



Generative AI DevOps in Real-Time Finance

Christopher Wilson
Asst Professor

Abstract

The emergence of Generative AI (GenAI) unlocks new capabilities in software development and operations. A related trend calls for the adoption of DevOps principles and practices. Furthermore, organizations across industries seek to process data in real time. In financial services, that imperative applies to transaction pipelines, where volumes peak at Black Swan events, yet the average workload is relatively modest. Ideas from GenAI-enabled DevOps can accelerate the design, implementation, and monitoring of such pipelines, but tensions remain. The landscape of existing work is therefore reviewed, gaps identified, Generalized Enterprise Architecture Theory summoned, and research questions posed. Addressing them will yield candidate solutions that reconcile GenAI-enabled DevOps principles with the requirements of real-time transaction processing. A foundation for that exploration is laid by: specifying the functional and nonfunctional requirements for such pipelines; providing an overview of GenAI-driven DevOps concepts, architectures, and operational paradigms amenable to financial Pipelines; and articulating suitable patterns, together with the associated challenges and trade-offs.

Real-time transaction processing pipelines in financial services must support Black Swan event volumes while meeting operational requirements at lower loads. Ideas from GenAI-enabled DevOps can accelerate the design, implementation, and monitoring of such pipelines, but tensions remain. Existing work is reviewed, gaps identified, Generalized Enterprise Architecture Theory summoned, and research questions posed. Addressing them will yield candidate solutions reconciling GenAI-enabled DevOps principles with the requirements of real-time financial transaction processing. A foundation for that exploration is laid by specifying the functional and nonfunctional requirements for such pipelines; providing an overview of GenAI-driven DevOps concepts, architectures, and operational paradigms amenable to financial pipelines; and articulating suitable patterns, together with the associated challenges and trade-offs.

Keywords: Generative Artificial Intelligence, GenAI Enabled DevOps, Real Time Transaction Processing, Financial Services Pipelines, Black Swan Event Handling, Enterprise Architecture Theory, Cloud Native Architectures, Transaction Pipeline Scalability, DevOps Automation Practices, Real Time Data Processing, Operational Resilience Engineering, Non Functional Requirements Analysis, Intelligent Pipeline Monitoring, Adaptive System Design, Financial Infrastructure Modernization, Event Driven Architectures, High Volume Transaction Systems, AI Assisted Software Operations, Architectural Patterns And Trade Offs, Enterprise Grade Transaction Systems.

1. Introduction

Advancements in Generative Artificial Intelligence (GenAI) and DevOps enable automation across the entire lifecycle of pipelines for online detection of fraudulent activities in real-time financial transaction processing. Historical lessons and challenges in implementing data-

driven GenAI models, organizations' needs for accelerated product delivery, and rising cybercrime costs demand automation in operationalization as functionally as possible. GenAI-Enabled DevOps for Real-Time Financial Transaction Pipelines addresses GenAI-enabled DevOps for real-time financial transaction pipelines, characterizing stakeholders, scoping DevOps automation in control,

scheduling, and security, and presenting architecture patterns. Considered deployment scenarios and architectures apply specifically to online detection of fraudulent financial transactions, with a broader lens on operationalization-as-a-function of organizations' data pipelines. Real-time data processing has seen increasingly adoption for multiple usage scenarios, especially by financial institutions, enabling them to detect and respond to events that require operational or business action explaining GenAI-driven DevOps concepts, architectures, and operational paradigms suitable for financial data pipelines. Automation capabilities across pipeline operations—model training and deployment, pipeline continuous monitoring, automated rollback, security-governed monitoring, and trade-offs—simultaneously support multiple use-cases, such as anomaly detection for security incidents, fraud detection for financial transactions, and prediction of user behaviours. Numerous architectures for secure, scalable, and performance-optimized implementation of data processing pipelines with controls on data access and governance address these challenges, focusing mainly on data processing.

1.1. Overview of the Study

Automation across the GenAI model lifecycle enables the construction of DevOps-driven data transaction pipelines for inference and decision-making tasks whose latency, operational knowledge, and expected service horizons correlate. Although it is rare to parallelize model prediction across multiple units with independent data, since the data must be routed through different branches according to its classification, high-frequency decision-making can also be achieved. These models are often known as Single-Task Models. Application in the banking and insurance industries requires special attention to Model Governance since profound consequences from decision errors or a breach in trade-off rules can occur; in those situations where a clear violation of the Business Policy occurs, the reason is usually also evident, so an adequate level of Explainability can be achieved and its adequate monitoring is less critical. The GenAI-DevOps principles generally

enable reasonably precise detection of potential errors in a GenAI-driven pipeline. Implementation in GenAI-DevOps is nevertheless sensitive to several factors. For missions demanding low latency and high throughput and, in particular, for those using recurrent models, the size of the model, the required speed of the inference rate, and the type and frequency of the monitoring checks introduce a complex trade-off. A parallelized model or an Externalized Model must avoid excessive size and complexity, while maintaining low cost and duration of inference, monitoring checks should preferably be executed during non-productive periods and by a less powerful and economical infantry non accessed by GenAI Decision-Making during production hours.




