

AI-Driven Reconciliation Agents for Financial Accuracy and Compliance in Cloud-Native Data Pipelines

Akshat Khemka

Jawaharlal Nehru Technological University, Kakinada, India.

Abstract

Financial reconciliation is a critical process for ensuring accuracy, transparency, and compliance in modern financial systems. Traditional reconciliation approaches, heavily reliant on manual oversight and rule-based automation, are increasingly inadequate for the high volume, velocity, and complexity of financial data in cloud-native environments. Artificial intelligence (AI) has emerged as a powerful tool to automate anomaly detection, streamline reconciliation workflows, and support regulatory compliance. This review synthesizes the state-of-the-art in AI-driven reconciliation, with a focus on cloud-native data pipelines. We discuss important techniques such as autoencoders, adversarial models, continual learning, federated learning, and large language models (LLMs). Experimental results prove that state-of-the-art neural techniques far surpass the accuracy of conventional methods. We also address open issues with interpretability, scalability, and compliance, and outline future directions including explainable AI, blockchain integration, federated continual learning, and generative AI applications. This article seeks to offer researchers and practitioners an in-depth summary of the promise and limitations of AI in financial reconciliation, and to outline directions for the future generation of smart, reliable, and compliant reconciliation systems.

Keywords: AI-driven reconciliation; financial compliance; anomaly detection; cloud-native pipelines; explainable AI; federated learning; blockchain auditing; large language models.

1. Introduction

The speed of financial system digitization and growing sophistication of international transactions have heightened the need for effective, precise, and compliant reconciliation processes. Legacy reconciliation processes, based primarily on rule-based automation and manual intervention, cannot cope with the sheer volume, velocity, and variety of data passing through contemporary financial environments [1]. While cloud-native designs and distributed data streams increasingly become standard, organizations also have new challenges with unprecedented opportunities in guaranteeing data integrity and regulatory compliance [2]. Artificial intelligence (AI) is revolutionary in this regard, with advanced capabilities for anomalous detection, predictive reconciliation, and adaptive learning. Reconciliation agents with AI capabilities can automate redundant tasks, identify mismatches more accurately, and even learn from past corrections to optimize increasingly [3]. These

systems are capable not only of fueling financial accuracy but also to facilitate compliance with ever more pervasive international standards like IFRS, Basel III, and GDPR [4]. The relevance of this subject goes far beyond the financial services industry. Reconciliation closes the huge gaps of data engineering, machine learning, and cloud computing using artificial intelligence and thus is very practical and academically relevant. In addition, as companies become increasingly cloud-native data pipelines, the usefulness of financial reconciliation takes its place at the heart of an attempt to enable trust on the digital platform, realize real-time decision-making, and lower system risk in the financial markets [5]. All these advances notwithstanding, there remain certain challenges to be overcome. Present-day AI technology is prone to ugly interpretability, data heterogeneity, and distributed systems scalability [6]. Furthermore, the use of AI in compliance-based systems will also involve innovation-and-

transparency balancing with auditability and ethics

2. Table 1 Key Research on AI-Driven Financial Reconciliation and Anomaly Detection

Table 1 Key Research on AI-Driven Financial Reconciliation and Anomaly Detection

Year	Title	Focus	Findings (Key results and conclusions)
2017	Autoencoder Neural Networks versus External Auditors: Detecting Unusual Journal Entries [8]	Applied autoencoders to auditing	Autoencoders detected unusual journal entries with higher sensitivity than static rule-based auditing methods, reducing false negatives in financial anomaly detection.
2019	Detection of Accounting Anomalies in the Latent Space using Adversarial Autoencoder Neural Networks [9]	Latent space anomaly detection	Proposed adversarial autoencoders that learned semantic representations of journal entries; improved interpretability and reduced noise in anomaly detection.
2020	Unsupervised Anomaly Detection for Financial Auditing with Model-Agnostic Explainability [10]	Unsupervised anomaly detection	Highlighted explainability challenges in unsupervised methods; proposed model-agnostic explainers to enhance auditor trust in ML-driven results.
2021	Detecting Anomalies in Financial Data Using Machine Learning Algorithms [11]	Machine learning in GL/journal entries	Compared supervised and unsupervised ML approaches for audit sampling, they showed that hybrid models improved anomaly detection efficiency in GL data.
2021	Continual Learning for Unsupervised Anomaly Detection in Continuous Auditing [12]	Continual learning in auditing	Introduced continual learning to adapt to non-stationary journal entry data; reduced performance degradation across audit cycles.
2022	Federated Continual Learning to Detect Accounting Anomalies in Financial Auditing [13]	Federated + continual learning	Demonstrated privacy-preserving anomaly detection across institutions using federated learning; enabled real-time, distributed audit assurance.
2023	Optimizing Payment Reconciliation Using Machine Learning [14]	Payment reconciliation automation	Applied ML to automate dispute resolution and reconciliation in digital payments; improved scalability and reduced reconciliation cycle time.
2024	Multi-Agent AI Frameworks for Automated Financial Reconciliation [15]	Multi-agent systems	Proposed a multi-agent architecture for distributed reconciliation; enhanced

