FIELDGUIDE

AI and Auditing: A Practical Guide for Firms

Introduction

Artificial Intelligence (AI) is rapidly becoming a cornerstone of modern business operations, promising unparalleled efficiency and innovation. However, alongside these benefits come a host of challenges, especially for firms who need to audit technologies and processes at those businesses. The evolving landscape of AI makes it particularly challenging to audit AI systems to ensure compliance, mitigate risks, and uphold ethical standards. Here are key best practices and findings in Fieldguide's conversations with Top 500 advisory and audit firms.

Growing Set of Standards & Guidelines

Policymakers at the state, national, and global levels are recognizing the imperative need for regulating AI technologies. With these regulations come the necessity for AI auditing to ensure compliance with specific legal standards and ethical guidelines. Some notable standards include the following:

- New York City's Local Law 144, commonly referred to as the "AI Audit Law," mandates independent audits of companies utilizing AI systems for employment hiring decisions. This law underscores the importance of transparency and accountability in AI-driven processes, particularly in sensitive areas such as hiring practices.
- In January 2023, the National Institute of Standards and Technology (NIST) introduced the AI Risk Management Framework (AI RMF 1.0), providing comprehensive guidelines for assessing risks associated with AI systems.
- In October 2023, the White House issued an Executive Order on AI, serving as a guideline for U.S. agencies and companies operating in the country. Emphasizing the importance of protecting both national interests and individual rights, this directive underscores the need for responsible AI development and deployment.
- In December 2023, the introduction of ISO/IEC 42001 marked a significant step towards global standardization in AI governance. These standards outline roles, policies, and processes that organizations should follow regarding AI, providing a common framework for AI management systems worldwide.
- Furthermore, the pending European Union AI Act, which started reviews in June 2023, focuses on protections within the EU while balancing innovation with citizen protection.

Key Elements of an ISO 42001 Audit

ISO 42001 provides a structured framework for organizations to manage AI systems, as well as for third parties to audit AI systems. Understanding the key components of an ISO 42001 audit is essential for ensuring compliance and promoting responsible AI implementation.

Performing an ISO 42001 audit shares conceptual similarities with an ISO 27001 audit, emphasizing the need for a systematic approach to managing AI systems. Key components of an ISO 42001 audit include:

- Scope and Statement of Applicability: Similar to ISO 27001, an ISO 42001 audit begins with defining the scope of the audit and preparing a statement of applicability. This helps delineate the boundaries of the AI management system and identify applicable requirements and controls.
- Documented Policies and Procedures: Organizations must develop and maintain documented policies and procedures governing the design and development of AI systems. These documents outline the principles, guidelines, and processes for ensuring the responsible and ethical use of AI technologies.
- Roles and Responsibilities: Clear delineation of roles and responsibilities is crucial for effective AI governance. Organizations must document roles related to AI design, development, deployment, and maintenance, ensuring accountability and oversight throughout the AI lifecycle.
- Risk Assessment and Risk Treatment: A fundamental aspect of ISO 42001 is the organization's risk assessment and risk treatment process. Assessing risks associated with AI technologies, including AI hallucinations or biases, requires a nuanced understanding of AI-specific risks and their potential impact on stakeholders. Organizations must develop strategies for mitigating identified risks to ensure the safe and ethical deployment of AI systems.
- Impact Assessment: One of the new control objectives within Annex A of the ISO 42001 standard is to assess AI system impacts on individuals, groups, and societies throughout its lifecycle. This entails evaluating the broader societal implications of AI technologies, considering ethical, social, and cultural factors.

- Resource and Document Management: Managing resources and documents related to AI systems can pose challenges due to the novelty and evolving nature of AI technologies. Organizations must ensure that resources, including data, algorithms, and models, are adequately managed, documented, and protected throughout their lifecycle. Additionally, organizations must stay informed about emerging standards and guidelines and adapt their documentation processes accordingly.
- Performance Evaluation and Continuous Improvement: Continuous evaluation and improvement are integral to the effective management of AI systems. Organizations should establish mechanisms for monitoring AI system performance, collecting feedback, and implementing corrective actions to address identified deficiencies.

ISO 42001 provides a structured framework for organizations to manage AI systems, as well as for third parties to audit AI systems. Understanding the key components of an ISO 42001 audit is essential for ensuring compliance and promoting responsible AI implementation.

Common Risks for IT and AI Audits

While the dynamic nature of AI presents its own challenges to even the most experienced auditors, there are several IT and security audit procedures that can also be applied to AI systems:

