The American Journal of Applied Sciences

ISSN 2689-0992 | Open Access

TYPE Original Research
PAGE NO. 56-65

AI-Powered Monitoring Systems for Liquidity Risk: Predicting Cash Flows, Liquidity Coverage Ratio, and Liquidity Shortfall with Deep Learning and Real-Time Analytics

Shaurya Shounik

MS in Finance, Brandeis University Independent Researcher, USA

OPEN ACCESS

SUBMITED 11 April 2024
ACCEPTED 27 May 2024
PUBLISHED 30 June 2024
VOLUME Vol.06 Issue 06 2024

CITATION
Shounik, S. (2024). AI-Powered Monitoring Systems for Liquidity Risk: Predicting Cash Flows, Liquidity Coverage Ratio, and Liquidity Shortfall with Deep Learning and Real-Time Analytics. The American Journal of Applied Sciences, 7(09), 56–65. Retrieved from https://theamericanjournals.com/index.php/tajas/article/view/6661

Abstract: As the COVID-19 pandemic and the financial crisis of 2008 have shown, liquidity risk is still an important factor in maintaining financial stability. Although buffers were enhanced by the Liquidity Coverage Ratio (LCR) and the Net Stable Funding Ratio (NSFR), day-to-day monitoring is still dependent on scenario analysis, which relies on historical assumptions, and stress testing. This research has gathered information from universities, government organizations, and companies to find out whether AI may help close that gap. Across studies, Long Short-Term Memory (LSTM), Gated Recurrent Unit (GRU), and Transformer models generally outperform statistical baselines such as ARIMA and GARCH, cutting forecast errors by roughly 20–40% and improving predictions of cash flows, LCR components, and liquidity shortfalls. Technically, continuous, low-latency monitoring is possible with event-driven data stacks like Apache Kafka with Flink, but there have been few practical installations for liquidity risk. With the use of AI-enabled monitoring, which seems to be more accurate and responsive, institutions are moving toward a unified, real-time view of financing risk. This development enables banks to efficiently manage risk and enhance systemic regulatory monitoring. Still, regulator-auditable compliance procedures, data-quality controls,

model-risk governance, and flexibility will likely be necessary for adoption.

Keywords: Liquidity risk, Liquidity Coverage Ratio (LCR), Real-time analytics, Risk management, Artificial intelligence in finance, Deep learning, Financial Stability, Liquidity shortfall

1. Introduction

1.1 Liquidity Risk and Financial Stability

The likelihood that a company would incur losses that are too high to justify its short-term commitments is known as liquidity risk (Safiyari & Nabati, 2023). Market liquidity may disappear in an instant, driving financially stable institutions to collapse (Heuver, 2020); ten years later, the COVID-19 shock demonstrated how sudden financing constraints and cash-flow disruptions can have far-reaching effects (Farooq et al., 2023). According to (Barongo & Mbelwa, 2023), Basel III imposed stricter regulations, with the Liquidity Coverage Ratio aiming for a stress horizon of 30 days and the Net Stable financial Ratio concentrating on financial resilience for a year. Despite their effective uses, these metrics have not been able to solve the issue entirely.

1.2 Limitations of Traditional Monitoring Approaches

Inspite of the changes and developments to traditional methods, limitations like stress testing, scenario exercises, and gap analyses still exist. They don't supply live feed but rather repose on predetermined assumptions and rely on delayed reporting cycles for updates (Heuver, 2020). There is minimal opportunity for early intervention or dynamic hedging in fast-moving markets since developing pressure is sometimes only identified after it is too late to move (Dionne et al., 2015).

Because of how fast technology is developing, the financial sector is becoming increasingly resilient and complex. Because of this, financial institutions should implement more flexible and real-time liquidity management strategies (Safiyari & Nabati, 2023). Banks and regulators were clearly exposed to a high risk of unforeseen liquidity crises due to the conventional methods' reliance on slow and outdated historical data. The employment of robust AI technology is now vital for proactive risk management and the stability of the financial system in the face of constant change, rather than just a creative exercise (Heuver, 2020).

1.3 Emergence of AI and the Research Gap

Better adaptive monitoring is possible because of advancements in AI, especially deep learning. Research depicts that when it comes to financial time-series applications, Long Short-Term Memory networks, Gated Recurrent Units, and Transformers often beat statistical baselines like as ARIMA and GARCH (Phien & Platoš, 2023). While conventional approaches often fail to account for non-linear, state-dependent correlations, these models may be able to include a wide variety of inputs. Concurrently, streaming infrastructures (such as Kafka for event handling or Flink for stateful processing) potentially allow for real-time score and intake (Gontarska et al., 2021).