Fig 1: Optimizing the GenAI-DevOps Lifecycle: A Framework for High-Frequency Inference, Latency Trade-offs, and Governance in Financial Decision Systems

2. Background and Related Work

Development and Operations (DevOps) shares with GenAI the ambition of accelerating release cycles while improving quality. Operational characteristics further inspire the adaptation of GenAI to gain efficiency, minimize downtime, and dynamically scale service resources. Integration of open-source GenAI into Hybrid Cloud production environments is viewed as a likely avenue: data ownership, residence, and sovereignty concerns associated with Third-Party Clouds can be ameliorated through a

Hybrid Cloud architecture that integrates the Cloud with localized Edge Machines for workloads sensitive to latency. More generally, Data Engineering and Event Sourcing for real-time processing are important domains that witness the emergence of GenAI tools and concepts. Milestones in their evolution into aligned domains are described, followed by an assessment of architectures, principles, and empirical realizations of DevOps and GenAI, as well as real-time processing. Real-time systems and their pipelines impose strong constraints on availability and response time; major incidents have starkly demonstrated the dire consequences of failure. Misuse of erroneous models can be catastrophic; automatic retraining models with new data and technologies that replace smart humans are, indeed, the path to disaster. Wariness of blithe acceptance of GenAI suggestions is warranted, and careful consideration is required to reduce their operation risk.



Equation 1: Latency & Throughput Targets — complete derivations

1.1 Transactions per second (TPS)

Let:

- N = number of transactions processed in time window T seconds

Definition

$$TPS = \frac{N}{T}$$

Visa up to 65,000 TPS with average latency 36 ms .

1.2 Throughput (MB/s) and conversion to TPS

Let:

- B = throughput in MB/s
- S = average message size in MB/transaction

Then transactions per second:

$$TPS = \frac{B}{S}$$

Step-by-step

1. MB per second divided by MB per transaction gives transactions per second:

$$\frac{\text{MB}}{\text{s}} \div \frac{\text{MB}}{\text{txn}} = \frac{\text{txn}}{\text{s}}$$

2.1. Historical Insights and Key Developments

The maturity of every social, economic, and technological aspect can be seen from the DevOps and GenAI literature is undeniable. The FiTS e-architecture is supported with continuously evolving industrial services. Although there is a lot of innovation happening at various fronts of business verticals, the journey of FiTS is the most appropriate to review for its complexity, integration, security and regulatory requirements. The key financial regulations like PCI-DSS, GDPR, BASEL III, Dodd-Frank, AML, BSA, the list goes on, affects every financial institution around the world. Many services required for these regulations are being offered based on Forms-of-subscribing. Data Processing-as-a-Service (DPaaS), Model-Lifecycle as a Service (MLaaS), Ground-motion-model as a Service (GMMaaS), Model as a Service (MaaS), Log-Monitor-as-a-Service (LMMaaS), Data-collection and Validation-as-a-Service (DCVaaS), Software as a Service (SaaS), Detect and Display (D&D) as a Service, Detect and Take Action (DTA) as a Service, Detect and Surface (D&S) as a

Service, etc., relieve SMEs from any capital investment and allow them to remain competitive with any high-cost subscription for complex services. Standards and tools allow integration with better security and reduced effort. Three architectural styles catering to specific transaction processing speed requirements have matured and have become reference models for Financial-as-a-Service (FaaS). Financial information exchange, the use of ATMs, and online services for transferring money, payments, stock-trading, etc., has integrated the services required to fulfil mandates. These service compositions are essentially integration of existing-monitoring-and-detection-and-display services. Retail a and b evolve from these methods. Special channels are also formed between the bank and designated customers based on the high volume of transactions in order to serve specialized needs. Financial institutions are offering relational database services, and point-of-interest authorisation, processing, and settlement services are moving towards Cloud. The SLAs of these services, and real-world patterns of market data motion have enabled the COTS products like CORBA/IIOP/DBPSRL demand-pull-algorithms.

3. Requirements for Real-Time Financial Transaction Pipelines

A comprehensive overview of the functional and nonfunctional requirements for real-time financial transaction pipelines is presented. Data correctness, integrity, security, and regulatory compliance are highlighted as pivotal nonfunctional requirements. A clear treatment of latency targets and throughput metrics aligns the discussion with concrete service-level objectives; requirements for data provenance, auditing, and observability establish criteria for external and internal governance. Together, requirements for consistency models, fault tolerance schemes, idempotency, replay protection, and end-to-end reliability guarantees address the preservation of nonfunctional properties during fault recovery.

