Volume **14** / **2026** ISSN 2344-102X
Issue **1** ISSN-L 2344-102X

FROM ALGORITHMIC OUTPUT TO AUDIT EVIDENCE: A CONDITIONAL FRAMEWORK FOR AI-ASSISTED AUDITING

Ionel BOSTAN
Ştefan cel Mare University of Suceava, 720229, Romania
ionel_bostan@yahoo.com
ORCID: 0000-0002-7653-7192

Abstract
The adoption of artificial intelligence (AI) in the audit has outpaced the empirical literature's ability to demonstrate where the benefits lie and where new risks are merely relocated. This work analyzes 15 primary articles in prominent accounting, auditing, finance, management, and information-systems journals. The empirical evidence is classified across the often blurred categories of AI adopted by audit firms, AI implemented by audit clients, and machine-learning applications to individual audit tasks. The results suggest that AI can lead to better risk assessment screening, fraudulent/erroneous reporting identification, analysis of internal controls, audit report timeliness, and certain aspects of audit quality. The results pertaining to audit fees and employment are less consistent as a result of technology investment potentially reducing the labor required for routine work while increasing audit scope, the use of specialists, and demand for tasks that involve higher cognitive demands. The study concludes that predictive ability only becomes economically significant when algorithms produce defensible audit evidence through the use of high-quality data, model validation, explanation capabilities, and auditor investigation, and that AI therefore must be understood as a conditional factor of the audit-production function and not as a substitute for the auditor's professional judgment.

Keywords: *artificial intelligence; auditing; machine learning; audit quality; audit fees; professional judgment*

JEL Classification: *C45; M42; O33.*

Received on: April 21, 2026 **Accepted on:** July 2, 2026 **Published on:** July 23, 2026

INTRODUCTION

The audit question around AI no longer is about technology that classifies documents, identifies anomalies, or produces risk scores - present-day systems are doing many of those functions - but about the extent to which technology use produces better audits, accounting for implementation costs, the quality of data, the necessity of professional judgment and algorithmic errors. This difference is significant, as audit quality can't simply be equated to speed of processing, data volume, or predictive power. Software may scour every journal entry, yet miss a deliberate misstatement if its model learned only on familiar data. A risk score may correctly assign most transactions to one group and inadvertently overlook a small group where all of the action is happening. The easiest assumption, that greater volumes of data result in more assurance, may be the most problematic assumption of them all, as coverage is a function of technique while confidence is a feature of evidence. A plethora of ongoing research includes empirical and applied work, though the results have been siloed with studies using different units of analysis: one stream measures AI investment at the firm level or local-office level and correlates it with financial restatements, fees, personnel, or audit opinion; another investigates client investment in AI as it relates to auditor work; and a third uses machine learning for fraud detection, misstatement, internal controls, text classifications or going-concern predictions. In all cases, the studies examine technology intervention at different points; document-extraction tools should not be compared to intelligent client production platforms and to systems for identifying fraud. The literature does not offer definitive conclusions despite an optimistic, techno-solutionist spin, with some work linking AI investment to fewer financial misstatements, lower fees, and declining accounting employment (Fedyk et al., 2022). Other work reports greater auditor hiring, particularly of junior and mid-career staff, stronger cognitive skills, more accurate going-concern opinions and internal controls (Law and Shen, 2024). More recently, evidence from China links AI implementation with reduced reporting lead times, fees, controls, and risk exposure - though with the nuances that adoption varies based on who implements the technology, auditors' backgrounds and client attributes (Liao et al., 2024; Rahman et al., 2024; Tan et al., 2025). These conflicting outcomes make for a peculiar, if practically relevant, academic puzzle: Why might AI lead to job reduction while creating new opportunities for auditor hires? Why would AI reduce audit fees in some instances but not in others? How can

DOI: 10.4316/EJAFB.2026.14113

software boost audit quality while simultaneously entailing new privacy and security risks, bias, and overreliance? Such contradictions cannot be explained away as "AI is good" versus "AI is bad" in accounting and auditing, but must be addressed by carefully examining the particular audit task that the technology changes, the organizational level at which it is implemented, the other resources it uses, and how it transmits technical output to evidential value. Although research so far has made valuable individual discoveries, none offers a simple framework that captures how technology, as measured in the applications, translates to evidence, quality and organizational performance. In this study, 15 empirical and applied articles are synthesized through an audit production framework to differentiate among AI adoption by auditors and clients, and among various tasks, identifying the conditions under which the use of technology produces different audit quality, timeliness, cost and staffing outcomes. It does not claim an exhaustive, holistic approach, nor does it aim to estimate an overall AI effect, but rather aims to provide a succinct understanding of why the observed findings are in agreement in some instances, in disagreement in others and support a conditional view of AI's role in the accounting world.