Even so, current applications are fragmented. Many studies focus on a single problem, such as cash-flow forecasting, deposit run-off, or LCR components, without tying results to shortfall detection or a full LCR trajectory. Operational issues also get limited treatment: interpretability for supervisors, model-risk governance, data quality, and the realities of production deployment are often left at the margins (Fritz-Morgenthal et al., 2022; Maple et al., 2023).

1.4 Objectives and Contributions of the Study

This study systematically reviews and synthesizes sources from academia, government, and business to examine the stated findings, compare AI performance to conventional approaches, and assesses the usability of real-time frameworks in banking situations. Contributions include:

Findings on cash-flow forecasting, LCR estimation, and shortfall detection were integrated into a single conceptual framework.

Assessing streaming analytics platforms and deployment patterns for proactive, low-latency monitoring.

Identifying interpretability, data governance, and model-risk controls as the key hurdles for regulatory acceptance.

2. Literature Review

2.1 Traditional Approaches to Liquidity Risk Management

Stress tests, scenario analyses, and gap analyses has been the toolkit for financial institutions to monitor their liquidity for a long time (Safiyari & Nabati, 2023). These tools map asset–liability mismatches, simulate adverse conditions, and size buffers for severe but plausible shocks (Tammenga & Haarman, 2020). Post-2008, Basel III raised the floor. The LCR and NSFR pushed institutions to hold larger, better-quality liquidity cushions (Barongo & Mbelwa, 2023). The baseline is stronger. Even so, the toolkit is mostly deterministic and backward-looking, which means it can lag fast-moving markets where risk builds intraday, not quarter to quarter (Heuver, 2020).

2.2 Quantitative Forecasting Models

Classical forecasting centers on ARIMA and GARCH for cash flows and liquidity ratios, with regressions and Monte Carlo simulations estimating HQLA and stressed outflows. These methods perform acceptably when relationships are stable. They are prone to de-calibration when non-linearities, volatility clustering, and regime transitions take front stage under stress (Rubio et al., 2023). Structural breaks, changing run-off elasticities, and feedback loops (e.g., margin calls tightening liquidity) are exactly where their predictive power appears to wane (Mariani et al., 2018).

2.3 AI and Machine Learning in Financial Applications

AI already underpins fraud screening (Kulatilleke, 2022), credit scoring (Aniceto et al., 2020), trading, and portfolio construction (Aziz & Dowling, 2018). Default prediction, volatility estimate, and operational loss detection are all areas of risk that have been enhanced by machine-learning algorithms (Ndikum, 2020). That said, many ML setups still rely on heavy feature engineering (Kim et al., 2019) and may struggle with long-range temporal dependence (Wittenbach et al., 2020), leakage risks, or drift (Jarquin et al., 2022). These are limitations that become visible precisely when conditions change fastest (Chen et al., 2023).

2.4 Deep Learning for Financial Forecasting

Deep learning addresses several of those pain points. LSTM and GRU networks learn sequential structure directly from time-series data (Liu & Long, 2020), while Transformers use attention to capture both short- and long-horizon relationships (Chen et al., 2023). Overall, these deep learning methods tend to outperform simpler statistical models on things like stock returns, volatility, and how likely a borrower is to default. In liquidity risk specifically, they've shown better accuracy for cash-flow (Weytjens et al., 2019) and deposit run-off forecasting than regressions or ARIMA. However, most treatments are still siloed with one model per task so LCR and shortfall dynamics are rarely integrated with cash-flow forecasts in a single, decision-ready view.

2.5 Real-Time Analytics in Financial Operations

Streaming stacks such as Kafka for events (Greco et al., 2018), Flink (Kalogerakis et al., 2022) or Spark Streaming (Carcillo et al., 2017) for stateful processing can let institutions ingest, transform, and score high-frequency data as it arrives (Fu & Soman, 2021). Where millisecond latencies come into effect, they are already supporting payment monitoring, fraud detection, and algorithmic trading (Kalogerakis et al., 2022). Extending the same architecture to liquidity monitoring seems feasible in principle, but documented production deployments are thin (Milojević & Redžepagić, 2021). Practical hurdles include stateful joins across ledgers and collateral systems (Scheinert et al., 2023), P99 latency and throughput targets (Xu et al., 2023), and feature pipelines that won't drift as booking practices evolve (Netti et al., 2022).

2.6 Identified Gaps in the Literature

Three gaps stand out.

