ORLANDO WILLIG
MARCO SCHREYER

BUILDING A FUTURE DIGITAL AUDIT WORKFORCE

Equipping artificial intelligence (AI) agents with computer-aided audit tools (CAATs)

Embedding autonomous artificial intelligence systems is no longer a distant prospect for auditing. Frameworks like the Model Context Protocol (MCP) and autonomous AI agent applications like OpenClaw provide the foundation for building an agentic audit workforce, which can extend the capabilities of human audit teams while preserving the independence, transparency, and scepticism that define oversight.*

1. INTRODUCTION

The auditing profession stands at a pivotal moment in its technological evolution. For decades, auditors have progressively integrated digital and software tools to gather evidence and perform analytical procedures. Today's advances in artificial intelligence mark the emergence of agentic AI, which refers to systems capable of reasoning, using tools and executing multi-step tasks autonomously. Recently, OpenClaw [1], an open-source agentic AI system, demonstrated how such technology can act on defined goals rather than merely respond to queries, thereby illustrating the shift from passive assistants to proactive digital collaborators. OpenClaw connects to everyday applications and executes multi-step tasks by dynamically selecting and combining available tools. While not designed for auditing, the underlying architecture closely mirrors the capabilities that an agentic audit system would require.

Against this backdrop, the developments in AI augmented auditing have increasingly shifted attention towards so-called agentic systems. Unlike earlier AI audit assistants, which typically support isolated tasks such as text analysis or data extraction, AI agents exhibit the capability to execute complex, multistep audit procedures [2]. In the context of financial audits, an AI agent could systematically review an entity's policies and procedures to identify key control objectives, query the IT-system for journal entries that violate controls, and prepare a draft audit report that summarises the exceptions noted and their impact on the control environment [3]. This development represents a significant leap forward from contemporary AI-enabled audit assistants often established by tailoring or adapting large language models [4] like GPT-4 [5], Llama [6] or Mistral [7].

In practical terms, agentic AI systems exhibit "agency", i. e. the capacity of operating autonomously, making independent decisions to achieve predefined objectives without constant human intervention [8]. As a result, AI agents can autonomously achieve audit-related goals over an extended period through planning, memory and tools [9] as illustrated in *Figure 1*:

→ *Planning:* AI agents break down large tasks into smaller, more manageable subgoals. They engage in self-criticism and reflect on past actions, learn from mistakes and refine future steps, thereby improving the quality of outputs [10].
→ *Memory:* AI agents use short-term memory techniques to adapt to new information during interactions [11]. They also have long-term memory capabilities that enable them to retain and recall extensive information over time [12].
→ *Tools:* AI agents learn to call external applications for real-time data, code execution and access to proprietary information sources. This tool use enables agents to extend beyond text generation and chatbot capabilities, accessing real-world data and systems [13].

Given these characteristics, AI agents excel in handling challenging, realistic tasks that often lack a single correct solution, offering real-world audit utility [14].

ORLANDO WILLIG,
MASTER OF ARTS,
JUNIOR AUDIT SPECIALIST,
SWISS FEDERAL AUDIT
OFFICE (SFAO)

MARCO SCHREYER,
PHD, AUDIT SPECIALIST,
SWISS FEDERAL AUDIT
OFFICE (SFAO),
POSTDOCTORAL DAAD IFI,
RESEARCH FELLOW,
INTERNATIONAL COMPUTER
SCIENCE, INSTITUTE (ICSI),
BERKELEY

