Demand for Governance in the Age of AI

Gregg Wishna, Will Shuman -- Protiviti Inc.

September 20, 2024

TAC

The Institute of Internal Auditors

With You Today

Gregg Wishna
Director

Internal Audit and Financial Advisory

- Director in Protiviti's Internal Audit and Financial Advisory practice and based out of Atlanta, GA.
- Earned a Master of Science as well as a BS in Decision and Information Sciences from the University of Florida.
- Certified as a CISA and has over 19 years consulting and audit experience with Protiviti; leading projects and working closely with Senior Management.
- Specializes in topics of Internal Audit innovation through advanced analytics and evaluating core Data Governance and Data Management functions.

Will Shuman
Director

Technology Consulting

- Director in Protiviti's Technology Consulting practice and based out of Atlanta, GA.
- Will is a leader within the Enterprise Data and Analytics team focusing primarily on Data Governance program establishment and supporting clients in ensuring data is used appropriately throughout their organization. Will also acts as a leader within the AI Governance space to support emerging technology and its impacts on data management.
- He serves as a board member on University of Georgia's MIS Advisory Board, a society board member for the Data Governance Society of TAG (Technology Associate of Georgia) and is a CISA (Certified Information Systems Auditor).

Overview of Today's Discussion

1

Key Concepts for Data Governance and AI

2

Explore

- AI Use Cases for Internal Audit
- Risks
- Key Considerations for Next Steps

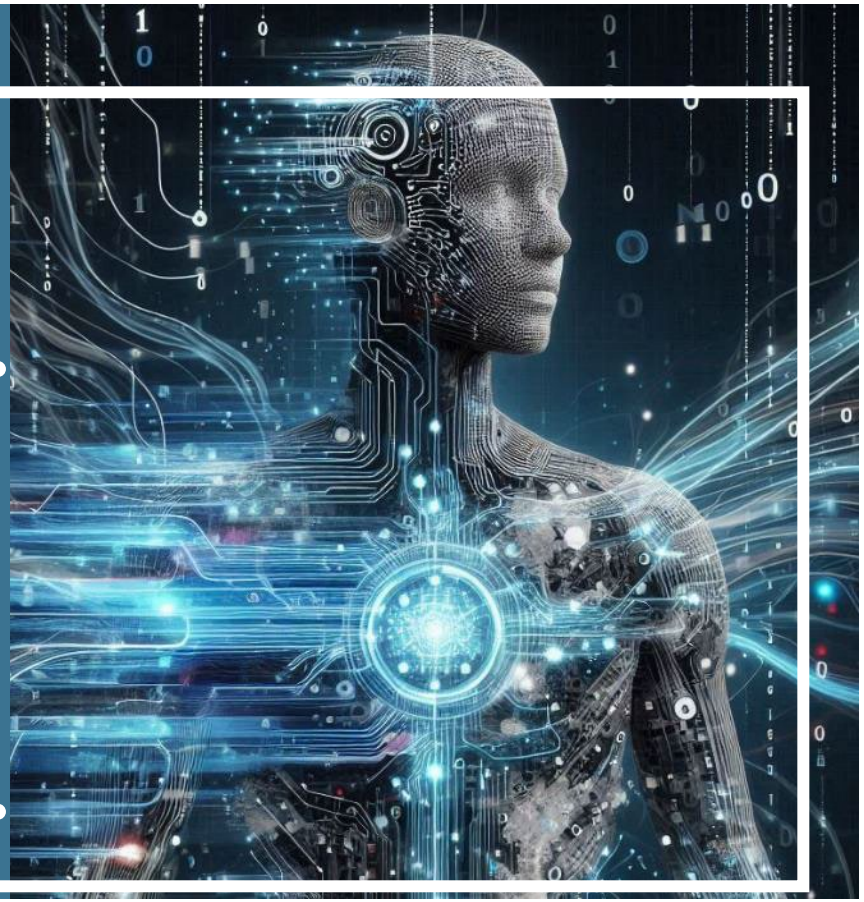

The AI Renaissance

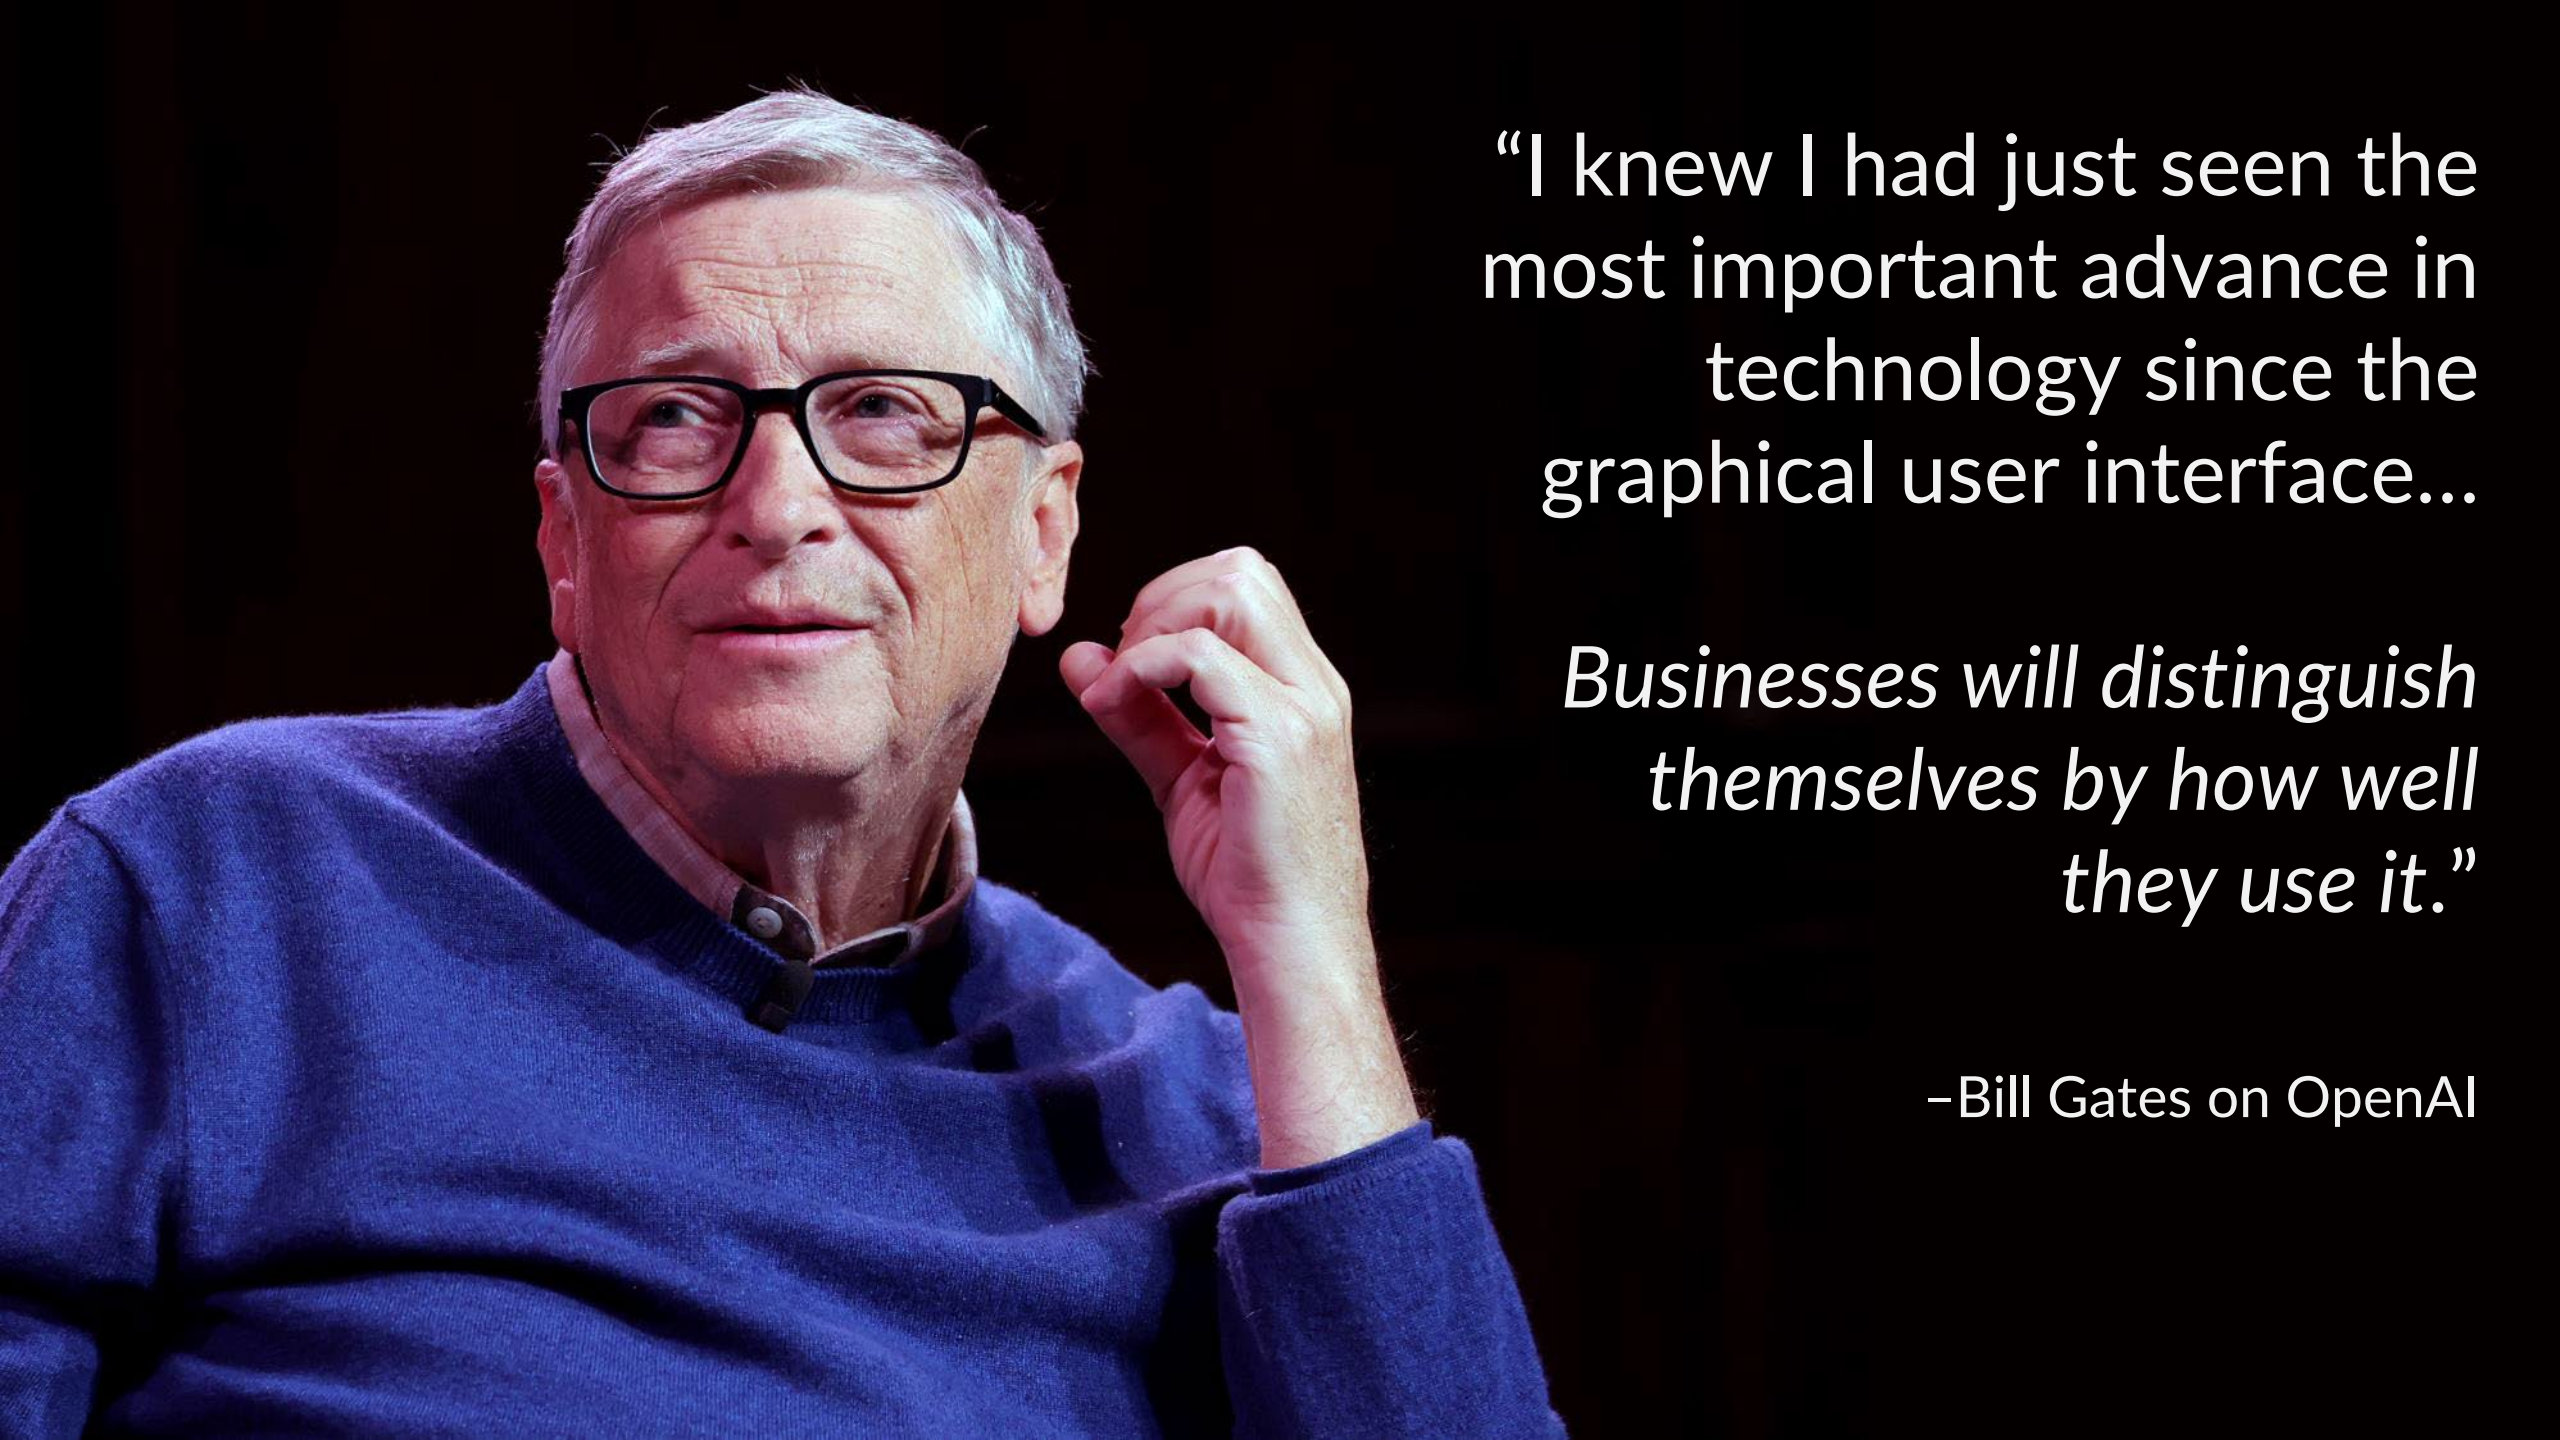
“I knew I had just seen the most important advance in technology since the graphical user interface…
Businesses will distinguish themselves by how well they use it.”
–Bill Gates on OpenAI


