Transforming Auditing with Artificial Intelligence: A Framework for Fraud Detection and Responsible Adoption

Hong Yang

Faculty of Business and Law, Curtin University, Perth, Australia
Email: yhzzywzy@gmail.com
Manuscript received January 21, 2026; accepted March 31, 2026; published April 15, 2026.

***Abstract*—This study investigates how Artificial Intelligence (AI) can transform auditing by improving fraud detection and audit quality. However, traditional methods, limited by sampling and manual inspection, failed to detect complex frauds such as Luckin Coffee. Using the Technology–Organization–Environment (TOE) framework and Socio-Technical Systems (STS) theory, the research adopted a qualitative case study design. Audit failures were reconstructed through triangulated evidence and counterfactual reasoning, assessing how Machine Learning (ML) and Natural Language Processing (NLP) could have detected anomalies in real time. Findings show AI provides broader coverage, timeliness, and granularity while preserving professional skepticism. The study contributes by situating AI adoption within socio-technical and organizational contexts and proposing a framework for responsible implementation, and it is crucial to position AI as a tool that complements, rather than substitutes, human auditors.**

***Keywords*—auditing, artificial intelligence, data analytics, machine learning, fraud detection, natural language processing**

I. Introduction

Auditing has long served as a foundation of financial accountability, as it upholds integrity in corporate reporting and maintains investor confidence.

The profession has in recent decades shifted to Computer Assisted Audit Techniques (CAATs) and data analytics, as opposed to the manual and sample-based testing (Appelbaum, Kogan, & Vasarhelyi, 2017). Despite these advancements, conventional auditing methods are still limited by the sampling, retrospective testing, and excessive dependence on auditor judgment. Consequently, these limitations contribute to major corporate scandals, such as the fraud at Luckin Coffee, which exposes structural vulnerabilities in audit strategies and eroded public trust (Carcello, 2020). Machine Learning (ML) and Natural Language Processing (NLP) are some methods that allow processing vast structured and unstructured data and have the potential to identify anomalies, improve risks evaluation, and facilitate continuous assurance (Kokina & Davenport, 2017). Professional services firms are testing various Artificial Intelligence (AI) tools. Examples include PwC's 'Halo' for transaction testing using NLP (PWC, 2018), EY's NLP-based contract review (EY, 2019), and Anti-Money Laundering (AML) systems in banks (Deloitte, 2020). These initiatives are helping to explore the opportunities and challenges of applying AI in assurance practices.

Nonetheless, there is more than technical deployment involved in the adoption of AI in auditing, which involves a need to align with organizational preparedness, regulatory expectations, ethical protections, and the core focus of the profession on judgment and independence (Li & Vasrhelyi, 2023).

There are two interrelated research questions addressed in this paper:

RQ1: How can AI and ML technologies improve the shortcomings of conventional auditing practices that led to high-profile audit failures?

RQ2: What is the likely trajectory of AI-enhanced auditing, and how can the profession prepare for its responsible adoption?

To address these questions, the study has three primary objectives:

- To assess the potential of AI in improving detection of fraud and evidential coverage in both structured and unstructured (in the form of contracts, disclosures, communications) sources (Sun & Vasarhelyi, 2018).
- To identify impediments and threats to AI adoption, which span issues such as algorithmic complexity, human capital shortages, and governance constraints (Kokina, Ma, & Davenport,2021).
- To suggest a cohesive system of the responsible AI adoption in auditing, to connect the technological innovation and the professional and regulatory aspects (Li & Vasarhelyi, 2023).

The remainder of the paper is structured as follows. Section II discusses the historical development of the audit technologies and the theoretical basis, relying on Technology-Organization-Environment (TOE) framework and Socio-Technical Systems (STS) theory. Section III discusses the research design, i.e. the design of a case study, case selection, and the analytical strategy. Section IV shows the result of case study, namely Luckin Coffee, along with counterfactual AI-enhanced auditing simulations. It also covers the challenges of technical, organizational, and ethical aspects and presents a comprehensive framework of responsible adoption of AI. Section V presents the conclusion summarizing the findings, providing theoretical and practical contributions, and future research directions.

II. Literature Review and Theoretical Foundations

A. *Evolution of Audit Technology: CAATs to AI*

The evolution of audit technology reflects the profession's ongoing effort to address increasingly complex business environments. Initially, CAATs automated routine procedures and enabled larger sample sizes, but did not fundamentally change audit methodologies (Minkkinen, Laine, & Mäntymäki, 2022). Subsequently, the adoption of Enterprise Resource Planning (ERP) systems facilitated continuous monitoring and real-time testing (de Oliveira, da Silva, & Neto, 2025). More recently, data analytics and visualization tools enhanced anomaly detection capabilities, though they still relied on sampling and manual interpretation (Hasan, 2022).

The current technological frontier involves AI, particularly ML and NLP, which represent a fundamental shift from earlier tools. ML algorithms can identify complex patterns across entire datasets rather than relying solely on samples (Bakumenko & Elragal, 2022). NLP enables systematic analysis of unstructured data such as contracts, emails, and management disclosures (KPMG, 2025). Thus, this evolution marks a critical transition from retrospective, sample-based verification toward proactive, comprehensive assurance (Leocádio, Malheiro, & Reis, 2024).