I. METHODOLOGY

The research follows a structured narrative evidence synthesis design and synthesizes evidence from 15 peer-reviewed publications issued between 2011 and 2026. The corpus was collected purposively around direct use and measurable consequences of AI, ML, text-analytics, XAI, and process mining in external auditing. All publication records and DOIs were cross-referenced against the respective publishers' platforms. Journal-based selection based on membership in relevant Web of Science categories and Q1 or Q2 positions was initially considered but quartiles did not replace judgement of each article's design and relevance at the selection stage. Included articles span original archival studies, machine-learning approaches, process-based models, and field studies. Excluded are literature reviews, bibliometric studies, purely conceptual articles and purely questionnaire surveys focusing primarily on perceptions or adoption intentions. The remaining articles had to focus on at least one identifiable audit outcome, namely: audit quality, audit risk, audit fees, reporting timeliness, employment and skills, fraud and misstatement identification, internal-control evaluation, restatement classification, explainability, or going-concern decision support. The search boundary was kept relatively narrow. The work does not include general literature on accounting automation, general AI governance, or auditing of AI systems unless the paper contributes directly to understanding AI-assisted auditing. The restriction on scope contributes to thematic cohesion at the cost of broader generalizability beyond external auditing, including for internal auditing, public-sector auditing, compliance auditing and performance auditing. Each selected article was coded along five dimensions: the level at which AI was observed, the type of technical or analytical capability concerned, the audit task or outcome under examination, the direction of the reported effect, and the major limitation impacting the interpretation of the findings. The 'level of AI' coding was distinguished as: adoption at the audit firm/office level, competence of the individual auditor, client adoption of AI systems, or task-specific use. This distinction is key because the term AI may refer to organisational investment, an individual's education, a client's operational technology, or a predictive model for a specific decision. Synthesis is qualitative and interpretive. Differences in samples, jurisdictions, adoption measures, dependent variables, and research designs preclude meta-analytic aggregation for a sample of this size and heterogeneity. The analysis therefore searches for mechanisms, tensions, and boundary conditions, rather than an average numerical effect. Table 1 reports the analytical profile of the 15 studies selected, summarizing the research context, method of study, outcome of interest and key limitation on interpretation.

Table 1. Profile of the Analysed Studies.

Study	Level and setting	Method or application	Audit-related outcome	Interpretive limitation
Perols (2011)	Task level; US financial statements	Comparison of statistical and machine-learning algorithms	Fraud classification	Performance depends on model choice and the costs assigned to classification errors.
Bao et al. (2020)	Task level; publicly traded US firms	Ensemble learning using raw accounting data	Accounting-fraud prediction	Historical fraud labels may not represent new concealment strategies.
Hayes and Boritz (2021)	Task level; restatement announcements	Machine learning and textual analytics	Classification of restatements	Textual signals can reflect disclosure strategy as well as underlying intent.

Volume 14 / 2026
Issue 1
ISSN 2344-102X
ISSN-L 2344-102X

Study	Level and setting	Method or application	Audit-related outcome	Interpretive limitation
Fedyk et al. (2022)	Audit-firm level; 36 large firms	Resume-based measure of AI employment and econometric analysis	Restatements, fees, employment	Centralised AI employment does not reveal the precise engagement-level application.
Zhang et al. (2022)	Task level	Application of LIME and SHAP	Explainability of misstatement-risk models	Explanations describe model behaviour but do not independently validate the model.
Liao et al. (2024)	Individual auditor; Chinese listed firms	Archival analysis of AI educational background	Audit-report timeliness	Educational background is an indirect proxy for actual technological use.
Law and Shen (2024)	Local audit offices; United States	Staggered AI hiring, archival tests, and partner interviews	Employment, skills, audit opinions	Local hiring captures deployment differently from firm-wide technological investment.
Rahman et al. (2024)	Auditor and client; China	Archival analysis of separate and joint AI adoption	Restatements and audit-report lag	Adoption proxies may not capture maturity, intensity, or system compatibility.
Duan et al. (2025)	Task and process level	Process mining combined with machine learning	Internal-control evaluation	Results depend on the completeness and integrity of event logs.
Kokina et al. (2025)	Large audit firms	Interviews with experienced professionals	Adoption maturity, opportunities, and risks	Field evidence offers depth but does not estimate causal effects.
Lai (2025)	Client level; Chinese listed firms	Archival panel analysis	Audit fees	Lower fees cannot by themselves establish improved audit quality.
Lee and Tahmoush (2025)	Task level; US archival data	Decision-tree model	Going-concern decision support	Cut-off points and interactions may not transfer across markets or periods.
Tan et al. (2025)	Client level; Chinese listed firms	Archival panel analysis	Audit quality, fees, and report lag	Client AI may proxy for broader managerial and governance capability.
Zhang Parker et al. (2025)	Task level	Machine-learning prediction models	Material-misstatement risk	Predictive gain must still be translated into assertion-level audit procedures.
Li (2026)	Client level; Chinese listed firms	Archival analysis	Corporate audit risk	The institutional concentration of the evidence constrains generalisation.

Source: Created by the author based on the analyzed studies.

It should be clear from the above that I deliberately created such a disparate corpus – since my question is how technological capability relates to the audit results, not the success of a particular algorithm. Such disparity does, however, necessitate an appropriate level of caution – results from the fraud-classification model cannot simply be extrapolated to the labour market for audit firms, and the evidence from Chinese listed companies must not be presumed to apply to every regulatory environment and audit market in existence.

