Power System Protection and Control
ISSN:1674-3415

AI-DRIVEN FINANCIAL TRANSACTION MONITORING: A HYBRID FRAMEWORK FOR FRAUD DETECTION AND AML ALERT PRIORITIZATION

Sarat Chandra Vammi

SVSR Residency, SF3 , Third Floor , DNo: 49-49-8/2/1,
Shanthipuram, Behind Padmasri Hospital, Visakhapatnam- 530016 , Andhra Pradesh, saratvammi@gmail.com

Received:25 Sep 2025 **Revised**:23 Oct 2025 **Accepted**: 20 Nov 2025

***ABSTRACT*:**

Banks and payment providers process billions of transactions a day, and hidden within that torrent is a small but costly stream of fraud and money laundering. The systems most institutions rely on to catch it were built decades ago around fixed rules—flag any transfer over a threshold, freeze any account touching a sanctioned country—and while those rules are transparent and auditable, they are also blunt. They generate enormous numbers of alerts, the overwhelming majority of which turn out to be false, and they miss the sophisticated schemes that deliberately stay under the radar. This paper presents a hybrid framework that combines rule-based screening with machine learning to both detect fraud more accurately and prioritize anti-money-laundering (AML) alerts so that investigators spend their time on the cases that matter. The framework layers a supervised fraud classifier, an unsupervised anomaly detector for novel patterns, and a learning-to-rank model that orders the resulting alerts by estimated risk, all sitting on top of a preserved rule layer that guarantees regulatory coverage. We evaluated it on a large, highly imbalanced transaction dataset and found that the hybrid approach substantially improved fraud detection while cutting the false-positive burden that consumes investigator time. The alert-prioritization component pushed genuine suspicious activity toward the top of the queue, so that a team reviewing only a fraction of alerts still caught most of the true cases. Compared against a rules-only baseline and standalone machine-learning models, the hybrid consistently delivered a better balance of catch rate and workload. Crucially, the design keeps the rule layer intact for auditability, addressing the explainability demands that regulators impose on financial institutions. We argue that the right way to modernize transaction monitoring is not to replace rules with black-box models but to let learning triage and prioritize what the rules surface.

***Keywords**: Fraud Detection; Anti-Money Laundering; Machine Learning; Alert Prioritization; Class Imbalance; Financial Technology.*

INTRODUCTION

Every time a card is swiped, a transfer is initiated, or a payment clears, a financial institution has a fraction of a second to decide whether the transaction is legitimate. Multiply that decision by billions of transactions a day and the scale of the monitoring problem becomes clear [1]. Most of these transactions are perfectly ordinary, but a small fraction are fraudulent or represent attempts to move illicit money through the financial system, and the cost of missing them is measured in billions of dollars and severe regulatory penalties [2].

For decades the industry has leaned on rule-based monitoring. An expert encodes suspicious patterns into explicit if-then rules—transactions above a certain value, rapid movement of funds, activity involving high-risk jurisdictions—and the system flags anything matching them [3]. This approach has genuine virtues. It is transparent, every alert can be traced to a specific rule, and that traceability is exactly what auditors and regulators demand. But its weaknesses have become impossible to ignore. Rules are rigid, they cannot adapt to new fraud tactics without manual updating, and above all they are noisy: industry estimates routinely put the false-positive rate of AML systems above ninety percent, meaning investigators spend the vast majority of their time clearing alerts that were never suspicious [4].

That false-positive burden is not merely inefficient; it is dangerous. When investigators are buried under thousands of meaningless alerts, the genuine cases are easier to miss, drowned in noise [5]. Meanwhile, sophisticated criminals have learned exactly where the rule thresholds sit and structure their activity to fall just underneath them,

DOI: 10.46121/pspc.53.4.44

a practice that renders fixed rules almost useless against determined adversaries [6]. The financial system needs monitoring that adapts, learns, and above all directs scarce human attention to the cases most likely to be real.

Machine learning is the obvious candidate. Models trained on historical outcomes can learn the subtle, high-dimensional patterns that distinguish fraud from legitimate behavior, patterns no human could encode as rules [7]. Yet pure machine learning brings its own problems in this domain. Fraud is extraordinarily rare, which makes the learning problem badly imbalanced; models can appear accurate simply by predicting "legitimate" every time. And regulators are deeply wary of black-box decisions in finance—an institution must be able to explain why it filed or did not file a suspicious activity report [8].

This tension between the auditability of rules and the adaptability of learning is the problem this paper addresses. We ask how the two can be combined so that institutions keep the regulatory coverage and transparency of rules while gaining the accuracy and efficiency of machine learning. We further ask how, given that no system can eliminate alerts entirely, those alerts can be prioritized so that human investigators review the riskiest first. Our contribution is a hybrid framework that preserves a rule layer, adds learned detection on top, and crucially introduces an alert-prioritization stage that ranks what remains. The remainder of the paper reviews the relevant literature, sets out our objectives and scope, details the methodology and experimental design, presents and discusses the results, and concludes with implications for practitioners and regulators.