B. Practical Applications of AI in Modern Auditing

Research demonstrates significant potential for AI in various auditing activities. In fraud detection, machine learning models generally outperform traditional statistical methods in identifying subtle anomalies within revenue patterns or vendor relationships (Gandhar, Gupta, Pandey, & Raj, 2024). Natural language processing tools are equally effective for analyzing unstructured text, such as detecting linguistic indicators of financial misrepresentation or obfuscation in management reports or regulatory filings (Wang, Zhang, & Li, 2023).

For contract review, major audit firms now employ AI to support activities like lease accounting analysis under IFRS 16 and transaction verification (Hamdan & Al Habashneh, 2024). These applications highlight AI's role in expanding evidence coverage, improving audit timeliness, and strengthening detection capabilities. By integrating ML and NLP, auditors can examine complete data populations and systematically incorporate qualitative information into their judgments (Leocádio *et al.,* 2024).

However, most existing research focuses on evaluating individual technologies' effectiveness rather than exploring how AI is integrated into end-to-end audit processes (Al-Qudah & Al-Hattami, 2023). This indicates a persistent lack of systematic understanding regarding how AI interacts with professional judgment and supports audit decision-making in real-world engagement contexts.

C. Overcoming Adoption Barriers: An Integrated Framework

A Successful AI adoption in auditing involves more than technological implementation. The TOE framework indicates that adoption decisions are shaped by technological capabilities, organizational readiness, and external pressures such as regulatory requirements and client expectations (Hamdan & Al Habashneh, 2024). For instance, even when AI tools are available, firms may lack the data infrastructure or specialized skills to utilize them effectively (de Oliveira *et al.,* 2025). Socio-Technical Systems (STS) theory further emphasizes that AI must align with human expertise, professional skepticism, and ethical standards (Leocádio *et al.,* 2024). AI should augment rather than replace auditor judgment (Hasan, 2022). Additionally, studies identify ethical concerns regarding data privacy, algorithm transparency, and accountability that may hinder widespread AI implementation. Collectively, existing literature reveals a significant gap between technological potential and practical application. Few studies provide comprehensive frameworks for integrating AI while maintaining compliance with established auditing standards and preserving professional responsibility and audit quality (Yang, Blount, & Amrollahi, 2024). This represents a critical area for future research and professional development.

III. Materials and Methods

The researcher utilized a qualitative case study design to address how AI might have prevented audit failures of the past, which fits best the how/why question of a complex organizational situation.

A. Case Study Design

The study combined the following elements:

- Triangulation between regulatory filings, judicial rulings, audit firm reports, and academic analyses;
- Transparency through a written coding procedure and evidence classification;
- Internal peer review of coding consistency and a traceable chain of evidence are validity checks.

B. Case Selection

The Luckin Coffee incident exposed the limitations of traditional audit processes, such as reliance on sampling, lack of real-time transaction data analysis, and excessive trust in management representations. It provided a natural experimental setting to test the effectiveness of traditional auditing versus AI-driven auditing. Moreover, as a listed company, Luckin's financial statements, news reports, regulatory penalty announcements, and internal investigation reports are relatively public. Its data is mainly digital transactions, which makes it suitable for big data analysis, ML, anomaly detection, and NLP.

C. Data Sources

The data was found in annual and quarterly financial announcements published on the company's official website, independent investigations conducted by third-party accounting firms such as Ruihua Certified Public Accountants, and investigation reports issued by the company's board of directors. Additional sources consist of penalties, investigations, or notices released by the U.S. Securities and Exchange Commission (SEC) and the China Securities Regulatory Commission (CSRC), as well as Bloomberg Terminal data and information from the company's mobile application.

D. Analytical Approach

A case-based failure analysis was conducted. Using the coded themes to guide the process, we identified failures in traditional procedures and performed counterfactual simulations using ML (e.g., random forests, clustering) to structured data and NLP (e.g., topic modeling, entity recognition) to unstructured texts to determine whether anomalies could have been detected sooner. Analysis was restricted to data that was available to auditors at the time to reduce the effects of hindsight bias.

E. Validity and Limitations

Construct validity: triangulation based on the multi-source, code-procedure alignment; Internal validity: cross-case comparison; External validity: application to other high-risk situations; Reliability: audit trail and transparent coding procedures.

Such shortcomings as the use of secondary data and simulations instead of actual interactions; future studies may involve experiments, or longitudinal fieldwork.

IV. Result and Discussion

A. *Audit failure Analysis of Luckin Coffee*

The audit failure at Luckin Coffee was primarily driven by inflated revenues and fabricated transactions. Under the direction of senior executives, employees created fictitious store sales and manipulated transaction records, amounting to approximately RMB 2.2 billion. The limitations of the audit engagement were evident in the auditors' heavy reliance on company-provided transaction data and written confirmations, with insufficient independent verification procedures such as on-site validation of store sales or systematic data analysis. Moreover, the audit continued to rely on traditional sampling and manual verification techniques, without leveraging ML or anomaly detection algorithms to process large transaction datasets. Unstructured data sources, including contracts, emails, coupon campaigns, and promotional records, were also not systematically examined, further constraining the ability to detect irregularities.