II. A CONDITIONAL FRAMEWORK FOR AI-ASSISTED AUDITING

AI comes into the audit at the stages of data retrieval, classification, anomaly identification, pattern discovery, forecasting, and judgment enhancement. None of these technical functions alone creates assurance. Assurance only comes to the fore once the AI's technical output is applied to a client assertion; supported by investigation; assessed for materiality; and brought to life in the context of a well-documented professional judgment. Figure 1 below represents the gated path whereby client information, using AI assistance and human judgment, gives rise to justifiable assurance results.

Volume **14** / **2026**
Issue **1**

ISSN 2344-102X
ISSN-L 2344-102X

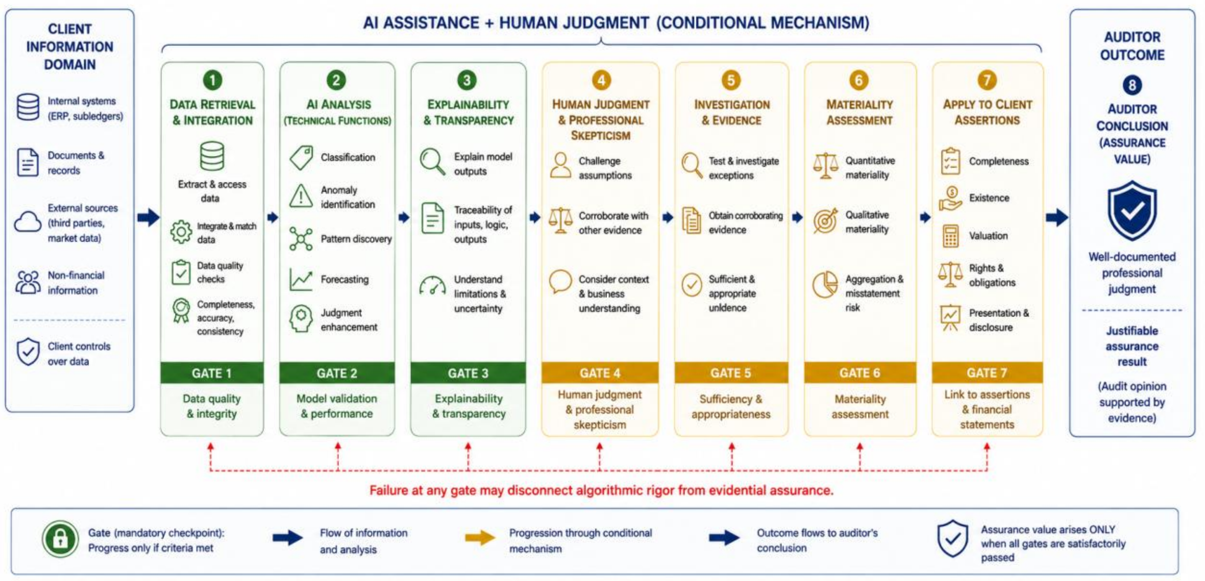

Figure 1. From data abundance to assurance value: a conditional mechanism. Source: Author's synthesis based on the analysed studies.

Figure 1 provides the analytical layer between the client's information domain and the auditor's conclusion that AI represents. It refutes any claim that the algorithm output is delivered directly to the audit opinion. There is an assumption, embedded in the framework, that data quality/integrity, model validation, explainability, and human judgment are between those two poles. A failure anywhere along this continuum may disconnect algorithmic rigor from evidential assurance. The framework helps to explain why studies of other forms of AI adoption can lead to varying results. Client AI could potentially increase quality/availability of accounting information (reducing information asymmetry and information costs of auditing). Firm AI might automate document retrieval or risk assessment; its economic impact depends upon whether saved costs result in reduced staff, more comprehensive sampling, deeper inspection of exceptions or new services. The performance gains might be detected without a change in the quality of audit services, when the task-level measures are not incorporated into the engagement method. The primary empirical test treats adoption as a homogeneous treatment. This is not true. A company that hires a team of data scientists, an auditor who got a CS degree and a company that mentions AI in annual reports are all very different things with very different causal effects on information production that supports assurance. It's not just a language issue; it defines what the research question really is.

III. WHERE AI PRODUCES AUDIT VALUE

3.1. Fraud and Material-Misstatement Risk

Fraud is one of the most enduring applications, owing to the classification of rare events, asymmetric error costs and massive amounts of financial data. Perols (2011) demonstrates that apparent model choice depends partly on how costs of false positives versus false negatives are weighed. No auditor is going to choose a model because it maximizes overall accuracy when an undiscovered fraudulent observation may cost far more than a wasteful inspection. Bao et al. (2020) apply ensemble learning to a set of raw accounting numbers, proving that machine learning can detect nonlinear combinations beyond what a conventional ratio-based specification captures effectively. Their important claim is not that the algorithm 'detects' fraud in the sense of conclusive evidence. The model can rank firms, or a specific portion of financial statements, based on patterns associated with firms involved in past fraud. This will serve to improve the selection of high-risk cases for audit. Treating the model as definitive proof of fraud, rather than an indicator that signals additional inquiry, would overreach what can be said from an algorithm built on past history. Zhang Parker et al. (2025) take predictive modeling one step further, and use it to identify the risk of a material misstatement. The application to practice would be during the planning stage of an audit, where more accurate ordering of firms, or reporting segments that pose a risk, would naturally inform the allocation of senior-accountant resources, as well as the amount and depth of audit testing. The predictive models will not automatically be smuggled into the function of audit evidence, however, because a calculated probability of a material misstatement provides no insight into the specific assertion and account(s) involved nor does it help to understand the economics of the possible error. Those will still require audit procedures grounded in the related