Figure 1: **OVERVIEW OF AN AI AUDIT AGENT INTEGRATING TOOLS, MEMORY AND REASONING CAPABILITIES**

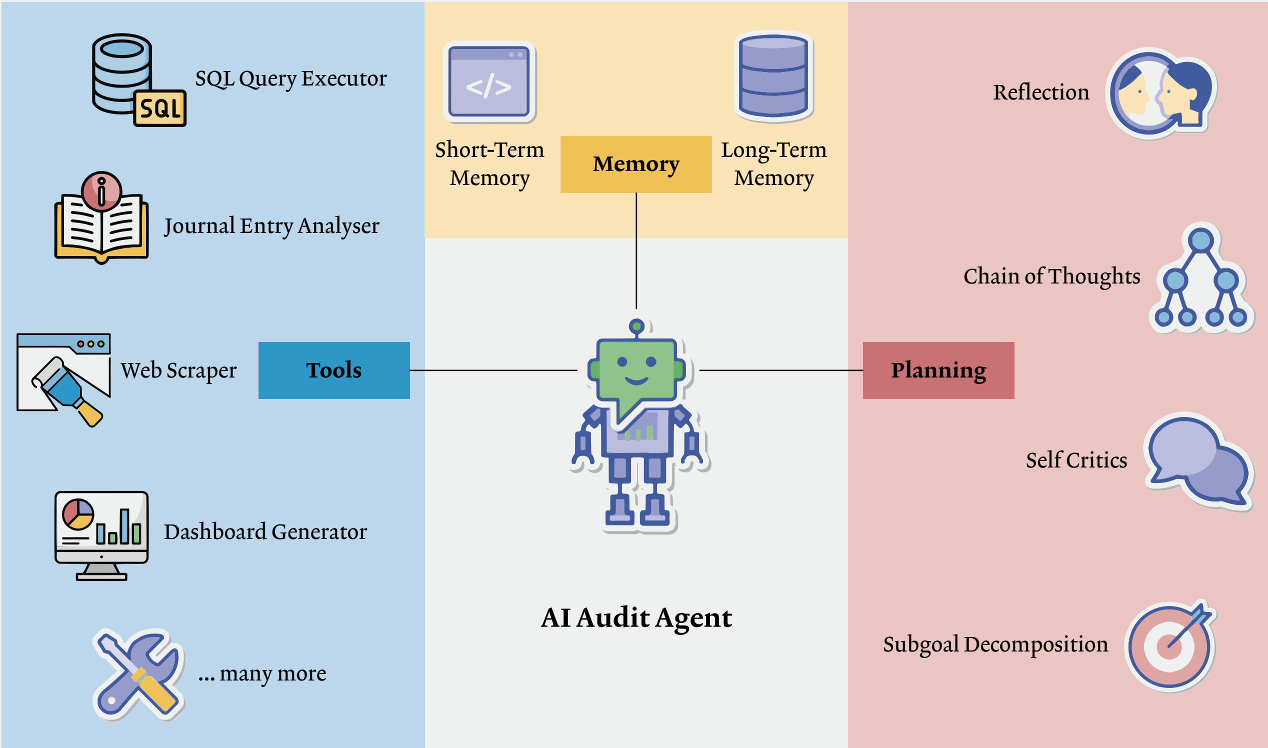

Over the years, auditors have developed a broad range of technology-enabled tools, commonly used in the context of computer-aided audit tools (CAATs) [15]. While such specialised CAATs have significantly improved audit coverage and testing efficiency, they are typically executed manually by auditors and remain tightly coupled to specific tasks.

In this article, we explore how CAATs can be designed, documented and exposed in a way that makes them accessible to AI agents. This approach aligns with broader developments in the field, where automation is increasingly used to support more streamlined analytical workflows [16]. Building on the Model Context Protocol (MCP) [17], we outline how traditional CAATs may be adapted into reusable, agent-ready audit tools and discuss how such approaches could contribute to the gradual development of a digital audit workforce. The article illustrates these considerations through selected business cases that indicate where such transformations may already be feasible in practice.

Consider an auditor instructing an AI agent to "identify journal entries that violate key controls." The agent would activate predefined audit tools – analogous to traditional CAATs – such as routines that translate policies into control rules, modules that extract journal-entry populations and scripts that test entries against approval, threshold or segregation-of-duties requirements. By orchestrating external audit tools, the agent performs the procedures end to end (see *Figure 2*). The auditor then evaluates these results, while much of the technical and analytical work is executed through the AI agent-operated audit tools.

2. FROM CAATS TO AI AGENTS

Computer-aided audit tools and techniques have significantly broadened the digital capabilities available to auditors. In general, the term CAAT denotes the practice of using computers to facilitate data analysis-related audit processes. Developed to address the limitations of traditional audit methods, CAATs allow auditors to build conclusions based on the examination of all available (or a large sample of) data, rather than upon a limited sample of a population.

Established in practice since around the 1960s, CAATs encompass a wide range of techniques, including rule-based queries for identifying control deviations, statistical sampling and estimation procedures, anomaly-detection analytics, process-mining methods, robotic process automation and visual analytics [18]. Commercial platforms such as ACL and IDEA have facilitated the adoption of these techniques and enabled full-population testing in many audit areas.

Despite their effectiveness, CAATs continue to rely on auditors to select, parameterise and execute each routine and to integrate the resulting outputs into their broader assessment. As the volume and heterogeneity of audit-relevant data continues to increase, the reliance on a manual application constrains the extent to which CAATs can support comprehensive analytical work.

Emerging advances in agentic AI suggest a further development in this trajectory through the introduction of agentic tools: audit procedures described in a structured format that enables intelligent systems to apply, sequence and document them under auditor supervision. Agentic tools build

Figure 2: **SCHEMATIC OVERVIEW OF AN AGENTIC AUDIT ARCHITECTURE DEMONSTRATING HOW AN AI AGENT ACCESSES DIVERSE AUDIT RESOURCES THROUGH A UNIFIED PROTOCOL**