B. *The article Discusses the Counterfactual Potential of AI in Auditing*

1) *Identifying structured data abnormalities*

Random forests and clustering algorithms are the type of ML models that allow auditors to examine large datasets as opposed to a limited sample.

In Luckin, transaction clusters with an abnormal frequency and circular flows that did not correspond to real customer behavior could have been identified using the method of clustering. Clustering analysis does not rely on predefined rules or labels. Instead, it identifies hidden patterns and anomalous groupings based on the intrinsic characteristics of the data. It can be used to detect clusters of fraudulent transactions, as such transactions are rarely isolated but often display quantifiable behavioral patterns that differ from genuine ones. In this context, clustering groups stores together, making it possible to identify those with abnormal performance. First, aggregated features were extracted for each store over a given period (e.g., one quarter) as part of data preparation and feature engineering. Second, clustering algorithms were applied to group stores with similar characteristics. Third, normal stores were found to typically fall into several clusters. For example: Mature shopping district stores were high transaction value, stable weekday traffic, increased weekend sales. Office district stores were strong morning peak, weak weekend sales.

The following image allows for a comparison between normal and abnormal stores, highlighting the characteristics exhibited by the abnormal stores.

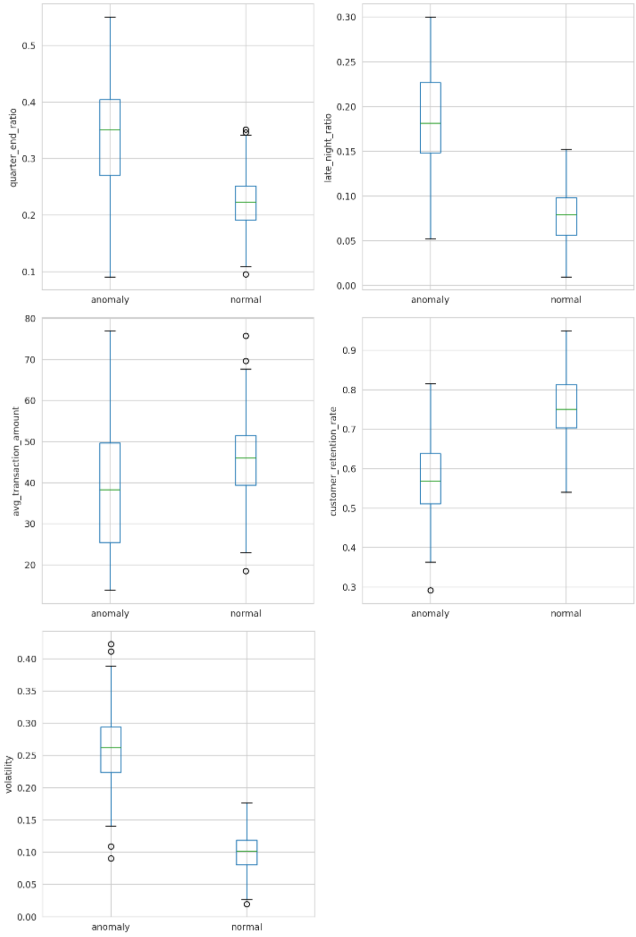

Fig. 1. Different types ratios between normal and abnormal stores.

The figure was based on data from 228 normal stores and 72 abnormal stores identified out of 300 stores using a clustering algorithm. As shown in Fig. 1, the quarter end ratio was 53.2%. In the validated cluster of normal stores, we observed that the Interquartile Range (IQR) of the quarter end ratio was 15% to 25%. And the median and overall distribution of the abnormal group were significantly higher than those of the normal group. Abnormal stores showed a stronger tendency to boost sales at the end of the quarter (e.g., in the last few days), indicating a quarter-end sprint signal. End-of-quarter sales share was excessively high. Late night ratio was 24.1% (Normal range was 5-10%), and the abnormal group was higher than the normal group, with both the median and upper quartile elevated. Abnormal stores recorded a larger share of transactions during late-night hours, which suggested that unconventional transactions may have been used in off-peak periods to inflate sales. Volatility was 0.156 (Normal range was 0.3–0.5), and the abnormal group was significantly higher than the normal group. Their sales were more volatile and unstable, often linked to short-term promotional bursts, end-of-month or end-of-quarter surges, or irregular transaction patterns. Regarding the average transaction amount, the differences between the two groups were minor, with the abnormal group slightly lower but not significant. Average transaction value was not a key feature distinguishing abnormal from normal stores; more variation appeared in timing patterns and volatility structures. The customer retention rate in the abnormal group was significantly lower than the normal group. Although abnormal stores may temporarily raise sales, they showed weaker customer loyalty, reflecting poor sustainability or excessive reliance on short-term promotions. Abnormal stores exhibited higher quarter-end ratios, late-night transactions, and sales volatility, displaying the behavioral traits of short-term sprinting, time concentration, and high fluctuation. At the same time, their lower customer retention indicated that such a model was unfavorable for long-term performance.