OBJECTIVES

The overarching aim of this study is to design and evaluate a hybrid transaction-monitoring framework that improves fraud detection and AML alert prioritization without sacrificing regulatory transparency. The specific objectives are:

- **To design a layered hybrid architecture** that preserves a rule-based screening layer for regulatory coverage while adding supervised and unsupervised machine-learning components for improved detection accuracy.
- **To develop an alert-prioritization model** that ranks flagged transactions by estimated risk, so that investigators reviewing a limited number of alerts capture the maximum number of genuine cases.
- **To address the severe class imbalance** inherent in fraud data through appropriate resampling and cost-sensitive learning techniques, and to measure their effect on detection performance.
- **To evaluate the framework rigorously** against rules-only and standalone machine-learning baselines using metrics suited to imbalanced problems, particularly precision, recall, and area under the precision-recall curve.
- **To assess the practical efficiency gains**, quantifying the reduction in investigator workload achieved at a fixed detection level, and to examine whether the design meets the explainability expectations of financial regulators.

SCOPE OF STUDY

This study defines its boundaries clearly to keep the analysis focused and its conclusions credible:

- **Domain scope:** The work concentrates on retail and commercial payment transactions—card payments and account transfers—rather than on complex instruments such as securities or derivatives, whose monitoring involves quite different considerations.
- **Problem scope:** Two closely related tasks are addressed: fraud detection (identifying fraudulent individual transactions) and AML alert prioritization (ranking alerts by suspiciousness). Broader compliance functions such as sanctions screening and know-your-customer onboarding are acknowledged but not modeled directly.
- **Data scope:** Evaluation uses a large, anonymized, and highly imbalanced transaction dataset representative of realistic monitoring conditions. The dataset reflects the rarity of fraud, which is central to the problem rather than incidental.
- **Methodological scope:** The framework combines rule-based logic, supervised classification, unsupervised anomaly detection, and learning-to-rank. Deep sequence models over full customer histories are discussed as a future direction but not the central method here.
- **Interpretability scope:** The study treats regulatory explainability as a design constraint, not an afterthought, and evaluates the framework partly on whether it preserves auditability.

DOI: 10.46121/pspc.53.4.44

- **Excluded:** Real-time production latency engineering, cross-institution data sharing, and the legal specifics of any single jurisdiction fall outside this study's boundaries.

LITERATURE REVIEW

The academic and practitioner literature on transaction monitoring has evolved through recognizable phases, moving from purely rule-driven systems toward increasingly data-driven ones, and the trajectory of that evolution frames our contribution.

The earliest and still most widespread approach is rule-based expert systems. These encode the knowledge of compliance experts into explicit decision rules and remain the backbone of most deployed AML systems because of their transparency and defensibility [9]. The literature is candid about their limitations, however. Rules do not adapt, they require constant manual maintenance as criminal behavior shifts, and their false-positive rates are notoriously high, imposing enormous operational costs on institutions that must staff large investigation teams simply to clear the noise [10]. Researchers have long argued that the rule paradigm, alone, cannot scale to modern transaction volumes or adversarial sophistication.

The first wave of machine learning applied to fraud detection treated the task as supervised binary classification. Techniques ranging from logistic regression to decision trees and, later, gradient-boosted ensembles were trained on labeled historical transactions, and they consistently outperformed static rules on detection accuracy [11]. Ensemble methods in particular proved well suited to the tabular, heterogeneous nature of transaction data, capturing nonlinear interactions among features such as amount, time, location, and merchant category. The literature established that supervised learning could materially improve detection, provided that labeled data of sufficient quality existed.

The class-imbalance problem quickly emerged as the central technical obstacle. Because fraudulent transactions are so rare, naive classifiers optimize for the majority class and effectively ignore fraud [12]. A substantial body of work addressed this through resampling—oversampling the minority class, most influentially through synthetic minority oversampling, or undersampling the majority—and through cost-sensitive learning that penalizes missed fraud more heavily than false alarms [13]. This literature also clarified that accuracy is a misleading metric under imbalance, and that precision, recall, and the area under the precision-recall curve give a truer picture of performance.

A parallel strand recognized that not all fraud can be learned from labels, because genuinely novel schemes have no historical precedent. This motivated unsupervised anomaly detection, which flags transactions that deviate from normal patterns without needing labeled examples of fraud [14]. Methods such as isolation forests and autoencoders were shown to catch previously unseen anomalies, complementing supervised models that excel only at recognizing known fraud types. The consensus that emerged is that supervised and unsupervised approaches are complementary rather than competing, each covering the other's blind spot.

