Closing the AI trust gap: The role of CPAs in strengthening AI governance and risk management

CPA CHARTERED PROFESSIONAL ACCOUNTANTS CANADA

AICPA®

24-3714

ABOUT THIS SERIES

In collaboration with the American Institute of CPAs (AICPA), CPA Canada has issued this publication as part of a series of resources for CPAs on artificial intelligence (AI) in the age of generative AI. This is the second installment of this series, with a third paper to follow: Closing the AI trust gap (Part 2): The role of CPAs in AI assurance.

Read more about this series and other AI resources.

WHO SHOULD READ THIS?

This paper aims to empower all CPAs, regardless of their background or specialization, to understand key components of an AI governance framework that will support building trust in the use of AI systems. This paper is primarily targeted towards CPAs and business leaders working in industry or advisory roles but also provides useful insights for those working in public practice and government.

- **business leaders:** Executives, managers and decision-makers in organizations are encouraged to engage with this paper to gain insights into their role in AI strategy, governance and risk management.
- **CPAs working in industry:** CPAs in finance and accounting, internal controls and audit, and other operational functions will find valuable information on how to play essential roles in supporting the governance and quality of the design, development, implementation and utilization of AI for the benefit of their stakeholders and functions.
- **CPAs in public practice:** CPAs in audit and advisory roles in public firms will find key insights into AI's governance and risk management best practices. For those operating in advisory roles, this paper provides insights that can be used to transform client services and compliance.
- **CPAs in government:** Government-employed CPAs will find this paper valuable for understanding AI's governance and risk implications, aiding in the enhancement of public accountability and financial integrity.

For more detailed insights, explore the first paper in this series, Navigating the AI Revolution: Key Updates for Today's CPA, and other resources in our broader AI series such as A CPA's Introduction to AI: From Algorithms to Deep Learning and The Data-Driven Audit: How Automation and AI are Changing the Audit and the Role of the Auditor. These papers provide foundational knowledge and practical guidance for CPAs on AI-related topics.

Table of Contents

Acknowledgments

About the authors

CPA Canada and the AICPA would like to express their gratitude to the authors, Cathy Cobey, FCPA, FCA, CISA and Pradeepa Ramu of Ernst & Young LLP (EY). The views reflected in this paper are the views of the authors and do not necessarily reflect the views of the global EY organization or its member firms.

About EY

EY exists to build a better working world, helping to create long-term value for clients, people and society and build trust in the capital markets. Working across assurance, consulting, law, strategy, tax and transactions, EY teams ask better questions to find new answers for the complex issues facing our world today.

CPA Canada and the AICPA would also like to express their gratitude to the following members of CPA Canada, the AICPA and the AI Project Advisory Group for sharing their perspectives and their contributions to this paper:

Olivier Blais	*Co-founder & Head of Decision Science, Moov AI*
Peter Hargitai	*Partner, Canada National Digital Risk Solutions Leader & Americas Enterprise Technology Capabilities Leader, PwC*
Tim Herrod	*Chief Executive Officer, InTension*
Tina Kim	*Deputy Comptroller for State Government Accountability, Office of the New York State Comptroller*
Erika Ordway	*Partner, Audit Quality & Transformation, Deloitte*
Eric Rae	*Partner, Risk Services, KPMG*
Theo Stratopoulos	*PwC Professor, School of Accounting & Finance, University of Waterloo*

CPA Canada

Melissa Robertson	*Principal, Research & Thought Leadership*
Taryn Abate	*Director, Research & Thought Leadership*

AICPA

Erin Mackler	*Senior Manager, Assurance, Advisory & Innovation*
Carrie Kostelec	*Senior Manager, SOC and Related Service*

CPA Canada and the AICPA would also like to express their gratitude to the AICPA's Assurance Services Executive Committee who assisted in the review of this paper.

Introduction

In the rapidly evolving landscape of advanced technologies, **artificial intelligence (AI)** is reshaping the way businesses operate and how employees engage with their work. As AI becomes more common across industries, there is a clear need for AI to be developed and used responsibly. Debates about the opacity of AI systems, the potential biases they may harbour, and the extent of their security, accuracy and transparency have intensified the urgency for governance and oversight of AI. One response to these concerns has been the development of voluntary standards, guidelines and regulations to help organizations navigate responsible adoption of AI. With the rapid advancement of AI capabilities, regulations and frameworks are embracing embedded governance and risk management approaches that can be derived from existing control systems.

As AI systems permeate financial systems, auditing processes and decision-making frameworks, Chartered Professional Accountants and Certified Public Accountants (collectively, CPAs) are uniquely positioned at the intersection of technological innovation and responsible AI governance. As stewards of financial integrity and control environments, CPAs in executive leadership positions and industry roles have an opportunity to lead the use of AI not only in financial processes and reporting, but throughout organizations' business, operational and internal control processes. CPAs can also drive digital transformation and AI adoption by translating technical and governance knowledge into actionable business cases and prioritization of use cases across all parts of their organizations. Working alongside the information technology (IT), legal, compliance and privacy groups, CPAs — with their blend of financial acumen, professional skepticism and commitment to integrity — can play an important role in navigating the challenges in the development and adoption of AI.

This paper is primarily directed towards CPAs in internal roles within organizations. CPAs acting or consulting in internal roles can have a direct impact on the AI initiatives of their organizations. It is important to distinguish that responsibilities and influence over AI systems differ significantly for CPAs in public practice who are responsible for providing assurance services and advising clients on AI governance and risk management. CPAs providing assurance services in public practice provide an external, objective evaluation of AI practices, assessing transparency and accountability. They do not implement AI systems, but assess and provide guidance on the systems implemented by their clients, taking the appropriate actions to safeguard objectivity and independence.

Building off our recent paper, **Navigating the AI Revolution: Key Updates for Today's CPA**, which highlighted recent advances in AI and potential use cases, this paper is intended to help CPAs understand the potential role they can play in building trust in AI systems and implementing AI governance processes that support responsible development and usage of AI systems.

Prepared as a guide for CPAs, this paper will explore:

- key components of a robust AI governance framework
- the interrelationship between corporate governance, voluntary guidelines, laws and regulations in the context of AI governance
- an industry perspective of AI in finance

Governing AI

A necessary foundation for organizations venturing into AI adoption is a robust AI governance program. AI governance encompasses the policies, practices and standards that ensure that AI systems are developed and operated in a manner that is trustworthy and aligned with societal, ethical and legal norms. It involves a systematic approach to managing AI-related risks and establishing accountability throughout the AI lifecycle. AI governance encompasses a broad spectrum of activities — from the oversight of individual AI systems to the management of enterprise-wide AI programs.[1] While many governance practices and principles are universally applicable, the scale and complexity of implementation will vary based on the specific context and scope of AI deployment within an organization.

An effective AI governance program begins with establishing a robust AI governance framework, including the guiding principles, rules and standards that direct the organization's development, deployment and use of AI. This framework is often informed and shaped by a variety of voluntary AI guidelines, industry standards and regulatory requirements, which serve as benchmarks for best practices and compliance. It is important to note that while a robust AI governance framework lays the groundwork for navigating the AI landscape, it must be complemented by a commitment to oversight and supervision, education, interdisciplinary collaboration, transparent decision-making, continuous monitoring and stakeholder engagement.

