10.48047/jocaaa.2024.33.08.323

Machine Learning-Driven Anomaly Detection in CI/CD Pipelines for Financial Applications

Avinash Reddy Segireddy
Lead DevOps Engineer
ORCID ID: 0009-0002-9912-0629

***Abstract*—The effects of machine learning-driven anomaly detection in CI/CD pipelines specific to financial applications are examined. Risk reduction while ensuring fault tolerance and corporate governance is paramount in the financial sector. To strike this balance, organizations rely on security, observability, and compliance frameworks that provide transparency and evidence of due diligence. It is through automation that organizations can achieve malware protection, anomalous transaction detection, incident response, and automated testing of non-production credentials. Automated anomaly detection reduces risk while maintaining velocity to market. The benefits extend to stakeholder buy-in, satisfying demand for auditability, traceability and governance. Operations in cloud platforms are auto-scaled to ensure proper functioning with changing workloads. Resource allocation for such auto-scaled operations is governed primarily by metrics that trigger scaling. In addition to making sure scaling operations do not hinder regular application functioning, investments are made to scale seamlessly and cost-effectively in response to such changes. CI/CD Pipelines in support of such operations must function correctly even with these auto-scaled resources. The pipelines need to have anomalies that indicate aberrations in regular functioning, the detection of which assists in reducing financial, operational, and reputational losses to the organization.**

***Index Terms*—Anomaly detection; financial applications; ML; CI/CD pipelines; integrity; compliance,Machine Learning–Based Anomaly Detection,CI/CD Security Monitoring,Financial Software Integrity,Behavioral Analytics for Pipelines,DevSecOps Automation,Pipeline Risk Scoring,Model-Driven Threat Detection,Real-Time Deployment Monitoring,Secure Continuous Delivery,Fraud and Compliance-Aware Pipelines.**

I. Introduction

Modern financial services organizations rely heavily on machine-learning (ML)-driven systems for risk detection, fraud prevention, and customer service. Responsibly deploying these models into production remains a challenge. An effective strategy is to reduce the risk of production ML models without sacrificing the delivery speed expected in a continuous integration/continuous delivery (CI/CD) environment. Automated ML-driven anomaly detection in CI/CD pipelines can provide these organizations with the confidence to deploy updates quickly while minimizing potential downstream impact. In doing so, development velocity is unshackled; the risk of loss caused by undetected failures is lowered; and the evidentiary burden of supporting an explainable, auditable, and governable operational ML system becomes easier to maintain. Anomalies in CI/CD pipelines for financial applications during design, implementation, testing, and deployment can have serious and costly repercussions in production. Traditional testing methods, despite being more thorough, cannot catch every unforeseen problem prior to deployment. Automated anomaly detection fills this gap by being continually applied throughout CI/CD, surfacing unusual conditions that an organization may be willing to treat as potential risk factors just above the operational level of noise but much lower than the disruptive level of incidents. Speed of deployment is further increased by providing the analysis required to maintain systems with diminished performance-benchmark confidence rather than ceasing deployment altogether.

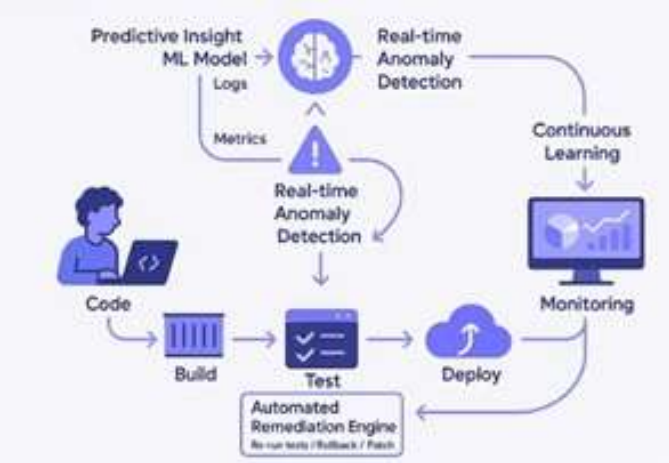

Fig. 1. Anomaly Detection in CI/CD Pipelines

A. Background and Significance

Every machine learning algorithm exhibits characteristics that delineate its behavior. The concepts of "Anomaly Detection" (AD) and "ML-driven AD" in the context of Continuous Integration/Continuous Delivery (CI/CD) workflows in financial applications are no exception. In the financial services domain, regulators or governing bodies impose stringent standards addressing risk and security. Violations can occur during any phase of development. An ML-driven AD framework incorporates data from logs, traces, metrics, and transactional sources to help identify anomalies, thereby accelerating detection, improving accuracy, and increasing trust. With auditable records of training and predictions, the evidence serves as a backstop during audits. Financial organizations are under relentless pressure to reduce the time it takes to deliver new offerings, products, and services without compromising quality

10.48047/jocaaa.2024.33.08.323

or security. Related elements in the CI/CD cycle that contribute to this pressure include cloud adoption and the decommission- ing of legacy applications. However, this drive has exacerbated incidents. Risk-related elements such as bid gobbling, broken communications links in trading systems, and unfortunate pricing on Initial Public Offerings validate this trend. Malware attacks and the increasing use of cryptocurrencies in ransom requests confirm this. Like all production systems, CI/CD pipelines cannot be immune to such risk. The aim is to achieve rapid delivery while incorporating a mechanism to detect potential risk during any phase of the CI/CD cycle. The early identification of these anomalies translates into significant savings.