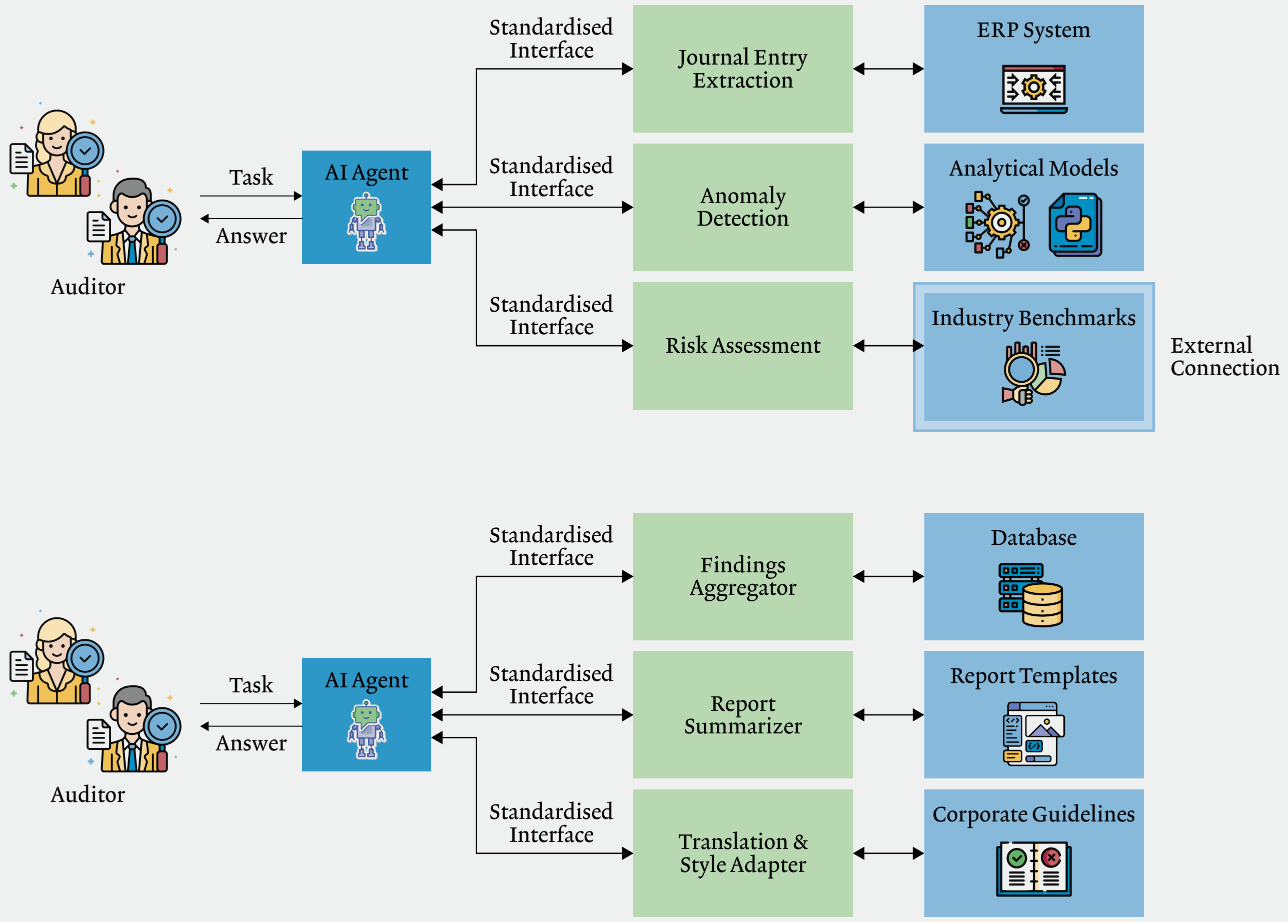

on the methodological foundations of CAATs, but express their functionality in a machine-interpretable form. This allows an AI agent to draw on an approved set of routines, such as control-rule mappings, journal-entry extraction modules or segregation-of-duties checks, and combine them as required to address the auditor's stated objective. The auditor retains responsibility for defining the task, setting the scope and evaluating the results, while the execution of procedural steps is carried out by the agent through the available tools.

The progression from CAATs to agent-operable tools establishes a conceptual basis for developing an agentic AI-enabled audit workforce. It extends established audit techniques in a controlled and transparent manner and provides a structured pathway for integrating intelligent systems into audit workflows.

3. BUILDING THE AI WORKFORCE

The progression from CAATs to agent-operable audit tools outlines what intelligent systems could do in an audit context. The remaining question is how these capabilities can be organised into a digital workforce that is reliable, controllable and reusable across engagements. This requires more than individual prototypes or isolated integrations. Audit teams need a structured environment in which AI systems can discover, invoke and combine approved tools under auditor supervision. The in 2025 introduced Model Context Protocol (MCP) provides such an environment [19], alongside similar alternative frameworks [20]. Although the frameworks differ in scope and design philosophy, they share a common principle: decoupling the AI model from the tools available to it, thereby enabling modular, governed access to external capabilities. In the following discussion, the MCP serves as a primary reference as an openly standardised communication protocol, whose explicit tool registration aligns well with audit requirements for transparency and traceability.

3.1 Establishing AI access to audit tools and techniques. From an audit perspective, frameworks like the MCP can be viewed as a structured approach to organising how AI systems access and use digital audit tools. Instead of establishing separate connections for each data source, analytical routine or CAAT, MCP provides a common framework through which approved tools can be made available to an AI system in a consistent and well-governed manner (see *Figure 2*).

This architecture enables flexible deployment, as a single agent can access multiple tools to complete complex audit tasks. Regardless of whether an audit team deploys one agent for substantive testing and another for compliance checks, or uses the same tools across multiple engagements, the tools remain reusable.

Within this framework, audit teams define which tools exist, the types of operations they perform and the conditions under which they may be used. These tools can reflect established elements of audit practice, for example: routines that extract and filter journal entries, procedures that evaluate control rules, access to policy or contract repositories, or scripts that re-perform calculations. MCP does not modify these tools; rather, it specifies a uniform way for an AI agent to discover them, request their execution, receive the resulting outputs and forward them to the auditor as illustrated in *Figure 2*.

A wide range of audit systems can be expressed as tools within the framework. The suite of frequently employed CAATs encompasses:

→ *data extraction and transformation tools,* to import, filter and prepare audit-relevant data, e. g. database connectors for ERP systems, Excel/CSV importers, PDF parsers, OCR engines for scanned documents and ETL pipelines for data consolidation [21].

→ *visualisation tools,* to facilitate data exploration and pattern recognition, e. g. journal entry dashboards [22], visual text analysis [23] and interactive continuous visualisations [24].

→ *process mining tools,* to analyse process conformity and discover deviations, e. g. process flow extractors from transaction logs, control violation detectors and sub-process discovery algorithms [25].

→ *anomaly detection tools,* to identify unusual patterns and internal inconsistencies, e. g. cluster analysis [26], extreme value detection [27] or financial ratio analysis [28].

This approach has implications for how digital audit work is coordinated. Instead of manually selecting and executing individual CAATs, auditors may formulate an analytical objective, and AI agents can identify and combine the relevant tools that have been made available (see *Figure 3*).

Looking ahead, recent developments indicate the potential for enhanced agent customisation through modular "skills" that can be, for example, integrated within Anthropic's framework [29]. This advancement further augments agentic capabilities by enabling models to access specialised expertise more efficiently, suggesting that the scope and sophistication of audit agent applications will continue to expand.