The opportunity in AI governance for CPAs

Although applied differently, AI governance draws upon many of the governance practices familiar to CPAs. This provides an opportunity for CPAs working within organizations to play a vital role in designing, implementing and monitoring governance practices and controls to manage the complexities of AI systems. CPAs who stay current with leading practices and emerging regulatory requirements will be better positioned to contribute to the construction of responsible AI governance frameworks and can act as facilitators for sound AI governance — connecting the dots between theoretical principles and practical applications.

1 Enterprise-wide AI programs may also be referred to as AI management systems (e.g., ISO 42001).

Why does AI require a different approach?

AI presents unique challenges and risks that distinguish it from traditional technologies, necessitating a governance framework tailored to its specific characteristics. Traditional IT governance frameworks are typically centered on well-defined, deterministic systems with predictable behaviours. In contrast, AI systems, particularly those trained using **machine learning**, can be non-deterministic, self-adaptive and capable of learning from data in ways that are not always transparent or predictable. Here are several examples of factors that support a governance approach tailored for AI:

1. **Complex decision-making processes:** AI systems, especially those utilizing deep learning, can make decisions based on complex, non-linear computations that are difficult for humans to interpret. A governance framework for AI must address the need for transparency and explainability in the AI systems' decision framework.
2. **Dynamic learning and adaptation:** Unlike traditional systems that follow static rules, some AI systems are designed to evolve and adapt over time as they are exposed to new data. This continuous learning can lead to changes in system behaviour that were not anticipated or designed by human operators. AI governance must therefore include mechanisms for ongoing monitoring and validation to ensure that AI systems continue to perform as intended and do not deviate from ethical or legal standards, whether due to continuous learning or other operational issues.
3. **Data dependency and quality:** AI systems are heavily reliant on the quality and quantity of the data they are trained on. Poor data quality or biased datasets can lead to inaccurate or unfair outcomes. Traditional governance frameworks may not fully account for the nuances of data management in AI, such as the need for diverse and representative training data.
4. **Scalability and integration:** AI technologies can be scaled and integrated into various aspects of an organization's operations, potentially amplifying their impact and introducing systematic risks. AI governance must consider the broader implications of AI integration across different business units and the potential for systemic failures.
5. **Ethical and societal implications:** AI systems can have a broad range of implications, such as influencing employment, privacy and human autonomy that traditionally were managed through human decision-making. AI governance frameworks need to incorporate ethical guidelines and safeguards for AI systems' involvement in these types of ethical decisions.
6. **Regulatory compliance:** The regulatory environment for AI is rapidly evolving, with new guidelines and laws being proposed to address the unique challenges posed by these systems. AI governance must be flexible and forward-looking to ensure compliance with current and future regulations.

7. **Enhancing return on investment:** AI technologies represent significant financial investments for organizations, both in terms of initial development and ongoing maintenance. Proper governance ensures that these systems are used efficiently and effectively, aligning their output with business objectives and maximizing the return on investment.

Understanding the difference between strategy, policy and governance

Before discussing the intricacies of AI governance, it is helpful to clarify the interconnected roles of strategy, policy and governance. These components form the backbone of an effective AI governance framework, ensuring that AI deployment aligns with organizational objectives, ethical standards and regulatory requirements.

- **Strategy:** The strategy is the high-level plan outlining an organization's direction and the actions required to achieve long-term objectives. It defines the overarching goals for AI deployment, aligning with the organization's mission and competitive advantage. For instance, an AI strategy may focus on enhancing customer experience through personalized recommendations or improving operational efficiency with predictive maintenance.
- **Policy:** Policies are the specific guidelines and principles that steer decision-making within an organization. It ensures actions are consistent with the AI strategy and compliant with legal and ethical standards. AI policies typically cover ethical AI use, data handling practices, transparency, accountability and compliance with relevant regulations.
- **Governance:** Governance refers to the systems of rules, practices and processes by which an organization is directed and controlled. It provides the structure through which strategies and policies are implemented, monitored and enforced. It encompasses the entire framework managing how AI systems are developed, deployed and monitored to maintain accountability and align with ethical, legal and societal expectations.

Key components of a robust AI governance framework

In recent years, there has been a significant increase in the publication of AI risk management frameworks and guideline documents. These publications often share common themes, reflecting a collective understanding of the core principles necessary for effective AI governance. A robust AI governance framework, when designed and implemented properly, serves as a roadmap for organizations to manage the complexities and potential risks associated with AI technologies. This framework must be versatile and comprehensive, operating at multiple levels within an organization. At the enterprise level, it ensures a cohesive strategy and consistent standards across the entire organization. At the AI systems level, it addresses the unique operational and technical risks of each AI deployment. AI governance must also be tailored to meet the unique needs of specific departments and functions, recognizing that areas such as marketing, finance and human resources may encounter distinct challenges and require customized governance solutions.

The key components of a robust AI framework described in this paper reflect many of the recurring themes of other AI risk management guidelines, including voluntary guidelines, standards, laws and regulations related to AI governance. Existing mechanisms, such as those **discussed further in this paper**, can be valuable tools for assessing the maturity and capabilities of an organization's AI governance structures and providing benchmarks for continuous improvement and alignment with global standards.

The following page contains a visual representation of the key components of an AI governance framework. The diagram contains the core components or activities of the framework, followed by the monitoring and assessment activities over the AI program and its components in the outer layer. Lastly, corporate governance practices, laws and regulations, and voluntary guidelines and standards influence how these AI governance practices are designed, monitored and executed.