II. THE IMPERATIVES OF FINANCIAL INTEGRITY AND RAPID DELIVERY

In sensitive domains like finance, the accuracy of code is paramount. However, the traditional testing paradigm cannot keep pace with CI/CD velocities, and the frequency of atypical code deployments increases the risk of financial errors. Automated anomaly detection in CI/CD pipelines leverages machine learning to identify anomalies in logs, traces, or metrics as soon as they occur. The consequent reduction in deployed anomalies mitigates the risk of financial errors while preserving deployment velocity. Organizations will therefore be able to cut their costs while moving faster, delighting customers, and building trust. Moreover, these predictions provide an added layer of compliance, automation, transparency, and auditability. Forecasting tests are trained on normal behavior modeled in historical data—distributed, labeled, and controlled by a data-management system so that quality can be monitored, lineage tracked, and drift detected, and that external teams can easily check data. Predictions can thus serve as part of evidentiary material and independently demonstrate whether the deemed "business requirement" has been fulfilled. As such, they offer the possibility of adhering to governance, audit, and compliance needs and of implementing robust business-rules structures.

Equation 01:Basic counts (from predictions vs. ground truth)

Let $y \in \{0, 1\}$ be the true label (0=normal, 1=anomaly) and

$\hat{y} \in \{0, 1\}$ be the predicted label at a threshold

τ on a continuous anomaly score s $TP = \sum 1[\hat{y} = 1 \wedge y = 1]$ $FP = \sum 1[\hat{y} = 1 \wedge y = 0]$

$FN = \sum 1[\hat{y} = 0 \wedge y = 1]$

$TN = \sum 1[\hat{y} = 0 \wedge y = 0]$

10.48047/jocaaa.2024.33.08.323

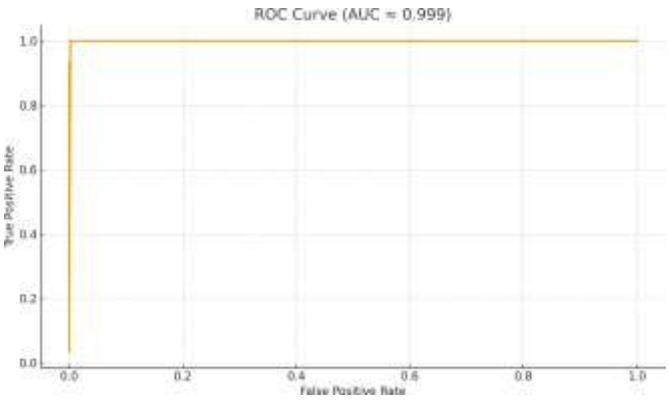

Fig. 2. ROC Curve (AUC ≈ 0.999)

threshold	precision	recall	F1
0.2246	0.06	1	0.1132
0.2827	0.0604	1	0.1139
0.2959	0.0607	1	0.1145
0.3281	0.061	1	0.1149
0.3413	0.0613	1	0.1156
0.3544	0.0616	1	0.1161
0.3684	0.062	1	0.1167
0.3851	0.0622	1	0.1172
0.3935	0.0626	1	0.1179

TABLE I
THRESHOLD SWEEP (FIRST 25 ROWS)

