UNMAS DENPASAR

**Jurnal Inovasi Akuntansi (JIA)**
**E-ISSN 2988-1536**
**https://e-journal.unmas.ac.id/index.php/jia/index**

# Auditing Artificial Intelligence-Driven Financial Systems: Accountability, Transparency, and Auditor Liability in Algorithm-Based Decision Making

**Dickson Mdhlalose**[1*]

[1]Department of Information and Communication Technology, National Electronic Media Institute of South Africa, Johannesburg, South Africa
[*]Corresponding author, e-mail: dsskosana@gmail.com

DOI: https://doi.org/10.36733/jia.v4i1.13758

**Abstract**

**Purpose:** This study examines how the growing use of Artificial Intelligence (AI) in financial reporting and auditing affects audit reliability, accountability, and transparency. It focuses on key challenges such as AI's "black box" nature, outdated auditing standards, limited auditor expertise, and unclear legal responsibility. **Method:** The study uses a conceptual and literature-based approach by reviewing prior research, auditing standards, and regulatory developments related to AI, financial reporting, and audit assurance. **Findings:** The study finds that although AI can improve risk assessment and audit efficiency, its complexity and lack of transparency may increase audit risk. Current standards, such as ISA 315 and ISA 500, are not fully suitable for algorithm-based decision-making. The study also highlights a shortage of auditors with data science skills and uncertainty over legal accountability between auditors, companies, and AI software providers. **Implications:** The study proposes the Assurance for Ethical and Governed AI Systems (AEGIS) framework, which emphasizes system review, AI explainability, and continuous monitoring. It recommends that standard setters, including the IAASB, develop AI-specific audit guidance, strengthen auditor training in data analytics and AI governance, and create a fairer legal responsibility framework. Without these changes, the audit profession may struggle to remain relevant in an AI-driven reporting environment.

**Keywords:** artificial intelligence, auditing, accountability, algorithmic transparency, auditor liability



## Introduction

Artificial intelligence's (AIs) increasing use in company finances completely changes how financial data is managed, examined, and reported. Now, machine learning handles important jobs such as determining income, forecasting loan defaults, and identifying sophisticated fraud (Brennan & Subramanyam, 2023; Zhang et al., 2022). Though switching from systems based on human decisions to intelligent algorithms should be quicker and offer improved understanding, it creates significant issues for the fundamental principles of auditing: being understandable, having proof, and fulfilling professional responsibilities. This research proposes that auditing is at a critical point, where the old ways of doing things just do not align with the technologically advanced environment that the company's auditors are meant to check. This mismatch is most evident in "black-box" AI, where even the developers do not understand how the system connects information to outcomes (Burrell, 2016). Financial institutions are using all kinds of "black box" models, for example, deep neural networks for evaluating how likely someone is to default on a loan, eXtreme Gradient Boosting (XGBoost) (a type of gradient boosting) for finding fraudulent activity, and natural language processing (NLP) to automatically look over financial agreements and reports (Kokina & Davenport, 2023). Random forest classifiers are commonly used to point out unusual entries in accounting journals, and recurrent neural networks work through lists of transactions to discover odd timing in how revenue comes in (Zhang et al., 2022). The important thing is that, unlike older systems, which use specific rules, each of these methods concludes by making millions of small changes to its internal settings. Because of this, it is pretty much impossible to follow the logic from the original data to the final financial decision in a way a person could understand. Auditing has always depended on a clear record of transactions following a demonstrable system. When a deep learning network makes a critical judgment about credit ratings or inventory valuation, that record vanishes into complicated calculations (Kokina & Davenport, 2023). This obscurity directly opposes the requirement to find "sufficient appropriate audit evidence" (International Standard on Auditing [ISA] 500, International Auditing and Assurance Standards Board [IAASB], 2022), resulting in what experts call a "responsibility void" in modern financial reporting (Brennan & Subramanyam, 2023).

This has caused a great deal of anxiety among researchers. Risks such as hidden biases, which can lead to repeated mistakes (Munoko et al., 2020), or the possibility of hacking with 'data poisoning' (Papernot et al., 2018) have been documented. There is also the potential for 'automation bias', where auditors might place too much faith in a machine and become less questioning (Alles & Gray, 2020). The Association of Certified Fraud Examiners (ACFE) 2024 worldwide fraud report has a very recent, clear example of this risk. It shows more fraudsters using generative AI to create fake financial documents and doing it on a large scale. They are making things like invented invoices, altered bank statements, and journal entries created by computer programmes; these entries look, statistically, exactly like proper records, and take advantage of the fact that AI-improved auditing systems are meant to be good at spotting patterns. In all these situations, people trusting too much in computers were directly involved. Because the AI did not point out anything unusual, human auditors did not do the thorough checking they normally would and would have discovered the fraud. This shows that fraud using AI and auditors believing what the AI tells them without question are issues that support each other, and we need specific steps in our processes to protect against both (ACFE, 2024; Munoko et al., 2020). Moreover, because AI continually learns, it is a continually changing target for assessment, leading to difficult questions regarding how long an audit opinion can be considered valid in such a rapidly evolving system (Vasarhelyi et al., 2015). Despite this increasing concern, the current research and professional guidance are missing something important. Most discussion is still based on theory, on 'what if' scenarios, and does not provide real-world examples or tested auditing procedures (Zhang et al., 2022). There is a significant shortage of detailed instructions to guide an auditor from the initial risk assessment of an AI model to the final verification of its results as credible evidence. This gap between what is discussed and what can be done leaves auditors without the necessary tools to do their jobs in a world where AI is heavily used.