WHAT WILL BE THE IMPACT OF CHATGPT ON OUR BUSINESS?
THERE'S A LOT WE DON'T KNOW FOR SURE...
LIKE HOW MUCH OF WHAT IT SAYS IS MADE UP...
OR IF IT WILL TAKE AWAY OUR JOBS...
OR THE SECURITY RISKS...
OR IF IT COULD DAMAGE OUR REPUTATION...
WHAT DO WE KNOW FOR SURE?
ONLY THAT WE WANT TO ADOPT IT EVERYWHERE AS FAST AS WE CAN.
TOM FISH BURNE
©marketoonist.com

Key Definitions - AI

Specific definitions and distinctions do not have academic and scientific consensus. For practical application, the broad term "AI-ML" is used to describe systems that emulate human thought without explicit programming.

Machine Learning ("ML")

The ability of algorithms and statistical models to learn and adapt without being explicitly programmed

Large Language Models ("LLMs")

A type of generative AI trained on vast amounts of data with a large number of parameters that generate novel text-based responses

Artificial Intelligence ("AI")

The ability of a computer system to emulate human-like cognitive abilities, such as learning and problem-solving

Generative AI ("GAI")

A type of AI capable of creating new and novel outputs, such as text, code, images, videos, and audio

ChatGPT

One specific LLM with a conversational interface, built by OpenAI. It is just one of many LLMs/GPT models created by OpenAI

Key Definitions – Data Governance

Data Governance Definition

Data Governance is the planning, oversight, and control over the management of data and the use of data-related resources[1]. Data Governance controls cover the full Data Lifecycle, from the initial Creation of Data (C), to the Reading of Data (R), Updating of Data (U), and eventual Deletion (D) of Data. Every Data Governance program has a uniquely defined scope, but common themes include data protection, data retention, data quality, proper use, and the management of data as a business asset.

Organizations rely on frameworks like the Data Management Association's DMBOK to outline objectives for Data Governance programs, but the actual implementations of Data Governance are very dependent upon the challenges they face. In simpler terms – ***there is no "one size fits all" answer to addressing data governance, though there are core foundational controls that are in most programs.***

Data Availability	Do I have the right data, available in the right format, to the right individuals, at the right time, in order to make required decisions?
Do I know enough about the data (e.g., sources, business definitions, technical constraints, proper usage) to ensure that I am using the 'right', or 'best', data for a given purpose?	**Data Knowledge**
Data Ownership	Are there appropriate controls – policies, procedures, architecture – in place to help manage data assets, and is there overall ownership and accountability to ensure these are operating as designed?
Do controls exist that help manage the day-to-day activities of storing and maintaining data within IT Systems, and are they well defined and well managed?	**Data Management**
Data Quality	Is the overall quality of data controlled, measured, and managed throughout its' lifecycle to ensure it is 'fit for purpose', or appropriate for given use cases? (e.g., can I trust the data?)

1. DAMA-DMBOK Guide

2023: The Year of Gen AI … What's Next?

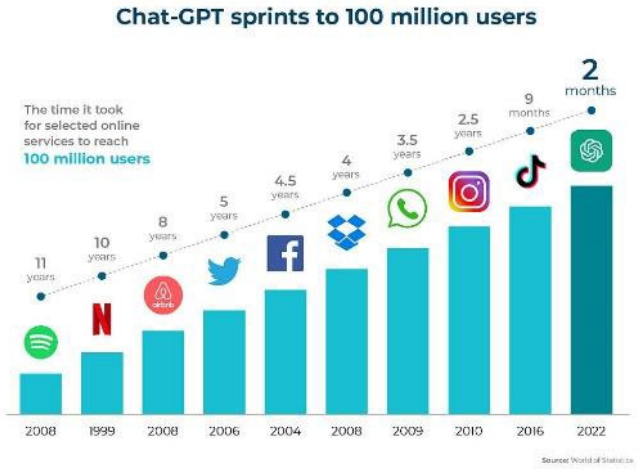

Executives are eager to capitalize on generative AI, with more than 6 in 10 planning to pilot or operate in some way by 2024.