A. Financial risk Mitigation through automated anomaly de- tection

Combined Machine Learning and DevOps techniques are increasingly being used to support rapid deployment of applications in many industries. Financial applications undergo much more scrutiny in these processes, though. Properly configured CI/CD pipelines should regularly pass audits when using, e.g., standard security control frameworks such as the NIST Cybersecurity Framework (CSF). Automated anomaly detection that incorporates Machine Learning approaches can help financial DevOps pipelines maintain audit-readiness and lessen the risk of exposed bugs in production systems, e.g., bugs leading to large financial losses or unauditable, illegiti- mate transactions. The need for rapid delivery remains, with study after study demonstrating the cost advantages of more frequent releases, by teams with fewer personnel. Machine Learning can be integrated into the CI/CD pipelines, of course, but like any other software component, the introduction of ML models can introduce new anomalies into the software, which can result in errant processing. Automated detection of such model errors will help reduce exposed risk. Another important determinant of speed of delivery remains achieving lower costs, both of creating and then maintaining CI/CD pro- cesses. Well-architected solutions that require less work at the deployment stage provide these advantages, especially in or- ganizations supporting many applications. Automated anomaly detection for ML models in the CI/CD pipelines undertaking financial processes is a powerful approach to lessening risk when striving for increased deployment frequencies. Many CI/CD pipelines do not yet provide the audit and compliance levels required for financial applications. Adding the ability to detect when a deployed model acts differently than has been experienced in the past will provide financial teams with more confidence that their deployments will pass scrutiny and comply with internal criteria. Financial services organizations, even their DevOps teams, are not often interested in the expense of training Artificial Intelligence models, particularly the large transformer-based architectures commonly used. But when speed and cost targets are established, those supporting financial applications can embrace a larger-than-usual appetite for risk, e.g., with respect to CI/CD pipeline anomaly detection. Research is a powerful conveyor of lessons learned and exploration results. Other research worlds will be served by lessons derived from this discussion.

B. Regulatory alignment and auditability

Recent incidents underscore financial institutions' double- edged nature, capable of sudden, substantial gains or catas- trophic failures. Although a specific outcome cannot be predicted, it is prudent to manage both the likelihood and scale of loss. Disturbances in the Continuous Integra- tion/Continuous Deployment (CI/CD) pipeline may be precur- sors of SI, just as minor

10.48047/jocaaa.2024.33.08.323

tremors can precede large earthquakes. Anomalies—unusual deviations from normal operational pat- terns—are often precursors of Incident failure. Automating the detection of these anomalies is key to risk management and speed. Automating this detection increases confidence, allowing the institution to accelerate its test-and-release sched- ule without increasing the risk of a serious failure. Such confidence is invaluable, especially where the situation is fully audited and must comply with legislation like Sarbanes- Oxley. The large amount of information generated in the CI/CD pipeline supports detailed post-mortems in the event of a breach. If CI/CD anomalies are detected automatically, information capture can be concentrated in areas of higher risk, resulting in even better incident reports.

III. CORE CONCEPTS OF ML-DRIVEN ANOMALY DETECTION IN CI/CD

Anomalies in Continuous Integration and Continuous De- ployment (CI/CD) pipelines for financial applications are desynchronized events that contribute to undesired results dur- ing the release process and make a financial transaction deviate from its intended effect. Anomalies differ from errors, which uncover defects in the code repository, and incidents, which cause an outage in the production environment. Three broad categories of anomalies can be distinguished within CI/CD pipelines: Code and Deployment Anomalies, Infrastructure Anomalies, and Monitoring and Testing Anomalies. A single event can belong to more than one category. For instance, an anomalously slow deployment can lead to anomalous monitoring results. A Code and Deployment Anomaly can be the cause of an Infrastructure Anomaly; for example, a database connection timeout caused by an unexpected change

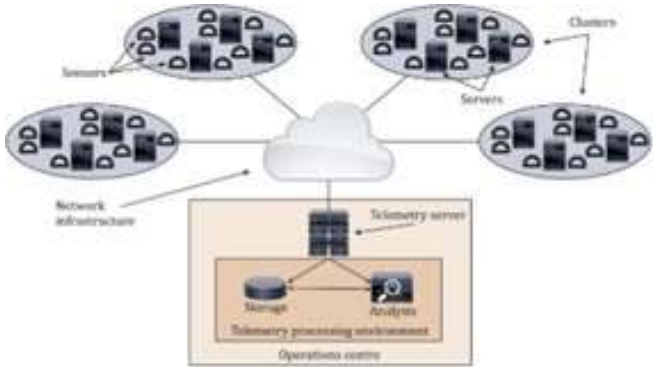

Fig. 3. ML-driven anomaly detection in CI/CD

in the number of available database connections can be the result of an anomalously high traffic volume or an anomalously low error rate in a web application. CI/CD pipelines for financial workstreams generate various signals: logs, traces, metrics, and transactional data. These signals can be processed individually or merged into one or more product feature stores, from which the features needed to train, test, and serve machine-learning models are built. Key aspects of feature engineering are feature alignment and the enrichment of events with temporal features (e.g., hour of the day, day of the week, month of the year) or domain-specific indicators of transactional volume, error rate, or latency. Features also require specific preprocessing or normalization, particularly for use in unsupervised models.