3.2 A gradual integration into audit processes and procedures. Once the auditor defines which tools the AI agent may access, the subsequent consideration is how these capabilities should be introduced into practice. Integrating agentic AI into audit workflows is typically incremental. Teams extend system access in stages, observe how individual tools are applied and refine governance arrangements as needed. This approach allows audit functions to expand digital support while maintaining clarity over system boundaries and supervisory processes.

An initial step may involve granting access to a single, well-defined tool, such as a document repository or a basic retrieval routine. In this configuration, the AI agent can locate and summarise relevant information, enabling audit teams to evaluate the quality and consistency of outputs under controlled conditions.

As experience is gained, additional tools – such as routines for extracting structured data or procedures that apply control rules – can be introduced. With multiple tools available, the AI agent can begin to combine them – for example by extracting figures from a document and comparing them with entries in a financial table. This supports more integrated analytical tasks while preserving a clear scope of operation.

Finally, broader toolchains may be enabled. Here, the AI agent can apply several tools in sequence, such as analysing transactions across units, applying validation rules, generating intermediate observations and preparing a consolidated summary for auditor review (see *Figure 3*). Thereby, each tool invocation can be logged, providing transparency into how the AI agent derived its outputs.

3.3 Examples of agentic AI-enabled audit processes. As audit teams expand the set of tools available to AI agents, the agents can begin to combine these capabilities into coherent workflows. The following examples illustrate how such toolchains may operate in practice.

3.3.1 Preparatory review. An AI agent could support the planning phase by synthesising existing information before fieldwork begins. Using a chain of document-retrieval, summarisation and cross-referencing tools, the system can assemble an initial briefing linking prior findings to current financial information. The auditor receives a structured overview that helps focus early risk discussions.

3.3.2 Substantive analytical procedures. An AI agent could support the execution of substantive analytical procedures by combining extraction, detection and assessment tools in sequence (see *Figure 3*). Using a journal-entry extraction tool, the agent retrieves the relevant transaction population from the entity's financial system. Finally, a risk-assessment tool prioritises the flagged anomalies based on materiality and risk indicators, producing a ranked set of entries for further investigation. The auditor can then concentrate on evaluating and corroborating the high-risk items rather than performing the underlying extraction and screening steps.

3.3.3 Drafting audit reports. Agents may also assist in preparing preliminary audit communications based on existing findings (see *Figure 3*). Such workflows might rely on findings-aggregation, summarisation and translation or style-guidance tools to generate initial drafts. Auditors begin with a structured document that they can refine, ensuring accuracy and appropriate framing.

Within the Swiss Federal Audit Office (SFAO), early experiments point toward similar possibilities. These visions are not science fiction; they are logical extensions of capabilities already being tested. But realising them demands careful governance, secure implementation, robust training and a culture that sees AI as a complement rather than a competitor.

Figure 3: **CONCEPTUAL END-TO-END AGENTIC AUDIT WORKFLOW: THE AI AGENT DECOMPOSES THE AUDITOR'S PROMPT, ORCHESTRATES THE SPECIALISED TOOLS, AND RETURNS STRUCTURED RESULTS FOR PROFESSIONAL JUDGMENT. THE SAME AGENT CAN ACCESS DIFFERENT TOOLS DEPENDING ON THE GIVEN TASK**