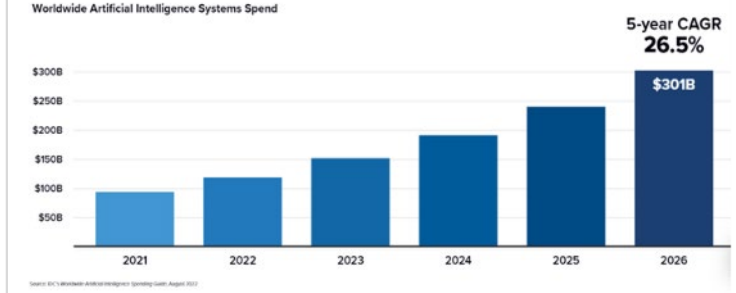

McKinsey & Company

Across the 63 use cases we analyzed, generative AI has the potential to generate $2.6 trillion to $4.4 trillion in value across industries.

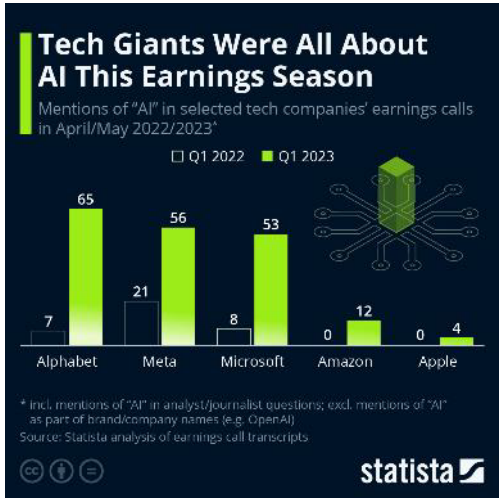

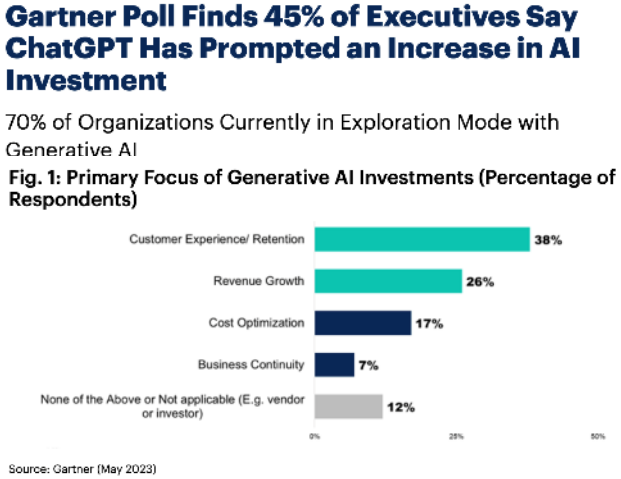

The Next Industrial Revolution?

"It's always easier to imagine the jobs that exist today and might be destroyed than it is to imagine the jobs that don't exist today and might be created."

JED KOLKO

Under Secretary for Economic Affairs at the U.S. Department of Commerce

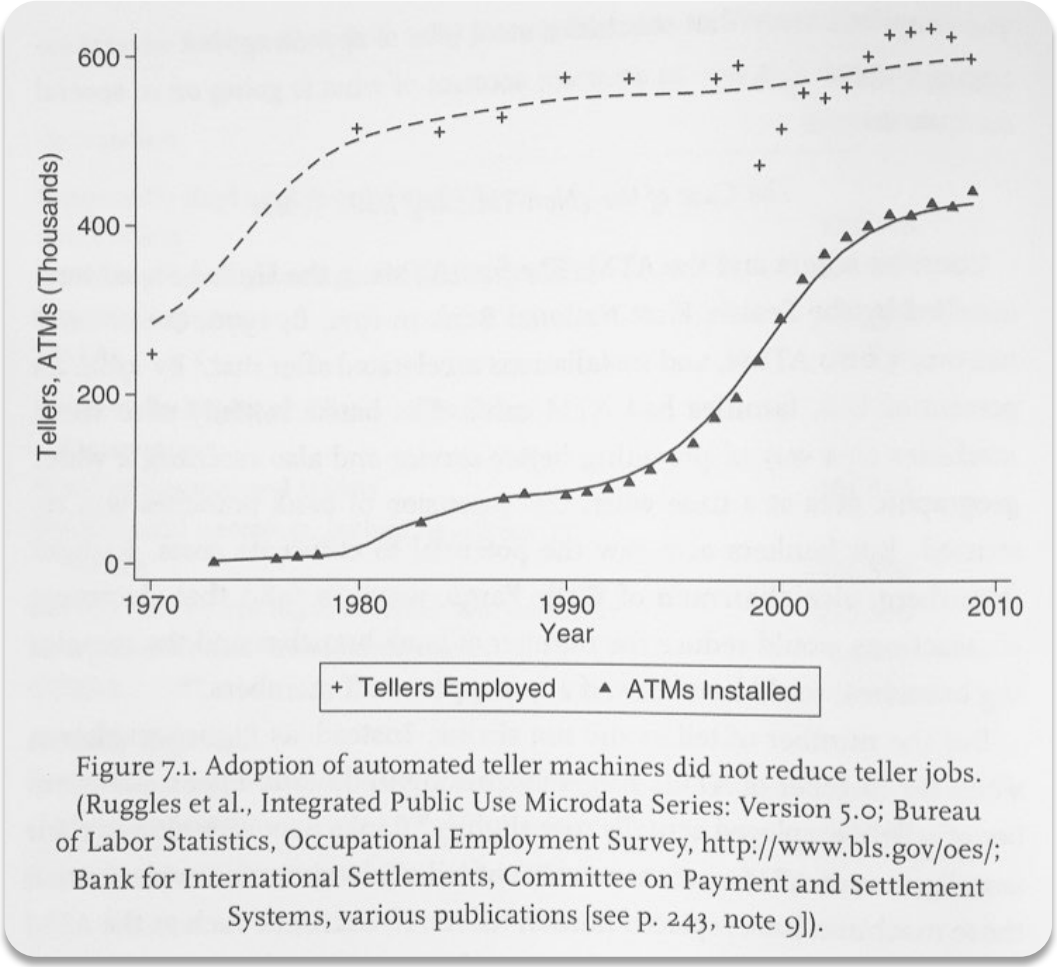

Figure 7.1. Adoption of automated teller machines did not reduce teller jobs. (Ruggles et al., Integrated Public Use Microdata Series: Version 5.0; Bureau of Labor Statistics, Occupational Employment Survey, http://www.bls.gov/oes/; Bank for International Settlements, Committee on Payment and Settlement Systems, various publications [see p. 243, note 9]).

AI Use Cases

Predominant AI Application Opportunities

Decision Support

Customer Experience

Knowledge Management

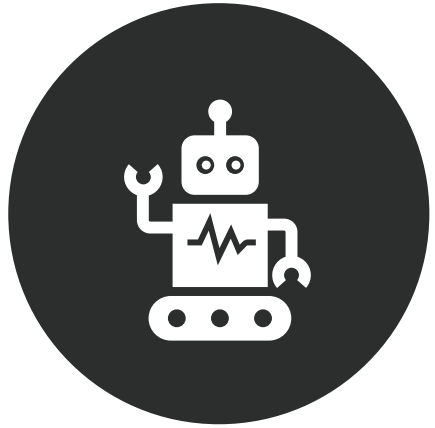

Process Efficiency & Automation

Tired of Creating Audit Notifications?