- Data Privacy and Confidentiality: In the context of AI, data privacy and confidentiality are paramount due to the processing of vast amounts of personal data. Organizations must prioritize a privacy-first approach when implementing AI systems, ensuring that adequate controls are in place to protect sensitive information. Auditors play a crucial role in assessing whether data collection, storage, and processing adhere to privacy laws and ethical guidelines. This includes evaluating mechanisms for obtaining consent, anonymizing data, and preventing unauthorized access. Additionally, audits should verify the organization's protocols for responding to data breaches and safeguarding data subjects' rights according to regulations such as GDPR or CCPA.
- Data Integrity: Data integrity, referring to the accuracy and consistency of data throughout its lifecycle, is essential for maintaining the trustworthiness of AI systems. Auditors must assess the quality of input data, validation methods, error detection and correction procedures, and overall management of AI models to prevent data alteration. Ensuring data integrity is particularly critical in sensitive applications such as healthcare, finance, and legal organizations.
- Security and IT Controls: AI systems introduce complex security challenges that extend beyond traditional IT controls. Auditors must evaluate security measures to protect against unauthorized access to AI systems and models. This includes examining encryption practices, authentication mechanisms, and access controls.
- Supply Chain Vulnerabilities: AI applications often rely on components, data, and services from third-party suppliers or vendors, making supply chain vulnerabilities a significant concern. Audits should assess the organization's practices for evaluating and managing the security and integrity of third-party components integrated into AI systems. This includes due diligence processes for vendor selection, ongoing monitoring of vendor compliance with security standards, and contingency plans for addressing supply chain disruptions.

What Makes an AI Audit Unique?

Auditing AI systems presents unique challenges and considerations compared to traditional IT audits. Understanding these differences is essential for auditors to effectively assess AI systems and mitigate associated risks.

- Distinction Between Compliance and Risk Audits: It is crucial to distinguish between compliance audits and risk audits. Compliance audits compare an entity's actions or properties to predefined standards or regulations, while risk audits focus on identifying and controlling risks through open-ended inquiries about system functionality and operation.
- Evolution of Standards: As AI technologies continue to advance, standards and regulations are likely to evolve, creating a challenge for auditors to keep pace with the latest developments. Additionally, clients may not be well-versed in the latest AI standards and guidelines, requiring auditors to provide education and guidance throughout the audit process.
- Fluidity of AI Technologies: Unlike traditional IT systems, AI technologies are highly fluid and dynamic. Traditional IT audits rely on pre-implementation and post-implementation testing to identify and mitigate risks. However, pre-implementation audits of AI systems may not fully capture all risks, as AI systems continue to "learn" and evolve over time. Similarly, post-implementation audits face challenges as AI systems lack standard change management processes typically found in traditional IT systems.
- Continuous Auditing: Effective AI auditing procedures must incorporate elements of continuous auditing to address the fluid nature of AI systems. Continuous auditing allows auditors to monitor and evaluate AI systems continuously, identifying and addressing risks in real-time. Techniques such as model fooling, functional testing, and template-based stress testing are commonly employed to ensure the integrity and reliability of AI systems.
- Scope and Pricing Challenges: AI audit engagement scope and pricing are not as established in the market compared to traditional IT audits. This lack of standardization makes it challenging for organizations to compare audit service offerings from different firms and for audit firms to determine their go-to-market approach. As a result, auditors must carefully define the scope of AI audits and establish transparent pricing structures to meet the needs of their clients effectively.

1. **Distinction between compliance and risk audits**
2. **Evolution of standards**
3. **Fluid AI technologies**
4. **Continuous auditing**
5. **Scope and pricing challenges**

AI Model Auditing

AI models are the foundation for AI technologies, so auditing AI models, particularly Large Language Models (LLMs), is critical. Evaluating their performance, robustness, security, and truthfulness is crucial to ensure their effectiveness and reliability in real-world applications.

- Model Performance: The efficiency of an AI model extends beyond its isolated performance and encompasses its functionality across a broad spectrum of tasks. Benchmarking AI models against human baselines provides a tangible measure of their capabilities. To accurately gauge AI performance, a comprehensive range of tasks and benchmarks must be employed.
- Robustness: Robustness refers to an AI model's ability to handle unexpected prompts or data effectively. Stress testing AI models under various conditions helps assess their robustness. For example, in Fieldguide's AI model, rigorous stress testing ensures its ability to execute testing procedures and document analysis reliably. The model's response in different scenarios, such as declining to provide irrelevant information or excluding sensitive data from responses, demonstrates its ability to maintain functionality and security under unusual conditions.
- Information Security: Information security is critical in auditing AI models, particularly concerning the protection of training data from extraction attacks. Hackers may attempt training data extraction attacks to recover personal information stored in training data. Defensive measures against these attacks are crucial to understanding and auditing an AI model's security effectively.
- Truthfulness: Truthfulness is a measure of an AI model's ability to differentiate between factual and false information accurately. An AI's tendency to generate false or misleading content can undermine trust in its applications significantly. Evaluating the truthfulness of AI models is a cornerstone of AI auditing, ensuring that they provide accurate and reliable information.

AI Audit Example

Deep Patient, an AI tool utilized at Mount Sinai Hospital for reviewing patient medical records, boasts the capability to predict a wide array of diseases in patients without any explicit instructions from experts. Its predictive abilities extend to anticipating the onset of psychiatric ailments. However, one notable risk associated is its inability to provide reasoning behind patient identifications.

Considerations for Audit Planning

Developing an audit plan for Deep Patient involves addressing several critical considerations. Firstly, data quality and bias must undergo rigorous scrutiny. Unlike traditional IT controls governed by established processes and logic, AI systems like Deep Patient may necessitate additional controls for monitoring and assessing both input data and output results. The integrity of the dataset is paramount, as biases within it can significantly impact AI predictions, potentially leading to discriminatory outcomes or misdiagnoses.