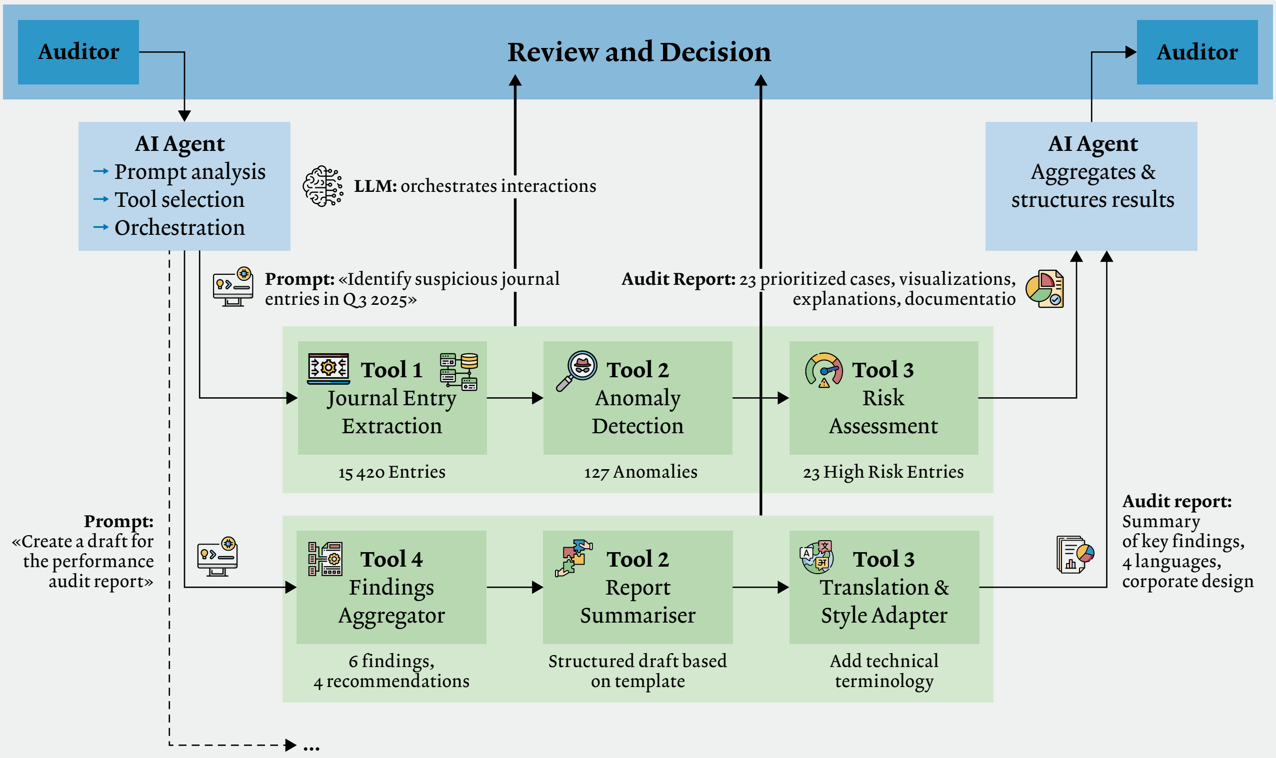

4. CHALLENGES

The transition from traditional CAATs to agentic AI auditing raises important considerations for how AI agents are granted access to, and make use of, these capabilities. AI agents can increasingly combine retrieval, validation, extraction and analytical tools to support audit tasks. Allowing such AI agents to autonomously operate tools introduces new forms of technical and organisational reliance. Understanding the associated risks, limitations and challenges remains essential.

4.1 Data quality. AI agents depend on the information made available to them, and incomplete records, inconsistent coding or incompatible formats can limit the reliability of outputs. Many public-sector systems were not designed for automated analysis, making data preparation and standardisation important prerequisites for deploying agentic tools [30].

4.2 Adoption. Auditors may be unfamiliar with the delegation of selected tasks to AI systems, and effective use requires clarity regarding system capabilities, limitations and potential failure modes. Training, pilot projects and transparent documentation can help support informed adoption and integration into existing practice [31].

4.3 Governance. Agentic AI requires clearly defined boundaries regarding tool access, autonomy and accountability. Agents should operate under least-privilege principles and within explicitly approved toolsets. Comprehensive logging and independent review of agent activity are essential to ensure traceability and detect unintended actions. Although frameworks such as MCP enhance transparency, organisations must define clear policies that preserve the auditor's responsibility for professional judgment. Agent outputs may inform decisions, but they cannot substitute them [32].

4.4 Resources. Deploying and maintaining agentic AI involves technical infrastructure, specialised expertise and ongoing operational investment. Institutions need to assess these demands relative to available capacity and long-term sustainability [33].

4.5 Bias. AI systems may reflect biases inherent in training data or design choices. Because impartiality is essential in auditing, outputs must remain explainable and subject to regular review. Structured logs can support such assessments, but continued monitoring and refinement remain necessary [34].

These challenges do not argue against the use of agentic AI; rather, they highlight the need for its careful and deliberate implementation. A successful deployment depends less on adopting advanced technology than on integrating it in ways that are consistent with institutional objectives, governance structures and operational conditions. As illustrated

Figure 4: **CRITICAL FACTORS FOR DEPLOYING AGENTIC AI SAFELY IN AN AUDITING CONTEXT**

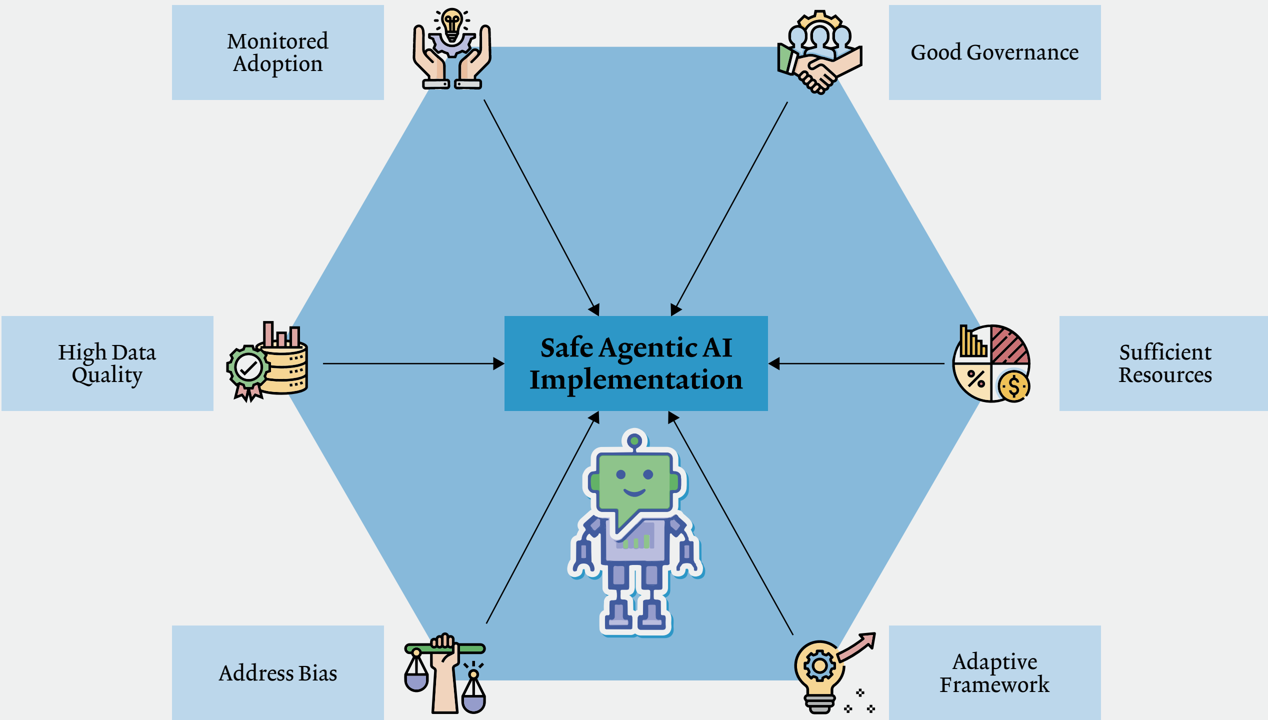

in *Figure 4*, the dimensions identified are not exhaustive. As agentic AI matures and its application in auditing broadens, additional considerations are likely to emerge. An adaptive approach that evolves alongside technological capabilities, regulatory expectations and institutional experience is therefore essential.