DOI: 10.4316/EJAFB.2026.14113

Volume **14 / 2026** ISSN 2344-102X
Issue **1** ISSN-L 2344-102X

economic activity and accounting and reporting standards. It has a weakness common to all such prediction algorithms in that they rely on historical patterns to predict the future and can be susceptible to manipulation as fraud evolves to take advantage of discovered patterns in control procedures. When this occurs, a model trained to discover patterns of a past generation of manipulation is unable to cope with that of a new one in a different industry, a different industry standard, or a different reporting framework. It is in this environment that human skepticism will be most needed, where there is no comfortable history available. Hayes and Boritz (2021) note that the wording of restatement announcements, analyzed using textual analysis, might shed light on managerial motivation. This will broaden the input information for an auditor beyond the financial statement numbers, while keeping in mind that the narrative part of the announcement is also controlled strategically. Any message about risk found by textual analysis should, therefore, prompt a deeper, closer look and corroboration rather than serving as a replacement for these fundamental auditing activities.

3.2. Internal-Control and Process Evaluation

Process mining takes a different path to value. It allows us to use digital event logs to remap actual workflows. Duan et al. (2025), using a combination of process mining and machine learning, apply this technique to the detection of control deviations and abnormal processes. The benefit for high-volume processes like procure-to-pay, sales, payroll and inventory management is identifying instances where an approver has repeatedly bypassed a control, where a set of users with no business need is jointly approving transactions, where a control is delayed in its execution, or where the transaction follows a path impossible to discover using random sampling. The financial value lies in allocating the auditors' time to process varieties that present a higher probability of control failure. Nevertheless, the seeming completeness achieved through analysing an event log may be deceptive because of missing manual entries, shared credentials and transactions outside the main system. A mapped process that has been digitally recreated may not be "the whole process". Prior to using this "map", the auditor must grasp the process of creating the map – an issue seldom stressed in technology-driven materials. Process mining can achieve coverage of the observed sample while ignoring the unobserved behaviour, and thereby offers an illusory view of an incomplete system.

3.3. Going-Concern Assessment and Explainability

These going-concern decisions exemplify a dilemma involving the trade-off between computational sophistication and auditor responsibility. Lee and Tahmoush (2025) utilize a decision tree approach to capture nonlinear interactions among financial and market variables for which going-concern opinions are issued. This model is valuable due to its white-box structure; users can readily observe how inputs and parameters combine in ways that are less transparent in models whose interior processes remain “black boxes”. It can, for example, help the auditor focus on areas requiring further attention but cannot assess the myriad of data that may impact the ultimate opinion, such as refinancing efforts, plans for the business, post-year-end events, waivers of loan covenants, government support, or the validity of cash-flow projections, all of which may not be present in databases. In effect, the tool is a supplement to, not a substitute for, the auditor’s conclusion. Zhang et al. (2022) explain how explainable artificial-intelligence approaches can aid in deciphering the impact of each variable in the calculation of the result, an important feature for documentation purposes because the audit team needs to be able to explain why an observation was classified as high risk. While useful, too much emphasis is often placed on explainability; an explanation doesn’t prove that the relationships on which it is based are economically significant, persistent, or unbiased. Interpretability and validity are indeed associated controls, but they are not the same things. Figure 2 provides an illustrative conceptual positioning of selected audit tasks according to their relative automation potential and dependence on professional judgment. The placement is intended to support analytical comparison and does not represent an empirically calibrated measurement.