Open Microsoft Word, Enter Prompt, Generate

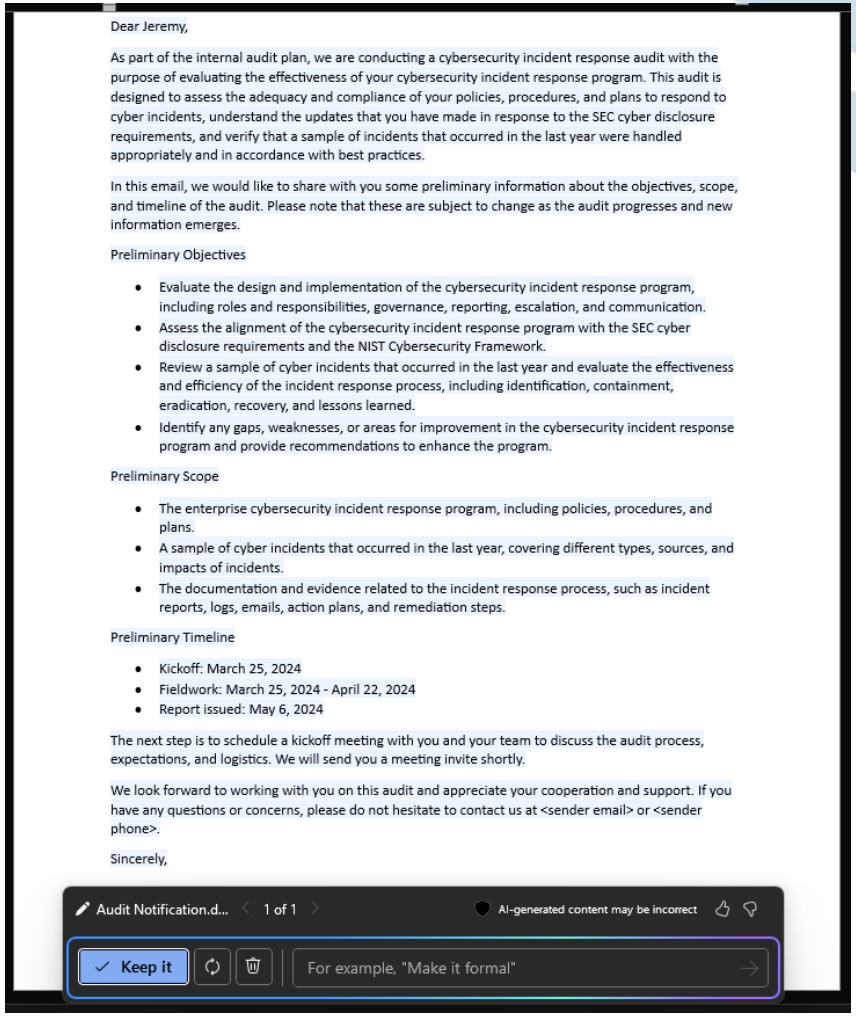

Leverage what's at your fingertips!