Functional and nonfunctional requirements for real-time financial transaction pipelines have far-reaching

implications yet are rarely articulated with sufficient care. Service-level objectives for latency, throughput, and rollback time influence the choice of architectural, operational, and governance design decisions. While functional and operational design considerations are often partitioned during design, effective implementation of the requirements requires integration awareness from both perspectives. Implementation trade-offs frequently necessitate intra- and cross-discipline compromises between security, data sensitivity, operational complexity, auditability, oversight, and environmental impact.

3.1. Latency and Throughput Targets

Demanding latency and throughput objectives are apparent from the governance policies and service-level agreements of many real-time financial processing environments. For example, Visa's network supports up to 65,000 transactions per second (TPS), with an average latency of 36 milliseconds. An internal study at JPMorgan Chase indicated that a latency of over 10 milliseconds during card authorization leads to a revenue drop. Data stream operators usually define their processing model and workflow based on the anticipated maximum rate of each data source, which serves as the upper input rate boundary for the processing workflow at scale.

Data streaming platforms, such as Apache Kafka and Confluent, define a specific target throughput of MB/s related to each broker. The end-to-end latency of data streams is usually associated with the message-processing delay of the various consumers and the producer acknowledgment time. The different parts of the network have different latencies, which leads to the concept of "worst-case" end-to-end latency that can be computed using the equations of series arrangement in logic circuits. Within the context of model reproduction, a key metric in the DevOps lifecycle is the amount of time it takes to complete a round of continuous training.




Fig 2: Quantifying the Latency-Revenue Nexus: Scalable Stream Architectures and Worst-Case Performance Modeling in High-Frequency Financial Networks

3.2. Consistency and Fault Tolerance

Real-time financial systems are severely affected by node or region failures. In a Service Level Agreement or Service Level Objective-driven environment, it is paramount to minimize recovery windows, and therefore advanced, nontrivial, and often customized solutions are usually adopted. As highlighted in prior work, solutions consist of data duplication, partitioning, and replication across regions and clouds to address region or cloud failures. However, it is also important to focus on the end-to-end implementation guaranteeing data consistency, completing the process without generating duplicates or packet drops, and achieving exactly-once semantics.


Managing in-flight data is critical for its correct management; therefore, mechanisms to identify the data processing path, such as sequence identifiers, and capabilities to replay in-flight data during a failure period must be implemented. Also, for proper consumption on the receiving side, it may be required that each message contains the complete state of a certain element (for instance, the last balance held by an account) or that

messages define a sequence that permits reconstructing the information on the other side.

4. GenAI-Driven DevOps: Concepts and Paradigms

This section explains concepts, architectural patterns, and operational paradigms of GenAI-enabled DevOps that adapt to real-time financial transaction pipelines. These elements support automation throughout the pipeline lifecycle, from model training and deployment to monitoring, rollback, and governance. GenAI capabilities enable a range of automation solutions for GenAI and non-GenAI models. Most tasks—except model training and some aspects of deployment—are executed in production environments; Model monitoring techniques are integrated with standard performance, security, and resiliency monitoring procedures.

GenAI- and non-GenAI-based business-specific models are trained and validated before being deployed in their runtime environments. Model inputs are continuously assessed by standard data pipelines. Underperformance triggers dedicated retraining pipelines to create new model versions. DevOps returns these models to production through automated deployment processes that perform the necessary rollout and rollback actions. Model validation is similarly automated; the processes detect functional segmentation and quality discrepancies between new and incumbent models and govern making them effective. These procedures facilitate continuous model validation with low overhead and support online learning when appropriate.



Equation 2: Queueing latency under peak load (Black Swan) — step-by-step

A standard minimal queueing model to show “why latency explodes near capacity” is **M/M/1**.

Let:

- arrival rate λ (TPS)
- service rate μ (TPS), where $\mu > \lambda$

Expected time in system

$$W = \frac{1}{\mu - \lambda}$$

Derivation sketch (stepwise)

- Utilization:

$$\rho = \frac{\lambda}{\mu}$$

- Expected number in system for M/M/1:

$$L = \frac{\rho}{1 - \rho}$$

- Little’s Law $L = \lambda W$ gives:

$$W = \frac{L}{\lambda} = \frac{\rho}{\lambda(1 - \rho)}$$

- Substitute $\rho = \lambda/\mu$:

$$W = \frac{\lambda/\mu}{\lambda(1 - \lambda/\mu)} = \frac{1}{\mu - \lambda}$$

4.1. Automation in Pipeline Lifecycle

Broad automation—covering the entire financial transaction pipeline lifecycle—is an indispensable element of GenAI-enabled DevOps. It encompasses the development, deployment, monitoring, rollback, and governance of the GenAI models that serve the pipeline. Automation of model training and deployment is relatively straightforward, with platforms like Amazon SageMaker Stream Processing facilitating both activities. Monitoring, however, is more complex; a monitoring solution should enable not only observability into model predictions, but also generation of alerts when production predictions begin to diverge from production labels or when confidence in predictions falls below a specified threshold. Sufficient confidence is critical, as inference costs usually increase with model size—and delayed attention to a failing model may compound operational risk. Where possible, such monitoring should also support early warning for other pipeline components, enabling corrective measures (like alternative processing paths) that minimize short-term disruption.