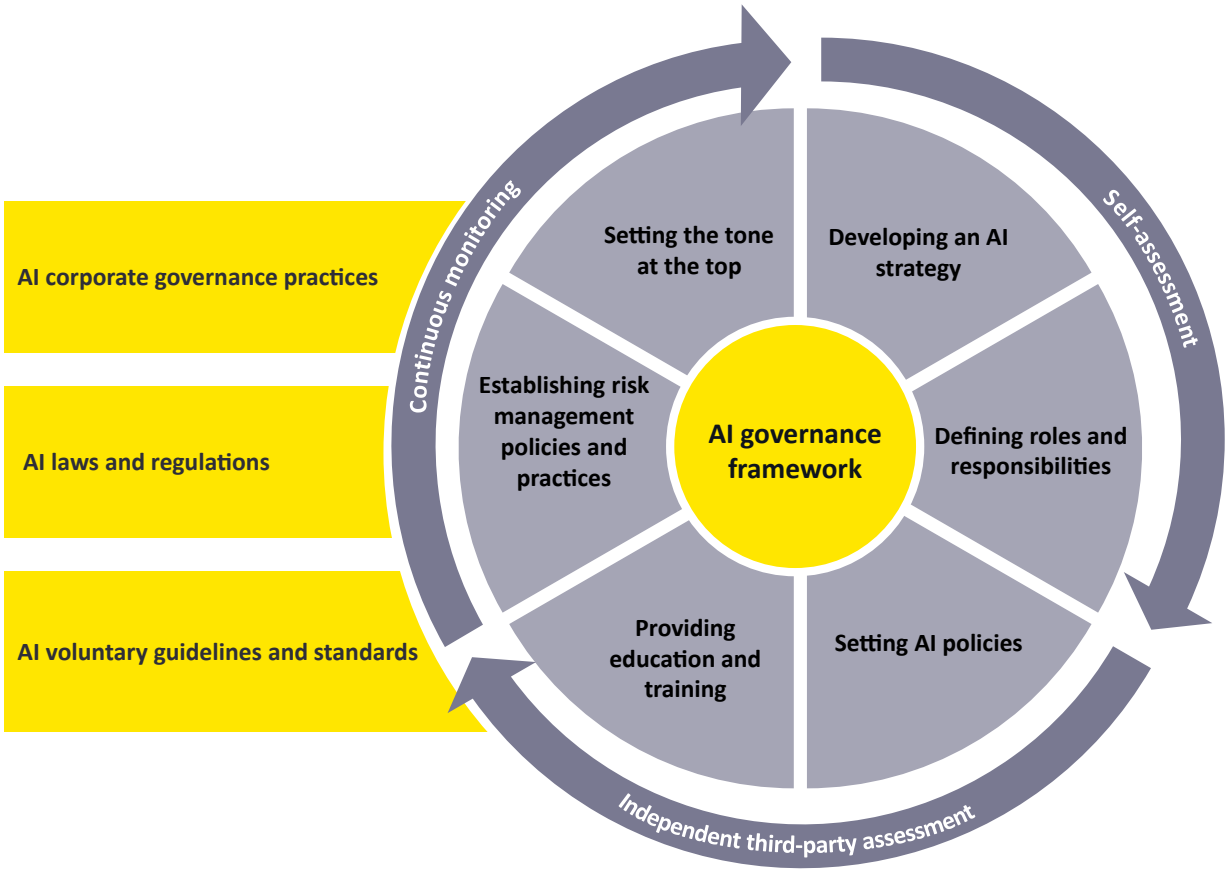

Source: EY

While not an exhaustive list of all the possible components of an AI governance framework, the above components set out a sound foundation for an AI governance program. Definitions for each of these AI governance framework components, and how CPAs might play a role within each area, are detailed below.

Setting the tone at the top

At the heart of a robust AI governance framework is the cultivation of a strong organizational culture rooted in ethical conduct. This begins with establishing a tone at the top, where leadership puts forth a vision and code of conduct that align with the organization's strategic goals. A crucial element in setting the right tone for AI includes adopting a common set of AI principles including fairness, explainability, data protection, transparency and accountability that guide all AI investment decisions from design to operation. Establishing responsibilities and enforcement mechanisms for these principles sends a message as to their importance.

Considering AI principles

The AI principles discussed in this paper are derived from the Organization for Economic Co-operation and Development (OECD)'s AI Principles, which include:

1. **Inclusive growth, sustainable development and well-being:** AI should benefit people and the planet by driving inclusive growth, sustainable development and well-being.
2. **Human-centred values and fairness:** AI systems should respect the rule of law, human rights, democratic values and diversity, and include appropriate safeguards.
3. **Transparency and explainability:** AI actors should commit to transparency and responsible disclosure to ensure that people understand AI-based outcomes and can challenge them.
4. **Robustness, security and safety:** AI systems should function in a robust, secure and safe way throughout their lifetimes, with continuous risk assessment and management.
5. **Accountability:** Organizations and individuals developing, deploying or operating AI systems should be held accountable for their proper functioning in line with these principles.

However, different frameworks may emphasize additional principles. For instance, the National Institute of Standards and Technology (NIST) AI Risk Management Framework (AI RMF) includes principles such as:

1. **Security:** Protect AI systems from adversarial threats and vulnerabilities.
2. **Privacy:** AI systems should adhere to privacy laws and respect individual privacy.
3. **Safety:** AI systems should not cause harm and operate safely.
4. **Reliability:** AI systems should perform consistently and accurately over time.

To avoid confusion and for a comprehensive coverage, organizations may consider adopting a holistic approach that incorporates principles from multiple frameworks, such as both the OECD and NIST, to guide their AI governance practices. Establishing clear responsibilities and enforcement mechanisms for these principles sends a strong message about their importance and ensures their integration into all AI-related activities.

Developing an AI strategy

AI strategy is a critical component that interlocks with the AI governance framework, guiding the organization in defining the who, what, when, why and how of AI investments. It outlines the strategic objectives for AI deployment, ensuring that the use of AI is purposeful and aligned with the organization's overarching goals. The strategy should consider the following key elements:

1. **Strategic value alignment:** The AI strategy must resonate with the organization's vision and contribute to its competitive advantage. It should identify areas where AI can add value, and enhance performance, innovation and customer experience. This would include a methodology for soliciting, capturing and prioritizing AI-related use cases.
2. **Principles and guidelines:** The strategy should be grounded in a set of ethical principles and guidelines that dictate the responsible use of AI. These principles ensure that AI systems are developed and used in a manner that is transparent, fair and respects privacy and human rights.
3. **Deployment roadmap:** The strategy should also include a clear roadmap for who will be involved in AI projects, what AI systems will be developed or acquired, when they will be implemented and how they will be integrated into existing processes and workflows.
4. **Success metrics:** The strategy must define how the organization will assess the success of its AI initiatives. This includes setting measurable goals, value-based performance indicators and regular review processes to evaluate the impact of AI on business outcomes and strategic objectives. At a more tactical level, it is essential for a strategy to consider and define precise success criteria for AI capabilities and systems.
5. **Risk management:** A thorough analysis of potential risks associated with AI, including ethical considerations, biases, security vulnerabilities, performance (e.g., reliability and accuracy) and compliance issues should be integrated into the strategy. This involves establishing protocols for risk assessment, mitigation and ongoing monitoring.
6. **Resource allocation:** The strategy should address the allocation of resources, including budget, talent, third-party partners and infrastructure, to support AI initiatives. It should also consider the need for upskilling and training to build AI literacy across the organization.
7. **Continuous learning and adaptation:** Recognizing the dynamic nature of AI, the strategy should include provisions for continuous learning and adaptation. This ensures that the organization remains agile and can adjust its AI initiatives in response to technological advancements and changing market conditions.

By incorporating these considerations into the AI strategy, organizations can ensure that their AI program is both robust and dynamic, and capable of steering AI initiatives towards delivering strategic value while adhering to ethical and operational standards. Another leading practice is to integrate AI into existing business strategies, viewing it as intrinsic to business operations rather than a distinct and separate strategic initiative. The approach emphasizes the seamless integration and alignment with overarching strategic goals, adapting AI initiatives as part of ongoing business strategy refinement. CPAs can be instrumental in driving AI investment decisions forward by supporting their organization in evaluating that the implications of AI investments and the economic outcomes of AI initiatives are in harmony with the organization's business and financial goals.

Strategy for selecting AI use cases

AI use cases are specific scenarios or applications where AI can be applied to achieve business objectives. It is essential that AI use case selection considers the full spectrum of costs, including design, development, operation and monitoring. CPAs can play a significant role in providing support in evaluating these use cases and their alignment with strategic goals and financial considerations. This involves not only traditional governance but also playing an active role in operational and capital allocation decisions, evaluating total benefits, costs and risks at both the micro (use case-specific) and macro (system-wide) levels.

Defining roles and responsibilities

Designing an effective AI governance framework necessitates the clear articulation of roles and responsibilities within an AI program, including a defined structure for oversight and decision-making. By delineating specific responsibilities, organizations can create a framework that promotes accountability at all levels, from the strategic direction set by senior leadership to the operational management by AI teams. CPAs involved in the design and oversight of governance structures for AI programs should ensure that roles and responsibilities are clearly defined, are appropriate, articulate areas of accountability and are managed with the same level of diligence and oversight as other critical business functions.