The lack of significant difference in the average transaction value suggests that abnormalities are more likely driven by unusual transaction timing and frequency rather than by the value of individual transaction.

Coupon usage was negatively correlated with sales ($r \approx -0.267$), meaning higher sales corresponded to fewer or irregular coupon uses, whereas the correlation was typically positive in normal stores, suggesting that fake orders did not rely on promotions to attract customers. While most normal stores reported 15%–25% in the final week, this cluster reached 40%–60%.

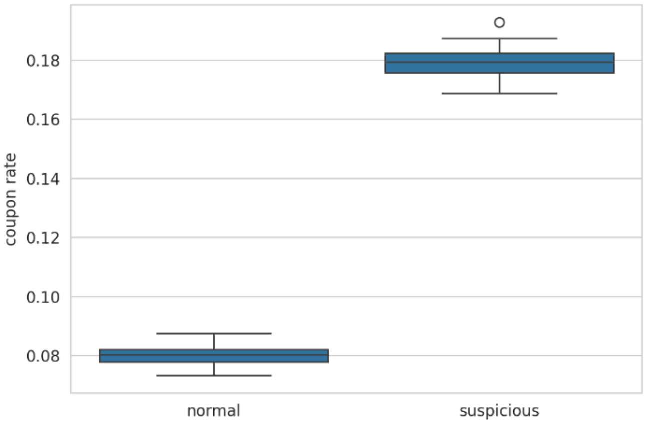

Fig. 2. Coupon rate between normal and suspicious stores.

Fig. 2 compared the distribution of coupon rates between the normal and suspicious groups. The median coupon rate of the normal group was about 0.08, with a concentrated distribution and a small range of variation. The suspicious group had a median coupon rate of about 0.18, also with a concentrated distribution, but at a significantly higher overall level than the normal group, with a few outliers. We used an independent samples test to compare coupon usage rates between the two groups and found that the suspicious store group had an average of 0.18, which was significantly higher than the normal store group. The elevated coupon rate in the suspicious group may indicate a link between suspicious behavior and high discount levels. A statistical test can be conducted to determine whether the difference was significant. Since coupon rate serves as a risk indicator, the higher level in the suspicious group warrants focused attention and further investigation.

Average sales compared the distribution of average sales between the normal and suspicious groups. The median average sales of the normal group were slightly higher than that of the suspicious group, with a more concentrated distribution. The suspicious group showed a wider distribution, greater volatility, a lower minimum (including extremely low values), and an upper limit like the normal group. The suspicious group had a lower median sales level and more extreme low values, indicating greater instability. The large fluctuations and lower median sales in the suspicious group suggested that suspicious behavior may be linked to sales anomalies. Further analysis can investigate the sources of extreme low values in the suspicious group to determine whether they were associated with specific periods, stores, or personnel. A joint analysis of average sales can be performed to identify anomalies with high coupon rates but low sales.

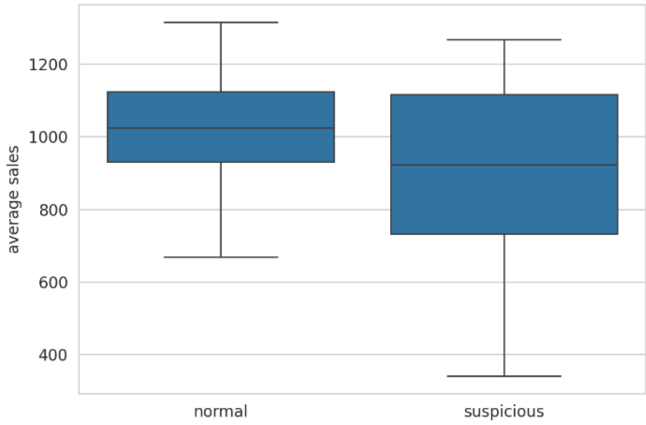

Fig. 3. Average sales rate between normal and suspicious stores.

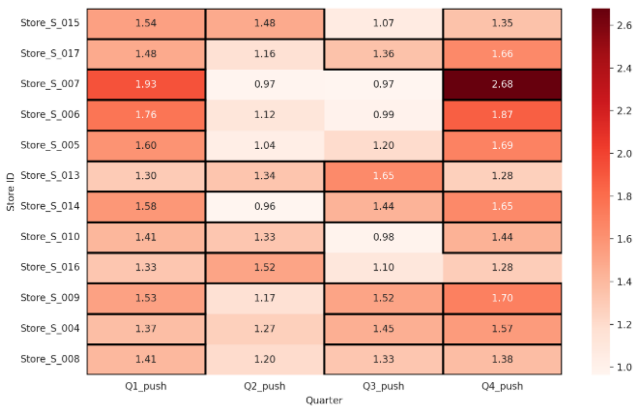

Fig. 4. Quarterly push index distribution by store.