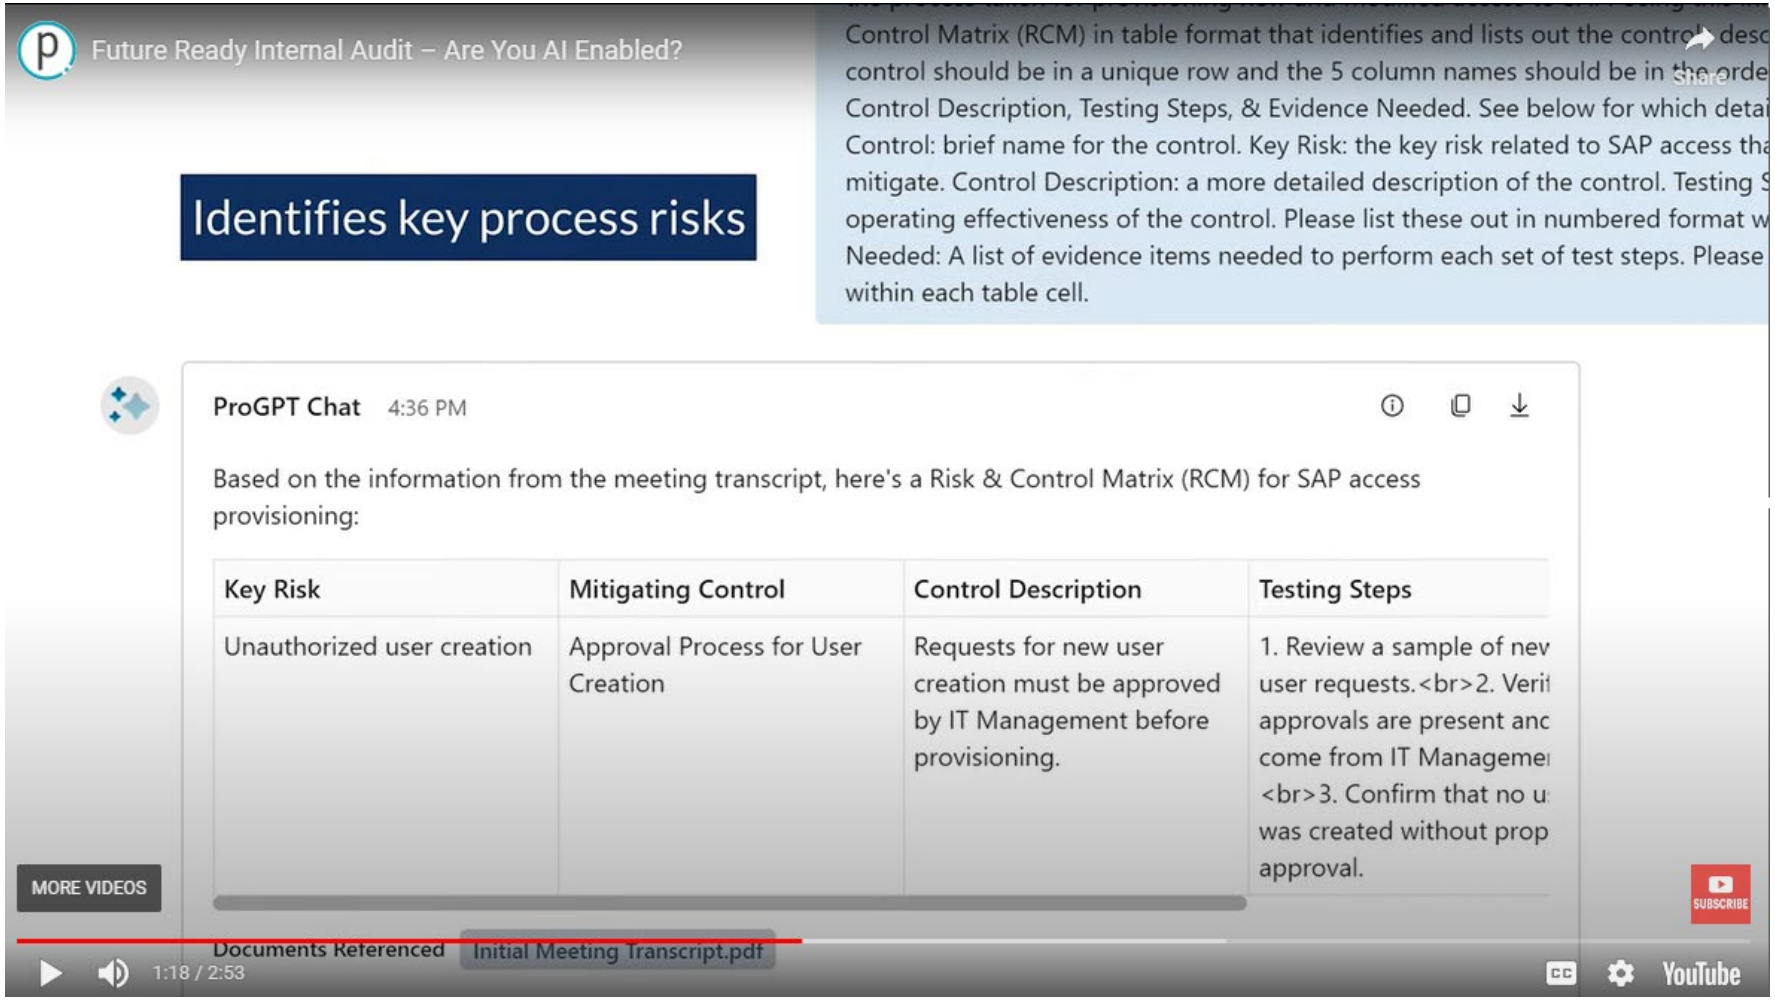

Protiviti Internal Audit AI Use Cases

Protiviti's AI roadmap was established to build the foundational for advancement by digitizing documentation, auto generate deliverables, and AI enabled reporting.

Use Case Name	Brief Description	Loading
AI Enabled Walkthroughs	Standardize delivery of RCM, process narratives, and process flow document from a walkthrough recording.	Completed on 7/10 NEW Add process narrative to process flow functionality
AI Enabled Audit Reports	Dynamically create audit reports with planning documentation and observations logs.	
AI Enabled Issues Management	Digest issue data and use AI to identify redundancies, deep insights, metrics for quarterly reporting.	
AI Enabled SOC Report Analysis	Dynamically identify key elements from SOC reports and automatically include them in an assessment template.	
AI Enabled ESG GHG Calculations	Dynamically digitize energy and utility consumption invoices into a structured format to auto calculate green house gas emissions.	
AI Unstructured Digitization	Submit any document(s) and receive a digitized and structured output in excel. Picture extracting key details from any invoice, contracts, and more.	

AI/ML Use Cases in Internal Audit

AI Risks

Examples of What Can Go Wrong

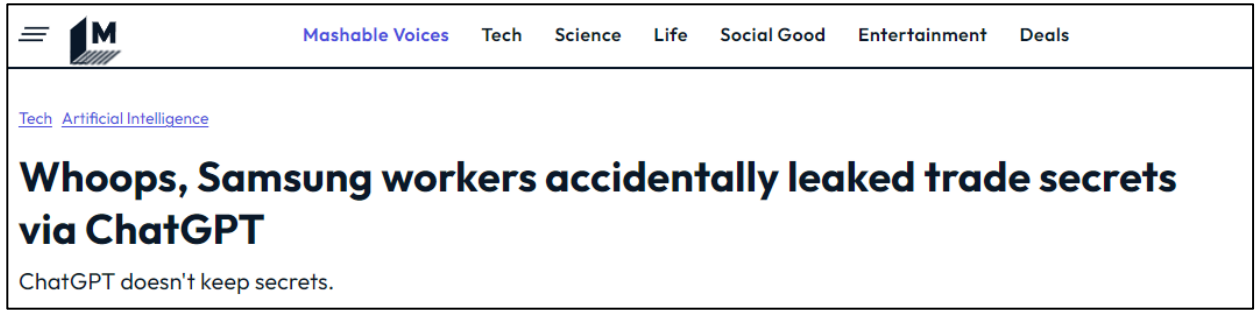

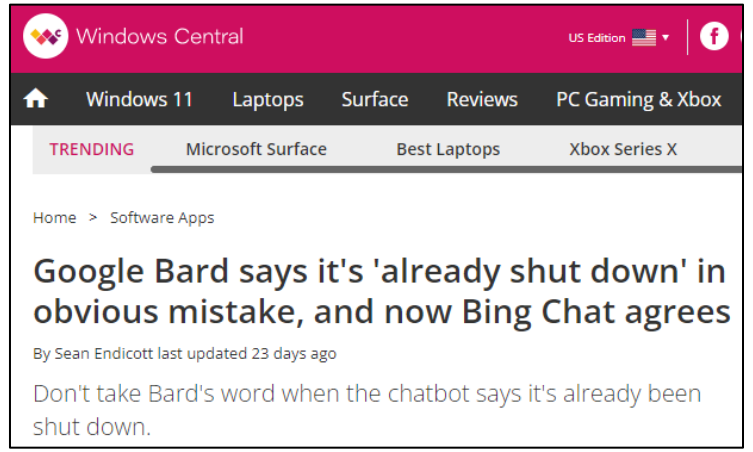

Why is AI Being Regulated?

New risks and threats for individuals' fundamental rights and freedoms

Need to build trust to keep the momentum in use and adoption of AI

Disruptive changes on business, economy, social life, environment, science and research

Ability to enable manipulation and advanced surveillance

Self-learning and self-evolving nature following the launch

AI-specific Frameworks and Regulations

i.e., EU AI Act, NIST Framework

Sectorial Regulations aligned to AI

i.e., UK Consumer Duty

Regulations on Data and Digital Infrastructure

i.e., DORA (Digital Operational Resilience Act), GDPR

Recognized Emerging Standards for AI Governance

NIST AI Risk Management Framework is designed to equip organizations and individuals 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.

ISO/IEC AI Framework provides guidance on managing risk associated with the development and use of AI. The document offers strategic guidance to organizations to assist in integrating risk management into significant activities and functions.