Defined accountabilities

When roles and responsibilities are well-defined, it becomes easier to identify who is responsible and accountable for each aspect of AI development and deployment, including ethical considerations, data management, model training and outcome monitoring. This clarity helps in assigning ownership for the performance and behaviour of AI systems, making it possible to trace decisions and actions back to individuals or teams. It facilitates prompt responses to any issues that may arise, such as biases in algorithms or data privacy concerns, and ensures that corrective measures are taken swiftly and effectively.

Segregation of duties

Clear roles and responsibilities also allow for the proper segregation of duties, whereby different individuals or groups are responsible for various stages of the AI lifecycle. Segregation of duties in the AI lifecycle reduces the risk of conflicts of interest and promotes objective evaluation of AI systems. This separation also encourages a culture of continuous improvement, as teams can independently assess and provide feedback on each other's work.

Targeted training

When roles are clearly established, it becomes possible to tailor training and development programs to the specific needs of each role, thereby enhancing the competencies required to manage AI responsibly. This targeted approach to skill-building supports the creation of a knowledgeable workforce that is equipped to handle the complexities of AI systems.

The case for a multi-disciplinary approach

The development of a competent, ethical and multidisciplinary workforce is integral to a good AI governance strategy. Collaborating with cross-functional teams allows organizations to obtain insights and perspectives from across the organization, including from those with a strong understanding of internal processes, resource availability, technology and legal and compliance requirements. Including diverse perspectives can also facilitate a multi-faceted examination of potential risks and rewards, allowing for strategic decision-making that balances innovation with ethical implications.

CPAs bring an understanding of existing business processes and internal controls and are knowledgeable of how to build trust in data, processes and technology. Including a CPA's expertise enables a more effective and holistic approach to AI governance and to the development and implementation of responsible AI practices within their respective organizations.

Setting AI policies

Defining clear and comprehensive AI policies is a key step when shaping the responsible parameters of AI adoption. These policies articulate the AI principles, scope and accountabilities within the organizational hierarchy. By defining a specific AI policy, organizations can provide a clear understanding of what is and is not considered AI and to which technologies the AI governance framework should be applied.

An AI policy should also serve as a complement to the AI strategy, outlining a selection and approval process that ensures only AI systems which are authorized and aligned with the organization's core values are pursued. This process should act as a gatekeeper, rigorously evaluating potential AI solutions for their ethical integrity, strategic fit, value generation, performance, cost effectiveness and contribution to the organization's overarching goals before granting approval for design and deployment.

What should an AI policy cover?

An AI policy document should succinctly outline the organization's approach to AI, including:

- **Purpose and scope:** Clarify the policy's intent and its coverage across the organization.
- **Definitions:** Offer clear explanations of AI-related terminology.
- **Acceptable use guidelines:** Guidance on what is approved or not approved for use.
- **Ethical principles:** Emphasize core ethical standards for AI use.
- **Governance structure:** Define roles and oversight mechanisms for AI management.
- **Compliance:** Address legal and regulatory adherence for AI applications.
- **Risk management:** Summarize procedures for identifying and mitigating AI risks.
- **Data management:** Set rules for data handling in compliance with privacy laws, fair use policies and other considerations and restrictions as needed.
- **Development and acquisition:** Establish criteria for creating or procuring AI systems.
- **Deployment:** Guide the integration of AI into operational processes.
- **Monitoring:** Outline methods for evaluating AI performance and impact.
- **Training:** Commit to educating staff on AI use and policy implications.
- **Review:** State the policy's review cycle for relevance and accuracy.
- **Accountability:** Detail enforcement and consequences for policy breaches.

Tools, such as the **Responsible AI Institute AI Policy Template,** are available to assist organizations with policy development.

Providing education and training

With AI continually evolving, ongoing education and employee upskilling are important components of governance. Integrating education into the governance structure enables the organization to remain adaptive and responsive to emerging challenges. AI education is also crucial in fostering understanding and adherence to AI policy, as it equips individuals across the organization with the knowledge and skills necessary to responsibly navigate the complexities of AI applications.

Upskilling a broad employee base in AI is crucial for organizations aiming to stay competitive in today's rapidly evolving digital landscape. It is not enough for only technologists to have a grasp of AI; business users across various departments must also understand its mechanisms and limitations to effectively adopt and utilize these technologies. Similarly, the education of data scientists and technical staff should not be limited to technical knowledge alone. Supplemental training in corporate governance, regulatory requirements, risk identification and management, ethical considerations, compliance and confidentiality is essential to ensure that AI applications align with legal standards and moral principles.

CPAs can play a pivotal role in this educational effort, leveraging their expertise in governance, risk management and compliance to lead training initiatives. By providing guidance on these critical areas, CPAs help promote AI applications that are not only technically sound but also ethically and legally compliant.

When employees comprehend how AI systems function and, most importantly, recognize the scenarios where AI may make mistakes, they can proactively mitigate risks and enhance decision-making processes. This familiarity not only fosters a culture of innovation but also accelerates the adoption and effectiveness of AI solutions within the company. This is also true of CPAs. As key stakeholders, CPAs must engage in continuous learning to stay abreast of AI technological advancements and responsible use considerations.

Establishing risk management processes

As AI systems introduce new or modified risks, it is important for organizations to enhance their existing risk-management processes or implement a new process tailored for AI that appropriately captures and assesses the impact of AI systems being used and the AI models that enable these systems to perform. While companies often apply their existing, principles-based risk assessment frameworks to AI, these processes must be comprehensive, encompassing the entire lifecycle of AI systems from design and development to deployment and maintenance. Such a process should be comprehensive, including the AI principles and risk considerations identified earlier on in this paper.

Risk assessments for AI systems

An AI-tailored risk assessment should consider the context in which AI is used, the sensitivity of the tasks it performs, and the impact of its decisions on stakeholders. With AI risk assessments, it is essential to align the evaluation process with the overarching strategy and business objectives, identifying and mitigating risks that could derail strategic plans or impede the achievement of business goals. By doing so, organizations can proactively address potential obstacles, ensuring that AI initiatives not only comply with governance protocols and regulatory requirements but also contribute effectively to the company's long-term vision and success.

An AI risk assessment typically encompasses several types of evaluations, each focusing on different facets of risk and impact, including:

1. **Impact assessment:** gauges the potential consequences of AI deployment on various stakeholders, including economic, social and environmental effects. It also considers the long-term implications for societal norms and structures.

Impact assessments are crucial to AI governance, as emphasized by standards such as International Organization for Standardization/International Electrotechnical Commission (ISO/IEC) 42001 and legislative measures such as the Artificial Intelligence and Data Act (AIDA) — proposed legislation in Canada. These assessments comprehensively analyze the potential effects of AI implementations across economic, social and environmental dimensions, as well as their lasting impacts on societal norms and structures. They evaluate economic, ethical, environmental, societal and legal implications for stakeholders, arising from both anticipated uses and potential misuses of AI systems. The goal is to systematically classify AI applications by their impact severity, pinpointing critical scenarios and areas of risk where adverse outcomes may occur. This structured approach empowers organizations to proactively mitigate risks and enhance the positive societal impacts of AI advancements.