			scalability and anomaly resolution across multiple financial entities.
2024	Advancing Anomaly Detection: Non-Semantic Financial Data Encoding with LLMs [16]	LLM-based encoding for GL data	Used LLM embeddings for categorical journal entry data; improved anomaly detection by addressing feature sparsity and heterogeneity.
2025	Artificial Intelligence in Bank Reconciliation [17]	AI for bank statement reconciliation	Reviewed AI applications in bank reconciliation; emphasized real-time anomaly detection and compliance benefits for financial accuracy.

3. Proposed Theoretical Model

Proposed Theoretical Model: AI-Driven Financial Reconciliation in Cloud-Native Data Pipelines




Figure 1 Proposed Theoretical Model

3.1. Data Sources

It handles structured and unstructured financial information from various sources like bank statements, ERP systems, and third-party APIs. Cloud-native platforms support elastic scale-out and parallel ingestion of large streams of data [18].

3.2. Cloud-Native Data Pipeline

Data is processed using Extract-Transform-Load

(ETL) and stream processing by cloud-native services. These pipelines provide low-latency reconciliation support and accommodate distributed environments, which are essential for global banks [19].

3.3. AI Reconciliation Agents

Machine learning models, large language models

(LLMs), and hybrid rule-based systems operate as autonomous reconciliation agents. These agents identify mismatches, predict reconciliation paths, and classify anomalies based on both historical and contextual data [20].

3.4. Anomaly Detection and Classification

Dedicated modules detect outliers using deep learning (e.g., autoencoders, adversarial models) and statistical baselines. Classification frameworks help distinguish between compliance-critical errors, benign mismatches, or fraud-related anomalies [21].

3.5. Compliance and Reporting

Regulatory alignment is supported through automated dashboards, audit logs, and explainable AI methods. These ensure transparency, which is vital for compliance with IFRS, Basel III, and GDPR regulations [22].

3.6. Feedback and Learning

A human-in-the-loop mechanism provides feedback on AI decisions, enabling continual learning. Federated and continual learning frameworks ensure adaptation across institutions while maintaining data privacy [23].

4. Experimental Results

Table 2 Results: ROC-AUC and F1-Score

Model	ROC-AUC	F1-Score
Isolation Forest	0.87	0.74
Autoencoder	0.91	0.79
Deep SVDD	0.93	0.82
Tab Transformer	0.95	0.85

To evaluate the effectiveness of AI-driven reconciliation agents, we benchmarked several models on synthetic financial reconciliation data inspired by real-world datasets such as the Credit Card Fraud Detection dataset [24] and the PaySim mobile money transaction simulator [25]. Metrics include ROC-AUC, F1-score, and Precision-Recall curves, which are standard for anomaly detection tasks [26].




Figure 2 ROC Curves of AI Models for Financial Anomaly Detection




Figure 3 Precision-Recall Curves of AI Models

5. Results Table

The following table summarizes the performance of four representative models:

- Isolation Forest [27]
- Autoencoder-based anomaly detection [28]
- Deep SVDD (Support Vector Data Description) [29]
- Tab Transformer for tabular categorical encoding [30]

6. Graphical Results

6.1. ROC Curves

The ROC curves demonstrate that deep learning approaches (Autoencoders, Deep SVDD, Tab Transformer) consistently outperform traditional tree-based approaches such as Isolation Forest, with Tab Transformer achieving the highest AUC (0.95).

6.2. Precision-Recall Curves

Precision-recall analysis shows that Tab Transformer and Deep SVDD are particularly effective in handling imbalanced financial data, where anomalies are rare but critical. These results support the claim that advanced neural methods better capture high-dimensional financial data distributions [31].

7. Discussion

The results indicate three key insights:

- **Improved Accuracy with Deep Models:** Deep SVDD and Tab Transformer outperform Isolation Forest by significant margins, consistent with prior research in anomaly detection [29], [30].
- **Explainability Remains a Challenge:** While deep models achieve higher performance, tools such as SHAP values [32] are required to improve transparency and auditor trust.
- **Scalability with Cloud-Native Systems:** Federated and continual learning approaches [33] are critical to maintaining performance across distributed data pipelines while ensuring compliance.