Adding to the lack of methodology is a considerable shortage of official regulations. Current auditing standards, such as ISA and 620 regarding risk and experts, were written before AI existed and are not suited to today's algorithmic complexity (Brown-Liburd et al., 2015). While organizations like the IAASB and Public Company Accounting Oversight Board

[PCAOB] have begun to discuss the issue (IAASB, 2022; PCAOB, 2019), they are being too slow to establish enforceable rules. This delay in regulation creates a risky situation where each firm approaches the issue differently, exposing auditors to legal problems they have not encountered before. The legal side of auditing AI is similarly confusing. Auditors are normally expected to exercise 'due care', but nobody has yet defined what this means when examining an algorithm that cannot be explained (Brennan & Subramanyam, 2023). If an AI error causes a financial mistake, responsibility could be shared between the auditor, the company, or the software creator, but our current laws are not equipped to fairly allocate blame (Munoko et al., 2020). This uncertainty increases audit risk and may cause firms to prioritize self-protection over effective auditing.

To complicate matters further, there is a large shortfall in relevant skills within the profession. Auditing AI requires a person with a combination of accounting, data science, and ethical understanding (Mikalef & Gupta, 2021). Current qualifications and certifications, for example, the Certified Public Accountant (CPA), do not yet create this 'audit data scientist' skillset. This results in a workforce that might not be ready for the new challenges they are facing (Institute of Chartered Accountants in England and Wales [ICAEW], 2019). Without changes to the training of future auditors, the difference between their abilities and what the technology demands will only grow. This study is designed to directly address these connected issues. It goes beyond simply identifying the problems to testing how well our current regulations, skills, and laws deal with AI in finance. The final aim is to combine theory with practical needs to develop a robust, evidence-based auditing framework. This framework must give a clear way to make AI transparent, determine who is responsible when things fail, and establish strict guidelines for evaluating AI data. Doing this will close the gap between rapidly changing technology and the reliable financial confidence we all need.

**Literature review**

*Current Audit Standards*

Currently, the rules for financial auditing do not really deal with how to review financial work done by AI or automated systems. These existing rules are based on how a person would make judgments and use a bit of checking here and there, and are far too slow for AI, which works with algorithms at great speed. ISA 315 and 330, for instance, require auditors to understand a company's data systems, but they presume a person is operating those systems, and we can understand and question their reasoning (IAASB, 2022). However, machine learning is a "black box", and this directly clashes with ISA 500's basic requirement to find "reliable evidence". Auditors are finding it almost impossible to follow the rules because they cannot determine why an AI highlighted a particular transaction or how it arrived at a complex valuation (Brennan & Subramanyam, 2023). This creates a huge problem: the rules say auditors need to understand, but not advise on, how to understand a machine whose reasoning is hidden and always changing. This absence of advice is particularly problematic with automated decisions and the controls meant to supervise them. Auditors are meant to use the Committee of Sponsoring Organizations of the Treadway Commission (COSO) frameworks to assess risks, but AI governance has not been included in their updates yet, so there is no obvious way to evaluate for biased data or how stable the model is (Zhang et al., 2022). Even ISA 620, which is about using an "expert", is not suitable; the "expert" is not a human professional with standards of behaviour, but a programme whose "judgement" is simply statistics (Kokina & Davenport, 2023). Because of this slow response in terms of regulations, auditors are currently treating sophisticated AI as if it were an old computer system or a human specialist, and this approach does not identify the unique dangers of a technology that can think and adjust itself.

Those who set standards have begun to discuss this widening gap, but we do not have any official, concrete rules yet. The IAASB has included technology in its "Future of Audit" plans, but is still researching with no changes to rules yet (IAASB, 2022). In the United States, the PCAOB and American Institute of Certified Public Accountants (AICPA) have had discussions and exchanged thoughts, but these are not rules that auditors are required to follow (PCAOB, 2019). This slowness is creating a risky, unregulated situation where auditing firms are inventing their own procedures with no common standards. And as many in the field point out, this inconsistency does not just increase the likelihood of errors; it also makes it difficult for the profession to justify its work on financial statements heavily based on AI (Alles & Gray, 2020). In effect, current rules barely touch AI, and we need a complete overhaul of how auditing standards are written. A little tinkering with old rules will not do; we require completely new, separate standards focused specifically on auditing automated decisions. These standards will need to give auditors a method for checking AI governance, confirming the quality of the data the AI learns from, and testing the AI’s internal workings with “explainable AI” methods (Mikalef & Gupta, 2021). Until these changes are officially in place, auditors are in a difficult, risky area. The public expects them to confirm finances that are driven by AI, but they have a set of rules that do not even acknowledge that the technology exists, and this puts the entire auditing profession at risk of becoming pointless.