5. FROM PILOTS TO PRACTICE

The integration of agentic AI into audit workflows is still in its early stages, but its trajectory suggests a transformation in how audits will be conducted in the near future [35]. Leading audit firms have forged strategic partnerships to establish competitive positioning in the emerging market for agentic AI audit solutions [36]. Concurrently, innovative startups have introduced agentic AI audit platforms [37] designed to enhance audit procedures [38] and complex accounting functions [39]. The rapid adoption of open-source agentic frameworks such as OpenClaw – which gathered over 100,000 GitHub stars within days of its release – underscores the pace at which agent-tool architectures are maturing and entering mainstream use [40]. It reflects a broad recognition that the paradigm of AI systems that *act* rather than merely *respond* has moved from research concept to practical reality.

As agentic AI moves from pilot applications towards broader use, it remains important to consider its integration beyond technical capability and in light of the foundational principles of the International Standards on Auditing (ISA). The profession may therefore reflect on how these developments align with its established standards and responsibilities. For example, three areas appear particularly worthy of consideration.

5.1 How is the auditor's professional responsibility preserved? ISA 200 requires the exercise of professional judgment and professional scepticism. Even if AI agents plan procedures and execute toolchains, the evaluation of results and the conclusions drawn from them remain the responsibility of the auditor.

5.2 How is sufficient appropriate audit evidence obtained and demonstrated? ISA 500 requires evidence that is both reliable and relevant. In an agentic environment, this increases the importance of transparent data sources, validated tool logic and a clear record of how outputs were generated.

5.3 How are quality management, supervision and documentation ensured? ISA 220 and ISA 230 anchor audit quality in direction, review and documentation. The integration of agentic systems therefore requires defined permissions, controlled autonomy and audit trails that enable an experienced auditor to understand what was performed and on what basis.

The contours of this debate are already visible in the most recent revisions to the ISA standards, which introduce explicit references to "automated tools and techniques". Although these provisions were not drafted specifically with agentic AI in mind, they reflect an awareness that audit methodologies are becoming increasingly technology enabled. The ongoing debate about agentic systems can therefore be seen as a natural continuation of this evolution.

In the meantime, the audit profession is already exploring these new possibilities and its related questions. Various AI technologies are being adopted in large firms – ranging from simpler machine learning tools to more nuanced ones like deep learning neural networks [41]. While comprehensive, public case studies remain limited, there are early signals from major firms indicating interest in agentic approaches to audit automation, particularly in areas like document analysis and risk assessment [42]. What these efforts share is a recognition that the future of auditing does not lie in isolated AI applications, but in integrated ecosystems where agents, data and auditors work in concert.

Over the medium term, standardised approaches to agentic audit tools could enable cross-institutional collaboration. If institutions adopt common frameworks for defining audit procedures as agent-accessible tools, they gain the ability to share capabilities and best practices more easily. An automated control testing routine developed by one audit team could be adapted by another for different contexts, without requiring deep technical expertise.

For public auditors, these industry trends offer both inspiration and caution. Large firms operate in competitive markets where early adoption can yield strategic advantages. Public institutions, by contrast, must prioritise transparency, accountability and long-term sustainability. Open, standardised approaches to agent-enabled audit tools are particularly well suited to this context: they allow public bodies to learn from industry innovations while maintaining independence and control over their analytical capabilities.

6. OUTLOOK

Looking ahead, agentic AI appears less as a short-term experiment and more as a significant development in audit technology. Its potential lies in supporting efficiency gains across selected audit workflows while preserving the principles of independence, transparency and professional judgment. The trajectory outlined in this article – from established CAATs to agent-operable tools, from isolated routines to orchestrated toolchains – represents a potentially fundamental shift in how audit work is conducted. Yet by building on established methodologies and introducing capabilities in defined stages, audit institutions can navigate this transition in a controlled and deliberate manner.

Realising this potential demands sustained institutional effort. The challenges identified are not arguments against agentic AI, but preconditions for its responsible deployment. AI frameworks such as MCP can support this transition by offering standardised and transparent interfaces between AI agents and audit tools, helping institutions retain control and oversight while exploring new forms of analytical support.

For public auditors such as the Swiss Federal Audit Office (SFAO), this development carries a dual mandate. On the one hand, agentic AI offers the opportunity to enhance the efficiency and depth of audit work. On the other hand, as a supreme audit institution that increasingly audits the use of AI across the federal administration, the SFAO must also develop a thorough understanding of the technology's capabilities, limitations and risks. Engaging with agentic AI in its own workflows is therefore not only a matter of operational improvement, but also a prerequisite for credible and informed oversight of AI systems deployed by the entities it audits.

The auditing profession has continuously evolved by integrating new technolgies into established practices. Agentic AI represents the next step in that journey. Recent developments, such as OpenClaw, demonstrate that AI systems can autonomously discover, select, and combine tools to achieve defined objectives. This offers a practical glimpse into how future audit teams may operate within increasingly complex data environments. With human expertise firmly at the center, these technologies can help to deliver greater assurance, improve efficiency, and navigate the growing complexity of modern auditing with confidence. ■

Notes: *The authors would like to thank Eric Serge Jeannet, Viola Sini and Daniel Hasler at the Swiss Federal Audit Office (SFAO) for their valuable comments and suggestions on earlier drafts of this article.