2. **Privacy assessment:** scrutinizes the AI system's adherence to data protection laws and regulations, evaluating how personal data is collected, processed and stored, as well as the risks of data breaches or misuse.
3. **Security assessment:** examines the system's vulnerability to cyber threats and the measures in place to safeguard against unauthorized access or malicious attacks.
4. **Ethical assessment:** analyzes the alignment of AI operations with ethical principles and values, to validate that the decisions made by AI are fair, transparent and accountable.

5. **Technical robustness assessment:** evaluates the system's reliability, accuracy and performance under varying conditions, as well as its resilience to errors or unexpected inputs. It also addresses the reliability of data used by the AI system, ensuring that data quality is sufficient for the intended use cases and that measures are in place to handle data uncertainties. A technical assessment should also include evaluations of relevant infrastructure components including foundational models, cloud and data storage and transmission components.
6. **Legal compliance assessment:** evaluates whether the AI system operates within the bounds of applicable laws and regulations, including those specific to AI and related technologies.

Collectively, these assessments form a multi-dimensional approach to AI risk management, aiming to identify and mitigate potential adverse outcomes while maximizing the benefits of AI innovations.

As standard-setting bodies and public policy setters work to incorporate AI principles into practice, they tend to encourage a risk-based approach to AI by aligning compliance obligations according to the risk profile and intended use of the AI system. For example, control expectations and compliance requirements are tailored based on the specific risks identified in AI systems. Systems categorized as presenting higher risks are often subject to more stringent control measures, reflecting the complexity and nature of the risks involved. Conversely, systems categorized as lower risk may require fewer controls. In both cases, control should be tailored to mitigate the specific risks identified. The focus is on implementing controls that are proportionate and effective in managing the identified risks.

Example 1: Risk assessment for an AI health monitoring system

To illustrate potential risk considerations for an AI system, the following example outlines potential risk considerations for an AI system deployed in a healthcare setting.

Scenario: A healthcare organization is implementing an AI monitoring system that predicts patient recovery progress and potential complications in real-time. This AI system analyses a vast array of variables, including patient medical history, current status indicators, treatments and recovery patterns, to provide recommendations for patient care plans and treatment.

Response: A risk assessment for this AI system should consist of several considerations, including:

1. **Data quality and privacy:** The underlying AI model may use sensitive patient data, requiring robust controls for data accuracy, privacy and security. In addition, as the AI is used to provide recommendations for patient care or treatment decisions, high data integrity and quality will be required.
2. **Algorithm bias and fairness:** With demographic information, including age, gender and ethnicity, being an important input into determining treatment plans, there is a risk of model bias if certain demographic or medical condition groups are underrepresented in the training data.
3. **Transparency and explainability:** The design of the AI system should accommodate the need for medical professionals and patients to be able to interpret and explain the AI predictions and decision framework for informed decision-making.
4. **Dependability and performance accuracy:** The underlying AI model's complex predictions could significantly affect patients' health outcomes, so it is vital that its predictions are trained and monitored for elevated levels of precision and confidence.
5. **Accountability:** An accountability structure should be clearly documented to determine responsibility if predictions are found to be inaccurate or diagnostics are misleading (i.e., the healthcare professional users, the AI operator or the AI developer/vendors).

An integrated, multi-faceted AI risk assessment process, regularly evaluated by an objective team including CPAs, plays a critical role in identifying, evaluating and responding to these risks.

Example 2: Risk assessment for an AI-powered revenue recognition system

The below is an illustration of an AI system used to support revenue recognition, a process relevant to financial reporting.

Scenario: A corporation operating in Canada and the United States is implementing an AI-powered system to automate its revenue recognition processes. This AI system analyzes contractual agreements, delivery milestones, customer payments and other data to recognize revenue in compliance with International Financial Reporting Standards (IFRS) and U.S. generally accepted accounting principles (GAAP).

Response: A comprehensive risk assessment for this AI system must consider multiple factors critical to financial accuracy and regulatory compliance in both Canada and the U.S.:

1. **Data quality and accuracy:** The underlying AI model may rely on historical sales data and current market trends to recognize revenue, necessitating stringent controls for data precision and reliability.
2. **Model bias and fairness:** The underlying AI model could inadvertently incorporate biases based on the data it is trained on, leading to unfair or skewed revenue recognition practices. For example, if the real estate portfolio has changed significantly and now has a much higher concentration of residential vs. corporate real estate.
3. **Reporting standards:** The AI system must be designed to comply with complex revenue recognition standards such as IFRS 15 or Accounting Standards Codification (ASC) 606. Failure to do so could result in misstated financials and potential penalties for non-compliance.
4. **Integration and interoperability:** The AI system must seamlessly integrate with existing financial systems and databases including the real estate subledger and general ledger. Poor integration between upstream and downstream technologies can lead to data silos, errors in revenue recognition and inefficiencies in financial reporting.
5. **Explainability and transparency:** The AI system's decision-making process must be transparent and explainable to stakeholders to ensure trust and accountability. Without clear insights into how revenue recognition decisions are made, there could be skepticism and resistance from users, auditors and regulators, potentially undermining the credibility of financial reports.

A comprehensive and dynamic AI risk management framework, periodically reviewed by an impartial multi-disciplined group, is essential for the proactive identification, assessment and mitigation of potential risks related to AI systems.

In both scenarios, there are additional risk considerations that should be considered that would apply across most AI use cases. These include:

1. **Regulatory compliance:** The AI system adheres to all relevant regulations including privacy, data residency, discrimination, cross-border transfers and intellectual property (IP) ownership laws.
2. **Cybersecurity and resilience:** AI systems and related infrastructure components are subject to robust cybersecurity measures to protect against threats that could compromise its integrity and reliability.
3. **Business continuity and contingency planning:** Develop contingency plans to maintain business continuity in case of AI system failures or disruptions, including protocols for fallback procedures and manual interventions to mitigate risks associated with system downtime.
4. **Operational complexity and management:** The implementation and ongoing operation of an AI system introduce complexities in business processes. Effective management of these complexities is crucial to ensure the AI system enhances rather than hinders operational efficiency.
5. **Talent and expertise:** The successful deployment and maintenance of AI systems require specialized skills and knowledge. Organizations must consider the availability of skilled personnel or the need for training to manage and oversee AI operations effectively.
6. **Ethical considerations and social impact:** AI systems should be designed and operated in a manner that is ethically responsible and considers the broader social impact, including potential job displacement and the effects on various stakeholders.
7. **Change management and user adoption:** Introducing AI systems into an organization's workflow requires careful change management to ensure user adoption and to minimize resistance to new technologies.
8. **Scalability and futureproofing:** AI systems should be scalable to accommodate growth, and flexible enough to adapt to future technological advancements or changes in business strategy.
9. **Intellectual property and proprietary data:** Safeguarding intellectual property related to AI systems and ensuring that proprietary data is not compromised or misused.

AI risk mitigation and control mechanisms

Organizations, guided by best practices in AI governance, will need to establish risk mitigation and control mechanisms that address the risks identified in the risk assessment. AI-related controls can be designed through a combination of oversight, monitoring and intervention strategies.