A. Anomaly definitions in CI/CD pipelines

Anomalies that may arise in the CI/CD pipelines of financial applications are detected through the application of machine learning models to operational transactional data (e.g., logs, traces, metrics). The operations of the financial applications supported by the CI/CD pipelines should, of course, be con- tinuously monitored, so that alerts can signal the occurrence of a known error or incident condition, which is like an anomaly but indicates a publicly known problem. If anomaly detection is effective in these applications, it can dramatically improve maintenance speed. Anomalies are defined here as unplanned disruptions in the quality of the operation of a CI/CD pipeline that may have negative consequences. This definition is careful to include only those conditions that cannot be captured by operational incident and error alerting, and thereby includes such situations as subtle performance changes and new failures before they are known public incidents or errors on the product. To illustrate the distinction between an anomaly and an incident or error, consider a CI/CD pipeline that had been returning a 404 page and is now returning an incorrect page: the first situation is an error alert, the second an incident alert, and the change from the former situation to the latter is an anomaly. Anomalies can also cover situations that, while perhaps becoming public knowledge, have had little exploration by the community.

10.48047/jocaaa.2024.33.08.323

B. Data sources and feature engineering for financial work- flows

CICD pipelines in financial applications produce vast amounts of heterogeneous data that may be the source of supervised or unsupervised training data for the automated detection of anomalies. CI/CD pipelines typically produce three main types of signals: - CI/CD execution logs and traces - Monitoring systems producing metrics on the CI/CD execution environment like SonarQube (development quality analysis), APM (newrelic, AppDynamics, etc..), . . . - Transactional data for the pipelines supporting the financial application itself Feature engineering and feature selection may heavily depend on the region of the CI/CD workflow where an anomaly is detected or in which detection is considered. Therefore, detailed research on preprocessing pipeline data for supervised or unsupervised training of ML models depend on the partic- ular area of interest. Some common considerations, however, are presented below. During preprocessing, log lines may be normalized to a standard set of columns like "timestamp", "level", "CI/CD identifier", "sub-component", "message"," . . . depending on the level of abstraction for the ML model. Domain/technical specific signals may also be produced like the timing since the last successful build or whether it is night time, blackout time, etc.. A specific temporal pattern can be generated from these signals to enforce temporal signals that allow models like LSTMs or ConvNets to trace temporal correlations.

IV. ARCHITECTURAL BLUEPRINT FOR INTEGRATION

The system architecture consists of two key parts: an information gain pipeline that collects and prepares data, and a monitoring pipeline that enables reliable anomaly detection. For the first part, the data ingestion layer should track the full lineage of applications. Lineage, however, can change quickly and might need time-varying dynamic schema adjustments. Therefore, it should first apply data-quality checks before being stored into long-term storage and made available for feature engineering. Feature engineering should prepare the raw input (i.e., logs, traces, metrics, transactional data) into the information gain required by the model and keep this level of preciseness across the monitored components to establish a comprehensive anomaly-detection system. The feature-engineering process should include the data-required normalizable steps, possible temporal aspects, and specific domain features of the components under test. The monitoring part of the architecture is responsible for serving the trained ML models in production and guaranteeing their correctness. It continually applies checks on the algorithm drift, monitors the deployed models, sets rollback plans when needed, and defines feedback loops in production to incorporate the detected issues into the retraining pipeline. The data flow of this part should align with the CI/CD approach, only adding the necessary elements for monitoring. As business continuity is essential, deployment patterns and observability must ensure the support of the tool as part of the production

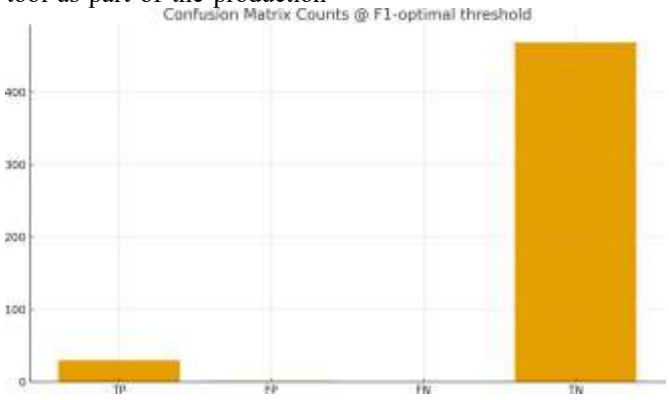

Fig. 4. Confusion Matrix Counts @ F1-optimal threshold

	feature	psi
1	error rate	0.1253
0	build time	0.104
3	mem	0.0981
2	cpu	0.0619

TABLE II
FEATURE DRIFT PSI

process.