Secondly, validation of predictions holds immense importance. Implementing controls to validate predictions and incorporating expert review of AI-generated diagnoses are essential steps in ensuring the accuracy and reliability of Deep Patient's outputs. Transparency challenges cannot be overlooked in auditing Deep Patient. Management's efforts to enhance the model's transparency are crucial, as transparency directly influences the trust and confidence of medical professionals and patients alike.

Lastly, ethical and privacy considerations loom large. Assessing how Deep Patient handles patient data privacy and confidentiality, and ensuring compliance with healthcare regulations such as HIPAA, are critical aspects of the audit process.

Implications and Lessons Learned

The Deep Patient case study underscores the transformative potential of AI in revolutionizing healthcare. However, it also underscores the indispensable need for AI audits to address data quality, validation, transparency, and ethical considerations. Management's ability to understand, explain, and justify Deep Patient's results is pivotal. Inability to do so may raise doubts about the completeness and accuracy of the subject matter, as well as the effectiveness of internal controls in mitigating risks. Similarly, auditors must be able to evaluate the results obtained from AI audit tools to form informed opinions. Failure to comprehend or evaluate these results compromises the sufficiency and appropriateness of audit evidence, potentially undermining the integrity of the audit process.

Next Steps

Top 500 firms have indicated to Fieldguide that there are three essential steps for firms embarking on AI audits. These steps are crucial for developing effective audit strategies and ensuring successful engagements.

- Developing In-House AI Expertise: Investing in in-house AI expertise is paramount for conducting thorough and insightful audits. Training your team on new terminologies and risks associated with AI systems is essential to ensure everyone is equipped with the knowledge needed to navigate complex audit engagements. Furthermore, embracing AI as a tool internally can help familiarize your team with its functionalities and applications, enhancing their effectiveness in conducting audits.
- Establishing Thought Leadership for AI: Staying abreast of the latest standards and regulations is crucial for maintaining thought leadership in the field of AI auditing. Actively monitoring developments in AI governance and participating in comment periods for new standards allows organizations to provide input and shape the regulatory landscape. Identifying and prioritizing clients that are best suited for AI engagements enables organizations to focus their resources effectively and maximize their impact. Creating a standard Statement of Work (SOW) and service offering streamlines the audit process, facilitating smoother engagements and ensuring consistency in deliverables.
- Navigating Uncertainties and Building for the Future: It is important to acknowledge that AI audits may involve uncertainties, especially in scoping and pricing engagements. Organizations should expect the unknown and approach audits with a mindset of learning and adaptation. By embracing uncertainties and leveraging them as opportunities for growth, organizations can build stronger teams and service offerings for the future. Investing in initial engagements to establish business and refine audit methodologies sets the foundation for long-term success in the evolving landscape of AI auditing.

Additional Resources

- The Data-Driven Audit: How Automation and AI are Changing the Audit and the Role of the Auditor, AICPA and Chartered Professional Accountants Canada
- AI Adoption Roadmap: Top Considerations, Hannes Hapke, Avani Desai, and Kevin Au
- Guidance on the AI Auditing Framework, Information Commissioner's Office
- Auditing Artificial Intelligence, ISACA
- LLM AI Security & Governance Checklist, OWASP
- Auditing Large Language Models: A Three-Layered Approach, Jakob Mökander, Jonas Schuett, Hannah Rose Kirk, and Luciano Floridi

While the dynamic nature of AI presents its own challenges to even the most experienced auditors, there are several IT and security audit procedures that can also be applied to AI systems, such as data privacy, data integrity, and IT controls.

FIELDGUIDE

About us

Fieldguide Powers the Future of Trust

Built by and for practitioners, Fieldguide is the category-defining and comprehensive AI Platform for Advisory and Audit Services that makes it amazingly easy for firms to automate the entire lifecycle of any type of engagement. Our cloud-based software helps firms take advantage of the increasing demand for advisory and audit services by addressing the challenges of scarce talent, remote collaboration, and modern client expectations.

Fieldguide provides end-to-end visibility across all engagements, streamlined collaboration with distributed clients and teams, and unparalleled staff productivity for any engagement, including SOC 1 & SOC 2, HITRUST, PCI DSS, HIPAA, and many more. Instead of juggling siloed legacy tools, practitioners can rely on Fieldguide as a single source of truth, with all the capabilities that they need to run engagements from kickoff to completion. Streamlined requests, collaborative document management, intelligent framework mapping, and one-click reports are just some of the unique Fieldguide features that help firms drive greater visibility across all engagements, while delivering a first-class client experience. Leading firms like Wipfli, Mazars, Aprio, and BerryDunn trust Fieldguide to increase revenue, boost client satisfaction, and improve profits.

Fieldguide has received multiple awards from industry associations, and the company is backed by top venture capital firms, including 8VC, Y Combinator, and Floodgate.

For more information or to schedule a free product demo, visit
fieldguide.io