We labeled stores with a sprint index above the mean plus two standard deviations (about 1.3) as highly suspicious. Stores with cell values greater than 1.3 in multiple quarters showed a consistent quarter-end sprint pattern. Store 9 was noticeably red in Q1, Q3, and Q4, indicating that the sprint was a fixed strategy rather than an isolated event. However, some individual cells had extremely high values (e.g., >2.0). Store 7, for example, showed 2.6–2.7 in Q4_push, meaning that sales in the last four days of the quarter were approximately 2.6–2.7 times the average of the earlier period. This represented suspicious sprint behavior and should be prioritized for review. We calculated the sprint index variance across four quarters; Stores 7 and 9 showed the lowest values, indicating persistent, strategic fraud rather than random fluctuations.

The figure showed a single highly suspicious store. The analysis highlighted abnormal fluctuations and an elevated coupon rate. Sales and coupon usage rate (Coupon Rate) increased around the end of each quarter, especially between Day 60 and Day 100. The raw data (light-colored line) showed noticeable spikes and troughs. Overall, sales rose first and then fell. From approximately Day 20 to Day 60, sales were relatively high (>1400), then decline to around 1000. Short-term increases were observed around the end of each quarter (near the dashed lines). The raw values fluctuated considerably, but the 7-day Moving Average (MA7) clearly rose between Day 60 and Day 100, peaking near 0.30, then slightly decreased. The increase in Coupon Rate partially overlapped with the rise in sales, indicating that promotions helped drive sales.

This AI model successfully identified 72 high-risk stores out of 300 by automatically detecting fraudulent patterns characterized by an abnormally high quarter-end share, unusual late-night orders, and anomalously low data volatility. It clustered 83.3% of previously known abnormal stores together and additionally identified 12 stores with the same anomalous features that had not been flagged before. This allowed audit resources to be concentrated on the top 11% of high-risk stores, rather than distributed evenly. Auditors conduct surprise inventory counts at these high-risk stores to verify whether the reported sales align with actual inventory consumption; perform 100% confirmation or telephone verification of large quarter-end orders; and review the performance evaluations of store and regional managers to assess whether sales figures are excessively tied to incentive structures.

The approach transforms auditors' work from "finding a needle in a haystack" to "targeted detection," greatly enhancing the ability to uncover systematic financial fraud.

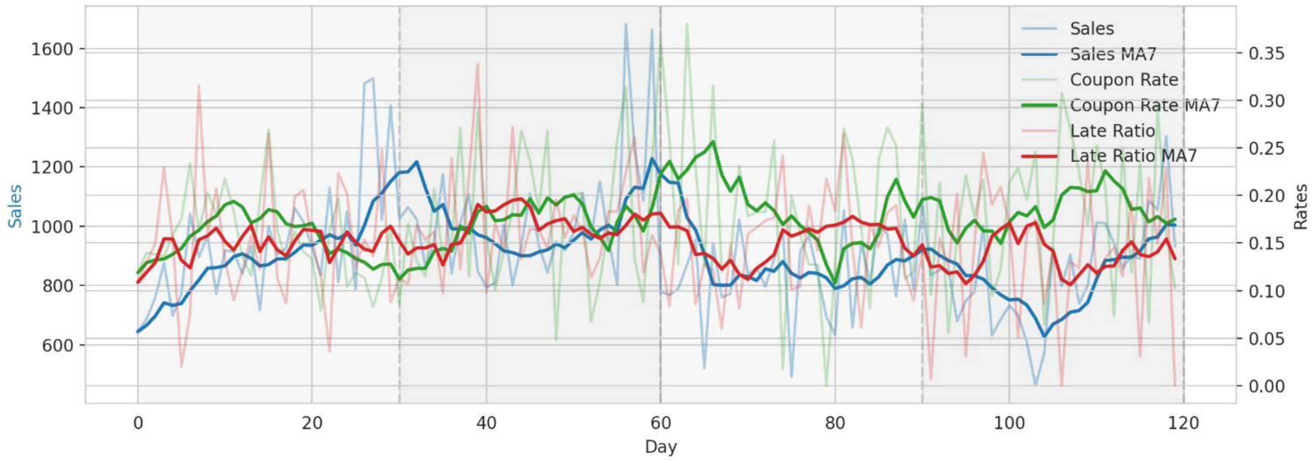

Fig. 5. Time series analysis of sales, coupon rate, and late ratio in store_S_008.

2) Uncovering insights in unstructured data

Natural Language Processing applied to contracts, regulatory filings, and email correspondence enables the extraction of signals missed by manual review. For Luckin Coffee, topic modeling of internal emails suggested collusive language patterns, while sentiment analysis of managerial disclosures revealed discrepancies between optimistic narratives and deteriorating operational metrics.

First, the data was preprocessed. Internal emails from Luckin Coffee were tokenized in Chinese, removing meaningless words and retaining only core content words, then converted into a numerical representation. Next, topic modeling was performed using Latent Dirichlet Allocation to uncover "hidden agendas". The model identified frequently co-occurring word groups, with terms such as assessment, pressure, completion rate, and private discussion appearing often in Luckin Coffee's emails.