Equation 02: Precision, recall, F1

10.48047/jocaaa.2024.33.08.323

These are the core metrics called out in your Section 5.1

$$\text{Precision} = \frac{TP}{TP + FP} \quad \text{———}$$

$$\text{Recall} = \frac{TP}{TP + FN} \quad \text{———}$$

Derivation of F1 (harmonic mean of precision μP and recall μR)

F1=P1+R1/2=P+R2PR

The "Summary @ F1-optimal threshold" table uses the τ that maximizes F1

A. Data ingestion, lineage, and quality controls

Reliable ML models for anomaly detection critically depend on the availability of trustworthy data. Anomaly detection solutions are expected to generate a substantial number of warnings, with only a portion being indeed anomalous and leading to incidents. This discrepancy triggers alert fatigue, negatively affecting the entire detection process and potentially rendering it ineffective. By integrating data ingestion pipelines in the CI/CD environment, the mean time to detect and resolve incidents can be significantly reduced. Consequently, ML- driven anomaly detection is strategically deployed upstream from the suspected failure points. However, enlarging the

input data with auxiliary information introduces new risks that must be mitigated. Such concerns are acknowledged by devising data lineage tracking, enabling data quality checks, and controlling access rights for the supporting data. Data consumption is monitored via provenance information, sup- porting auditable records of the data paths and the underlying motivations for each data source. Missing or low-quality data need to be detected before the model evaluation stage; otherwise, the acceptance criteria can be violated. To address these needs, a dedicated strategy can facilitate quality controls on the supporting data. Auxiliary data are certainly prone to schema evolution and may need to be checked for duplicates, unascribed values, ethics, and bias. Data quality patterns can be implemented to provide a focused diagnosis of the auxiliary datasets. Implementing data quality checks is crucial when the input data originate from DevOps teams not directly responsible for model training and adoption. It may also be relevant when operational data from production are employed in a semi-supervised way. Before considering a deviation from the learned behavior, the existence of missing data must be taken into account and provision made for its treatment. Finally, appropriate access controls play a pivotal role in upholding data confidentiality and compliance with company policy and external regulations.

B. Model serving, monitoring, and feedback loops

Design the model-serving and monitoring strategy in tan- dem with the application-development cycle's deployment pattern. Deployment in containers and Kubernetes facilitates scale-out for both resources and access. To enable rapid rollback, recent versions must be retained passively or actively for recovery; a candidate version should be retrained on observational data before rollout. Observe model performance to detect potential issues requiring retraining for sustained efficacy; use new labels to monitor developers' reactivity and feedback. Carefully assess predicted labels of new sam- ples for temporal shifts; for modeling drift, adapt features, training intervals, and refresh workflows based on data-from- production patterns. Specify an alerting approach for drift and monitoring-level nonresponsiveness. Integrate the monitoring- and-feedback-loop refinement with existing CI/CD channels to provide developers effective ML-driven production coverage at scale.

V. Efficacy, risk, and governance

Automated anomaly detection in CI/CD pipelines for finan- cial workflows necessitates a comprehensive strategy to ensure efficacy, risk mitigation, and governance. Evaluation criteria comprise deployment-time guarantees and a rigorous accep- tance process. Validation metrics include precision, recall, F1, AUROC, and MCC, measured against appropriate benchmarks and employing cross-validation. A considered governance framework addresses bias, fairness, robustness, and adversarial resilience. A meticulous evaluation and deployment checklist facilitates acceptance, while a robust governance regime mit- igates risk, encourages auditability, and establishes account-ability. Efficacy hinges on detecting real anomalies while minimizing false positives—capturing problems that mandate direct intervention. Insights from exploratory milestone CI/CD runs can enhance the underlying model and its deployment- time guarantees. Projected metrics guide acceptance: pre- cision conveys candidness in alerts, F1 balances precision and recall, AUROC indicates performance across thresholds, and Matthews correlation coefficient

10.48047/jocaaa.2024.33.08.323

quantifies correlations between anomalies and alerts. Financial CI/CD contains tem- poral correlations, enabling spatio-temporal models and key performance indicator predictions via landmark datasets. Such datasets also facilitate cross-validation between folds across logical time—intuitively a sounder partition than network- or user-based splits. However, any automated system intro- duces risk, especially in high-stakes domains such as finance. Standard data-science conventions—testing against multiple models and problem formulations, exhaustively exploring hy- perparameter space, and executing causative testing—partially mitigate risk. A more robust methodology addresses bias, fairness, and robustness. Final service acceptance further con- siders adversarial resilience.

A. Evaluation metrics and performance guarantees

The effectiveness of ML-driven anomaly detection is best understood through precision, recall, and related measures, all of which can be computed automatically using classi- fication engagements within the standard machine-learning routine. Stable production performance in these dimensions is a necessary condition for, but does not entail, long-term business impact. Meaningful business impact can be expressed as a business rating—a composite measure that combines risk-based revenue, cost-of-investment considerations, and the implied cost of rollback failures diminishable by a deployment guarantee. Therefore, along with links to operational data and continuous cross-validation across geographically sepa- rate cluster installations, development teams should explicitly engage with risk and governance materials to anticipate the challenges of getting and staying on production. Operational and business risk factors also need to be considered for each model being deployed. Supervisors should be alert to intrusion threats facing bank operations, in which adversaries aim to create, inject, or modify deployment artefacts in ways that will introduce exploitable backward shifts in the risk-return profile. Such risks are most apparent in cases where collateral damage triggers trade- or corruption-related issues some distance away from the actual insertion point. Both companies' tax and finance departments face unique external threats that diminish the independent internal reproduction defences provided by the normal business processes and automated ML audit systems deployed across the remainder of CI/CD.