Integration: Cash flows, LCR, and shortfall risk are commonly modeled in isolation; few works propose a unified framework that links them end-to-end (Cont et al., 2020).

Real time: Despite mature streaming tools, evidence on genuinely low-latency liquidity monitoring (with live features and on-demand scoring) is limited (Zhang et al., 2009).

Governance: Explainability (Deza et al., 2021), data lineage/quality (Chen, 2021), and model-risk controls (Fritz-Morgenthal et al., 2022) remain underexplored relative to what supervisors are likely to require (Sarker & Bhowmik, 2021).

3. Methodology

3.1 Research Approach

This study took a structured literature-review route (Snyder, 2019) with secondary data analysis (Johnston, 2017). We study scholarly articles, reports from businesses, and publications that deal with regulations side by side to compile our evidence. Wherever possible, we favored sources that reported concrete metrics so performance claims could be compared. We did not clean or link bank-internal data. The goal was breadth and external validity, not institution-specific depth. That choice broadens coverage but may sacrifice some granularity.

3.2 Data Sources

We derived data from four main sources:

Writings for academic journals and books covering topics including liquidity risk, financial time series forecasting, and artificial intelligence/deep learning in the financial sector (Özbayoğlu et al., 2020).

Papers pertaining to regulation, such as those pertaining to Basel III, the European Banking Authority, the Federal Reserve, and other studies published by central banks (KV, 2023).

Case studies, reference designs, and industry white papers on artificial intelligence (AI) for risk functions have been compiled by banks, consultants, and vendors (Jáuregui-Velarde et al., 2023).

The following datasets are available to the public: LCR disclosures, summaries of regulatory stress tests, and market data used to show liquidity trends.

3.3 Data Collection Process

We used a three-stage process:

Identification (2010–present). Searches on SSRN, ScienceDirect, and official regulatory repositories combined terms such as liquidity risk, LCR/NSFR, cash-flow forecasting, run-off, deep learning/LSTM/GRU/Transformer, and real-time/Kafka/Flink.

Screening. Titles/abstracts were checked for relevance to (i) liquidity measurement or forecasting, (ii) AI/ML methods for financial time series, or (iii) real-time/streaming architectures (Hazel et al., 2021; Pérez-Moure et al., 2023). Items lacking empirical content or method detail were set aside; duplicates and superseded drafts were removed.

Categorization. There were five topics that were used to for the sources: (i) classic liquidity models, (ii) statistical forecasting, (iii) applications of AI and ML, (iv) architectures of real-time systems, and (v) governance and interpretability. Within each theme we noted model families, features, validation windows, and evaluation metrics to enable cross-study comparisons.

3.4 Analytical Method

We used comparative synthesis (Schick-Makaroff et al., 2016) to benchmark findings across sources (Bartz–Beielstein et al., 2020). Reported metrics (e.g., LSTM vs. ARIMA error reductions, classification performance for shortfall detection, latency for streaming pipelines) were extracted and viewed against regulatory expectations for liquidity monitoring (e.g., timeliness for LCR oversight, early-warning utility). Evidence from case studies and technical reports was analyzed to judge operational feasibility and to surface regulatory implications around explainability, governance, and model-risk management. Where results depended heavily on design choices (feature sets, retraining cadence, backtest windows), we note that sensitivity explicitly.

Step	Description
Research Approach	Systematic literature review & secondary data analysis
Data Sources	Academic studies, regulatory reports, industry whitepapers, public liquidity disclosures
Data Collection	Identification, screening, categorization into themes (traditional, AI models, real-time, governance)
Analysis	Comparative synthesis of findings; benchmarking AI vs traditional models

Figure 1. Flow chart representing research methodology

4. Results and Discussion

4.1 Predictive Accuracy from Existing Studies

Across the literature, AI models generally beat traditional statistical baselines. For cash-flow forecasting, LSTM and GRU architectures cut error rates by roughly 20–40% relative to ARIMA and GARCH. In LCR estimation, deep learning models more reliably captured drivers like deposit run-off and contingent drawdowns, which translated into tighter forward-looking ratio calculations. Exactly when classical models have a tendency to lose calibration—in the face of turbulent conditions—these benefits seem to expand. That said, performance still hinges on sensible feature design (e.g., segmenting retail vs. wholesale flows), data quality, and retraining cadence (Artificial Intelligence, Machine Learning, and Deep Learning Models for Risk Management, 2022). One of the main reasons why deep learning models outperform conventional approaches when it comes to prediction accuracy is their remarkable capacity to automatically learn and extract complicated characteristics from input data.

