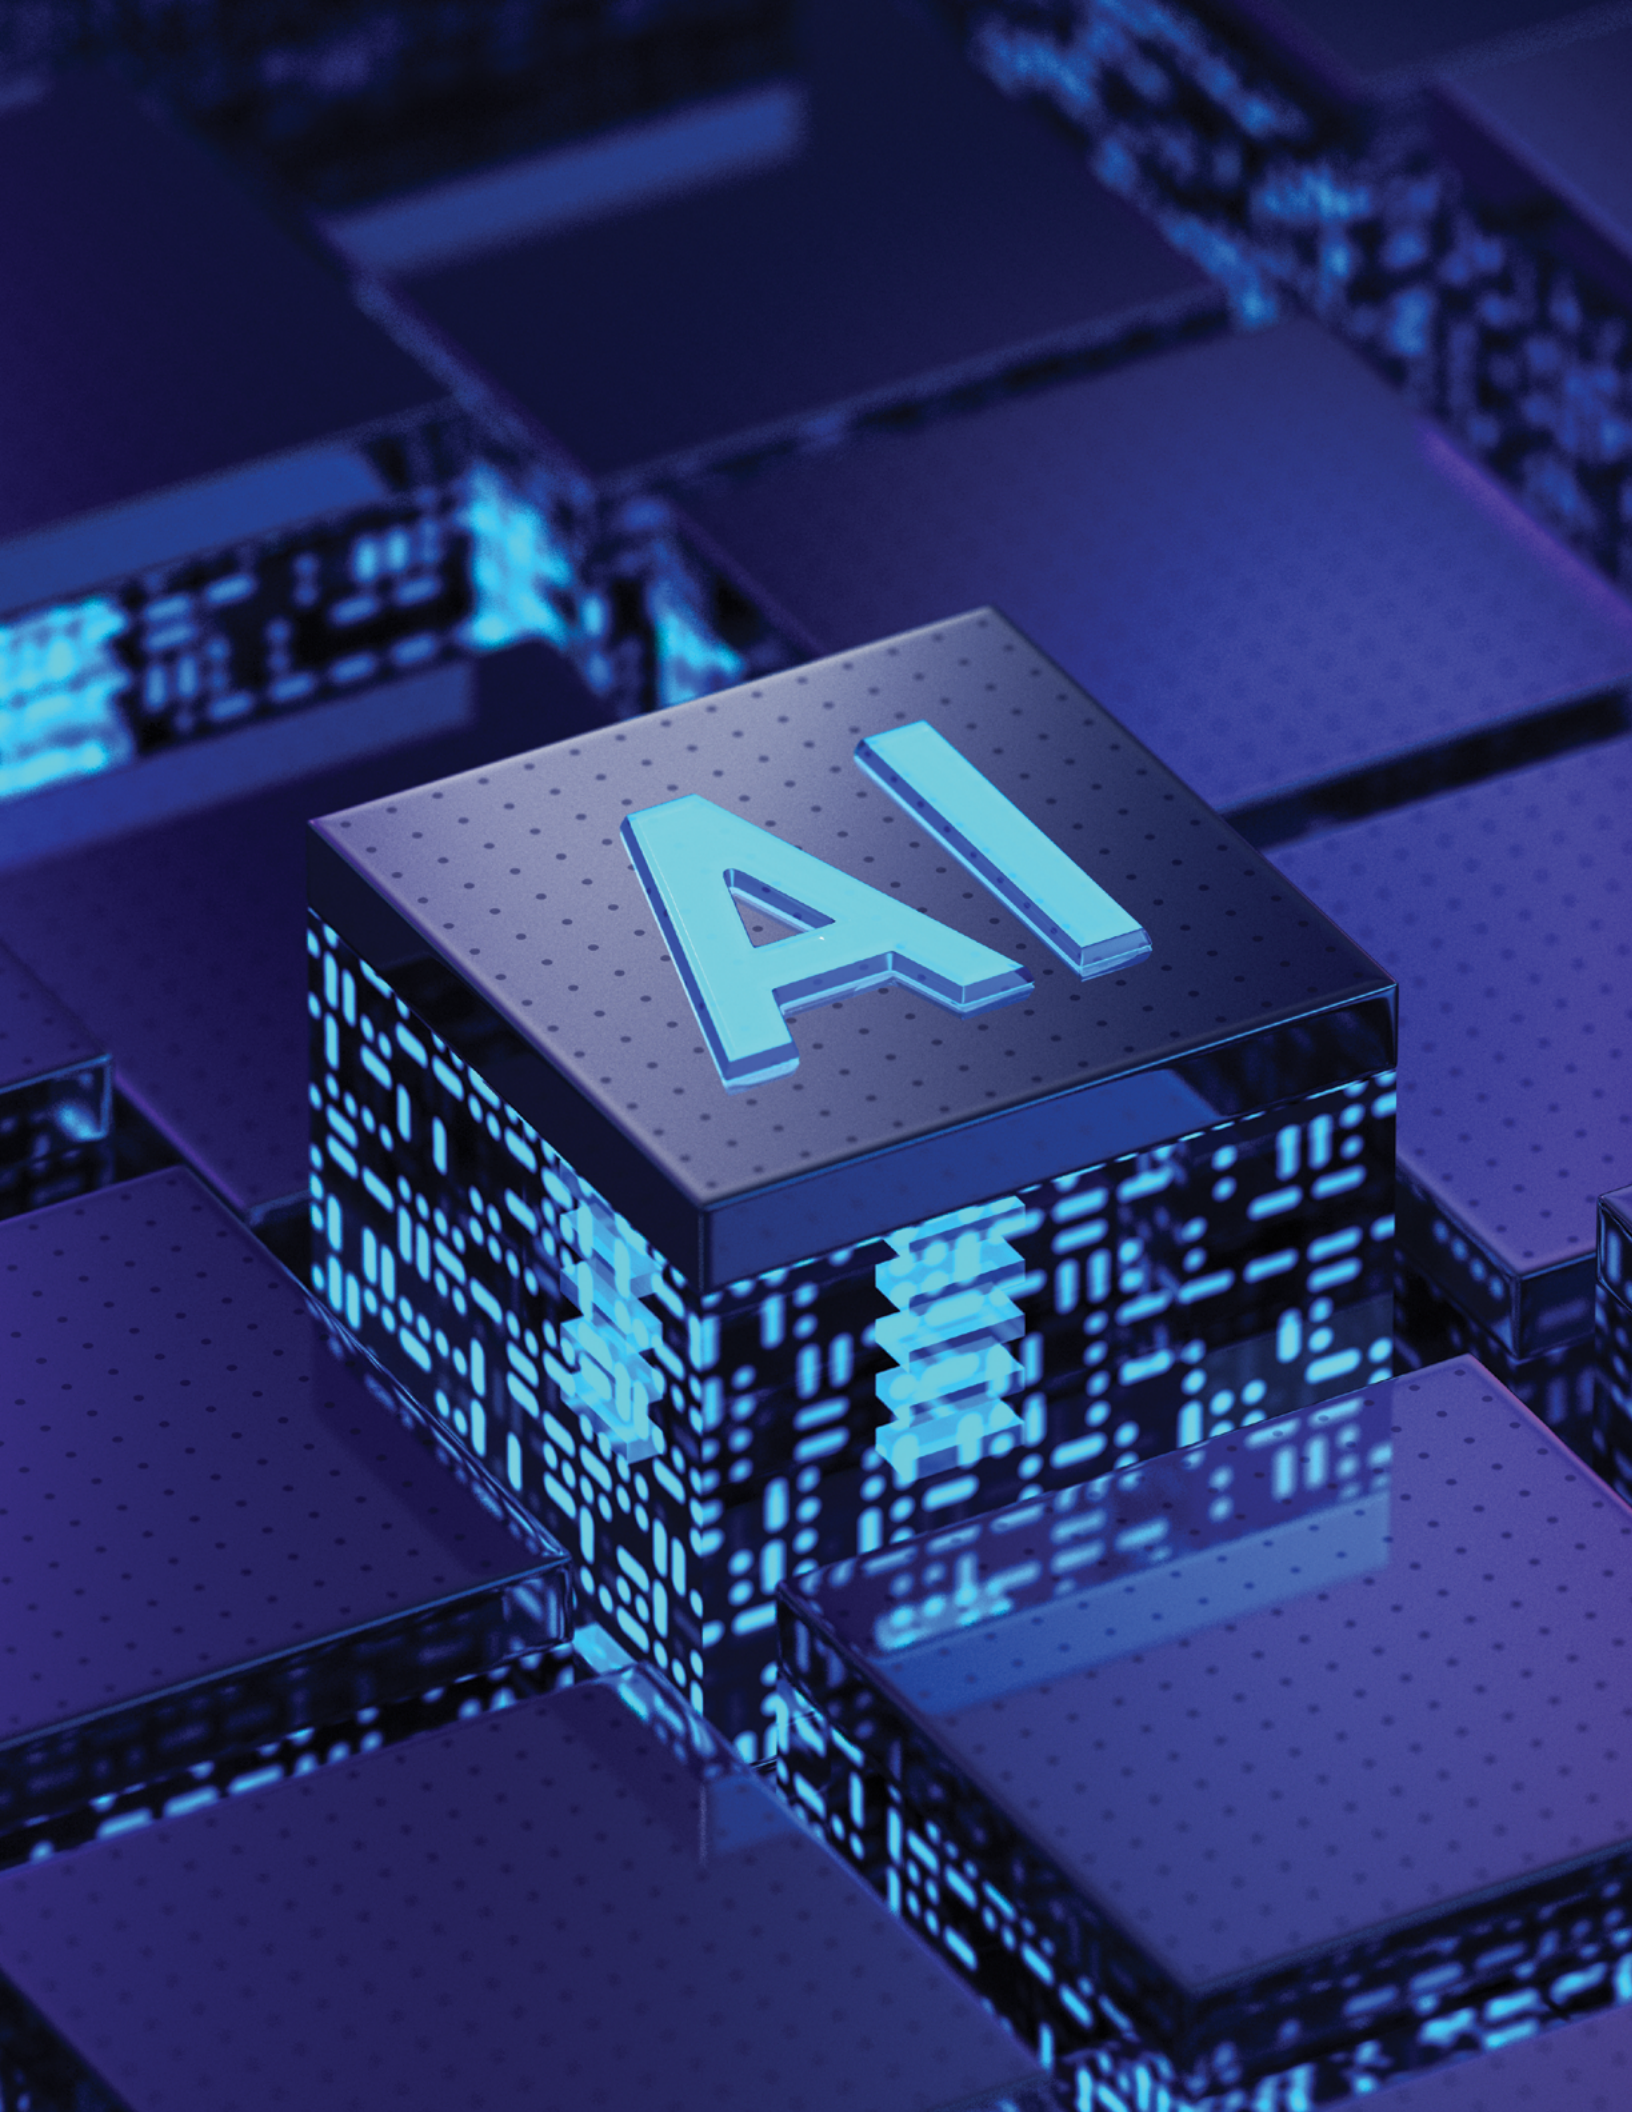
AI

Framework Development Team
George Barham, Allison Banzon, Anne Mercer, Kat Seeuws, Geoff Nordhoff.

Contributors
Andrew Cook, Pam Stroebel Powers, Robert Perez, Jim Enstrom, Scott Moore.

About The IIA
The Institute of Internal Auditors (IIA) is the internal audit profession's most widely recognized advocate, educator, and provider of standards, guidance, and certifications. Established in 1941, The IIA today serves more than 245,000 members from more than 170 countries and territories. The association's global headquarters is in Lake Mary, Fla., USA. For more information, visit **theiia.org**.

Disclaimer
The IIA publishes this document for informational and educational purposes. This material is not intended to provide definitive answers to specific individual circumstances and as such is only intended to be used as a guide. The IIA recommends seeking independent expert advice relating directly to any specific situation. The IIA accepts no responsibility for anyone placing sole reliance on this material.

IIA Global Headquarters
1035 Greenwood Blvd., Suite 401
Lake Mary, FL 32746 USA

The Institute of
Internal Auditors