DOI: 10.4316/EJAFB.2026.14113

Volume **14** / **2026**
Issue **1**
ISSN 2344-102X
ISSN-L 2344-102X

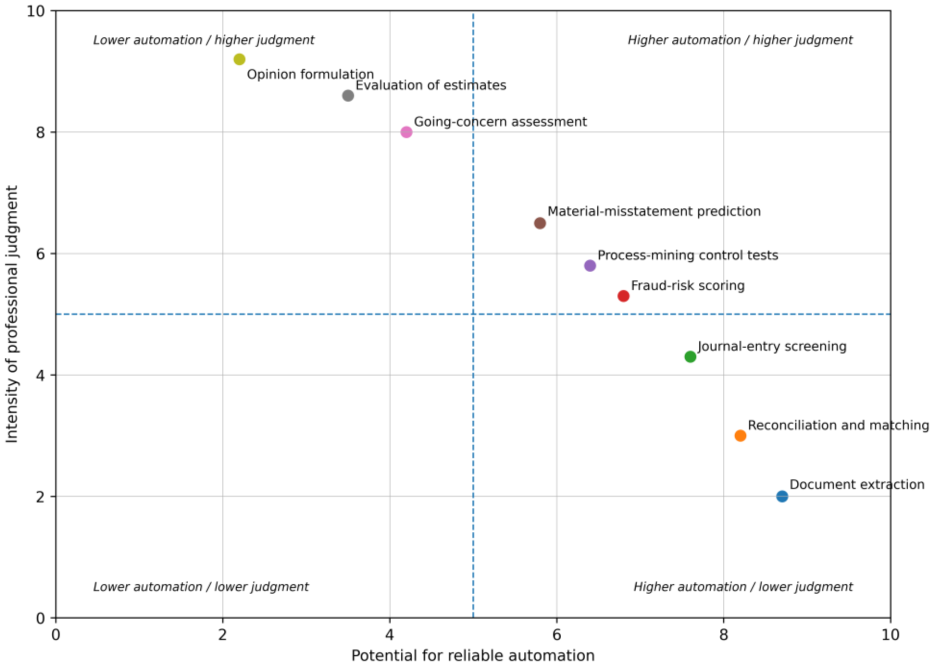

Figure 2. A conceptual positioning of audit tasks by relative automation potential and professional-judgment dependence. Source: Author's synthesis based on the analysed studies.
Note: Task positions are illustrative analytical assessments derived from the reviewed literature. The numerical axes facilitate comparison and should not be interpreted as empirically estimated, statistically validated, or expert-consensus scores.

Figure 2 conceptually differentiates the selected tasks according to their relative potential for automation and their dependence on professional judgment. The matrix should be read as an interpretive synthesis of the reviewed evidence, rather than as a validated quantitative classification. The area with higher automation potential but lower reliance on judgment includes data capture and matching, as well as routine reconciliation. Areas with significant analytical capacity but higher judgment input include fraud risk assessment and control testing. Those areas with higher levels of dependence on judgment, regardless of whether AI can bring insight, include going-concern assessment, accounting estimates and opinion-formation. Note that figure 2 does not indicate a static technological threshold – this threshold will likely evolve, but it should move when a new capability produces more reliable evidence, not simply because it is newer.

IV. ECONOMIC AND ORGANISATIONAL EFFECTS

4.1. Audit Quality

Strong evidence that AI investment results in better observable audit outcomes is provided by Fedyk et al. (2022), whose study demonstrates that a larger fraction of AI workers is associated with fewer restatements and identifies improvement as an explicit goal stated by audit partners. Local deployment of AI in audit offices has also been associated with more accurate going-concern and internal-control opinions (Law and Shen, 2024). These results are significant, yet the outcome variables measured – restatements and reported errors – capture only a limited part of the picture of audit quality. Absence of restatements or reporting errors does not mean that no material risk remained undisclosed, just that certain detectable anomalies or standard errors are fewer. AI may enhance detection of familiar errors while being less effective in estimating forecasts, dealing with complex transactions, assessing managerial integrity or anticipating new risks. There is another way to document an impact through the client side. Firms that adopted AI have higher audit quality through better internal controls and disclosure transparency, as presented by Tan et al. (2025). Li (2026) documents reduced corporate audit risk associated with firms' AI adoption. This association makes sense when AI in client firms better facilitates integration of information, monitoring, and accessibility. It is less obvious when firms utilize opaque models, automated forecasting or a multitude of complex data dependencies that are not accessible to auditors. Selection could also be the reason for the observed link between AI on the client-side and better outcomes, because firms ready to adopt cutting-edge technology may be better managed, governed and financially supported irrespective of the technological adoption. Existing empirical evidence includes control variables, but this does not fully alleviate the concern that the benefits derived should not necessarily be credited solely to AI rather than the inherent

DOI: 10.4316/EJAFB.2026.14113

Volume **14** / **2026** ISSN 2344-102X
Issue **1** ISSN-L 2344-102X

capabilities of firms choosing to adopt AI technology.

4.2. Timeliness and Efficiency

Auditors with an education in AI skills are associated with shorter audit-report lags, especially for more complex organisations and for engagements where the perceived risk of misstatement is high (Liao et al., 2024). This shows that technological skill has value beyond software ownership. Auditors with skills in programming logic, data structures, and analysis techniques might better diagnose data problems and communicate with other experts. Rahman et al. (2024) explain why efficiency may not appear instantaneously and symmetrically: individual adopters—auditors or clients—may face compatibility issues, training costs, and further validation, whereas common technological capability helps to increase timeliness and reduce restatements. This result challenges a general proposition within the digital-transformation literature: that buying or announcing technology is equivalent to deploying that technology as part of an inter-organisational workflow. Kokina et al. (2025) draw a stark contrast between common applications of AI such as document extraction and OCR, and more complex applications that are still under construction. The real-world data here keep the debate from devolving into a fantastical scenario of an entirely self-directed audit—an idea for which the profession is not currently ready. The more direct impacts on productivity seem to be confined to information gathering and processing of repeated tasks, while the sophisticated usage of AI in judgment-enhancing roles still suffers from concerns about reliability, privacy, transparency and compliance.