Model	Strengths	Limitations	Reported Performance
ARIMA	Simple, interpretable, widely used	Assumes linearity, weak under volatility	Baseline; higher forecast error
GARCH	Captures volatility clustering	Focus on volatility, limited in structural shifts	Baseline; higher forecast error
LSTM	Captures long-term temporal dependencies	Black-box, data intensive, requires tuning	20–40% error reduction vs ARIMA/GARCH
GRU	Computationally efficient with similar benefits to LSTM	Still black-box, less mature in literature	Similar to LSTM; faster training
Transformer	Handles long-range dependencies via attention	Computationally demanding, less adoption in liquidity risk	Promising but less empirical liquidity evidence

Figure 2. Model Comparison

4.2 Evidence from Case Studies

Case studies point to a broader pattern which included techniques that work for credit risk, default prediction (Du & Shu, 2023), or volatility which translated to liquidity monitoring (Heuver, 2020). Several implementations reported detecting likely LCR breaches one to two weeks before traditional stress-testing workflows would have flagged them. For example, an LCRisk measure demonstrated a 48% probability of illiquidity a few days before a bank underwent resolution, providing a crucial early warning signal (Lusby & Stidsen, 2022). The early signal seems to come from models that update continuously as new transactions and market quotes arrive, rather than waiting for end-of-day batches (Heuver, 2020). Transferability isn't automatic, though; label definitions, horizons, and cost of false alarms need to be re-tuned for liquidity use. These models leverage broad sets of liquidity measures and classification techniques for higher predictive accuracy, adapting insights from high-frequency trading to the liquidity domain.

4.3 Real-Time Analytics Infrastructure

Public reports suggest the pipeline is in place. Flink (also known as Spark Streaming) allows for low-latency feature creation and stateful joins, while Apache Kafka is capable of streaming events with a high frequency. These stacks already power fraud screening (Artificial Intelligence, Machine Learning, and Deep Learning Models for Risk Management, 2022) and high-frequency trading, so extending them to liquidity monitoring looks technically feasible (Arner et al., 2017). AI technologies can be integrated into existing fintech systems, such as high-frequency trading platforms and derivative trading platforms, to incorporate liquidity effects into volatility forecasting (Ding et al., 2021). Furthermore, the adoption of cloud technology provides significant benefits in data processing, real-time analytics, and scalability, fundamentally altering how banks manage liquidity risk by centralizing vast amounts of data and enabling sophisticated AI/ML models on demand. The harder parts are operational - guaranteeing P99 latency targets, keeping feature pipelines stable as booking practices evolve, and scaling model serving without accuracy drift. Production-grade examples for liquidity remain scarce (Milojević & Redžepagić, 2021), but nothing in the stack appears to be a showstopper. New treasury ecosystems backed by AI are envisioned to connect banks, suppliers, customers, and regulators in a seamless network of financial intelligence, enabling real-time cash flow forecasting across multiple entities.

4.4 Regulatory and Governance Considerations

Model opacity is the recurring concern. Complex neural nets can feel like “black boxes,” which complicates supervisory review, especially when outputs inform LCR decisions (Maple et al., 2023). Post-hoc tools (e.g., SHAP, LIME) help explain feature influence but may not fully satisfy expectations for traceable, auditable logic (Fritz-Morgenthal et al., 2022). Data governance matters just as much.

To prevent undetected deterioration, it is crucial to have documented retraining processes, evaluate quality, monitor drift, and look at lineage (Kurshan et al., 2020; Maple et al., 2023) Fewer false positives and more fruitful oversight conversations were observed by institutions that used a combination of main and challenger models, alert levels, and transparent escalation playbooks (Fritz-Morgenthal et al., 2022).

4.5 Synthesis of Findings

Enhanced predictive accuracy. Deep learning models tend to outperform classical statistics, with the advantage most visible in volatile regimes. Quantitative evidence indicates improvements ranging from 20-40% in forecast error reduction and significant gains in predictive accuracy. Results, however, depend on data quality, feature choices, and retraining discipline. Real-time is feasible. Modern streaming infrastructure can support continuous liquidity surveillance when tightly integrated with predictive models. Cloud platforms are often used in this integration to analyze high-frequency data and give real-time insights. This is achieved by linking AI models to current finance and treasury management systems. Problems often arise with dependability and integration rather than with raw technology. Deficiencies persist. It is unusual to find completely integrated, real-time frameworks in the literature that simulate cash flows, LCR, and shortfall risk simultaneously. Still, the most significant obstacles to regulatory acceptance are model-risk practices (Kurshan et al., 2020), governance, and explainability (Fritz-Morgenthal et al., 2022).