Automation of model rollback may also be of interest, especially in scenarios where a failed prediction causes higher downstream costs than simply waiting for a return to service. Additionally, prediction monitoring may yield sufficient lead time for supervised model retraining; any such retraining should incorporate the labels accompanying the production predictions. Governance automation facilitates quality-assured model updates by operationalizing traditional data science governance, testing, and monitoring procedures. A pivotal aspect here is the dimensionality of the training data: larger data volumes lead to larger models, which add cost and risk to inference.

5. Architectural Patterns for GenAI-Enabled Pipelines

A set of architectural patterns emerges from the preceding analysis and can accommodate GenAI-driven automation of the full pipeline lifecycle: (i) Hybrid Cloud and Edge Architecture Patterns define the integration of Hybrid and Edge-Cloud infrastructures with the application workload, (ii) Streaming Ingestion and Processing Patterns outline the design aspects of streaming data ingestion and processing for the pipeline.

Hybrid Cloud and Edge Integration Patterns establish the topology and placement of workload execution. Depending on the traffic flows, processing workloads and data regulatory and sovereignty requirements, ingestion and processing may be executed using an edge-local solution or engineered in collaboration with the Hybrid Cloud, using the edge as an accelerator that moves data close to the Cloud resources that actually process it. The Streaming Ingestion Pattern governs the real-time ingestion of streaming data flows, focusing on the definition of the input event schema, control and coordination mechanisms for event windowing, any required event enrichment, stateful processing with updates to internal state stores, idempotent writes to external systems, and requirements for exactly-once processing guarantees.





Fig 3: Integrated Architectural Patterns for GenAI Lifecycle Management: Synergizing Hybrid Edge-Cloud Topologies with Stateful Stream Orchestration

5.1. Hybrid Cloud and Edge Integration

Hybrid Cloud infrastructures enable the best of both worlds by combining the benefits of Public Clouds (unlimited scalability, low cost, diversity of services) with the advantages of Private Clouds (cost control, exclusivity, Security, and Compliance). Many financial institutions leverage a Hybrid Cloud setup to store non-sensitive data on the Public Cloud or leverage services such as Natural Language Processing, Optical Character Recognition, or SenseMaking that have a high potential cloud cost due to large Resource Management requirements. Workloads are offloaded to the Public Cloud when it is cheaper than using the on-Premise Infrastructure. Placing the Transaction Pipeline either fully in the Private Cloud or Partitioned with parts on the Public Cloud makes sense from a Cost Efficiency and Risk Compliance point of View. Certain Data Paths, such as Data Classification, Money Laundering Monitoring, Fraud Detection, are however Priority Cloud Paths for these Financial Institutions and are required to execute locally to Safeguard Sensitive Data or to Ensure Low Latency Operation. Executing on Private Edge Sites leverages the low-latency and low-continuous cost Benefits of Edge Processing while maintaining the Integrity of the Data within the Organization. Thus, a Subset of Workload is Specified to run on the Edge.



5.2. Streaming Data Ingestion and Processing

A typical architectural pattern for continuously ingesting streaming data and performing real-time analysis and processing is shown in . Such architectures form the foundation of most GenAI-enabled application scenarios that are not one-shot transactions but require continuous interaction with external users or systems and, consequently, have a large number of transactions not all requiring GenAI model inference capabilities. Data is continuously ingested from one or multiple sources, analyzed, transformed, and possibly enriched using model inference along the way. The results from these analyses, transformations, and enrichments are streamed to one or more systems. While the simplest instantiation of this pattern treats all incoming data as atomic transactions, the pattern becomes far more useful when support for evolving schemas and throwing away, modifying, or adding some of the ingested messages over time is incorporated. Streaming ingestion can be driven by a variety of events and trigger different types of processing. Very common schema types for ingestion are messages from message buses, traffic and sensor data, logs, and clickstreams. Data can also be ingested in a batched manner while still being processed with millisecond resolutions. Every ingesting message has a schema associated with it. The processing component can be internally merged into one component or distributed over many for both performance and scaling reasons. It can process different windows of the data, re-ingest messages into a message bus once processed, or

route them to different processing components. Different processing components can operate on the same data using different windows or separated amounts of data. Windows can be generated using sliding time windows, tumbling time windows, or session-driven windows. Windows processing with complete and update models is also commonplace.