*Accountability and liability challenges*

AI's increasing use in preparing financial reports is fundamentally changing how we determine who is responsible for accuracy. Auditors are traditionally expected to carefully consider and use their own judgment to confirm that financial reports are correct (AICPA, 2014). However, when AI systems – belonging to either the company being audited, or the auditors themselves – are choosing accounting entries or determining values, that human judgment is being transferred to the machine. Because these sophisticated systems are complex and opaque ("black boxes"), auditors may be unable to say why an AI made an error, even though they are the ones formally approving the report (Brennan & Subramanyam, 2023). This creates a worrying “accountability gap”: the auditor is legally liable for their opinion, but realistically cannot understand the machine’s reasoning for the evidence it provided. As a direct result, audit firms face significantly higher legal risks. Currently, auditors can be sued for being negligent if they have not been careful enough, and lawyers will soon be claiming that this carefulness absolutely must include a thorough understanding of AI (Restatement of Torts, 1965). A “skill gap” liability arises if auditors do not have the technical skills to review these systems (Kokina & Davenport, 2023), and they might even be blamed for “automation bias” – that is, simply accepting the machine’s answer without enough questioning (Munoko et al., 2020). It will be very difficult in court to determine exactly who is to blame when an autonomous system makes a mistake, and auditors could end up being held responsible for failures they could not have stopped.

The legal complications get even worse because of all the people involved in creating and operating AI: software companies, data scientists, and the IT department at the company being audited (Zhang et al., 2022) could all be at fault when something goes wrong with the data or the algorithm itself. However, most software agreements have very strong protections for developers against lawsuits (Issa et al., 2016). This frequently means the audit firm is the only one left that can be sued, the 'last person standing' with the money and professional liability insurance. Without revised guidance from organizations like the IAASB to give a definite defence, auditors have no solid way to demonstrate they performed their jobs as required. To deal with these serious dangers, the profession needs a large, multi-faceted overhaul. Auditors should start receiving education in “algorithmic awareness” and insist on “explainable AI” – AI that provides a clear, readable record of how it reached its important conclusions (Kokina & Davenport, 2023). The PCAOB and IAASB regulatory bodies need to quickly publish specific rules for how much investigation is needed when assessing accounts that AI is running (Brown-Liburd et al., 2015). We might even require new legal approaches, such as “enterprise liability”, to make sure everyone involved, including the developers, shares the blame equitably (Brennan & Subramanyam, 2023). Otherwise, auditors are in a difficult spot: they are pushed to employ fast AI tools but will be assessed against old standards that do not allow for errors made by machines.

*The AEGIS audit framework*

This study aims to develop a way to audit, and therefore make more understandable and trustworthy, financial data created by artificial intelligence, as you can see in Figure 1. AEGIS is intended to be a constantly repeating cycle of human review and machine checking. The main part of this is a progression that begins with the AI system and financial data themselves, goes through testing and evaluation, and ultimately delivers a final audit decision. Three important things support this: ensuring auditors have the skills to work with AI, a "3-Party Responsibility Matrix" to clarify who is responsible for what, and making sure everything is current with regulations. This means auditing AI is not simply completing a list; it is a system that will evolve as skills and regulations, and the audit itself, are used. The core of the system is verifying things in two ways. Auditors are not just looking at what the AI calculates; they are also looking at how it calculates it. The "AI System Audit" is about the model's technical health and how it is managed, while “Output Statistical Validation” is another set of tests on the results. This double check makes sure auditors do not just accept the AI's answers but actively check its reliability and correctness. Auditors will be able to determine how much confidence to have in what they are looking at using an Evidence Reliability Pyramid (ERP), which ranks evidence from fully human checked (Level 1) to incomprehensible math (Level 4), and gives a clear way to audit according to how much risk is involved.

AEGIS also has a "Sentinel System" that sends instant warnings for issues as they arise. Because it is always looking, auditors can identify "model drift" or data breaches during the audit, instead of discovering them later after something has failed. Imagine a store that uses AI to figure out how much money it has made. Let us say around halfway through the year, the type of things they sell changes significantly, and because of this, the AI now puts transactions into categories in noticeably different amounts. Sentinel would instantly notice this big statistical difference from how the AI usually performs and send an immediate warning to an auditor to look at it before those numbers get into the official financial reports (Vasarhelyi et al., 2015). And, if the information going into the AI suddenly has numbers far outside what the AI is used to - perhaps incredibly large changes due to currency rates or unusually large markdowns on a lot of old stock - the system will flag it if the value passes a certain point, and send the unusual issue to an auditor. This stops incorrect AI results from finding their way into the financial results the company shares. Furthermore, at important points in the process, the system has “branching points”; if an AI result is unusually large or complicated, a human must review it. This ensures that anything with a lot of risk will not be missed by a person. This forward-thinking method changes auditing from looking at what has happened to a continuing check on a business’s technology. Lastly, the entire structure is to produce an audit conclusion that openly discusses what AI can do and what it cannot do. The final decision is either a typical opinion with a disclosure about AI, or a restricted opinion if the AI’s workings were too complex to be verified. By continually watching, checking twice, and having clear evidence standards, AEGIS goes beyond discussing the problems of AI and provides a practical, organized audit method. It offers a way to balance modern technology with traditional thoroughness and regulatory adherence for trustworthy and sound financial markets.