5. Contributions and Limitations

5.1 Contributions of the Study

This paper adds to the liquidity-risk and fintech literature in four ways:

Bringing the pieces together. Prior work tends to treat cash-flow forecasting, LCR estimation, and shortfall detection as separate problems. This review pulls them into a single lens, showing how an AI-enabled framework could address all three at once and why that matters for day-to-day monitoring (Cont et al., 2020).

Real-time feasibility, not just theory. Drawing on technical reports and case evidence, we show that streaming stacks (e.g., Kafka + Flink) can plausibly support low-latency model serving for liquidity use cases (Heuver, 2020). The open questions are integration and ownership rather than raw capability.

Regulatory and governance through-lines. The review connects technical gains to supervisory expectations highlighting that adoption will likely hinge as much on governance design as on model accuracy (Fritz-Morgenthal et al., 2022).

Input to the stability debate. By benchmarking AI results against traditional approaches, the paper offers a starting point for practitioners and regulators weighing systemic implications of moving from static ratios to adaptive, model-driven monitoring (Maple et al., 2023).

5.2 Limitations of the Study

The limitations of the study are:

Reliance on secondary sources. Published academic, regulatory, and business articles were used to gather the data. That option increases coverage overall but restricts findings that are particular to certain institutions (Achter et al., 2023).

Heterogeneous study designs. Metrics, horizons, and datasets vary across sources, which challenges like-for-like comparisons. We therefore emphasize directional results and effect sizes rather than exact point estimates (Chow et al., 2023).

Limited operational validation. While real-time stacks are well documented in adjacent domains (fraud, trading), there are few production-grade examples for liquidity monitoring itself. The gap between technical feasibility and sustained live performance remains to be closed (Heuver, 2020).

Changing regulatory policies. Changes are ongoing in the expectations around model-risk management, explainability, and governance. Current recommendations are included in this assessment; however, future requirements for supervision or enforcement techniques cannot be entirely anticipated (Maple et al., 2023).

6. Conclusion

6.1 Summary of Findings

On balance, the literature points in the same direction: deep learning tends to beat classical baselines on financial time-series tasks that matter for liquidity (Michańków et al., 2023). Studies report meaningful error reductions on cash-flow forecasts (Dadteev et al., 2020) and better estimates of LCR components such as deposit run-offs and contingent outflows. Several case reports also show earlier flags for potential shortfalls—sometimes a week or two ahead of what static stress tests would have caught (KV, 2023). On the plumbing side, event-driven stacks like Kafka with stateful processors such as Flink appear technically capable of supporting continuous monitoring (Kontaxakis et al., 2021), even if production use for liquidity remains relatively rare (Lee et al., 2020).

6.2 Implications for Practice

The takeaway for financial institutions is the need to unify their disparate models into a unified, future-oriented perspective. When fed low-latency data (such as intraday deposits, credit-line drawdowns, collateral haircuts), merging cash-flow projections, LCR trajectories, and shortfall alarms into a single process might improve accuracy and accelerate decision-making. Having said that, a handful of less glamorous aspects are likely to determine success: feature pipelines that remain stable (Bohlke-Schneider et al., 2020), thresholds that are adjusted according to the monetary consequences of false alarms for the company, and dashboards that enable risk teams to easily go from top-line LCR to segment-level drivers with a single click. Interpretability remains the biggest sticking point; without auditable reasoning behind alerts, supervisors and risk committees will be hesitant (Fritz-Morgenthal et al., 2022).

6.3 Implications for Policy and Regulation

For policymakers, the promise is earlier, more reliable warning signals that could improve system resilience. The challenge, however, is governance: complex models can look like black boxes (Mirestean et al., 2021). Standards may need to evolve to spell out documentation, challenger-model expectations, periodic back-testing, and explainability sufficient for decisions tied to LCR compliance (Fritz-Morgenthal et al., 2022). Clear guidance on data lineage, retraining cadence, and escalation playbooks would lower adoption friction while keeping safeguards intact.