4.3. Audit Fees

Evidence consistent with lower audit fees is reported in Fedyk et al. (2022), Lai (2025) and Tan et al. (2025). Proposed explanations involve lower information asymmetry, higher-quality accounting information, increased internal control effectiveness, and automation of repetitive tasks. All these would move the audit-production frontier outward as it takes less time to extract, reconcile and filter. A lower fee per audit is not, of itself, a proxy for a better audit. It could be a result of a more productive audit process, but could also result from competition, narrowed audit scope, increased use of client systems or inadequate pricing for new model risk. The quality impact turns on what the audit firm does with resources freed by automation. Law and Shen (2024) view technology less as simple labour displacement. Their research suggests audit firms may add to their teams, complementing the technology with additional auditors with advanced cognitive abilities. This type of technological adoption expands either the volume of tasks carried out or their complexity rather than their cost. Audit firms could also appropriate part of their productivity gains for investment in platforms, expertise, security systems, model risk governance and staff training. This allows the fee-based evidence to align with a range of possible production choices. An audit firm might use technology to do the same audit for a reduced hour input, to increase the scope of work undertaken without an increase in hour input, or to offer new or more complex audit services. Without data at the engagement level, it is difficult to see what is done in relation to this, based solely on changes in fees.

4.4. Employment and Skills

The labor results only appear contradictory when viewed as a simple, static concept of employment. Fedyk et al. (2022) found a slow, gradual decline in accounting employment after investment in AI with the lag effect kicking in several years down the road. In contrast, Law and Shen (2024) observed a 4.3% rise in auditor roles in the local offices of firms with AI employees—these were primarily concentrated in entry- and mid-level positions—and an increased demand for cognitive and interpersonal skills. Several explanations can reconcile the diverging results. Fedyk et al. (2022) focus on AI employment at large firms and a long-term displacement of labor, whereas Law and Shen (2024) assess phased local office implementation and the total number of audit roles. Increased output from centralized automation might be associated with reduced local office labor inputs as these offices acquire clients, increase audit testing, or employ additional auditors to interface with the automated systems. Also, slower job growth differs from an actual reduction in numbers. More likely, we observe a recomposition of occupational structure: rote information retrieval, information matching, and standardized documentation will increasingly be automated; exception-based reasoning, modeling skills, client interaction, complex estimation challenges, and ultimate accountability for the audit opinion will continue to require a human touch. What, therefore, appears most at stake for the profession is less outright job destruction and more the diminishing importance of the junior audit experience. Firms that automate at the entry level without re-envisioning the pipeline may make it cheaper to complete certain task outputs in the short term, but at the cost of eroding the supply of future professional judgment. Table 2 summarizes the findings of our review regarding those areas on which evidence is coalescing, those that are in conflict, and when the benefits from the use of AI in audit are likely to be most potent.

DOI: 10.4316/EJAFB.2026.14113

Volume 14 / 2026 ISSN 2344-102X
Issue 1 ISSN-L 2344-102X

Table 2. Areas of Agreement, Contradiction, and Conditional Interpretation.

Outcome	Evidence tending toward improvement	Evidence or issue preventing a simple conclusion	Conditional interpretation
Audit quality	Fewer restatements and more accurate opinions are associated with audit-firm or office AI adoption.	Quality proxies capture only selected failures and may omit complex judgment areas.	AI is most likely to improve quality when outputs are linked to assertion-level procedures and reviewed by competent auditors.
Reporting timeliness	AI-educated auditors and coordinated auditor-client adoption are associated with shorter lags.	New systems can initially increase validation, learning, and compatibility work.	Efficiency depends on complementary skills and integration, rather than adoption in isolation.
Audit fees	Several client- and firm-level studies report lower fees.	Lower fees may reflect productivity, competition, scope reduction, or unpriced risk.	Fee reduction is informative only when assessed together with scope and quality indicators.
Employment	Firm-level evidence indicates delayed displacement of some accounting labour.	Office-level evidence shows more junior and mid-level auditor jobs and stronger skill demand.	AI automates tasks while potentially expanding services and changing occupational composition.
Adoption maturity	Basic extraction and automation tools are already used in practice.	Complex, judgment-oriented AI remains constrained by transparency and reliability concerns.	Near-term value lies mainly in augmentation and screening, not autonomous opinion formation.
Client AI	Improved controls and transparency may reduce audit risk and cost.	Technological complexity can create new model, access, and data-governance risks.	Client AI is beneficial where governance capability develops alongside operational technology.

Source: Created by the author based on the analyzed studies.

V. GOVERNANCE REQUIREMENTS AND RESIDUAL RISK

Our evidence base thus gives credence to the use of AI within audit methodology but also makes it apparent that capability must be subject to the same rigorous discipline as any other form of audit evidence. Key safeguards focus on data, models, people and accountability. Data governance must address provenance, completeness, accuracy, authorized access and the trustworthiness of transformations carried out prior to analysis. Models should be tested against the population of interest, monitored for drift, have their assumptions documented, and include explicit analysis of the impact of false positives and negatives. Explainability needs to be coupled with substantive review, not simply a veneer applied over an unquestioned black box. Review by human beings must be framed to resist automation bias; for an auditor to merely have to click a button to 'accept' a model output is not effective oversight; the reviewer needs to understand the goal of the model, what it hasn't considered, where it could fail and what steps should be taken to verify a model alert or challenge a reassuring false-negative result. Accountability is indivisible. A partner cannot delegate ultimate responsibility for their opinion to a software vendor, data scientist, client system or to a generative model; where specialists are brought to assist the audit, their output should be incorporated into the evidence trail and related to the financial statement assertions in question. Table 3 provides an illustration of key uses of AI within audit along with their residual risks and the required audit procedures to maintain evidential reliability and accountability.