B. Bias, fairness, and robustness considerations

Machine learning is rarely free from bias, and it would be na¨ıve to assume that model predictions are impartial, fair, robust, adversarial-resilient, and governed by a rigorous process. Striving to minimize bias is crucial; traversing the

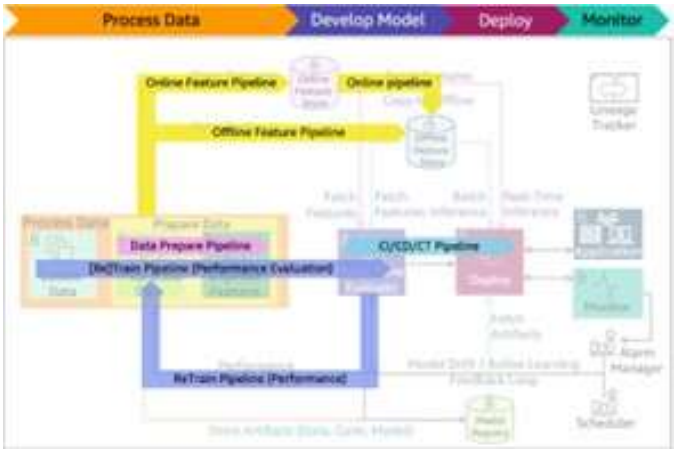

full pipeline—from data accumulation and preparation to de- velopment, deployment, and reconstruction—requires special attention at each juncture. Throughout, humans must remain accountable for supervising the actions attributable to their creations. Auditing procedures applied to predictions provide another layer of responsibility. Although distinct from techni- cal considerations, these human aspects cannot be overlooked without jeopardizing the predictive process. Bias can creep into predictions for various reasons. For example, in a CI/CD pipeline, training-data pre-processing and feature engineering may contain subjectivity—that is, a statistical model may

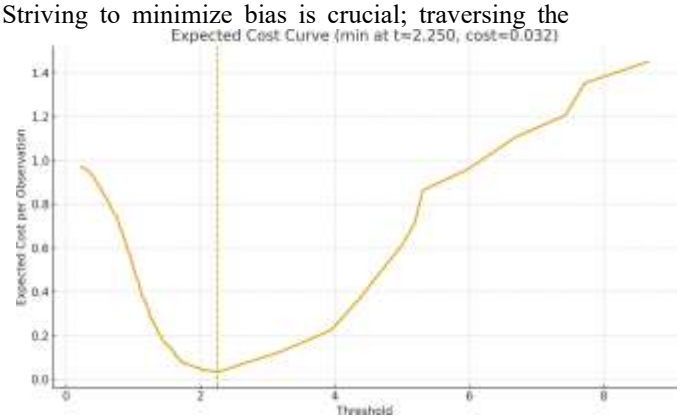

use a specific feature, while another service may not use it even though the data are present. Certainly, the choice of training data will also determine the resulting model's fairness. These requirements can be mitigated during deployment by continuously examining prediction drift and performance as new observations arrive. A labelled dataset of model outputs over time makes it feasible to assess trends in the model's capacity to fulfil its purpose.

VI. CONCLUSION

Four trends are shaping future development in ML- driven anomaly detection for CI/CD automation in financial services: 1. Scaling pipeline complexity. The proliferation of microservices is increasing the number of stakeholders contributing into the CI/CD process throughout the software

10.48047/jocaaa.2024.33.08.323

development lifecycle. 2. Business process and risk-awareness. An increasingly precise mapping of the business process supported by CI/CD pipelines together with new metrics and governance processes to track risk at a pipeline level are driving a more agile detection of anomalies and divergent changes with respect to the controlled pipeline "template". 3. Integration into broader business process and risk-assurance frameworks. Past monitoring has primarily focused within the computational box of the CI/CD system. The next step is to connect it wider in the context of the CI/CD business process, its dependencies and KPIs, upstream and downstream. These KPIs are heavily oriented toward risk, the dimension of

Fig. 6. Expected Cost Curve (min at t≈2.250, cost≈0.032)

threshold	expected cost
0.2246	0.97
0.2827	0.964
0.2959	0.958
0.3281	0.954
0.3413	0.948
0.3544	0.944
0.3684	0.938
0.3851	0.934
0.3935	0.928

TABLE III
EXPECTED COST VS THRESHOLD (FIRST 25 ROWS)

main concern when modernizing core banking systems.