6.4 Directions for Future Research

Three research gaps stand out, but they differ in urgency. The most immediate need is for integrated frameworks that jointly forecast cash flows, Liquidity Coverage Ratios (LCR), and shortfalls, as most studies still treat these dimensions separately; testing such frameworks head-to-head against current supervisory practice would provide the clearest evidence of added value (Heuver, 2020). The second priority is to generate operational evidence on real-time deployment. While conceptual diagrams of streaming architectures abound, rigorous demonstrations of latency budgets, stateful joins across systems, and drift monitoring in regulated environments remain scarce, even though liquidity crises can escalate within hours (Heuver, 2020). The third area of focus is explainable AI. This field has to go beyond post-hoc visualizations to approaches that supervisors can put into practice, including reliable local and global attributions, diagnostics that are in line with policy triggers, and so on (Fritz-Morgenthal et al., 2022). Public standards, such as shared datasets and defined prediction horizons, are essential to all three goals because they provide credible technique comparisons and speed up the process of translating research into practice (Arner et al., 2017).

References

1. Achter, S., Borit, M., Cottineau, C., Meyer, M., Polhill, G., & Radchuk, V. (2023). How to conduct more systematic reviews of agent-based models and foster theory development: Taking stock and looking ahead. Environmental Modelling & Software, 173, 105867. https://doi.org/10.1016/j.envsoft.2023.105867
2. Aniceto, M. C., Barboza, F., & Kimura, H. (2020). Machine learning predictivity applied to consumer creditworthiness. Future Business Journal, 6(1), 1–12. https://doi.org/10.1186/s43093-020-00041-w

3. Arner, D., Barberis, J., & Buckley, R. (2017). FinTech, RegTech, and the reconceptualization of financial regulation. Northwestern Journal of International Law & Business, 37(3), 371–413. http://hub.hku.hk/bitstream/10722/235796/1/content.pdf
4. Artificial intelligence, machine learning, and deep learning models for risk management (pp. 31–53). (2022). In Handbook of Financial Data Science. Wiley. https://doi.org/10.1002/9781119824961.ch3
5. Aziz, S., & Dowling, M. (2018). Machine learning and AI for risk management. In Palgrave studies in digital business & enabling technologies (pp. 33–55). Springer. https://doi.org/10.1007/978-3-030-02330-0_3
6. Barongo, R. I., & Mbelwa, J. (2023). Using machine learning for detecting liquidity risk in banks. Machine Learning with Applications, 15, 100511. https://doi.org/10.1016/j.mlwa.2023.100511
7. Bartz-Beielstein, T., Doerr, C., Bossek, J., Chandrasekaran, S., Eftimov, T., Fischbach, A., Kerschke, P., López-Ibáñez, M., Malan, K. M., Moore, J. H., Naujoks, B., Orzechowski, P., Volz, V., Wagner, M., & Weise, T. (2020). Benchmarking in optimization: Best practice and open issues. arXiv. https://doi.org/10.48550/arXiv.2007.03488
8. Bohlke-Schneider, M., Kapoor, S., & Januschowski, T. (2020). Resilient neural forecasting systems. Proceedings of the 4th International Workshop on Data Management for End-to-End Machine Learning (DEEM'20), 4, 1–5. ACM. https://doi.org/10.1145/3399579.3399869
9. Carcillo, F., Pozzolo, A. D., Borgne, Y. L., Caelen, O., Mazzer, Y., & Bontempi, G. (2017). SCARFF: A scalable framework for streaming credit card fraud detection with Spark. Information Fusion, 41, 182–194. https://doi.org/10.1016/j.inffus.2017.09.005
10. Chen, J. (2021). Ontology drift is a challenge for explainable data governance. arXiv. https://doi.org/10.48550/arXiv.2108.05401
11. Chen, W., Hussain, W., Cauteruccio, F., & Zhang, X. (2023). Deep learning for financial time series prediction: A state-of-the-art review of standalone and hybrid models. Computer Modeling in Engineering & Sciences, 139(1), 187–218. https://doi.org/10.32604/cmes.2023.031388
12. Cont, R., Kotlicki, A., & Valderrama, L. (2020). Liquidity at risk: Joint stress testing of solvency and liquidity. Journal of Banking & Finance, 118, 105871. https://doi.org/10.1016/j.jbankfin.2020.105871
13. Dadteev, K., Shchukin, B., & Nemeshaev, S. (2020). Using artificial intelligence technologies to predict cash flow. Procedia Computer Science, 169, 264–270. https://doi.org/10.1016/j.procs.2020.02.163
14. Deza, G., Travers, A., Rowat, C., & Papernot, N. (2021). Interpretability in safety-critical financial trading systems. arXiv. https://doi.org/10.48550/arXiv.2109.15112
15. Ding, S., Cui, T., Zhang, Y., & Li, J. (2021). Liquidity effects on oil volatility forecasting: From fintech perspective. PLoS ONE, 16(11), e0260289. https://doi.org/10.1371/journal.pone.0260289
16. Dionne, G., Pacurar, M., & Zhou, X. (2015). Liquidity-adjusted intraday value at risk modeling and risk management: An application to data from Deutsche Börse. Journal of Banking & Finance, 59, 202–219. https://doi.org/10.1016/j.jbankfin.2015.06.005
17. Du, P., & Shu, H. (2023). Design and implementation of China financial risk monitoring and early warning system based on deep learning. IEEE Access, 11, 78052–78065. https://doi.org/10.1109/access.2023.3280934
18. Farooq, M. B., Hunjra, A. I., Ullah, S., & Al-Faryan, M. A. S. (2023). The determinants of financial distress cost: A case of emerging market. Cogent Economics & Finance, 11(1), 2186038. https://doi.org/10.1080/23322039.2023.2186038
19. Fritz-Morgenthal, S., Hein, B., & Papenbrock, J. (2022). Financial risk management and explainable, trustworthy, responsible AI. Frontiers in Artificial Intelligence, 5, 779799. https://doi.org/10.3389/frai.2022.779799
20. Fu, Y., & Soman, C. (2021). Real-time data infrastructure at Uber. Proceedings of the 2021 International Conference on Management of Data (SIGMOD'21), 2503–2509. ACM. https://doi.org/10.1145/3448016.3457552
21. Gontarska, K., Geldenhuys, M. K., Scheinert, D., Wiesner, P., Polze, A., & Thamsen, L. (2021). Evaluation of load prediction techniques for distributed stream processing. Proceedings of the 2021 IEEE International Conference on Cloud