*Professional competencies and regulatory reforms*

Auditing financial systems that use AI needs a big change in what auditors are able to do. We are moving beyond standard accounting towards something of a combined role - a lot of people are now calling this person an ‘audit data scientist’. If teachers and those designing courses want a real idea of what an audit data scientist can do, the most important technical skills fall into five areas. Firstly, they need to be good at programming: Python is key for getting data out, looking inside models, and automating jobs, and Structured Query Language (SQL) is for asking questions of large financial databases and following a transaction’s history. Secondly, they must understand machine learning. That means knowing how models are built, how they are tested and improved, how to spot when they are too closely fitted to the data, and how to find prejudice in the supervised classifiers (those are the types of models often used for predicting things in finance). Thirdly, they need to know how to use explainability tools. Specifically, being able to use things like SHapley Additive exPlanations (SHAP) and Local Interpretable Model-agnostic Explanations (LIME) is important; these allow an auditor to find out why a model gave a particular answer and to decide if that reasoning is good enough to be used as proof during an audit. Statistical analysis is fourth, and ideally, they did use R (or something similar) to do tests of assumptions, look at differences in values, and check regression analysis to check the risk scores created by the AI on their own. Finally, they must be able to manage and be sure about the data and how it is brought to the AI. This means checking the data’s history documentation, the security of who can access it, and that the data going into the AI has not been damaged, changed in a dishonest way, or significantly altered before the AI uses it (Mikalef & Gupta, 2021; ICAEW, 2019).

Auditors cannot just treat AI as if it is a black box anymore; they need to have enough technical understanding to thoroughly investigate every stage of a model's lifecycle, from the initial data selection to its final application (Brennan & Subramanyam, 2023). That includes understanding ideas like ‘overfitting’ or hidden biases to determine if a system is suitable for its intended purpose (Kokina & Davenport, 2023). Without this knowledge, they will not be able to effectively question the design or the way a system operates. So, professional qualifications and examinations such as the CPA need to start including data science and AI ethics to get the next group of auditors ready. But knowing the technology is not the whole story, stronger ethics and governance skills are also essential for auditing AI. Auditors need to be very good at assessing a company's AI rules, so how they manage data quality and the dangers of the model. This involves looking at the processes that generate the figures, where the data comes from, and how the model’s performance is tracked over time (Sun, 2019). They must also bring long-established principles, like ‘completeness’, to these new technological situations. A decision about whether a biased model led to a financial mistake, for instance, requires both technical and ethical thought. And to get to this point, we need more than classroom learning; we need practical training, led by the professional organizations.

Currently, developing these skills is being slowed down by a lack of current regulations. Existing auditing guidelines like ISA 315 and 500 were written for processes done by people, and do not cover AI at all (IAASB, 2022). Regulators like the PCAOB and the IAASB need to provide detailed guidance that states clearly what an auditor is expected to do when evaluating AI controls or when using models from another company (Brown-Liburd et al., 2015). Without these revisions, auditors will be left to make things up based on outdated rules, which are unsuitable, resulting in inconsistent work and the possibility of legal issues, because they will have no clear professional standard to refer to in court to justify their decisions. Ultimately, to make AI auditing work, a partnership between universities, auditing businesses, and regulators is required. Universities must bring their courses up to date to create these ‘hybrid’ professionals. Audit firms need to invest more in staff training and create positions like ‘AI Audit Specialists’ (Deloitte, 2021). And crucially, regulators need to accelerate the updating of the regulations. That could mean a requirement for new reports on AI governance, or rules for 'explainable AI' tools, which would make the reasoning behind machine decisions easier to follow (Kokina & Davenport, 2023). If we use 20th-century rules on 21st-century technology, the audits will fail. The profession will only be able to genuinely assure people of a system's reliability in an economy driven by AI if everyone moves forward together.