Although there is currently heavy reliance on human-in-the-loop oversight, as AI systems become more complex and widespread, greater investment will be required in automated, machine-based supervision and monitoring. This automated monitoring goes beyond monitoring of system availability and performance to encompass behavioural monitoring – assessing whether the AI system is operating as intended and within acceptable boundaries. Common areas to monitor include fairness, reliability and data quality. Although rules-based behavioural monitoring is currently the norm, **AI agents** are emerging as an alternative approach. AI agents are autonomous software entities equipped AI capabilities. They are deployed within AI ecosystems to autonomously monitor, evaluate and verify the outputs of primary AI models. For example, utilizing advanced technologies such as natural language processing and machine learning, AI agents can interpret and assess the quality, accuracy and reliability of information generated by **large language models** (LLMs).

Real-time continuous monitoring of algorithmic behaviour is also needed for models that continue to learn after deployment and dynamically adjust their decision framework as they encounter new data. Given that AI threats are evolving rapidly, and formerly unidentified vulnerabilities are being regularly discovered, it is imperative that controls are sufficiently evaluated to address new threats and vulnerabilities.

As organizations increasingly incorporate AI into their operations, there is an opportunity for CPAs working within organizations to oversee and advise on the refinement and implementation of effective control mechanisms over AI. With an understanding of financial systems, data analysis and risk management principles, CPAs can contribute to the development of robust control frameworks that mitigate potential risks associated with AI systems.

CPAs in public practice may also be tasked with providing assurance to stakeholders regarding the reliability and integrity of AI-driven processes, and the completeness and accuracy of the outputs. Providing assurance over these aspects may involve obtaining an understanding and testing of the design and operation of the controls.

The following page contains an example of the controls that may be considered for a potential AI system. These controls should be considered on top of IT general controls and business process controls that might be considered for any new IT applications or changes.

AI control considerations: A credit scoring system

In this illustration, an AI system will soon be deployed to assess the creditworthiness of applicants based on a myriad of variables, from financial history to behavioural patterns. Given the impact on an individual's financial health of these decisions, the potential for model bias or inaccuracy raises ethical, legal and financial risks.

Potential controls for this AI system include:

- **interpretability tools**, enabling lenders to explain credit decisions transparently to applicants, facilitating understanding and trust in the decision-making process
- **bias detection algorithms** to flag potential discriminatory patterns in decision-making, enabling fairness and compliance with regulatory standards
- **data shift monitoring** that raises alerts about changes in the underlying data that might affect the underlying AI model's performance or introduce biases over time
- **feedback loops and continuous monitoring** to allow for real-time adjustments to align decision-making with ethical considerations and mitigate unforeseen risks.

Here, CPAs can play a role in overseeing or providing assurance over these controls, verifying data integrity and supporting ethical decision-making, demonstrating how they help bridge the complex world of AI with established governance principles. CPAs can specifically contribute by:

- **Data integrity assessments:** Verifying the quality and reliability of data fed into the underlying AI model, ensuring it aligns with accounting standards and regulatory requirements.
- **Ethical decision support:** Supporting ethical decision-making by validating the alignment of AI outputs with established governance principles and regulatory guidelines.
- **Impact on credit reserves:** Assessing how the underlying AI model's decisions influence the calculation and recording of credit reserves in financial statements.
- **Controls over data and model outputs:** Verifying that robust controls are in place to manage data quality, validate model outputs and incorporate them appropriately into financial reporting processes.

Model testing and validation

AI models are the core components of AI systems, serving as the algorithms that process data, learn from it and make decisions or predictions, thereby enabling the system to perform intelligent tasks. Model testing and validations, typically carried out by an objective model validation team consisting of data scientists, machine learning (ML) engineers and/or AI developers, involve testing the conceptual soundness, accuracy and reliability of AI models during the development and training phases. Conducted through rigorous testing

and benchmarking against predefined criteria, validations help with determining if AI models perform as intended. Once placed into production, AI models should be periodically re-validated to determine that they continue to operate within expected parameters. The frequency of re-validation is determined based on the identified risks of **data drift** or **model drift**, which may favour more continuous monitoring.

The outcomes of validations primarily cater to technical stakeholders involved in the model development process; however, they can also be relied upon by business stakeholders to gain comfort that AI models are ready to be moved into, or stay in, production. This is particularly crucial in contexts where AI is utilized for high-impact decision-making. For instance, validations performed by the model validation team can be leveraged by CPAs to build trust that AI systems underpinning any financial calculations and predications are performing as designed and are aligned with stakeholder expectations.

Managing risks of third-party AI systems

The changing landscape of generative AI, coupled with the escalating costs of building and maintaining AI infrastructure, has steered organizations toward increased reliance on third-party AI systems. For CPAs navigating this shift, effective risk management becomes paramount in enforcing responsible use, compliance and data security by vendors.

Here are some best practices to manage the risks of third-party systems and improve vendor management processes:

- Advocate for independent assessments of AI systems that rigorously assess the data protection, accuracy, reliability, fairness, explainability, transparency and accountability of algorithms bought or licensed from third-party vendors.
- Scrutinize the assessment process, emphasizing benchmark testing against industry standards and real-world scenarios to garner insights into the reliability and limitations of third-party AI applications.
- When independent assessments are not available, CPAs should take an active role in designing and executing AI system assessments to verify the AI systems maintain their integrity, accuracy and compliance with ethical standards over time.
- Evaluate the completeness and sufficiency of contractual and service level agreements, data handling practices, vendor transparency, incident management and overall alignment with organizational policies.
- Champion continuous monitoring mechanisms for third-party AI systems against changing standards, guidelines and regulations.
- Conduct **red teaming** and simulation testing to identify vulnerabilities and gaps in third-party AI system and processes.

Third-party assurance

As organizations mature in their AI governance practices, obtaining independent third-party assurance based on established guidelines can demonstrate an organization's commitment to accountable AI practices. While independent assurance over AI practices and systems, including AI management systems, is still emerging, it is a promising mechanism to instill trust across the AI ecosystem. Whether seeking assurance for the AI governance practices applied across the organization and/or over a specific AI system, independent assurance can provide confidence that an organization adheres to selected criteria. As assurance gains prominence, it is poised to become an integral part of the broader strategy for engendering trust in the deployment of AI across its ecosystem.

The third paper in this AI series will address assurance over AI systems in more detail, outlining why assurance is needed and how management can use it to gain an objective evaluation of their operations, which in turn aids in making informed decisions, improving internal controls and enhancing operational performance. Third-party assurance can also assist in building trust in AI systems for external stakeholders including regulators and customers.

The role of voluntary guidelines, standards, laws and regulations in AI governance

The interrelationship between corporate governance, voluntary guidelines and laws and regulations creates a synergistic framework for governing AI. As presented in the Governing AI section of this paper, AI corporate governance practices are the policies, governance frameworks and procedures put in place by an organization to promote and enforce responsible AI behaviour.

Voluntary standards and guidelines provide useful guidance in developing AI corporate governance practices and promoting standardization across industries and geographies. Developed by international organizations, industry bodies and expert consortia, these guidelines offer a global perspective on responsible AI design and deployment. They often encompass principles such as fairness, accountability, transparency and explainability, providing organizations with a benchmark against which to measure their AI practices. The voluntary nature of these guidelines allows for flexibility in implementation, empowering organizations to tailor their AI governance strategies to align with their specific contexts.