1) Steinberger, P. (2025), OpenClaw: Your own personal AI assistant, GitHub,. retrieved February 10, 2026. https://github.com/openclaw/openclaw; Béchard, D.E. (January 30, 2026), OpenClaw – what happens when AI stops chatting and starts doing. Scientific American. Retrieved February 10, 2026. https://www.scientificamerican.com/article/moltbot-is-an-open-source-ai-agent-that-runs-your-computer/. **2)** Schreyer, M., Mäder, T., Ruud, T. (June 2025), A Novel Digital Audit Workforce: Embedding Agentic Artificial Intelligence in Internal Auditing, in: Expert Focus 2025/June. 260–267. **3)** Gu, H., Schreyer, M., Moffitt, K., Vasarhelyi, M. (2024). Artificial intelligence co-piloted auditing, International Journal of Accounting Information Systems, 54, 100698. **4)** Yee, L., Chui, M., Roberts, R., with Xu, S., Why Agents are the next Frontier of Generative AI, McKinsey Quarterly (July 24, 2024), https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/why-agents-are-the-next-frontier-of-generative-ai. **5)** OpenAI. GPT-4. (March 14, 2023), Retrieved November 27, 2025, https://openai.com/index/gpt-4-research/. **6)** Meta. Llama 4: Leading intelligence. (n.d.), retrieved November 27, 2025, https://www.llama.com/. **7)** Mistral. Frontier AI In Your Hands, (2025). retrieved November 27, 2025. https://mistral.ai/. **8)** Russell, S., Norvig, P. (2021), Artificial Intelligence: A Modern Approach, Global Edition, Pearson Higher Ed. **9)** NVIDIA, (2025, March 30), Retrieved November 24, 2025, AI Agents. NVIDIA Glossary, https://www.nvidia.com/en-us/glossary/ai-agents/. **10)** Huang, X., Liu, W., Chen, X., Wang, X., Wang, H., Lian, D., Chen, E. (2024), Understanding the Planning of LLM Agents: A Survey, arXiv Preprint arXiv:2402.02716. **11)** Wei, J., Bosma, M., Zhao, V.Y., Guu, K., Yu, A.W., Lester, B., Le, Q.V. (2021), Fine-tuned Language Models are Zero-Shot Learners, arXiv Preprint arXiv:2109.01652. **12)** Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W., Rocktäschel, T., and others. (2020), Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks, Advances in Neural Information Processing Systems, 33, 9459–9474. **13)** Masterman, T., Besen, S., Sawtell, M., Chao, A. (2024), The Landscape of Emerging AI Agent Architectures for Reasoning, Planning, and Tool Calling: A Survey. arXiv Preprint arXiv:2404.11584. **14)** Eulerich, M., Sanatizadeh, A., Vakilzadeh, H., Wood, D.A. (2024), Is it All Hype? ChatGPT's Performance and Disruptive Potential in the Accounting and Auditing Industries, Review of Accounting Studies. **15)** Coderre, D. (2009). Computer-aided fraud prevention and detection: A step-by-step guide. John Wiley & Sons. **16)** n8n. Product overview, (2025), retrieved January 9, 2026. https://n8n.io/features/. **17)** Anthropic. Introducing the Model Context Protocol. (November 25, 2024), retrieved November 4, 2025. https://www.anthropic.com/news/model-context-protocol. **18)** Schreyer, M., (2023), Representation Learning in Financial Auditing Using Deep Neural Networks, Ph.D. **19)** Anthropic. Architecture overview. (n.d.), retrieved November 4, 2025, https://modelcontextprotocol.io/docs/learn/architecture. **20)** Radulescu, R. and Garcia, J.V., Universal Tool Calling Protocol, (2026), retrieved January 14, 2026. https://github.com/universal-tool-calling-protocol; LangChain, Build agents faster, your way. (n.d.), retrieved February 12, 2026, https://www.langchain.com/langchain; Google. Announcing the Agent2Agent Protocol (A2A). (April 9, 2025). Retrieved February 12, 2026). https://developers.googleblog.com/en/a2a-a-new-era-of-agent-interoperability/. **21)** Alteryx, Where Analytics, Automation, and AI Come Together (n.d.), retrieved December 10, 2025. https://www.alteryx.com/de/; R Core Team (2021), R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria, https://www.R-project.org/; Wickham, H., Bryan, J., (2025), readxl: Read Excel Files, R package version 1.4.5, https://github.com/tidyverse/readxl; https://readxl.tidyverse.org; Wickham, H., Bryan J., (2025), readxl: Read Excel Files. R package version 1.4.5, https://github.com/tidyverse/readxl, https://readxl.tidyverse.org; Wickham, H., François, R., Henry, L., Müller, K., Vaughan, D., (2025), dplyr: A Grammar of Data Manipulation. R package version 1.1.4, https://dplyr.tidyverse.org; The pandas development team, Pandas, (February 2020), Zenodo (latest), 10.5281/zenodo.3509134. **22)** Tatu, A., Schreyer, M., Hagelauer, J., and Wang, J., Visual Exploration of Journal Entries to Detect Accounting Irregularities and Fraud, In IEEE VIS 2014 Workshop – Business VIS 14 (2014). **23)** Lee, H., Zhang, L., Liu, Q., and Vasarhelyi M., Text Visual Analysis in Auditing: Data Analytics for Journal Entries Testing, International Journal of Accounting Information Systems 46 (2022), 100571. **24)** Zhang, G., Atasoy, H., and Vasarhelyi M., A. Continuous Monitoring with Machine Learning and Interactive Data Visualization: An Application to a Healthcare Payroll