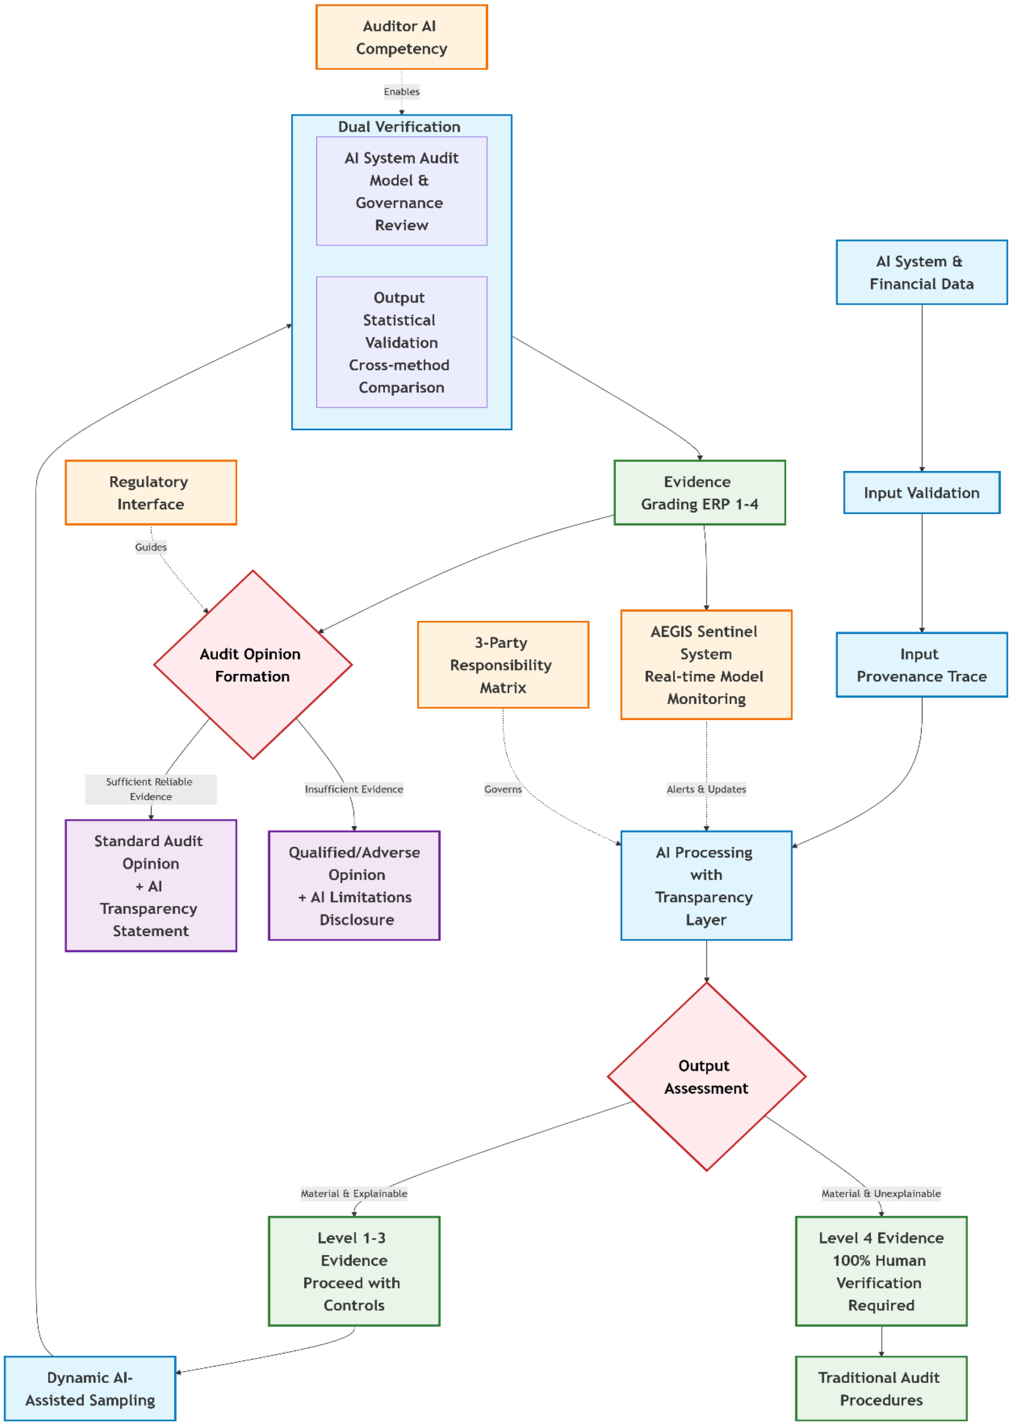


Figure 1. Audit model
Source: developed by the author

*Impact of AI adoption on audit risk assessment*

AI is fundamentally altering how audits assess risk. Rather than looking at just a selection of transactions, AI can now examine every single one as it happens, and as Zhang et al. (2020) point out, this leads to much more accurate identification of unusual trends or potential risks. For example, AI can simultaneously monitor for fraud at every branch of a company worldwide, changing risk assessment from an annual event into a continuous one. This does, though, have a drawback: auditors might develop "automation bias" and over-rely on the machine, losing their professional scepticism (Munoko et al., 2020). Although AI offers a more thorough view of risk, human evaluation remains crucial to ensure the risk assessment tool itself works as it should. And when collecting evidence, AI allows a shift from testing a small amount to looking at the entirety of the data. NLP can read through thousands of contracts in seconds for compliance, and other applications can chart connections to uncover concealed fraud (Kokina & Davenport, 2023). Because it is based on all the data available, this move makes audit evidence much stronger. However, this creates a very significant “garbage in, gospel out” issue; biased data or errors in the AI’s programming will lead to huge amounts of inaccurate evidence (Brennan & Subramanyam, 2023). Auditors now need to concentrate on validating the "data pipelines" which provide information to the AI, instead of selecting samples.