6. Challenges and Trade-offs

Concrete challenges and trade-offs surface throughout pipeline design and operation. Different design aspects must be carefully balanced to avoid misguided decision-making. For instance, supporting strict latency and high throughput in an automated high-volume fraud detection application using 10 challenger models covered twenty slices of the prediction phase, adding significant overhead. Motivation for the additional latency and cost stemmed from the fact that a single successful fraud attempt, typically demanding at least in the low five-part sum, could have high impact on the business.

User acceptance, maintenance, monitoring, and continuous assessment of model quality become increasingly challenging at scale, and the need for governance, security, and explainability increases with the size of the generative models. Manpower-demanding complex models, such as text-to-code or text-to-image systems, enable business transformation. However, negligence of all three aspects of MLOps could result in high operational risk or worse business outcomes than with simpler, well-governed, and comprehensively monitored solution models.

Equation 3: Consistency, replay protection, exactly-once — equations/logic

3.1 Idempotent processing condition

Let:

- message has unique sequence id s
- state store keeps processed ids set D

A processing function is idempotent under duplicates if:

$$\forall s: \text{Apply}(s) \text{ multiple times} \equiv \text{Apply}(s) \text{ once}$$

Implementation rule:


$$\text{if } s \in D: \text{ skip output} \quad \text{else: process and add } s \rightarrow D$$

This is the mathematical basis for “no duplicates” even under replay.

3.2 Replay correctness

If failure causes messages $s_a, \dots, s_b$ to be replayed, correctness requires:

Final state after replay = Final state if no failure occurred



6.1. Performance vs. Model Complexity

Automating the development and deployment of GenAI models can reduce risks associated with human intervention but may also increase it, especially for low-latency environments. Long-running Pipeline Enhancements deliver pre-trained models that potentially lower operational risk, but they accelerate the flywheel in the opposite direction, since unMonitored or poorly placed models can exert significant operational pressure. As with all changes in open systems, care must be taken to keep entry standards high. Human review and patience are just as important as the automation of mundane operational tasks. Although unMonitored models can relieve pressure

from the backlog, they place additional pressure on operations. The addition of an unMonitored model should surface operational pressure in the monitoring. The accepted complexity of a model is not a simple median: model performance and operational pressure are often highly correlated. Balancing model size and parameter count with latency requirements, supported data, and monitoring is important to strike the right operational complexity.

The tension between latency and throughput is multi-dimensional. A single low-latency model placed at ingress can often be replaced with several larger models with lower per-inference latency and more complex processing pipelines with larger windowing bursts. Reducing the frequency that GenAI calls are made can reduce the pressure on monitoring and repoussé systems that track and redirect the model traffic and populate the training data needed for long-running pipelines.

7. Conclusion

GenAI, banks, and other financial market participants are keenly exploring the potential of engines like OpenAI's ChatGPT for a variety of use cases. This paper narrows the focus to the area of DevOps applied to Financial Pipelines needing a real-time capability. The paper synthesizes the active and important area of DevOps with GenAI techniques and their ramifications for architecture and pipelines in real-time financial transaction processing systems. GenAI-enabled DevOps—a formalization of both concepts in this application domain—offers the promise of lessening the burden of these, often manual and random, activities in the operational lifecycle and governance of such systems, enabling pipeline owners to operate their systems with a greater degree of confidence.

GenAI, banks, and other financial market participants are keenly exploring the potential of engines like OpenAI's ChatGPT for a variety of use cases. This paper narrows the focus to the area of DevOps applied to Financial Pipelines needing a real-time capability. The paper synthesizes the active and important area of DevOps with GenAI techniques and their ramifications for architecture and

pipelines in real-time financial transaction processing systems. GenAI-enabled DevOps—a formalization of both concepts in this application domain—offers the promise of lessening the burden of these, often manual and random, activities in the operational lifecycle and governance of such systems, enabling pipeline owners to operate their systems with a greater degree of confidence.




Fig 4: Formalization of GenAI-Enabled DevOps

7.1. Final Thoughts and Future Directions

Despite the operational challenges and risks, the potential rewards warrant continued experimentation with GenAI-enabled DevOps in real-time financial transaction pipelines. These systems automate the management of the pipeline model lifecycle, from training and deployment to monitoring, rollback, and governance. Such automation allows organizations to decisively reap the benefits of real-time systems—reduced latency and increased business value—while offloading continuous operations to a controlled, automated process that requires limited manual intervention.