Engineering (IC2E), 91–101. https://doi.org/10.1109/IC2E52221.2021.00023

22. Greco, L., Ritrovato, P., & Xhafa, F. (2018). An edge-stream computing infrastructure for real-time analysis of wearable sensors data. Future Generation Computer Systems, 93, 515–528. https://doi.org/10.1016/j.future.2018.10.058

23. Hazel, C. A., Bull, S., Greenwell, E., Bunik, M., Puma, J., & Perraillon, M. C. (2021). Systematic review of cost-effectiveness analysis of behavior changes communication apps: Assessment of key methods. Digital Health, 7, 20552076211000559. https://doi.org/10.1177/20552076211000559

24. Heuver, R. (2020). Applications of liquidity risk discovery using financial market infrastructures transaction archives. Other Publications TiSEM. Tilburg University. https://doi.org/10.26116/center-lis-2007

25. Jarquin, D. V., Sharp, A. J., Zhang, R., Nguyen, T., Hoang, A. T., & Pianykh, O. S. (2022). Temporal quality degradation in AI models. Scientific Reports, 12(1), 12834. https://doi.org/10.1038/s41598-022-15245-z

26. Jáuregui-Velarde, R., Andrade-Arenas, L., Molina-Velarde, P., & Yactayo-Arias, C. (2023). Financial revolution: A systemic analysis of artificial intelligence and machine learning in the banking sector. International Journal of Electrical and Computer Engineering, 14(1), 1079–1090. https://doi.org/10.11591/ijece.v14i1.pp1079-1090

27. Johnston, M. P. (2014). Secondary data analysis: A method of which the time has come. Qualitative and Quantitative Methods in Libraries, 3(3), 619–626. http://www.qqml.net/papers/September_2014_Issue/336QQML_Journal_2014_Johnston_Sept_619-626.pdf

28. Kalogerakis, S., Papaioannou, A., & Magoutis, K. (2022). Efficient processing of high-volume tick data with Apache Flink for the DEBS 2022 grand challenge. Proceedings of the 16th ACM International Conference on Distributed and Event-Based Systems (DEBS'22), 133–138. https://doi.org/10.1145/3524860.3539649

29. Kim, A., Yang, Y., Lessmann, S., Ma, T., Sung, M., & Johnson, J. E. V. (2020). Can deep learning predict risky retail investors? European Journal of Operational Research, 283(1), 217–234. https://doi.org/10.1016/j.ejor.2019.11.007

30. Kontaxakis, A., Deligiannakis, A., Arndt, H., Burkard, S., Kettner, C.-P., Pelikan, E., & Noack, K. (2021). Real-time processing of geo-distributed financial data. Proceedings of the 2021 ACM Symposium on Cloud Computing, 190–204. https://doi.org/10.1145/3465480.3467842

31. Kulatilleke, G. K. (2022). Credit card fraud detection: Classifier selection strategy. arXiv. https://doi.org/10.48550/arXiv.2208.11900