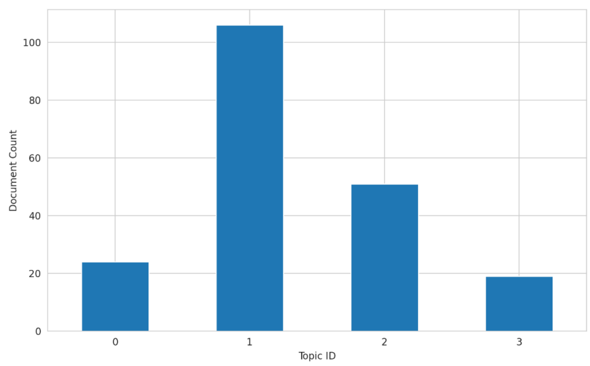

Fig. 6. Topic distribution of internal message.

According to the analysis of Fig. 6, Topic 1: Daily operations and resources (e.g., coffee beans, new products, training) had the highest volume of email exchanges. Topic 2: Suspicious/avoidance activities (e.g., delete, private, virtual) ranked second in email volume, indicating potential issues that warrant attention. Topic 3: Performance pressure (e.g., complete, performance, assessment) and Topic 0: Store services (e.g., store, service, service quality) accounted for similar proportions of emails.

Then, discourse analysis was conducted by combining the topic modeling results with communication metadata to identify key personnel. In the Luckin Coffee model, middle managers' communications were highly concentrated around suspicious operations. They occupied central positions, connecting senior management and frontline employees.

Finally, signals were captured by tracking changes in topic intensity over time. In Luckin Coffee, the intensity of the suspicious operations topic noticeably increased at the end of each quarter.

According to the analysis of Fig. 7, the highest value occurred in May, which was the end of the first quarter. The line chart also showed that in the following year, May data were slightly higher than in April, indicating that the proportion of suspicious operations increased noticeably at the end of the quarter. We applied the Cox-Stuart trend test to the time series of this topic share and found a significant upward trend in the quarter-end months ($p < 0.05$).

It was noteworthy that the 'suspicious operations' topic peaked at the end of Q1 in 2022 (May), which fully coincided with the period when large-scale fabricated transactions were later revealed by investigations.

These results showed that NLP, even with limitations, significantly expanded the auditor's evidence base by systematically processing volumes of unstructured data that human auditors could not reasonably cover.

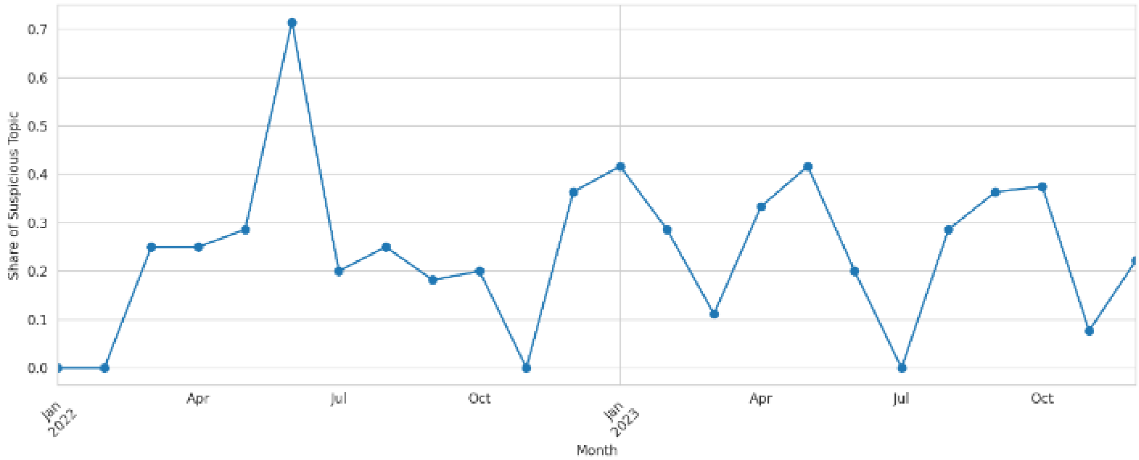

Fig. 7. Monthly share of suspicious topic.

C. *Comparative Advantage: AI vs Traditional Methods*

A systematic comparison highlights the systemic advantages of AI as compared to conventional methods. Table 1 gives eight evaluative dimensions: coverage, timeliness, evidence diversity, interpretability, cost efficiency, impact on auditor independence, granularity of risk localization, and traceability of working papers. Although conventional techniques prove to be relatively effective in terms of interpretability and professional judgment, AI proves to show a real benefit in terms of coverage, timeliness, and granularity, thus providing support to human auditors and not replacing them.

Table 1. Comparison of traditional vs. AI-enhanced auditing

Dimension	Traditional Audit	AI-Enhanced Audit
Coverage of Transactions	Sampled subset	Full population
Timeliness of Detection	Delayed (post-year-end)	Near real-time
Evidence Type	Structured only	Structured + Unstructured
Interpretability	High (manual judgment)	Moderate, requires model explanation
Cost Efficiency	Labor-intensive	Higher upfront, scalable thereafter
Independence Impact	High reliance on management	Lower, data-driven verification
Risk Localization Granularity	Low (aggregate tests)	High (account-/transaction-level)
Working Paper Traceability	Manual narratives	Automated logs + audit trails

In conclusion, this comparison demonstrates that AI can improve the coverage, timeliness, and richness of evidence, while also posing challenges regarding interpretability, data governance, and competence of audit.