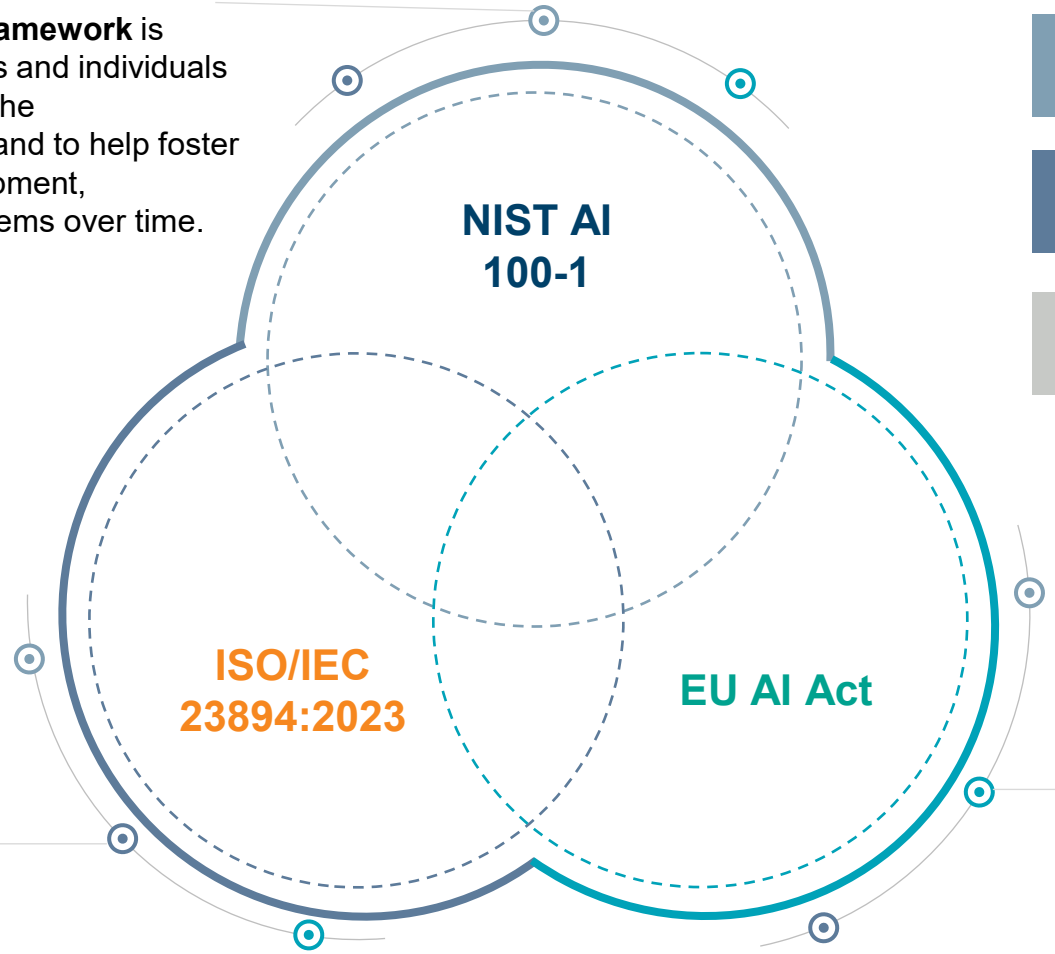

Predominant Themes

- Safety, Security & Resilience
- Transparency & Accountability
- Privacy
- Ethics & Fairness
- Validity / Reliability
- Explainability & Interpretability
- Data Quality & Integrity

European Union AI Act focuses primarily on strengthening rules around data quality, transparency, human oversight and accountability. It also aims to address ethical questions and implementation challenges.

Generative AI Risk Areas

01 **Data Privacy / Confidentiality / Security –** Data entered into an LLM such as sensitive information or trade secrets may be utilized for training and could potentially be inappropriately stored, shared or disclosed.

02 **Reliability and Accuracy –** LLMs may experience “hallucinations” or provide inaccurate or outdated information due to limitations in how it was trained

03 **Lack of Professional Judgement –** Relying on LLMs for audit conclusions may limit the use of professional judgement, which is critical in audit

04 **Adversarial Attacks / Malicious Use –** LLMs may be the target of attacks resulting in inaccurate information, used to generate misinformation, or used in other cybersecurity attacks (e.g., phishing)

05 **Bias –** The LLM may reflect biases in training data, which could lead to biased results when utilizing the LLM

06 **Overreliance / Impact to Skills –** Overreliance on LLMs may lead to a lack of human oversight and introduction of errors as well as degradation of critical thinking and problem-solving skills amongst users

07 **Non-compliance –** LLMs may not be familiar with the latest compliance requirements, or their usage may be restricted, resulting in non-compliance with laws and regulations

AI Observations Across Organizations

Observation	Impact	Recommendations
Lack of standardized AI definitions	• Misalignment of AI projects with business objectives • Over characterization of AI use cases	• Create an AI governance steering committee with cross-functional representation • Conduct workshops to define AI for the organization
Tension between innovation and risk	• Conflicting priorities among departments leading to missed opportunities for competitive advantage	• Develop an AI strategy document that align with business and technology strategies • Develop Key Performance Indicators (KPI) and Key Risk Indicators (KRI) to monitor the design and implementation lifecycle of AI applications
Non-alignment with traditional risk frameworks (i.e., Model Risk Management v NIST AI)	• Ineffective management of risks associated with AI applications • Bifurcation of model & AI policies	• Update risk management policies to include AI-specific risks • Conduct regular model validation and bias assessments • Create unified or mapped MRM & AI Risk Framework
Uncertainty on aligning AI practices globally or locally	• Inconsistent AI practices and difficulty complying across different regions • Unequal treatment of acquisitions	• Develop global AI policy and standards to align with ethical principles and legal requirements • Allow regional teams to customize policy and standards based on local regulations
Shadow AI currently being used	• Deployment of noncompliant AI applications that are biased or unfit for use	• Create intake and approval process for all AI initiatives • Develop training to help employees understand AI technologies, associated implications, and best practices
Lack of centralized tracking for AI use cases	• Challenges governing and managing applications • Unable to report to regulators or third parties	• Deploy a structured tool to inventory AI applications and use cases • Mandate intake and approval process for all AI initiatives (including initiatives with third parties)
Conflating perceived high risk with actual legal risk	• Increased administrative burden and reduced efficiencies	• Establish AI risk tolerance with risk considerations attributed to all phases of the AI lifecycle • Develop and leverage risk control matrix to differentiate between perceived and actual risks
Explainability	• Opaque black box decision making challenge regulators, consumers and companies in meeting explainability and transparency obligations	• Develop Process and outcome-based rationales for input and output variables, to better understand roles of components and impact on decisioning to enable explanations for various parties

Main Categories of Bias

Data Bias

- **Measurement bias** how data is assessed based on preconceived notions
- **Representation bias** when process of data collection does not consider outliers, diversity of the population
- **Sampling bias** happens when a non-random sample of the data is used to train the model
- **Aggregation bias** means assuming that conclusions that are true for the whole group, must be true for a sub-group or individual
- **Omitted features/variable bias** happens when one or more variables or features are missing from dataset

Algorithmic Bias

- **Algorithmic bias** is the systematic and repeatable error during the model development that results in bias towards one group versus another group
- **Bias- Variance tradeoff**
  - Models that are mathematically complex, are highly sensitive to “noise”, overfitted leading to high variance
  - Models that are mathematically less complex, lead to underfitting and hence high risk of bias
- **Statistical unfairness** is when the model uses the average statistics in a dataset to represent and judge individuals in that group

User Interaction Bias

- **User interaction** bias is a type of bias, which is mostly observant on the Web and can get triggered by user interface and self-selection bias, behavior and interaction (clicks) from user
- **Ranking bias** is the idea that high-ranked results will be more relevant and therefore will get more clicks in online search results
- **Popularity bias** is where items that are more popular tend to gets more clicks
- **Emergent bias** arises from change in population and cultural values after completion of model design

Hallucinations in Generative AI

Hallucination

Instances where Generative AI systems or models generate inaccurate or nonsensical outputs.