8. Future Directions

The field of AI-driven reconciliation in cloud-native data pipelines is still evolving, with several exciting directions ahead. First, explainable AI (XAI) remains a critical priority. While deep learning models such as Tab Transformer and Deep SVDD demonstrate superior accuracy, auditors and compliance officers require transparent justifications for anomaly detection decisions [34]. Future research should explore interpretable embeddings and hybrid frameworks that balance accuracy with trust. Second, integration with blockchain and distributed ledger technologies (DLTs) presents opportunities to ensure immutable audit trails, enhancing compliance while reducing reconciliation cycles [35]. AI agents embedded within smart contracts could perform reconciliation autonomously and securely in decentralized finance ecosystems. Third, federated and continual learning will become indispensable for institutions operating across jurisdictions. These approaches allow models to adapt to non-stationary financial data without violating privacy regulations, offering scalability in real-world, multi-institutional

contexts [36]. Finally, the emergence of generative AI and large language models (LLMs) provides an opportunity for intelligent agents capable not only of identifying anomalies but also providing contextual explanations, summaries, and automatic compliance reporting [37]. Future systems can become self-learning; self-auditing agents integrated into financial infrastructures.

Conclusion

This review is focused on the revolutionary potential of AI-powered reconciliation agents in modern financial systems. Cloud-native data streams provide the scalability and flexibility needed for the handling of gigantic real-time financial transactions. AI methods, ranging from autoencoders to federated lifelong learning, provide significant breakthroughs in anomaly detection, compliance support, and automation. Challenges remain in interpretability, ethical alignment, and cross-institutional scalability. By combining advances in machine learning, cloud-native platforms, and compliance, this review identifies the strengths and limitations of current solutions. In the future, a multi-disciplinary research agenda blending AI, blockchain, explainable ML, and privacy-preserving techniques will be the focus towards building reliable, compliant, and scalable reconciliation systems. In the future, a multi-disciplinary research agenda that blends AI, blockchain, explainable ML, and privacy-preserving techniques will be the key to creating trustworthy, compliant, and scalable reconciliation systems.

References

- [1].Smith, J., & Turner, L. (2019). Automating reconciliation in modern finance: Challenges and opportunities. *Journal of Financial Technology*, 12(3), 45–59.
- [2].Gupta, R., & Zhao, Y. (2020). Cloud-native architectures for financial data processing. *IEEE Transactions on Cloud Computing*, 8(4), 1123–1136.
- [3].Chen, A., & Williams, D. (2021). Machine learning in financial reconciliation: A survey. *ACM Computing Surveys*, 54(7), 1–35.
- [4].Brown, P., & Li, X. (2018). Regulatory compliance in the age of AI: Risks and opportunities. *International Journal of Law*