D. *Discussion: Interpreting Findings Through a Socio-Technical Lens*

This study demonstrates that AI's value in auditing extends beyond simple automation. My findings indicate that AI fundamentally augments audit quality by addressing systemic blind spots inherent in traditional, sample-based approaches. The Luckin Coffee case illustrates this point clearly: AI not only expands coverage but, more importantly, identifies systemic behavioral patterns—like the coordinated quarter-end sales surges and anomalous late-night transactions—that human sampling would almost certainly miss. This aligns with the core tenet of Socio-Technical Systems (STS) theory: it reveals a powerful complementarity where AI excels at detecting patterns across vast datasets, while human auditors excel at interpreting the business substance and intent behind those patterns. The synergy transforms the audit from a retrospective exercise into a more proactive, diagnostic process.

1) *The interdependence of human expertise and algorithmic output*

The effective use of AI relies critically on human capital. As our simulations show, algorithmic outputs, such as the outlier clusters in Fig. 2 or the suspicious patterns in Fig. 3, require human interpretation and contextual understanding. Algorithms can flag a store as high-risk based on statistical anomalies, but only a skilled auditor can determine whether this indicates fraud, aggressive accounting, or an unusual but legitimate business practice. This necessity directly supports the STS perspective, underscoring that AI cannot function in isolation. Auditors must view these tools as instruments to support decision-making, not as replacements for their judgment. Reconciling AI with professional skepticism means fostering digital skepticism—the ability to question and validate algorithmic findings—which must become a core component of professional training and evaluation.

2) *Reframing adoption challenges using the TOE framework*

The challenges of AI adoption can be systematically understood through the TOE framework. Our analysis confirms that success depends on more than just technological capability.

Technologically, while tools like ML and NLP are powerful, issues of data quality, model interpretability (black box concerns), and seamless integration with existing ERP systems remain significant hurdles. The proposed use of Explainable AI (XAI) in audit workpapers is a direct response to this challenge. Organizationally, firms must bridge substantial human capital and workflow gaps. This involves reskilling auditors, forming multidisciplinary teams (auditors, data scientists, IT specialists), and formally updating audit methodologies to incorporate AI-driven insights. The resource constraints faced by smaller firms threaten to create a significant AI adoption gap within the profession. Environmentally, regulatory uncertainty, evolving professional standards, and client data readiness exert strong external pressures. A collaborative effort between firms, regulators, and standard-setters is essential to establish clear guidelines for accountability, independence, and for the auditing of algorithms themselves.

3) A governance framework for responsible integration

Moving from technical potential to responsible practice requires a robust governance structure. Building on the TOE and STS foundations, our proposed four-pillar framework ensures that AI adoption is balanced and accountable:

- Technology: Include XAI outputs in audit workpapers, define interoperability standards with ERP systems, and ensure model drift and robustness monitoring is carried out continuously.
- Organization: Reskill auditors, institutionalize multidisciplinary audit teams, and update audit methodologies in a manner that incorporates AI insights into sampling, testing, and documentation.
- Environment: Work with regulators, standard setters, and clients to ensure AI practices work within the professional requirements and legal mandates; encourage industry standards of detection accuracy and error levels.
- Governance: establish AI oversight committees, implement exception-handling workflows, conduct regular algorithms audits, and mandate third-party evaluations to ensure accountability.

This integrated model positions AI as an augmentation tool that enhances assurance quality while preserving auditor independence and professional skepticism.

4) Future trajectory: toward continuous, integrated assurance

The future of AI-enhanced auditing lies in its convergence with other technologies like blockchain and the Internet of Things (IoT). This will accelerate the shift from periodic audits to continuous, real-time assurance. Smart contracts on blockchain could provide immutable audit evidence, while IoT sensors could automatically verify physical assets and supply chain transactions. However, this future amplifies risks related to data dependence and privacy. A dual-track strategy is therefore essential: a technical track focused on integrating AI detection and explainability into continuous monitoring systems, and an institutional track dedicated to establishing the standards, governance, and compliance frameworks for the responsible adoption of AI.

V. Conclusion

This paper discusses the role of AI fundamentally enhances audit quality by overcoming the systemic limitations of traditional, sample-based methods. The counterfactual analysis of the Luckin Coffee case reveals that:

- ML clustering techniques successfully identified anomalous store behaviors—such as abnormally high quarter-end sales ratios (53.2% vs. a normal range of 15–25%), elevated late-night transactions (24.1% vs. 5–10%), and high sales volatility—that were indicative of coordinated fraud. The model pinpointed 72 high-risk stores out of 300, transforming the audit approach from finding a needle in a haystack to targeted detection.
- NLP and topic modeling uncovered latent risks within internal communications. It detected a significant proportion of emails related to suspicious activities and performance pressure, with the intensity of suspicious topics peaking at quarter-ends, providing a crucial, previously untapped evidence source.
- A systematic comparison confirms AI's superior coverage, timeliness, and granularity over traditional audits.