The main benefit expected is a substantial improvement in audit quality, fewer mistakes, and improved fraud detection. In an ideal situation, AI would do the tedious, repetitive parts of the job, allowing auditors to concentrate on more challenging problem-solving (ICAEW, 2017). AI, because it does not tire and is perfectly consistent, has the potential to lower human error and make audits more reliable (Frey & Osborne, 2017). However, this isn't automatic. Audit quality relies on scepticism and good ethics, which machines do not possess. If auditors become simply “tick-box” approvers of the AI’s outputs, then the quality of their judgment will decline (Alles & Gray, 2020). Achieving success requires audit firms to build a cooperative relationship where humans and AI both assess things carefully. Ultimately, getting more effective audits from these advanced tools will require surmounting a number of significant challenges. One major one is the “black box” issue: if an auditor can't explain why the AI flagged something as a risk, it is difficult to defend their conclusion under standards like ISA 500 (IAASB, 2022). Regulations are also lagging, and many firms are working without a clear definition of what a 'good' AI-driven audit looks like (Klynveld Peat Marwick Goerdeler [KPMG], 2019). Real success will depend on three things: “explainable AI” (XAI) that shows its working, speedy revisions to professional guidelines from bodies like the PCAOB, and a large-scale initiative to give auditors new skills. Without this kind of balance, AI could simply add to the complication of audits without improving them.

## Method

The research involved a detailed qualitative approach, carefully examining and combining ideas from what is already known. Its primary purpose was to carefully look at and assess academic research, professional auditing standards, and official statements by regulators to address what’s missing from auditing financial systems that use artificial intelligence. This approach is very suitable for initial exploration with the goal of forming new theories and identifying the essential difficulties in a relatively new subject where concrete data is scarce (Creswell & Poth, 2018). The research was structured to first identify and deconstruct the biggest obstacles, such as the ‘black box’ nature of AI, complicated questions of legal responsibility, and the slow pace of relevant laws, then to integrate the understanding gained from this to develop a completely new auditing structure.

Data was collected by way of a large amount of documentation and was not based on conducting experiments. This meant a rigorous look at published academic work, important books on auditing AI (for example, Brennan & Subramanyam, 2023), and current and proposed regulations from key organizations such as the IAASB and PCAOB, as well as reports by leading accountancy firms (ICAEW, 2019; Deloitte, 2021). A thematic content analysis was employed, and many texts were labelled according to repeated concepts related to responsibility, openness, needed skills, and the dangers involved. This allowed for the merging of different perspectives from accounting, the legal profession, data science, and ethics into a unified view of how prepared auditing is for AI, and it used the accepted best practices for research that crosses multiple disciplines (Snyder, 2019).

This work is fundamentally about developing theory and proposing practical improvements. The research method goes from a thorough criticism of the weaknesses in current auditing approaches and auditor training, to the development of the Assurance for Ethical and Governed (AEGIS) AI Systems framework. Constructing this framework is the key focus of the research, changing abstractly identified shortcomings into a firm, structured guide for auditors. Consequently, the research is not about proving complicated mathematical ideas, but providing a sensible, thoroughly supported basis for directing future experiments, lawmaking, and professional development. It satisfies an important requirement in an area which has, up to now, largely consisted of assumptions and conjecture (Bailey, 2017).

**Results and discussion**

This analysis shows an important and pervasive problem: what AI can do in auditing is a lot more than what we have rules and ways to officially manage it. Brown-Liburd and others (2015) pointed out that AI could really change how we assess risk by looking at everything, all the time. However, the results here confirm that the fact that we cannot see how complicated AI makes its decisions (its "black-box" nature) prevents auditors from getting enough proper proof – that is what ISA 500 (IAASB, 2022) says they need. This leads to something of a contradiction. AI can find more risks than ever before, yet auditors very often cannot follow or check the reasoning an algorithm used to find them, and that puts the proof supporting their opinion at risk (Brennan & Subramanyam, 2023; Kokina & Davenport, 2023). Because of this, the better quality of assurance we were promised by AI is, for now, only a thought. We absolutely need Explainable AI (XAI) as the base for being able to audit, and the discussion should focus on that. The findings show auditors cannot properly do their job when dealing with completely unreadable systems. Burrell (2016) says that being unable to see how these advanced machine learning models work is not a temporary problem; it's how they are built. So, using AI in important areas of preparing financial reports must depend on using XAI methods that provide a traceable audit trail that a human can understand. This is not just a technical improvement; it is a professional obligation to fill the accountability gap Munoko and colleagues (2020) identified. The audit framework in this study uses ERP and ranks evidence on how easy it is to explain, creating a clear path for an auditor to rely on things, and fitting with professional scepticism.