32. Kurshan, E., Shen, H., & Chen, J. (2020). Towards self-regulating AI: Challenges and opportunities of AI model governance in financial services. arXiv. https://arxiv.org/abs/2009.09052

33. Lee, M. S. A., Floridi, L., & Denev, A. (2020). Innovating with confidence: Embedding AI governance and fairness in a financial services risk management framework. SSRN Electronic Journal. https://doi.org/10.2139/ssrn.3830579

34. Liu, H., & Long, Z. (2020). An improved deep learning model for predicting stock market price time series. Digital Signal Processing, 102, 102741. https://doi.org/10.1016/j.dsp.2020.102741

35. Lusby, R. M., & Stidsen, T. (2022). Mathematical programming with Julia: New open-source book. Research Portal Denmark. https://local.forskningsportal.dk/local/dki-cgi/ws/cris-link?src=dtu&id=dtu-9165028e-8bfa-42d3-911d-2f31b993a117

36. Maple, C., Szpruch, Ł., Epiphaniou, G., Staykova, K., Singh, S. B., Penwarden, W., Wen, Y., Wang, Z., Hariharan, J., & Avramović, P. (2023). The AI revolution: Opportunities and challenges for the finance sector. arXiv. https://doi.org/10.48550/arXiv.2308.16538

37. Mariani, M. C., Bhuiyan, M. A. M., Tweneboah, O. K., González-Huízar, H., & Florescu, I. (2018). Volatility models applied to geophysics and high frequency financial market data. Physica A: Statistical Mechanics and Its Applications, 503, 304–316. https://doi.org/10.1016/j.physa.2018.02.167

38. Michańków, J., Kwiatkowski, Ł., & Morajda, J. (2023). Combining deep learning and GARCH models for financial volatility and risk forecasting. SSRN Electronic Journal. https://doi.org/10.2139/ssrn.4589950

39. Milojević, N., & Redžepagić, S. (2021). Prospects of artificial intelligence and machine learning application in banking risk management. Journal of Central Banking Theory and Practice, 10(3), 41–62. https://doi.org/10.2478/jcbtp-2021-0023

40. Mirestean, A., Farias, A., Deodoro, J., Boukherouaa, E. B., AlAjmi, K., Iskender, E., Ravikumar, R., & Shabsigh, G. (2021). Powering the digital economy: Opportunities and risks of artificial intelligence in finance. IMF Departmental Paper, 2021(24), 1–42. https://doi.org/10.5089/9781589063952.087

41. Ndikum, P. (2020). Machine learning algorithms for financial asset price forecasting. arXiv. https://doi.org/10.48550/arXiv.2004.01504

42. Netti, A., Ott, M., Guillén, C., Tafani, D., & Schulz, M. (2022). Operational data analytics in practice: Experiences from design to deployment in production HPC environments. Parallel Computing, 113, 102950. https://doi.org/10.1016/j.parco.2022.102950

43. Özbayoğlu, A. M., Gudelek, M. U., & Sezer, Ö. B. (2020). Deep learning for financial applications: A survey. Applied Soft Computing, 93, 106384. https://doi.org/10.1016/j.asoc.2020.106384

44. Pérez-Moure, H., Lampón, J. F., Velando-Rodriguez, M.-E., & Rodríguez-Comesaña, L. (2023). Revolutionizing the road: How sustainable, autonomous, and connected vehicles are changing digital mobility business models. European Research on Management and Business Economics, 29(3), 100230. https://doi.org/10.1016/j.iedeen.2023.100230

45. Phien, N. N., & Platoš, J. (2023). The PSR-transformer nexus: A deep dive into stock time series forecasting. International Journal of Advanced Computer Science and Applications, 14(12), 68–76. https://doi.org/10.14569/ijacsa.2023.0141292

46. Rubio, L., Pinedo, A. P., Castaño, A. M., & Ramos, F. (2023). Forecasting volatility by using wavelet transform, ARIMA and GARCH models. Eurasian Economic Review, 13, 803–823. https://doi.org/10.1007/s40822-023-00243-x

47. Safiyari, M. H., & Nabati, H. (2023). Liquidity risk in banking: Integrating two hybrid machine learning models to introduce a comprehensive risk assessment system. SSRN Electronic Journal. https://doi.org/10.2139/ssrn.4677575

48. Sarker, N., & Bhowmik, P. K. (2021). Bank liquidity risk: Significance of financial disclosure and governance practice. Asian Economic and Financial Review, 11(9), 724–744. https://doi.org/10.18488/journal.aefr.2021.119.724