4. Supporting regulations and compliance. Policy-makers and regulatory authorities perceive the potential benefits of Digital Transformation as positive. Therefore, appropriate recommendations and directives are being formulated in parallel, such as Capgemini's 2020 report for the European banking authority. In the U.S., various agencies will have to support Digital Transformation while maintaining soundness and solvency. The combination of strategy and governance has major implications for speeding up Global IT. These trends point a road map for ML-driven anomaly detection. Its presence will be systematically integrated into the CI/CD architecture and scope. The proper consideration of these aspects will enable all Financial Services Organizations, large or small, to effectively implement ML anomaly detection in a scalable and effective way, allowing it to become a permanent safety surveillance during facilitation and transformation projects.

Equation 03: ROC curve and AUROC Define the True Positive Rate and False Positive Rate at threshold τ

$$TPR(\tau) = \frac{TP}{TP + FN} \qquad \text{———}$$

$$FPR(\tau) = \frac{FP}{FP + TN} \qquad \text{———}$$

$$\text{AUROC} = \int_0^1 TPR(FPR)\, d(FPR)$$

10.48047/jocaaa.2024.33.08.323

A. Future Trends

ML-driven anomaly detection in CI/CD pipelines for fi- nancial applications will increasingly use consolidated mod- els rather than a myriad of specialized models. Such ap- proaches simplify operational overhead across a vast range of applications and CI/CD operations. Alternative architectures will enable distributed and semi-autonomous implementations. Scaling will leverage concurrent pipeline executions along with consolidated models or cloudbursting architectures, with the latter relying solely on inferencing when resource demand peaks. Transparent operation will support data privacy and ownership concerns while providing CI/CD controls. Trends in ML-driven anomaly detection will impact CI/CD pipeline providers by granting ML plug-and-play capability and mon- itoring services capable of assessing data and model quality. New ML platform will be fully automated by merging its oper- ations into a single job offering an end-to-end semi-automated, ML-ready pipeline. Self-evident data quality checks, train- ing requantification, and concurrent multiactivity scaling will shield the ML deployment from interest differences arising from nearness to planned ilgible activity intervals and their expected increased volumetric pace. Risk assessment will address these non-detection normative factors, along with model development time, by introducing auxiliary CI/CD anomaly control without specific data rules.

REFERENCES

[1] Jia, F., & Zhang, Y. (2024). Supply chain transparency: Opportunities, challenges and risks. International Journal of Operations & Production Management, 44(9), 1525-1549.

[2] Nikookar, E., & et al. (2024). Supply chain resilience: When the recipe is more important than the ingredients. International Journal of Production Economics, 290, 108155.

[3] Zhang, S., & et al. (2024). Digital supply chain: Literature review of seven related trends. Manufacturing Review, 11, 1-23.

[4] Lakkarasu, P. (2024). From Model to Value: Engineering End-to-End AI Systems with Scalable Data Infrastructure and Continuous ML Delivery. European Journal of Analytics and Artificial Intelligence (EJAAI) p- ISSN 3050-9556 en e-ISSN 3050-9564, 2(1).

[5] Vasconcelos, L. F., & et al. (2024). Supply chain 4.0: A multi-sector grey systems maturity assessment in Brazilian industry. Grey Systems: Theory and Application, 14, 123-145.

[6] Nakayima, F., Namagembe, S., Kabagambe, L., Ntayi, J. M., & Muh- wezi, M. (2024). Asset specificity, inter-firm ecosystem, firm adaptability and supply chain integration. Modern Supply Chain Research and Applications, 6(4), 1-17.

[7] Yellanki, S. K. (2024). Co-Creation and Connectivity: The Role of Con- sumers in Digital Ecosystem Evolution. Journal of Artificial Intelligence and Big Data Disciplines, 1(1).

[8] Siddiqui, M. F., Khalid, W., & Arsalan, M. (2024). Reshoring concepts: Definitions and a structured bibliometric review. Modern Supply Chain Research and Applications, 6(4), 18-35.

[9] Gan, Y., Gao, X., Zhou, W., Ke, S., Lu, Y., & Zhang, S. (2024). On the benefit of developing customer profile analysis to implement personalized pricing in a supply chain. Modern Supply Chain Research and Applications, 6(3), 55-72.

[10] Motamary, S. (2024). Data Engineering Strategies for Scaling AI-Driven OSS/BSS Platforms in Retail Manufacturing. BSS Platforms in Retail Manufacturing(December 10, 2024).

[11] Tetteh, F. K., Atiki, G., Kyeremeh, A., Degbe, F. D., & Apanye, P. (2024). Linking business analytics capability and sustainability perfor- mance: The mediating role of circular economy implementation. Modern Supply Chain Research and Applications, 6(2), 101-120.