As GenAI improves and stabilizes, areas to monitor and explore include support for dispersed Hybrid Cloud-Edge pipelines with strict data sovereignty constraints; reliable streaming data ingestion and processing; and continuously supervised, multi-objective, long-term training of models for multiple and evolving objectives. The need for oversight and management of the auto-generation, in-context training, and execution risk fosters a second axis for monitoring and experimentation—establishing and

controlling an appropriate balance between simplicity and transparency of GenAI usage and the associated risk of blind trust in the output. A third direction for focus is the definition of a suitable standard operation generator and the exploration of the trade-offs between operational complexity, optimal operation, and exploitation of operation explainability.

8. References

1. Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., ... Amodei, D. (2020). Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33, 1877–1901.
2. Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W.-T., Rocktäschel, T., Riedel, S., & Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. *Advances in Neural Information Processing Systems*, 33, 9459–9474.
3. Jogi, N. L., Kumar, K. R., Mangala, N., M.Govindarajan, Gamini, P., & Bundhe, S. (2026). Machine Learning-Based Random Forest Model for Sustainable Supply Chain Sales Forecasting. In 2026 6th International Conference on Intelligent Technologies (CONIT) (pp. 1–6). IEEE. 2026 6th International Conference on Intelligent Technologies (CONIT). <https://doi.org/10.1109/conit69683.2026.11620788>
4. Goodell, J. W., Kumar, S., Lim, W. M., & Pattnaik, D. (2021). Artificial intelligence and machine learning in finance: Identifying foundations, themes, and research clusters from bibliometric analysis. *Journal of Behavioral and Experimental Finance*, 32, 100577.
5. Warin, T., & Stojkov, A. (2021). Machine learning in finance: A metadata-based systematic review of the literature. *Journal of Risk and Financial Management*, 14(7), 302.

6. Recharla, M., Garapati, R. S., Bandi, V. D. V. K., Yandamuri, U. S., & Mangalampalli, B. M. (2026, March). Generative AI-Enhanced Data Engineering Pipelines for Predictive Biomarker Discovery in Alzheimer's and Kidney Disease. In 2026 IEEE International Conference on AI Engineering and Innovations (AIEI) (pp. 1-5). IEEE.
7. Ashta, A. (2021). Artificial intelligence and fintech: An overview of opportunities and risks for banking, investments, and microfinance. *Strategic Change*, 30(3), 211–222.
8. Ryll, L., & Treleaven, P. (2021). Algorithmic trading and artificial intelligence in finance: A literature review. *International Journal of Data Science and Analytics*, 11, 1–15.
9. Kolla, S. H., Nagabhyru, K. C., & Kummari, D. N. (2026). Letter to the Editor re: "CurvAssist: An AI assisted pipeline for penile shaft segmentation/curvature measurement in children with hypospadias". *Journal of Pediatric Urology*.
10. Wamba-Taguimdje, S.-L., Wamba, S. F., Kamdjoug, J. R. K., & Wanko, C. E. T. (2020). Influence of artificial intelligence on firm performance: The business value of AI-based transformation projects. *Business Process Management Journal*, 26(7), 1893–1917.
11. Sarker, I. H. (2021). Machine learning: Algorithms, real-world applications and research directions. *SN Computer Science*, 2, 160.
12. Lakhanpal, S., Tunjungsari, H. K., Raj, S., Chukka, A., Dawadi, D., & Mangalampalli, B. M. (2026, April). The Digital Transformation of Healthcare: Balancing Opportunities, Risks, and Ethical Considerations. In 2026 International Conference on Multidisciplinary Innovations For Smart & Sustainable Future (MISSF) (pp. 1-6). IEEE.
13. Kaur, H., & Rani, R. (2022). Artificial intelligence and machine learning in financial services: A systematic review. *Journal of Risk and Financial Management*, 15(11), 523.
14. Ahmed, S., Alshater, M. M., El Ammari, A., & Hammami, H. (2022). Artificial intelligence and machine learning in finance: A bibliometric review. *Research in International Business and Finance*, 61, 101646.
15. Dinh, T. H. T., & Raza, M. (2021). Data-driven approaches in FinTech: A survey. *Information Discovery and Delivery*, 49(2), 123–135.
16. Kshetri, N. (2021). The economics of artificial intelligence in the financial sector. *IT Professional*, 23(5), 26–31.
17. Ozili, P. K. (2020). Financial inclusion and FinTech during COVID-19 crisis: Policy solutions. *SSRN Electronic Journal*.
18. Nandan, B. P., Kumar, M. V. K., Garapati, R. S., Bandi, V. D. V. K., Davuluri, P. N., & Mangalampalli, B. M. (2026, March). AI-Enhanced Semiconductor Yield Optimization Using Hybrid Deep Learning and Edge Data Analytics. In 2026 IEEE International Conference on AI Engineering and Innovations (AIEI) (pp. 1-6). IEEE.
19. Philippon, T. (2020). On fintech and financial inclusion. *BIS Working Papers*, 841, 1–24.
20. Fuster, A., Goldsmith-Pinkham, P., Ramadorai, T., & Walther, A. (2022). Predictably unequal? The effects of machine learning on credit markets. *The Journal of Finance*, 77(1), 5–47.
21. Gambacorta, L., Huang, Y., Qiu, H., & Wang, J. (2022). Data versus privacy: A comparison of cloud computing and artificial intelligence in finance. *BIS Working Papers*, 956, 1–35.
22. Lopez-Lira, A., & Tang, Y. (2023). Can ChatGPT forecast stock price movements? Return predictability and large language models. *SSRN Electronic Journal*.
23. Dowling, M., & Lucey, B. (2023). ChatGPT for (finance) research: The Bananarama conjecture. *Finance Research Letters*, 53, 103662.
24. Wu, S., Irsoy, O., Lu, S., Dabrovolski, V., Dredze, M., Gehrmann, S., Kambadur, P., Rosenberg, D., & Mann, G. (2023). BloombergGPT: A large language model for finance. *arXiv preprint arXiv:2303.17564*.
25. Li, Y., Wang, S., Ding, H., & Chen, H. (2023). Large language models in finance: A survey. *Proceedings of*