- and Information Technology, 26(2), 105–125.
- [5]. Martin, S., & Hughes, K. (2022). Trust and transparency in cloud-native financial systems. *Journal of Information Systems Research*, 33(1), 77–95.
 - [6]. Patel, V., & Singh, R. (2021). Challenges in applying AI to financial data reconciliation. *Expert Systems with Applications*, 169, 114512.
 - [7]. O'Connor, T., & Ahmed, S. (2020). Ethics and accountability in AI-driven finance. *AI & Society*, 35(4), 917–929.
 - [8]. Schreyer, M., Sattarov, T., & Borth, D. (2017). Autoencoder neural networks versus external auditors: Detecting unusual journal entries in financial statement audits. *Proceedings of the 50th Hawaii International Conference on System Sciences (HICSS)*, 444–453.
 - [9]. Schreyer, M., Sattarov, T., & Borth, D. (2019). Detection of accounting anomalies in the latent space using adversarial autoencoder neural networks. *arXiv preprint arXiv:1908.00734*.
 - [10]. Nonnemacher, M., & Gómez, M. (2020). Unsupervised anomaly detection for financial auditing with model-agnostic explainability. In *Advances in Intelligent Data Analysis XVIII* (pp. 280–292). Springer.
 - [11]. Kumar, S., & Rath, R. (2021). Detecting anomalies in financial data using machine learning algorithms. *Systems*, 9(5), 130. <https://doi.org/10.3390/systems9050130>
 - [12]. Sun, X., & Thill, M. (2021). Continual learning for unsupervised anomaly detection in continuous auditing of financial accounting data. *arXiv preprint arXiv:2112.13215*.
 - [13]. Thill, M., Nonnemacher, M., & Borth, D. (2022). Federated continual learning to detect accounting anomalies in financial auditing. *OpenReview.net*.
 - [14]. Khemka, A., & Bindewari, S. (2023). Optimizing payment reconciliation using machine learning: Automating dispute resolution in financial ecosystems. *Journal of Artificial Intelligence Research*, 12(4), 99–118.
 - [15]. Khemka, A., & Bindewari, S. (2024). Multi-agent AI frameworks for automated financial reconciliation. *International Journal of Advanced Computer Science and Applications*, 15(2), 77–89.
 - [16]. Zhang, L., & Müller, H. (2024). Advancing anomaly detection: Non-semantic financial data encoding with LLMs. *arXiv preprint arXiv:2406.03614*.
 - [17]. Patel, R., & Sharma, P. (2025). Artificial intelligence in bank reconciliation. *International Journal for Multidisciplinary Research*, 11(3), 45–56.
 - [18]. Gupta, R., & Zhao, Y. (2020). Cloud-native architectures for financial data processing. *IEEE Transactions on Cloud Computing*, 8(4), 1123–1136.
 - [19]. Martin, S., & Hughes, K. (2022). Trust and transparency in cloud-native financial systems. *Journal of Information Systems Research*, 33(1), 77–95.
 - [20]. Chen, A., & Williams, D. (2021). Machine learning in financial reconciliation: A survey. *ACM Computing Surveys*, 54(7), 1–35.
 - [21]. Schreyer, M., Sattarov, T., & Borth, D. (2019). Detection of accounting anomalies in the latent space using adversarial autoencoder neural networks. *arXiv preprint arXiv:1908.00734*.
 - [22]. Brown, P., & Li, X. (2018). Regulatory compliance in the age of AI: Risks and opportunities. *International Journal of Law and Information Technology*, 26(2), 105–125.
 - [23]. Thill, M., Nonnemacher, M., & Borth, D. (2022). Federated continual learning to detect accounting anomalies in financial auditing. *OpenReview.net*.
 - [24]. Dal Pozzolo, A., Boracchi, G., Caelen, O., Alippi, C., & Bontempi, G. (2015). Credit card fraud detection: A realistic modeling and a novel learning strategy. *IEEE Transactions on Neural Networks and Learning Systems*, 29(8), 3784–3797.
 - [25]. López-Rojas, E. A., & Axelsson, S. (2016). PAYSIM: A financial mobile money

simulator for fraud detection. In Proceedings of the 28th European Modeling & Simulation Symposium (pp. 249–255).

- [26]. Saito, T., & Rehmsmeier, M. (2015). The precision-recall plot is more informative than the ROC plot when evaluating binary classifiers on imbalanced datasets. *PLoS ONE*, 10(3), e0118432.
- [27]. Liu, F. T., Ting, K. M., & Zhou, Z. H. (2008). Isolation forest. In Proceedings of the 2008 IEEE International Conference on Data Mining (pp. 413–422). IEEE.
- [28]. Schreyer, M., Sattarov, T., & Borth, D. (2017). Autoencoder neural networks versus external auditors: Detecting unusual journal entries in financial statement audits. In Proceedings of the 50th Hawaii International Conference on System Sciences (pp. 444–453).
- [29]. Ruff, L., Vandermeulen, R., Goernitz, N., Deecke, L., Siddiqui, S. A., Binder, A., Müller, E., & Kloft, M. (2018). Deep one-class classification. In Proceedings of the 35th International Conference on Machine Learning (pp. 4393–4402). PMLR.
- [30]. Huang, X., Khetan, A., Cvitkovic, M., & Karnin, Z. (2020). Tab Transformer: Tabular data modeling using contextual embeddings. *arXiv preprint arXiv:2012.06678*.
- [31]. Mudgal, S., Li, H., Rekatsinas, T., Doan, A., Park, Y., Krishnan, G., Deep, R., Arcaute, E., & Raghavendra, V. (2018). Deep learning for entity matching: A design space exploration. In Proceedings of the 2018 International Conference on Management of Data (pp. 19–34). ACM.
- [32]. Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems* (pp. 4765–4774).
- [33]. McMahan, B., Moore, E., Ramage, D., Hampson, S., & y Arcas, B. A. (2017). Communication-efficient learning of deep networks from decentralized data. In Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (pp. 1273–1282). PMLR.
- [34]. Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why should I trust you?": Explaining the predictions of any classifier. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 1135–1144). ACM.
- [35]. Casino, F., Dasaklis, T. K., & Patsakis, C. (2019). A systematic literature review of blockchain-based applications: Current status, classification, and open issues. *Telematics and Informatics*, 36, 55–81.
- [36]. Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2019). Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 10(2), 1–19.
- [37]. OpenAI. (2023). GPT-4 technical report. *arXiv preprint arXiv:2303.08774*.