[12] Kravchenko, K., Gruchmann, T., Ivanova, M., & Ivanov, D. O. (2024). Responding to the ripple effect from systemic disruptions: Empirical evidence from the semiconductor shortage during COVID-19. Modern Supply Chain Research and Applications, 6(3), 130-150.

[13] Inala, R., & Somu, B. (2024). Agentic AI in Retail Banking: Redefining Customer Service and Financial Decision-Making. Journal of Artificial Intelligence and Big Data Disciplines, 1(1).

[14] Usman Khalid, R., Sadiq Jajja, M. S., & Ahsan, M.-B. (2024). Supply chain sustainability and risk management in food cold chains – a literature review. Modern Supply Chain Research and Applications, 6(2), 82-100.

[15] Moleme, L. O., Omoruyi, O., & Quayson, M. (2024). Supply chain visibility and integration in the age of the Internet of Things: A retail perspective. Modern Supply Chain Research and Applications, 6(3), 73- 89.

[16] Somu, B. (2024). Agentic AI and Financial Compliance: Autonomous Systems for Regulatory Monitoring in Banking. European Data Science Journal (EDSJ) p-ISSN 3050-9572 en e-ISSN 3050-9580, 2(1).

[17] Chen, W., Kang, Z., Yang, H., & Shang, Y. (2024). Research on interregional oil cooperation-sanctions with evolutionary game. Modern Supply Chain Research and Applications, 6(2), 20-38.

[18] Mehmood, S., Nazir, S., Fan, J., & Nazir, Z. (2024). Achieving supply chain sustainability: Enhancing supply chain resilience, organizational performance, innovation and information sharing: Empirical evidence from Chinese SMEs. Modern Supply Chain Research and Applications, 6(1), 1-19.

[19] Meda, R. (2024). Predictive Maintenance of Spray Equipment Using Machine Learning in Paint Application Services. European Data Science Journal (EDSJ) p-ISSN 3050-9572 en e-ISSN 3050-9580, 2(1).

[20] Urmston, A., Song, D., & Lyons, A. (2024). The development of risk assessments and supplier resilience models for military industrial supply chains considering rare disruptions. Logistics, 8(2), 57.

[21] vThakker, S., Rane, S. B., & Narwane, V. S. (2024). Implementation of blockchain–IoT-based integrated architecture in green supply chain. Modern Supply Chain Research and Applications, 6(1), 45-64.

[22] Pandiri, L., & Chitta, S. (2024). Machine Learning-Powered Actuarial Science: Revolutionizing Underwriting and Policy Pricing for Enhanced Predictive Analytics in Life and Health Insurance.

[23] Quayson, M., Avornu, E. K., & Bediako, A. K. (2024). Modeling the enablers of blockchain technology implementation for information man- agement in healthcare supply chains. Modern Supply Chain Research and Applications, 6(1), 28-44.

[24] Tang, C., Hou, Q., & He, T. (2024). Research on closed-loop supply chain decision-making of power battery echelon utilization under the scenario of trade-in. Modern Supply Chain Research and Applications, 6(2), 121-140.

[25] Nandan, B. P. (2024). Revolutionizing Semiconductor Chip Design through Generative AI and Reinforcement Learning: A Novel Approach to Mask Patterning and Resolution Enhancement. International Journal of Medical Toxicology and Legal Medicine, 27(5), 759-772.

10.48047/jocaaa.2024.33.08.323

[26] Ren, L., Zhou, Z., Fu, Y., Liu, A., & Ma, Y. (2024). Integrated optimization of logistics routing problem considering chance preference. Modern Supply Chain Research and Applications, 6(3), 95-113.
[27] Ge, X.-L., Xu, M., Wang, B., & Yin, Z. (2024). Service equilibrium of urban transportation energy supply station based on cooperative game. Modern Supply Chain Research and Applications, 6(1), 11-27.
[28] Sheelam, G. K. (2024). AI-Driven Spectrum Management: Using Ma- chine Learning and Agentic Intelligence for Dynamic Wireless Op- timization. European Advanced Journal for Emerging Technologies (EAJET)-p-ISSN 3050-9734 en e-ISSN 3050-9742, 2(1).
[29] Hasan, M. R., Khan, M. A., & Wuest, T. (2024). Towards Industry 5.0: A systematic literature review on sustainable and green composite materials supply chains. [Pre-print].
[30] Fernandez-Carames, T. M., Blanco-Novoa, O., Froiz-Miguez, I., & Fraga-Lamas, P. (2024). Towards an autonomous Industry 4.0 ware- house: A UAV and blockchain-based system for inventory and trace- ability applications in big-data-driven supply chain management. [Pre- print].
[31] Singireddy, J. (2024). Deep Learning Architectures for Automated Fraud Detection in Payroll and Financial Management Services: Towards Safer Small Business Transactions. Journal of Artificial Intelligence and Big Data Disciplines, 1(1), 75-85.