Critically, AI serves as a powerful complement to—not a replacement for—human auditors, augmenting professional skepticism by extending the evidential foundation.

From a theoretical perspective, this research offers a systematic explanation of the adoption of AI in the field of auditing as a complex process influenced by the technological capacity, organizational preparedness, and environmental forces through the use of the TOE framework and the STS theory. It also contributes to the discussion on professional skepticism as it demonstrates the relationship between human judgment and the product of algorithms to create a hybrid audit model.

The findings offer clear guidance for stakeholders: Audit firms should invest in XAI and integrate AI tools into methodologies, supported by reskilling and multidisciplinary teams. Regulators and standard-setters can utilize the proposed four-pillar governance framework to update standards for the AI era.

Several limitations point to important avenues for future research. First, the counterfactual design relies on historical data and simulations. Second, Detection metrics require empirical validation in live audit settings. Lastly, the single-case focus limits generalizability. Future research may therefore benefit from longitudinal field studies, controlled experiments on auditor–AI collaboration, and comparative investigations of governance models across different jurisdictions.

Overall, AI will not eliminate fraud but can transform auditing into a more proactive and risk-sensitive profession. In light of these findings, the pivotal challenge is no longer if AI can be adopted, but how to implement it responsibly. This necessitates a collaborative effort to ensure AI enhances, rather than undermines, audit integrity in the digital economy.

Conflict of Interest

The authors declare no conflict of interest.

References

Al-Qudah, A., & Al-Hattami, H. (2023). Artificial intelligence in auditing: A conceptual framework for enhancing assurance and ethics. ***Journal of Accounting and Organizational Change***, 19(4): 657–674.

Appelbaum, D., Kogan, A., & Vasarhelyi, M. A. (2017). Big data and analytics in the audit profession. ***Journal of Accounting Literature***, 36: 19–50.

Bakumenko, A., & Elragal, A. (2022). Detecting anomalies in financial data using machine learning algorithms. ***Systems***, 10(5): 130.

Carcello, J. V. (2020). Luckin Coffee and lessons for audit oversight. ***Accounting Today***.

De Oliveira, F. R., da Silva, E. P., & Neto, R. L. (2025). The increased role of advanced technology and automation in audit. ***Journal of Accounting Literature***, 54: 101–118.

Deloitte. (2020). AI in financial crime prevention: The future of AML and KYC. ***Deloitte Insights.***

EY. (2019). EY uses artificial intelligence for contract review. ***EY Global.***

Gandhar, A., Gupta, K., Pandey, A. K., & Raj, D. (2024). Fraud detection using machine learning and deep learning. ***SN Computer Science***, 5: 453.

Hamdan, S. A. R., & Al Habashneh, A. K. (2024). The advantages and difficulties of using AI and BT in auditing procedures: A literature review. ***In Artificial Intelligence-Augmented Digital Twins:*** 111–126. Springer Nature.

Hasan, A. R. (2022). Artificial intelligence (AI) in accounting and auditing: A review**. *Journal of Accounting and Business Research*,** 42(2): 88–105.

Kokina, J., & Davenport, T. H. (2017). The emergence of artificial intelligence: How automation is changing auditing. ***Journal of Emerging Technologies in Accounting***, 14(1): 115–122.

Kokina, J., Ma, Y., & Davenport, T. H. (2021). The adoption of artificial intelligence in auditing: Challenges and opportunities. ***Accounting Horizons***, 35(2): 33–56.

KPMG. (2025). A more fluent audit with natural language processing. ***KPMG Insights***

Leocádio, D., Malheiro, L., & Reis, J. C. G. (2024). Auditors in the digital age: A systematic literature review. ***Digital Transformation and Society***, 4(1): 5–20.

Li, Y., & Vasarhelyi, M. A. (2023). Responsible AI adoption in auditing: A framework for ethical assurance. ***Journal of Information Systems***, 37(1): 89–104.

Minkkinen, M., Laine, T., & Mäntymäki, M. (2022). AI-driven auditing and the transformation of professional logic: A socio-technical perspective. ***Accounting, Auditing & Accountability Journal***, 35(8): 1987–2009.

PwC. (2018). PwC Halo: Transforming audit through technology. ***PwC Global.***

Sun, T., & Vasarhelyi, M. A. (2018). The effect of data analytics and AI on audit evidence and assurance. ***Accounting Horizons***, 32(4): 29–45.

Wang, Y., Zhang, X., & Li, C. (2023). Text mining and linguistic cues in financial reporting fraud detection: Evidence from NLP-based auditing analytics. ***Expert Systems with Applications***, 227: 120325.

Yang, J., Blount, Y., & Amrollahi, A. (2024). Artificial intelligence adoption in a professional service industry: A TOE framework perspective. ***International Journal of Accounting Information Systems***, 55: 101739.