Table 3. AI Applications, Residual Risks, and Required Audit Responses.

Application	Output supplied to the audit team	Residual risk	Required audit response
Document extraction and classification	Structured information obtained from contracts, invoices, or reports	Extraction error, omitted context, incorrect document version	Test accuracy on representative items and retain access to original documents.
Journal-entry and transaction screening	Ranked anomalies or unusual clusters	Excessive false positives, missed novel schemes, inappropriate thresholds	Calibrate error costs and investigate alerts through independent evidence.

DOI: 10.4316/EJAFB.2026.14113

Volume 14 / 2026 ISSN 2344-102X
Issue 1 ISSN-L 2344-102X

Application	Output supplied to the audit team	Residual risk	Required audit response
Fraud and misstatement prediction	Entity, account, or period risk score	Confusion of prediction with proof; unstable historical relationships	Translate risk signals into assertion-specific procedures and reconsider contradictory evidence.
Process mining	Visualisation of actual system-recorded workflows	Missing off-system activity, shared credentials, incomplete logs	Test log integrity and reconcile digital events with organisational practice.
Going-concern decision support	Interacting risk indicators and probability classifications	Omission of forward-looking, qualitative, or post-reporting information	Evaluate forecasts, financing, subsequent events, and management plans independently.
Explainable AI	Variable-level reasons for a model output	Plausible explanation of an invalid or biased model	Separate interpretability testing from model validation and data-quality assessment.
Client-side AI and automated controls	Faster data access, monitoring, and exception reports	Increased system complexity, opaque estimates, unauthorised model changes	Understand governance, access controls, change management, and model ownership.
Generative assistance in documentation	Drafted summaries, working-paper text, or proposed procedures	Hallucination, confidentiality breach, loss of professional authorship	Restrict approved use, verify every substantive statement, and document human responsibility.

Source: Created by the author based on the analyzed studies.

The table also highlights a significant asymmetry: the marginal cost of generating thousands of AI-generated alerts is near zero, yet it's expensive to analyze each one. A poorly calibrated model can thus shift rather than solve a bottleneck, and an AI-assisted audit will only be as effective as the capacity an engagement team has to deal with the exception reports that arise from using the AI.

VI. DISCUSSION

These findings align with a vision of AI as an conditional input in the audit-production function. Technology is introduced through a variety of channels affecting the quality of client information, routine task automation, risk prioritization, auditor allocation, and detection of interactions that traditional models miss. These channels have varying costs and risks. Such a framework reconciles conflicting results. Firms investing in AI might achieve fewer hours spent on tasks, but greater total employment if the audit is extended to other areas of client operations, or a greater audit scope for a larger client. Client technology may reduce information asymmetry but may increase system complexity requiring audits. Hours may drop, remain flat (if the firm retains all gains), or rise (if more use of specialists and system controls become necessary). These studies also highlight that the most successful applications are those that assist, not substitute for, judgmental work. Document extraction, reconciliation, process analysis, and anomaly identification appear natural candidates for AI given that their results are subject to validation from source material. Opinion formulation and the assessment of volatile estimates are more challenging since the relevant evidence is typically context dependent, predictive, and contested. A related conclusion has implications for professional competency. Part of AI's value derives from auditor skepticism. Understanding data lineage, classification errors, model drift, proxies, validation results and explanations should be part of technological competency, not just how to operate a dashboard. Failure to do so might turn a theoretically advanced audit into one that is less, not more skeptical, as management's unsupported assertions might be rejected whereas an automated system is given due respect. More work remains to be done. Almost all archival research relies on measures of technology adoption instead of firm-level data about technology deployment within individual engagement teams. The geographic coverage of the literature is largely restricted to U.S. and Chinese evidence. Public-sector auditing, internal auditing, compliance engagements, and performance auditing are generally neglected even though large administrative data files and public accountability standards apply. The next stage of the research needs to match individual AI technologies with specific audit procedures, hours, and results. Furthermore, a clearer understanding of what happens during failures, overconfidence, and review is necessary, along with cases where algorithm alerts are triggered, either appropriately or inappropriately.