Process. International Journal of Accounting Information Systems 46 (2022), 100570. **25)** Jans, M., Alles, M.G., and Vasarhelyi, M.A., A Field Study on the Use of Process Mining of Event Logs as an Analytical Procedure in Auditing, The Accounting Review 89, 5 (2014), 1751–1773; Jans, M., Van Der Werf, J.M., Lybaert, N., and Vanhoof, K., A Business Process Mining Application for Internal Transaction Fraud Mitigation. Expert Systems with Applications 38, 10 (2011), 13351–13359; Laghmouch, M., Jans, M., and Depaire, B., Classifying Process Deviations with Weak Supervision, in 2020 2nd International Conference on Process Mining (ICPM) (2020), IEEE, pp. 89–96; Werner, M., and Gehrke, N., Multilevel Process Mining for Financial Audits, IEEE Transactions on Services Computing 8, 6 (2015), 820–832; Werner, M., and Gehrke, N., Multilevel Process Mining for Financial Audits, IEEE Transactions on Services Computing 8, 6 (2015), 820–832. **26)** Jolly, D., KPMG Joins MindBridge in Global Launch of AI-Assisted Audit Tech. https://news.bloombergtax.com/financial-accounting/kpmg-joins-mindbridge-in-global-launch-of-ai-assisted-audit-tech, 2023. Accessed: 2023-07-08; Thiprungsri S., and Vasarhelyi M.A., Cluster Analysis for Anomaly Detection in Accounting Data: An Audit Approach. International Journal of Digital Accounting Research 11 (2011). **27)** Argyrou, A., Auditing Journal Entries using Extreme Value Theory, in 21st European Conference on Information Systems, ECIS 2013 (2013), Association for Information Systems. **28)** Spathis, C.T., Detecting False Financial Statements using Published Data: Some Evidence from Greece, Managerial Auditing Journal 17, 4 (2002), 179–191. **29)** Anthropic. Introducing Agent Skills, (October 16, 2025), retrieved December 11, 2025. https://www.claude.com/blog/skills. **30)** 6 key challenges of auditing AI – and how to approach them. KPMG. (May 2025), retrieved December 16, 2025. https://kpmg.com/sg/en/insights/ai-and-innovation/ai-in-internal-audit/download-guide.html. **31)** OECD (2025), Governing with Artificial Intelligence: The State of Play and Way Forward in Core Government Functions, OECD Publishing, Paris, https://doi.org/10.1787/795de142-en. **32)** Kokina, J., Blanchette, S., Davenport, T.H., Pachamanova, D., (2025), P. 10–11, Challenges and opportunities for artificial intelligence in auditing: Evidence from the field. International Journal of Accounting Information Systems, 56, 100734; Jim Reavis, (January 28, 2026), Leveling Up Autonomy in Agentic AI, retrieved February 9, 2026, https://cloudsecurityalliance.org/blog/2026/01/28/levels-of-autonomy; Meijer, E. (2025), Guardians of the Agents, Commun. ACM 69, 1 (January 2026), 46–52. https://doi.org/10.1145/3777544. **33)** The 5 biggest AI adoption challenges for 2025. Stryker, C. (February 14, 2025). Retrieved December 16, 2025. https://www.ibm.com/think/insights/ai-adoption-challenges. **34)** Murikah, W., Nthenge, J.K., Musyoka, F.M., (2024), Bias and ethics of AI systems applied in auditing-A systematic review, Scientific African, 25, e02281; Lacmanovic, S., Skare, M., (2025), Artificial intelligence bias auditing–current approaches, challenges and lessons from practice, Review of Accounting and Finance, (ahead-of-print). **35)** Kokina, J., Blanchette, S., Davenport, T.H., Pachamanova, D., (2025), Challenges and opportunities for artificial intelligence in auditing: Evidence from the field. International Journal of Accounting Information Systems, 56, 100734. **36)** Datasnipper, What Is the Future of Audit? Agentic AI in Excel. (2025), retrieved December 9, 2025, https://www.datasnipper.com/resources/agentic-ai-in-audit-and-finance. **37)** neuland.ai, Die KI-Plattform für Ihr Unternehmen, (n.d.). retrieved January 9, 2026. https://new.neuland.ai/#integrationen. **38)** Trullion, AI Meets Accounting. (2025), retrieved December 9, 2025, https://trullion.com/. **39)** Tellen, The AI Agent Workforce for Accounting Firms, (2025), retrieved December 10, 2025, https://www.tellen.ai/. **40)** Steinberger, P., (2025), OpenClaw: Your own personal AI assistant, GitHub, retrieved February 10, 2026, https://github.com/openclaw/openclaw. **41)** Kokina, J., Blanchette, S., Davenport, T.H., Pachamanova, D., (2025), Challenges and opportunities for artificial intelligence in auditing: Evidence from the field. International Journal of Accounting Information Systems, 56, 100734. **42)** Chandrasekaran, S., Edin, P., Lohr, T., (June 2025), Agentic AI advantage: Unlocking next-level value, https://kpmg.com/kpmg-us/content/dam/kpmg/pdf/2025/kpmg-agentic-ai-advantage.pdf; Wood, M., (May 16, 2025), PwC adds support for Model Context Protocol in agent OS, https://www.pwc.com/us/en/about-us/newsroom/press-releases/pwc-adds-support-for-mcp-in-agent-os.html; Sharma, P., Srinivas, V., Chauhan, A., (August 14, 2025), How banks can supercharge intelligent automation with agentic AI, https://www.deloitte.com/us/en/insights/industry/financial-services/agentic-ai-banking.html.