What is important is to use tools for XAI as a starting point for investigation, not as a substitute for an auditor's professional opinion. Auditors should be actively questioning and carefully checking the explanations given by XAI against their usual audit work (like looking at trends and detailed testing of figures) and not just believe what the AI says it did. For any important decision made by AI, auditors need to write down both what the XAI tool said and their own evaluation of whether that makes sense, considering what they know about the client's business and its industry. If the XAI explanation does not match other audit information or the auditor's understanding of the company, that difference should lead to more detailed testing, not to blindly trusting the AI. Auditors also must be suspicious of the XAI tool itself: the method used to make it understandable (whether it is SHAP values, LIME, or visuals of what the AI is focusing on) can potentially bend the truth and should be tested for accuracy before being used as audit proof. This means that with AI helping with audits, professional scepticism must be applied to both the financial numbers the AI is working with and the tools explaining them. Ultimately, a human auditor must be the final decision-maker in an audit.

Moreover, the results emphasise that the biggest obstacle to moving forward is not the technology itself, but regulations and training. The analysis makes it clear that current auditing standards are not prepared for a world where AI is in charge, as they are based on the idea of human-controlled processes and logical steps that can be traced (Alles & Gray, 2020). The delay in regulation, as the IAASB (2022) and PCAOB (2019) have shown, is creating a dangerous situation with lots of different ways of doing things. And the shortage of people with the right skills - "audit data scientists" - is a direct result of accounting courses and professional qualifications not keeping up (ICAEW, 2019). This leads to the conclusion that small changes will not be enough. We need a complete overhaul at the same time: standard-setting bodies need to create specific auditing guidance for AI, universities and professional organisations need to completely revise education to create people with both technical and professional skills, and firms need to invest in strong ways of managing and auditing AI (Mikalef & Gupta, 2021). Ultimately, the results back up the main idea of this study: without a well-defined system like the AEGIS model suggested, using AI in auditing will increase the chances of being held liable and damage reputations, rather than improving quality. Because of the patchy way liability is allocated (auditors could be left as the only ones to blame for AI errors because of how vendors are protected, Zhang et al., 2022), a more forward-thinking definition of professional responsibility is needed. We must discuss and push for a new way of auditing, which includes continually checking that AI systems are working correctly, watching for changes in how the model performs in real time, and openly stating how much the audit has depended on AI in the audit opinion. Only by making these kinds of systems standard practice can the profession deal with the difference between the new technology and its legal responsibility to give meaningful and justifiable assurance about financial statements produced by AI.

**Conclusions**

The research has carefully looked at the big problems artificial intelligence causes for the very basics of how financial audits are done. The fact that we cannot easily see inside complex “black-box” AI systems, that current audit rules are not enough, and that responsibility for errors is all over the place, all combine to make a hole in who is held to account. AI could hugely improve how we assess risk and get evidence, but now using it in important financial reporting increases audit risk instead of making us surer that things are right. The AEGIS system, which focuses on checking things twice, making sure AI can explain itself, and always looking at what the AI is doing, is a needed change to auditing methods. This is to bridge the gap between how complicated algorithms are and what auditors need: proof that is enough and can be confirmed. In the end, just having AI in financial systems is not enough to be confident in them. A complete and coordinated overhaul of the whole audit system is needed. This means the IAASB, PCAOB, and similar organizations must quickly create AI auditing standards that can be enforced, a massive change to how auditors are trained to create professionals who are both auditors and data scientists, and a rethinking of the law about who is liable when many AI programmes are working together. Auditing is at a crucial point; if it is to remain important, it must actively update its methods, the skills auditors have, and the rules that govern it. Auditors can only properly do their job for the public – providing assurance that is meaningful, clear, and can be defended - in this increasingly computer-controlled economy if they fully adapt.

The limitations of this study are that, for a start, AEGIS is currently just an idea and has not been used in real businesses or by smaller auditing companies, which may not be able to afford the more advanced AI it requires. The research also concentrates on major international regulators like the IAASB and PCAOB, so it might overlook local regulations or the practical problems of getting all countries to agree to the same rules. And because AI evolves so rapidly, the legal and liability points are based on current situations. These legal guidelines could change quickly as newer, more complicated "multi-agent" AI systems become available.

Audit firms are best off beginning with a trial of the AEGIS Framework. They can use its ‘double-check’ and ‘logic-sharing’ features on financial information that does not pose much of a risk. This will help get everything working well in an actual setting before they completely implement the system. Organizations that make regulations should create secure “sandboxes” for auditors to experiment with sophisticated AI. Auditors can find out how to create proper, formal rules without being in danger of lawsuits during their testing. Colleges and accounting professional organizations must quickly update what is taught to include data science, the moral implications of technology, and managing AI. Future auditors will need this combination of abilities to remain current. Future studies should investigate how to use things like ChatGPT responsibly and in a way that respects privacy to examine complicated information - things such as legal agreements, or what was said in meetings - to find proof for our audits. The study must be done without inadvertently revealing confidential details of our clients. Future studies need to monitor the difference in technological ability within the auditing world. Specifically, they should look at how mid-sized firms are doing with the expense and complicated regulations of using AI for audits and compare that to the enormous amount of resources the “Big Four” (the biggest firms) have. Future studies should explore what laws are needed to figure out who is to blame when a fully independent AI agent, working without a human looking over its shoulder, makes a major error in a financial report.