As AI technologies advance, governments and regulatory bodies are introducing legislated rules to address the societal and legal implications of AI. These laws and regulations set mandatory requirements for AI governance, imposing legal obligations on organizations to adhere to specific standards for areas that matter most to public policy setters and society at large. While varying across jurisdictions, common elements include data protection, bias mitigation, protection of minors, digital identity rights, intellectual property protection and accountability measures. New and expanded oversight functions are being created to enforce adherence to these legal frameworks, and organizations must align their AI practices with these regulatory mandates to avoid legal consequences.

Together, voluntary standards, guidelines, laws and regulations provide useful criteria for CPAs to assess the completeness, appropriateness and maturity of their organization's AI corporate governance practices.

Illustrated below is a visual representation adapted and updated as of June 2024, from a 2023 diagram crafted by the Boston Consulting Group (BCG) and the Responsible AI Institute,[2] which encapsulates the foremost guidance aimed at aiding organizations in shaping their AI governance practices. These valuable publications referred in the illustration not only serve as a resource for internal stakeholders, but also play a crucial role in building trust across a wide range of stakeholders.

AI corporate governance practices

These refer to the set of rules, controls, policies, and resolutions put in place to dictate corporate behaviour. These practices are designed to balance the interests of a company's many stakeholders, such as shareholders, management, customers, suppliers, financiers, government, and the community.

- OECD AI Principles
- OECD Framework for the Classification of AI Systems
- NIST AI Risk Management Framework
- HTI-1 Final Rule
- Coalition of Health AI
- Asilomar AI Principles
- IEEE's Ethically Aligned Design
- Montreal Declaration for a Responsible Development of Artificial Intelligence
- Artificial Intelligence | COSO

AI voluntary guidelines & standards

These are voluntary commitments that organizations choose to adhere to in order to maintain a level of operational or ethical standards beyond what is required by law. These can include codes of conduct or ethical guidelines that help shape the culture and operations within an organization.

- Biden Administration Executive Order
- White House's voluntary commitments from leading AI companies
- Canada's generative AI code of conduct
- ISO/IEC 22989, 42001
- ISO/IEC JTC 1/SC 42
- IEEE P7000 series of standards projects
- CEN/CENELEC standards development
- RAI Institute's Certification Program for AI Systems

AI laws & regulations

These are legal requirements established by governmental bodies that must be followed by organizations. They serve as a formal framework for ensuring compliance with minimum standards of practice in various areas such as finance, labour, environment, and corporate governance.

- AI Act (EU)
- The Artificial Intelligence and Data Act (Proposed) (Canada)
- NYC Local Law 144 of 2021 (regulating automated employment decision tools)
- Colorado Final Rule SB24-205-Consumer Protections for Artificial Intelligence
- American 'Data Privacy and Protection Act, Section 207' (U.S.)
- Digital Services Act/Digital Markets Act
- Internet Information Service Algorithm Recommendation Management Regulations (China)
- Future of Artificial Intelligence Innovation Act of 2024
- Secure Artificial Intelligence Act of 2024

Original image from BCG, Responsible AI Institute; updated by EY as of June 2024.

2 BCG, Responsible AI Institute, A Guide to AI Governance for Business Leaders, November 23, 2023.

For internal stakeholders, such as project sponsors, business operators and risk and control partners, these publications serve as a starting point to evaluate the maturity and sufficiency of AI governance practices within an organization. They provide a benchmark against which the effectiveness of existing frameworks can be measured. Although there are ongoing efforts to evolve the above publications in response to public comments and multi-stakeholder consultations, in their current forms, they may provide objective criteria for CPAs working within organizations to evaluate the design and implementation of their organization's AI governance program.

Common themes across AI voluntary guidelines, standards and regulations

As international and regional bodies intensify their efforts to formulate guidelines, standards and regulations for the responsible advancement of AI, the landscape for compliance is becoming increasingly intricate. Amidst the diversity in approach, terminology and objectives across jurisdictions, there is a shared commitment to interoperability. Despite the complexity, there is a common thread among these initiatives — an impetus to protect the public and foster the development of safe, trustworthy AI while encouraging investment in AI.

In analyzing the evolving AI standard-setting and regulatory landscape, seven common themes emerge:

1. There is consensus in the fundamental principles that AI systems should uphold, including respect for human rights, sustainability, transparency and strong risk management.
2. Most jurisdictions are taking a risk-based approach, tailoring their AI governance and control requirements to the perceived risks of AI such as privacy, non-discrimination, transparency and security.
3. Public policy bodies are working towards an approach that balances the need for sector, technology and use case specific rules. Building upon general, agnostic approaches, policy setters are starting to develop complementary use cases and sector-based specific rules.
4. AI rulemaking considers other digital policy priorities and existing standards, laws and regulations in the areas of digital identity, cybersecurity, data privacy and intellectual property protection.
5. Many jurisdictions are using regulatory sandboxes as a tool for the private sector to collaborate with policymakers to promote safe and responsible AI, as well as to consider the implications of higher-risk innovations where closer oversight may be needed.
6. Jurisdictions are defining different roles involved in the design, development and operation of AI and are specifying accountabilities and obligations for each (e.g., AI developer, provider and user).

7. Public policy setters are working collaboratively to create laws, regulations and standards that are interoperable across borders and different AI use cases with international standards serving as the foundation on which laws and regulations are based.

While each jurisdiction is focused on translating AI principles into practice, their approaches range from voluntary guidance to mandatory rules. Many began with guidelines which are now being encoded into concrete laws and regulations as the understanding and governance of AI matures. As AI technologies continue to evolve and expand in their applications and use, continuous enhancement of guidance by public policy setters will persist. This trajectory, as is common with other emerging technologies, indicates that AI guidelines and regulations will continue to evolve to become more prescriptive and use case specific as the technology advances.

An industry perspective on AI in finance

Tim Herrod is the CEO and Co-Founder of InTension Inc., a consultancy firm dedicated to driving transformative strategies and innovative solutions for clients. Tim currently serves as CPA Canada's representative on the International Federation of Accountants (IFAC) Professional Accountants in Business Advisory Group.

AI presents a transformative opportunity for finance functions, enabling significant strides towards the operational excellence business expects, and talent demands. Marrying top-down expectations of "AI everywhere" to bottom-up realities (e.g., a series of micro-improvements via activity-by-activity endeavours) is essential in leveraging these amazing new tools for redirecting time and effort to the highest value-creating uses while preserving governance and trust responsibilities. The following suggestions do just that.

Leading an AI-driven finance strategy

Ownership and strategy: As finance leaders, CPAs must take ownership of AI initiatives. Develop and drive the finance strategy with AI at its core, ensuring alignment with business goals.

Bottom-up use cases: AI value is driven bottom up, use-case-by-use-case. Identify and prioritize AI use cases within the finance function and then support other functions in doing the same. Present these use cases to IT and other departments to ensure they are supported and integrated effectively.

Optimizing finance operations

Automate low-value tasks: Deploy AI to handle repetitive tasks such as data entry, transaction processing and reconciliations to free up capacity for high-value activities, enhancing efficiency and accuracy.

Enhance data and insights and improve decision making: Utilize generative AI tools to analyze and integrate large datasets in ways not previously possible, providing deeper insights about the past and forecasts of the future and supporting faster, more accurate[3] decision-making for and by the business. Provide business users with "finance-endorsed" AI-powered interfaces that

3 Accuracy of AI is dependent on appropriate controls as without such, AI may be susceptible to inaccuracies or hallucinations.

offer real-time insights and predictive analytics. Enabling DIY users (they will find a way on their own) with internally sanctioned and properly governed data, processes and tools will enable better decision-making at all levels of the organization.