DOI: 10.4316/EJAFB.2026.14113

Volume **14** / **2026** ISSN 2344-102X
Issue **1** ISSN-L 2344-102X

VII. CONCLUSION

This work explored how AI is impacting audits at four levels – those of the firm, the client, the auditor and the analytical procedure. Evidence is most compelling in document management, anomaly identification, screening for fraud and misstatement, process analysis and the targeting of audit efforts. These applications offer opportunities to improve certain quality indicators and reduce effort associated with repetitive tasks. The economic impact is more nuanced. Fees might fall in line with enhanced efficiency, but they are inextricably linked with quality in the absence of more detailed audit evidence. Employment may decrease in roles associated with basic tasks while it increases in firms using technology to service more clients or perform richer analyses. The use of AI by the client may lead to improved quality of information and internal control, but also new risks surrounding model integrity, data governance and complexity. A central tension in the paper, between algorithmic output and audit evidence, holds for the industry. An output is meaningful only once the auditor verifies the input data quality, appreciates the models used, understands their purposes and limitations, inquires about the output, and then relates it to relevant assertions and reporting judgments. Explainability is an aid to these actions but not a substitute; automation can expand the scope of audits but it cannot transfer the auditor's professional responsibility. In terms of practice, an approach based on selective augmentation is the most defensible. Firms should automate reliable procedures for verifying inputs and outputs; ensure human oversight in subjective matters affected by uncertainty and intention; and adjust training so that a pursuit of efficiency does not compromise the formation of professional judgment. Regulators should prioritize evaluating whether an AI tool is governed and auditable as part of a traceable chain of evidence, rather than whether a firm simply "uses" it.

To be of value, the artificial intelligence employed in auditing must not silence professional skepticism but rather render it better informed, more precisely focused, and ultimately, more rigorously accountable.

Funding: This research received no external funding.
Conflicts of interest: The author declares no conflicts of interest.
Ethics statement: This study is based exclusively on publicly available legal, institutional, policy and secondary academic sources. It does not involve human participants, interviews, surveys, experiments, personal data collection or clinical/biomedical material. Therefore, formal ethical approval was not required.

REFERENCES

1. Bao, Y., Ke, B., Li, B., Yu, Y. J., & Zhang, J. (2020). Detecting accounting fraud in publicly traded U.S. firms using a machine learning approach. Journal of Accounting Research, 58(1), 199-235. https://doi.org/10.1111/1475-679X.12292.
2. Duan, H. K., Vasarhelyi, M. A., & Codesso, M. (2025). Integrating process mining and machine learning for advanced internal control evaluation in auditing. Journal of Information Systems, 39(1), 55-75. https://doi.org/10.2308/ISYS-2022-028.
3. Fedyk, A., Hodson, J., Khimich, N., & Fedyk, T. (2022). Is artificial intelligence improving the audit process? Review of Accounting Studies, 27(3), 938-985. https://doi.org/10.1007/s11142-022-09697-x.
4. Hayes, L., & Boritz, J. E. (2021). Classifying restatements: An application of machine learning and textual analytics. Journal of Information Systems, 35(3), 107-131. https://doi.org/10.2308/ISYS-19-003.
5. 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, Article 100734. https://doi.org/10.1016/j.accinf.2025.100734.
6. Lai, J. (2025). Artificial intelligence applications and audit fees: An empirical study. International Review of Economics & Finance, 103, Article 104421. https://doi.org/10.1016/j.iref.2025.104421.
7. Law, K. K. F., & Shen, M. (2024). How does artificial intelligence shape audit firms? Management Science, 71(5), 3641-3666. https://doi.org/10.1287/mnsc.2022.04040.
8. Lee, E. J., & Tahmoush, D. (2025). Auditors' decision-making aid for going concern audit opinions through machine learning analysis. International Journal of Accounting Information Systems, 56, Article 100732. https://doi.org/10.1016/j.accinf.2025.100732.
9. Li, Y. (2026). Can artificial intelligence reduce corporate audit risk? Evidence from Chinese listed companies. International Review of Economics & Finance, 108, Article 105209. https://doi.org/10.1016/j.iref.2026.105209.
10. Liao, F.-N., Zhang, C., Zhang, J.-J., Yan, X., & Chen, T.-X. (2024). Hyperbole or reality? The effect of auditors' AI education on audit report timeliness. International Review of Financial Analysis, 91, Article 103050. https://doi.org/10.1016/j.irfa.2023.103050.
11. Perols, J. (2011). Financial statement fraud detection: An analysis of statistical and machine learning algorithms. Auditing: A Journal of Practice & Theory, 30(2), 19-50. https://doi.org/10.2308/AJPT-50009.
12. Rahman, M. J., Zhu, H., & Yue, L. (2024). Does the adoption of artificial intelligence by audit firms and their clients affect audit quality and efficiency? Evidence from China. Managerial Auditing Journal, 39(6), 668-699. https://doi.org/10.1108/MAJ-03-2023-3846.
13. Tan, J., Chang, S., Zheng, Y., & Chan, K. C. (2025). Effects of artificial intelligence in the modern business: Client artificial intelligence application and audit quality. International Review of Financial Analysis, 104, Article 104271. https://doi.org/10.1016/j.irfa.2025.104271.
14. Zhang, C. A., Cho, S., & Vasarhelyi, M. A. (2022). Explainable artificial intelligence (XAI) in auditing. International Journal of Accounting Information Systems, 46, Article 100572. https://doi.org/10.1016/j.accinf.2022.100572.
15. Zhang Parker, C., Jiang, L., Cho, S., & Vasarhelyi, M. A. (2025). Predicting material misstatements using machine learning. The Accounting Review, 100(6), 225-262. https://doi.org/10.2308/TAR-2024-0035.