**References**


Alles, M. G., & Gray, G. L. (2020). Will the AI auditor please stand up? A manifesto for the future of AI in accounting. Journal of Emerging Technologies in Accounting, 17(1), 13-20. https://doi.org/10.2308/jeta-10717

American Institute of Certified Public Accountants (AICPA). (2014). Statement on Auditing Standards No. 122, Clarification and Recodification. AICPA.

American Institute of Certified Public Accountants (AICPA). (2017). Audit analytics and continuous audit: Looking toward the future. AICPA.

Association of Certified Fraud Examiners (ACFE). (2024). Occupational fraud 2024: A report to the nations. Association of Certified Fraud Examiners. https://www.acfe.com/report-to-the-nations/2024

Bailey, J. (2017). The construction of auditing expertise: Measuring the meaning of experience. Accounting, Auditing & Accountability Journal, 30(5), 1165–1187. https://doi.org/10.1108/AAAJ-03-2016-2483

Brennan, N. M., & Subramanyam, J. (2023). Auditing artificial intelligence: The future of assurance. Routledge.

Brown-Liburd, H., Issa, H., & Lombardi, D. (2015). Behavioral implications of big data’s impact on audit judgment and decision making and future research directions. Accounting Horizons, 29(2), 451–468. https://doi.org/10.2308/acch-51023

Burrell, J. (2016). How the machine ‘thinks’: Understanding opacity in machine learning algorithms. Big Data & Society, 3(1). https://doi.org/10.1177/2053951715622512

Creswell, J. W., & Poth, C. N. (2018). Qualitative inquiry and research design: Choosing among five approaches (4th ed.). SAGE Publications.

Deloitte. (2021). The future of audit: How advanced technologies are reshaping the profession. Deloitte Insights.

Frey, C. B., & Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? Technological Forecasting and Social Change, 114, 254–280. https://doi.org/10.1016/j.techfore.2016.08.019

Institute of Chartered Accountants in England and Wales (ICAEW). (2019). AI and the future of accountancy: The impact of artificial intelligence on accounting and audit skills. ICAEW.
Institute of Chartered Accountants in England and Wales (ICAEW). (2017). Artificial intelligence and the future of accountancy. ICAEW.
International Auditing and Assurance Standards Board (IAASB). (2022). The IAASB's future work: Enhancing audit quality in a digital world. https://www.iaasb.org
International Auditing and Assurance Standards Board (IAASB). (2022). International Standard on Auditing (ISA) 500, Audit Evidence. https://www.iaasb.org
Issa, H., Sun, T., & Vasarhelyi, M. A. (2016). Research ideas for artificial intelligence in auditing: The formalization of audit and workforce supplementation. Journal of Emerging Technologies in Accounting, 13(2), 1–20. https://doi.org/10.2308/jeta-10511
Kokina, J., & Davenport, T. H. (2023). The emergence of explainable AI in audit. Journal of Information Systems, 37(2), 1–14. https://doi.org/10.2308/isys-2023-001
Klynveld Peat Marwick Goerdeler (KPMG). (2019). Audit 2025: The future of audit. KPMG International.
KPMG. (2020). Audit of the future: Reimagined by data and analytics. KPMG International.
Mikalef, P., & Gupta, M. (2021). Artificial intelligence capability: Conceptualization, measurement calibration, and empirical study on its impact on organizational creativity and firm performance. Information & Management, 58(3), 103434. https://doi.org/10.1016/j.im.2021.103434
Munoko, I., Brown-Liburd, H. L., & Vasarhelyi, M. (2020). The ethical implications of using artificial intelligence in auditing. Journal of Business Ethics, 167(2), 209–234. https://doi.org/10.1007/s10551-019-04407-1
Papernot, N., McDaniel, P., Sinha, A., & Wellman, M. P. (2018). SoK: Security and privacy in machine learning. 2018 IEEE European Symposium on Security and Privacy (EuroS&P), 399–414. https://doi.org/10.1109/EuroSP.2018.00035
Public Company Accounting Oversight Board (PCAOB). (2019). Spotlight: Data and technology research project. PCAOB.
Restatement (Second) of Torts § 552. (1965). American Law Institute.
Snyder, H. (2019). Literature review as a research methodology: An overview and guidelines. Journal of Business Research, 104, 333–339. https://doi.org/10.1016/j.jbusres.2019.07.039
Sun, T. (2019). Applying deep learning to audit procedures: An illustrative framework. Accounting Horizons, 33(3), 89–105. https://doi.org/10.2308/acch-52455
Vasarhelyi, M. A., Kogan, A., & Tuttle, B. M. (2015). Big data in accounting: An overview. Accounting Horizons, 29(2), 381–396. https://doi.org/10.2308/acch-51071
Zhang, Y., Xiong, Y., Lee, C., & Li, H. (2022). Artificial intelligence and audit risk: A review and research agenda. International Journal of Accounting Information Systems, 46, 100571. https://doi.org/10.1016/j.accinf.2022.100571