Empowering teams and functions

Redefine roles and responsibilities: Reorganize teams to focus on strategic cognitive vs. repetitive tasks, fostering an environment where individuals can contribute to higher-value activities visible to the business, and higher job satisfaction.

Skill development: Invest in training for AI literacy, advanced analytics and prompt engineering to equip teams with the skills needed to leverage AI effectively.

Integrated collaboration: Work closely with IT and other functions to ensure seamless integration of AI solutions, enhancing finance data sharing and overall decision quality and efficiency.

Delivering value to external stakeholders

Enhance customer, operations and supplier relationships: Use AI to uncover insights, streamline integrated business planning, enhance interactions and improve service delivery. CPAs leveraging AI can help non-finance professionals better analyze and understand customer needs, translate them into financial and risk decisions, and optimize end-to-end supply chain operations.

Sustainable practices: Leverage AI to support sustainable business practices, optimizing resource use and enhancing relevant environmental, social and corporate governance factors in financial decision-making about opportunities, risks and compliance.

Managing risks and ensuring trust

Data governance and security: Own and maintain stringent data governance practices to ensure data quality and security. This is critical for building and preserving trust in AI-driven processes.

Ethical and responsible AI use: Establish ethical and responsible guidelines and accountability for AI applications, ensuring they align with organizational values and regulatory requirements.

Final thoughts

CPAs are uniquely positioned to lead the transformation of finance functions through AI. By focusing on optimization, team empowerment, enhanced decision-making and stakeholder value, they can drive significant improvements in efficiency and strategic impact. This is a call to action for CPAs to lead the charge on responsible AI adoption by embracing an integrated mindset, raising literacy, leveraging AI effectively and positioning the finance function at the forefront of technological innovation. And, critically, demonstrating that this is a great – and essential – place to work for amazing finance talent looking for broad experience and career growth in a rapidly changing world.

Conclusion

CPAs, including those in executive roles such as CFOs, stand at the forefront of addressing the multifaceted challenges presented by AI and developing corporate strategies and responses to balancing digital transformation and innovation responsibly. Navigating the evolving landscape of AI necessitates an understanding of new risks introduced by AI, as well as an ability to steer organizations towards effective AI governance programs. By actively leveraging expertise in governance, risk management and assurance mechanisms, CPAs can play a critical role in promoting transparent and responsible conduct in AI adoption.

To prepare for this role, CPAs should take the following actions:

- **Deepen their AI and data expertise:** Engage in advanced studies in AI, machine learning and big data, focusing on applications in financial contexts.
- **Craft comprehensive AI policies:** Develop detailed governance frameworks that address AI ethics, data security and algorithmic transparency to mitigate risks and ensure accountability in automated systems.
- **Commit to ongoing education and collaboration:** Regularly update skills through professional courses and work closely with experts in technology, ethics and data science to build robust AI assurance strategies.

In this three-part series on AI, our first paper *Navigating the AI Revolution: Key Updates for Today's CPA* explored the advances in AI and the opportunities for CPAs to understand and use this emerging technology. In this paper, the second in this series, we highlighted the opportunity for CPAs to build trust in AI through playing a leading role in the design and operation of governance and controls over AI systems. The third and final paper will explore the role of CPAs in building trust in AI systems through third-party assurance engagements over AI systems.

Appendix: Glossary of terminology

Artificial intelligence (AI) system[4]**:** The OECD defines an artificial intelligence (AI) system as a machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations or decisions that can influence physical or virtual environments. Different AI systems vary in their levels of autonomy and adaptiveness after deployment.

AI agent: AI agent refers to a software entity that acts autonomously, taking actions on behalf of or in collaboration with a user or another program. AI agents are designed to perceive their environment through data acquisition, interpret the collected data and act to achieve specific goals. They often utilize techniques such as machine learning, natural language processing and decision-making algorithms to perform tasks ranging from simple to complex, all while adapting to changing circumstances and learning from experience.

Data drift: Data drift in AI refers to the phenomenon where the statistical properties of the input data used to train a machine learning model change over time, leading to a degradation in the model's performance. This change in data distribution can occur due to several factors such as shifts in user behaviour, changes in the environment, or updates to the underlying systems generating the data. Data drift poses a significant challenge for machine learning models as they may become less accurate or reliable over time if not properly monitored and addressed. Detecting and mitigating data drift is essential for maintaining the effectiveness and performance of AI systems in real-world applications.

Generative AI: Generative AI refers to a class of artificial intelligence models and algorithms that are designed to generate new content, often in the form of images, text or other data types. These models are trained on large datasets and learn patterns, structures and styles from the input data. Once trained, they can generate content that shares similarities with the training data.

Large language model: An LLM is a deep learning algorithm that can perform various natural language processing tasks like language generation and classification. These models use transformer architectures and are trained on massive datasets, which makes them capable of recognizing, translating, predicting or generating text and other content. When given a prompt or question, an LLM uses neural networks to predict the next logical word, producing coherent output.

4 Explanatory memorandum on the updated OECD definition of an AI system | OECD

Machine learning: Machine learning is a subset of AI that focuses on the development of algorithms and statistical models that enable computers to perform specific tasks without explicit instructions, relying on patterns and inference instead. It involves the use of data to train algorithms to learn and improve over time, allowing machines to make predictions, decisions and automate tasks based on past experiences without being explicitly programmed for each task.

Model drift: Model drift in AI refers to the degradation in the performance of a machine learning model over time, even when the input data remains consistent. Unlike data drift, which occurs due to changes in the input data distribution, model drift occurs when the underlying relationships between the input features and the target output change or when the model's assumptions no longer hold true. This can happen as a result of shifts in the environment, changes in user behaviour, or other external factors that were not accounted for during model training. Detecting and addressing model drift is crucial for maintaining the accuracy and relevance of machine learning models in production systems.

Red teaming: An advanced form of network penetration testing where a contracted or in-house red team (as opposed to a defending blue team) emulates an advanced threat actor using physical, digital and human vectors to identify gaps in the organization's defensive strategy.

Third-party assurance: Third-party assurance refers to independent verification to assess and confirm assertions made by management in relation to a company's processes, systems or financial information. This form of assurance is typically sought by organizations to validate their assertions to customers, regulators or other stakeholders to meet certain standards or regulatory requirements or instil trust. The third-party, often an audit firm, conducts evaluations and provides a report or opinion that the organization can use to demonstrate due diligence, manage risks and build trust with external parties.

Validations: In the context of AI, validations refer to the process of rigorously testing and assessing the conceptual soundness, accuracy and reliability of machine learning models during their development, training and deployment phases. Validations involve benchmarking the performance of AI models against predefined criteria and evaluating their ability to produce accurate and reliable predictions or outputs. These assessments help verify that AI models perform as intended and meet the desired performance standards before being deployed in real-world applications. Additionally, validations may be periodically conducted post-deployment to verify that the models continue to operate within expected parameters and to identify any potential degradation in performance over time.

CHARTERED
PROFESSIONAL
ACCOUNTANTS
CANADA

277 WELLINGTON STREET WEST
TORONTO, ON CANADA M5V 3H2
T. 416 977.3222 F. 416 977.8585
WWW.CPACANADA.CA