Causes

- **Training Data Sources:** Generative AI models are trained on vast amounts of data containing both accurate and inaccurate content, as well as societal and cultural biases. Since AI models mimic patterns in their training data without discerning truth, they can reproduce any falsehoods or biases present in the data set.
- **Limitations of Generative Models**: Generative AI models are designed to predict the next word or sequence based on observed patterns to generate plausible content.
- **Inherent Challenges in AI Design:** The technology behind generative AI tools isn't designed to differentiate between what's true and what's not true. Even if generative AI models were trained solely on accurate data, their generative nature would mean they could still produce new, potentially inaccurate content by combining patterns in unexpected ways.

Mitigation

- **Use High-Quality Training Data:** Train on diverse, balanced and well-structured data
- **Define Purpose:** Spell out how you will use the AI model, as well as any limitations on the use of the model
- **Use Data Templates: T**emplates provide teams a predefined format, increasing the likelihood that an AI model will generate outputs that align with prescribed guideline
- **Limit Responses:** Define boundaries that limit possible outcomes
- **Test and Refine:** Evaluate on an ongoing basis to adjust or retrain the model
- **Human Oversight:** Make sure a human being is validating and reviewing AI outputs is a final backstop measure

Key Considerations for Internal Auditors

Key Distinctions that Matter!

How it's Trained

- Supervised
- Unsupervised
- Reinforcement

What Data it's Trained On

- Public Domain
- Subscription Data Feeds
- Enterprise-Owned
- Industry-Specific

How it's Architected

- Public Cloud
- Private Cloud
- On Premises

Identifying Value & Evaluating Readiness of AI

Value Identification	Data	Skills	Ecosystem	Experimentation	Change Management
Where should we apply AI? **Key Focus:** • Do we have a strong understanding of the strengths of AI? • What areas of the business offer the greatest potential? • What job roles & How many? Cost Basis? • What levers of value will we pull? (productivity improvement, etc)? How will we measure and how fast could we deliver results?	**Do we have the data and is it ready?** **Key Focus:** • What data will be needed? • Where is it located? • Is it organized for usage and scalable? • What governance/security procedures do we have in place?	**Do we have the talent we need to unlock the value?** **Key Focus:** • What AI skills do we have in our organization? • Where are they and how many? • Is there alignment between our skills and our identified value hypothesis? • What is our training/dev plan?	**Where to build/buy/partner?** **Key Focus:** • How closely does our value hypothesis align to existing or emerging solutions? • Are we actively engaged with the AI ecosystem? Big & small? • Do we have a framework on how to evaluate build, buy or partner decisions?	**Do we have a capacity to test & learn quickly?** **Key Focus:** • Do we have a technology environment available? • Do we have a Rapid Prototype team/approach? • Are we leveraging best practices in Design Thinking / Agile / Lean / Innovation?	**Do we have an organized plan to execute?** **Key Focus:** • What is the Change Readiness of our Org? • What is most important to achieve buy in? • How do we communicate our plan? • How will we collect, evaluate, and act on feedback?

Quick Steps: AI Governance & IA Considerations

1

Governance

Promotes cross-functional representation and leadership buy-in, fostering cohesive decision-making and accountability in AI governance initiatives

2

Policy, Standards, and Controls

Provide guidelines and standards to employees for AI development and deployment, ensuring alignment with ethical principles, legal requirements, and organizational values

3

Training

Helps employees understand AI technologies, their implications, and best practices for effective governance, enhancing the organization's capacity to responsibly manage AI initiatives

4

AI Inventory & Intake Process

Enables organizations to gain awareness of different AI initiatives in flight, AI platforms and capabilities (including third parties), facilitating risk assessment, compliance monitoring, and resource allocation for governance efforts

5

Risk-Prioritized Approach

Allows organizations to allocate resources effectively, focusing on addressing the most critical concerns first and mitigating potential harms associated with AI deployments

Internal Audit Considerations

- ✓ Explore participation in Governance steering committees or working groups to drive awareness for audit purposes, provide risk and control guidance, and assess the impact of AI implementation on audit plans and the broader audit universe
- ✓ Engage in advisory and consultancy roles early and often to evaluate the design of AI program implementation, helping to avoid gaps and limit risk exposure
- ✓ Understand the risk assessment process and consider the business's use of AI when planning future audit assessments

Closing Considerations

As LLMs continue to gain prevalence and demonstrate utility across various industries and business functions, it's crucial for Internal Audit to leverage these advanced technologies benefits and stay progressive.

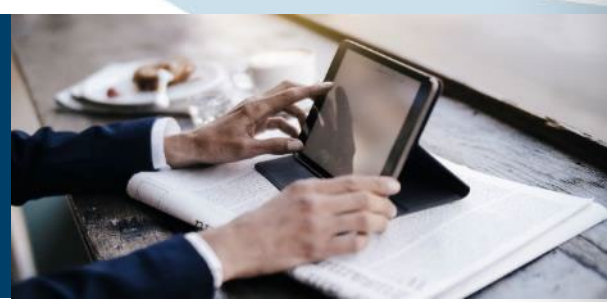

Familiarize Yourself with LLMs

Invest time in learning the capabilities, risks, and best practices associated with LLMs to understand their impact on both the Internal Audit function as well as the business.

Identify Internal Audit LLM Opportunities

Explore opportunities to safely and effectively incorporate LLMs into the audit process, such as enhancing risk assessments, supporting audit execution, and streamlining reporting.

Partner with Stakeholders on LLM Risks and Controls

Work closely with stakeholders to help them understand the risks associated with LLM usage and provide guidance on implementing appropriate controls to mitigate those risks and drive compliance with relevant regulations and organizational standards.


Gregg Wishna
Director at Protiviti | Expert in IT Internal Audit, Analytics, and Data Governance | Published Thoug...
Will Shuman
Director at Protiviti
THANK YOU!