- the Fourth ACM International Conference on AI in Finance, 374–382.
26. Gopinath, A., Davuluri, P. N., Kumar, U. B., MP, S., & Yamini, G. (2026, April). Communication Aware Malware Detection Using Network Flow Bytecode Fusion With Adaptive Focal Optimization. In 2026 2nd International Conference on Intelligent Systems and Computational Networks (ICISCN) (pp. 1-8). IEEE.
 27. Huang, A. H., Wang, H., & Yang, Y. (2023). FinBERT: A large language model for extracting information from financial text. *Contemporary Accounting Research*, 40(2), 806–841.
 28. Yang, H., Liu, X.-Y., & Wang, C. D. (2023). FinGPT: Open-source financial large language models. *Proceedings of the Fourth ACM International Conference on AI in Finance*, 492–500.
 29. Chen, Z., Chen, W., Lin, Z., Zhang, Q., Wang, X., & Yang, Y. (2023). PIXIU: A large language model, instruction data and evaluation benchmark for finance. *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 10590–10614.
 30. Chen, Z., Li, L., & Smola, A. J. (2021). FinQA: A dataset of numerical reasoning over financial data. *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 3697–3711.
 31. Kolla, S. K., Bandi, V. D. V. K., & Meda, R. (2026). Comment on “From laboratory promise to decision-grade practice: strengthening reproducibility and casework relevance in AI-assisted bloodstain ageing”. *Forensic Science, Medicine and Pathology*, 1-2.
 32. Zhu, F., Lei, W., Huang, Y., Wang, C., Zhang, S., Lv, J., Feng, F., & Chua, T.-S. (2021). TAT-QA: A question answering benchmark on a hybrid of tabular and textual content in finance. *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics*, 3277–3287.
 33. Chen, W., Ma, X., Wang, X., & Cohen, W. W. (2022). FinQA: A dataset for numerical reasoning over financial data. *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics*, 644–657.
 34. Wu, Y., Raghavan, P., & others. (2022). FinQA and financial question answering: Recent advances in financial NLP. *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, 1–12.
 35. Kreuzberger, D., Kühl, N., & Hirschl, S. (2023). Machine learning operations (MLOps): Overview, definition, and architecture. *IEEE Access*, 11, 31866–31879.
 36. Symeonidis, G., Nerantzis, E., Kazakis, A., & Papakostas, G. A. (2022). MLOps: Definitions, tools and challenges. *Proceedings of the 2022 IEEE 12th Annual Computing and Communication Workshop and Conference*, 453–460.
 37. John, M. M., Olsson, H. H., & Bosch, J. (2021). Towards MLOps: A framework and maturity model. *Proceedings of the 2021 International Conference on Software Engineering and Knowledge Engineering*, 384–389.
 38. Davuluri, P. S. L. (2023). AI-Augmented Sanctions Screening: Enhancing Accuracy and Latency in Real Time Compliance Systems. *AI-Augmented Sanctions Screening: Enhancing Accuracy and Latency in Real Time Compliance Systems* (December 15, 2023).
 39. Kumari, S., & Kaur, R. (2021). MLOps: A systematic review of machine learning operations. *Journal of Systems and Software*, 182, 111086.
 40. Amershi, S., Begel, A., Bird, C., DeLine, R., Gall, H., Kamar, E., Nagappan, N., Nushi, B., & Zimmermann, T. (2020). Software engineering for machine learning: A case study. *Proceedings of the 41st International Conference on Software Engineering: Software Engineering in Practice*, 291–300.
 41. Sculley, D., Phillips, T., Ebner, D., Chaudhary, V., Young, M., Crespo, J.-F., & Dennison, D. (2020). Hidden technical debt in machine learning systems. *Advances in Neural Information Processing Systems*, 33, 2503–2515.
 42. Zhang, T., Li, J., & others. (2021). Engineering machine learning systems: Challenges and opportunities. *IEEE Software*, 38(4), 47–54.

43. Tamburri, D. A. (2020). Design principles for machine learning pipelines. *ACM Computing Surveys*, 53(1), 1–37.
44. van der Aalst, W. M. P. (2021). Process mining and RPA: A new partnership. *Business & Information Systems Engineering*, 63, 585–595.
45. Forsgren, N., Humble, J., & Kim, G. (2021). *Accelerate: The science of lean software and DevOps. IT Revolution*.
46. Erich, F. M. A., Amrit, C., & Daneva, M. (2020). A qualitative study of DevOps usage in practice. *Journal of Software: Evolution and Process*, 32(6), e2246.
47. Leite, L., Rocha, C., Kon, F., Milojevic, D., & Meirelles, P. (2020). A survey of DevOps concepts and challenges. *ACM Computing Surveys*, 52(6), 1–35.
48. Ebert, C., Gallardo, G., Hernantes, J., & Serrano, N. (2020). DevOps. *IEEE Software*, 37(2), 94–99.
49. Luz, W. M., Pinto, G. H. L., & Steinmacher, I. (2021). What is DevOps? A systematic mapping study. *Information and Software Technology*, 123, 106304.
50. Shahin, M., Zahedi, M., Babar, M. A., & Zhu, L. (2021). An empirical study of DevOps practices in software development organizations. *Empirical Software Engineering*, 26, 1–44.
51. Subramanian, N., & Kolla, T. Equity in Healthcare Access Policy Lessons from Universal Coverage Models.
52. Katal, A., Wazid, M., & Goudar, R. H. (2021). Big data: Issues, challenges, tools and good practices. *Proceedings of the 2021 International Conference on Intelligent Computing and Control Systems*, 1–8.
53. Mäkinen, S., Skogström, H., Laaksonen, E., & Mikkonen, T. (2020). Who needs MLOps: What data scientists seek to accomplish and how can DevOps help? *Proceedings of the 2020 IEEE/ACM 1st International Workshop on Software Engineering for AI in Autonomous Systems*, 37–44.
54. John, M. M., Hotti, V., & others. (2022). Machine learning operations: Challenges, practices, and future directions. *IEEE Software*, 39(5), 56–63.
55. Fan, A., Gokkaya, B., Harman, M., Lyubarskiy, M., Sengupta, S., Yoo, S., & Zhang, J. M. (2023). Large language models for software engineering: Survey and open problems. *Proceedings of the 2023 IEEE/ACM International Conference on Automated Software Engineering*, 1–12.
56. Hou, X., Zhao, Y., Liu, Y., Yang, Z., Wang, K., Li, L., Luo, X., Lo, D., Grundy, J., & Wang, H. (2024). Large language models for software engineering: A systematic literature review. *ACM Transactions on Software Engineering and Methodology*, 33(8), Article 220.
57. Liu, J., Wang, K., Chen, Y., Peng, X., Chen, Z., Zhang, L., & Lou, Y. (2024). Large language model-based agents for software engineering: A survey. *arXiv preprint arXiv:2409.02977*.
58. Gao, C., Hu, X., Gao, S., Xia, X., Jin, Z., & others. (2024). The current challenges of software engineering in the era of large language models. *arXiv preprint arXiv:2412.14554*.
59. Peng, S., Kalliamvakou, E., Cihon, P., & Demirer, M. (2023). The impact of AI on developer productivity: Evidence from GitHub Copilot. *arXiv preprint arXiv:2302.06590*.
60. Di Grazia, L., & Pradel, M. (2023). Code generation by large language models: Evaluation and limitations. *Proceedings of the 2023 IEEE/ACM International Conference on Software Engineering*, 1–12.
61. Chatterjee, S., Liu, C. L., Rowland, G., & Hogarth, T. (2024). The impact of AI tool on engineering at ANZ Bank: An empirical study on GitHub Copilot within corporate environment. *arXiv preprint arXiv:2402.05636*.
62. Hou, X., Zhao, Y., Liu, Y., Yang, Z., Wang, K., Li, L., Luo, X., Lo, D., Grundy, J., & Wang, H. (2024). Large language models for software engineering: A systematic literature review. *ACM Transactions on Software Engineering and Methodology*, 33(8), Article 220.
63. Chiu, I.-C., & Hung, M.-W. (2025). Finance-specific large language models: Advancing sentiment analysis



and return prediction with LLaMA 2. Pacific-Basin Finance Journal, 90, 102632.

64. Baird, M. (2026). Firms' GitHub Copilot adoption and labor market outcomes for software engineers. Contemporary Economic Policy.