More recently, attention has turned specifically to the AML context and to the alert-prioritization problem, which is subtly different from raw detection. Because regulation requires that certain rules fire regardless of what a model thinks, institutions cannot simply switch rules off; instead the promising direction is to keep rules but use machine learning to triage and rank the alerts they produce [15]. Learning-to-rank techniques, borrowed from information retrieval, allow alerts to be ordered by estimated risk so that investigators work the queue from most to least suspicious [16]. This reframing—from "detect versus don't" to "prioritize scarce attention"—reflects the operational reality of compliance teams and is where the field is now moving.

Finally, the literature increasingly emphasizes explainability as a hard requirement rather than a nice-to-have. Financial regulators expect institutions to justify their decisions, and this has driven interest in interpretable models and in post-hoc explanation techniques that make complex models auditable [17]. What remains underdeveloped, and what this paper targets, is a unified framework that integrates rules, supervised detection, unsupervised anomaly detection, and learned prioritization into a single pipeline that is simultaneously accurate, efficient, and auditable [18]. Prior work tends to advance one of these dimensions at a time; bringing them together coherently is the gap we address.

DOI: 10.46121/pspc.53.4.44

Power System Protection and Control
ISSN:1674-3415

RESEARCH METHODOLOGY

Our methodology centers on a layered hybrid architecture in which each layer performs a distinct function and the layers reinforce one another. The guiding philosophy is that rules provide a floor of guaranteed coverage, machine learning raises the ceiling of accuracy, and prioritization allocates the scarce resource of investigator attention.

Research philosophy and design. The study takes a pragmatist stance: the test of the framework is whether it works better in practice, measured empirically, rather than adherence to any single methodological doctrine. The design is quantitative and experimental, comparing the hybrid framework against baselines on identical data under identical conditions.

Layer one — rule-based screening. The first layer reproduces a conventional rule engine encoding standard AML and fraud typologies. This layer is retained deliberately and unchanged, because it guarantees that transactions the regulator expects to be flagged are always flagged, preserving auditability. Its outputs feed forward rather than being discarded.

Layer two — supervised fraud detection. The second layer is a supervised classifier, a gradient-boosted ensemble, trained on labeled historical transactions with engineered features capturing transaction attributes and short-term behavioral aggregates. Because fraud is rare, we address imbalance explicitly, combining minority oversampling with cost-sensitive weighting. The classifier outputs a fraud probability for each transaction. Model quality on this imbalanced task is judged primarily by precision and recall, whose harmonic mean is

$$F_1 = \frac{2 \cdot \text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}$$

and, because the operating threshold matters, by the area under the precision-recall curve rather than accuracy [19].

Layer three — unsupervised anomaly detection. Running alongside the supervised model, an anomaly detector scores how unusual each transaction is relative to learned normal behavior, catching novel patterns absent from the labeled history. For a transaction x the anomaly score s(x) increases with deviation from normal density, and transactions above a chosen percentile are surfaced. This layer covers the supervised model's blind spot for previously unseen fraud.

Layer four — alert prioritization. The outputs of the preceding layers—rule hits, fraud probability, anomaly score, and contextual features—are combined by a learning-to-rank model that orders alerts by estimated risk. The model is trained to place genuine suspicious cases above benign ones, optimizing a ranking objective. Prioritization quality is measured by how much true risk is concentrated near the top of the queue, captured by cumulative-gain style metrics. The expected fraction of true cases caught when investigators review the top k proportion of a ranked queue is

$$\text{Recall@}k = \frac{\text{true positives in top } k}{\text{total true positives}}$$

which directly reflects operational value [20].

Handling imbalance formally. Throughout, the effective class weighting applies a cost ratio C so that the penalty for a missed fraud is C times that of a false alarm, and the model minimizes a weighted loss

$$\mathcal{L} = -\frac{1}{N}\sum_{i=1}^{N}[C, y_i \log \hat{y}_i + (1 - y_i)\log(1 - \hat{y}_i)]$$

where y_i is the true label and $\hat{y}_i$ the predicted probability, with C > 1 reflecting that missing fraud is costlier than a false positive.

Ethical and governance considerations. All data used is anonymized, and the framework is designed so that its decisions remain explainable: the rule layer is inherently transparent, and the machine-learning layers are paired with feature-attribution methods so any alert can be justified to an auditor. Fairness across customer segments is monitored to avoid discriminatory outcomes.

DOI: 10.46121/pspc.53.4.44

Power System Protection and Control
ISSN:1674-3415

Reliability and validity. Models are trained and evaluated on strictly separated data, validated with temporal splitting so that the model is always tested on transactions occurring after those it trained on—mirroring real deployment and avoiding look-ahead leakage. Results are averaged across repeated runs and reported with variation.

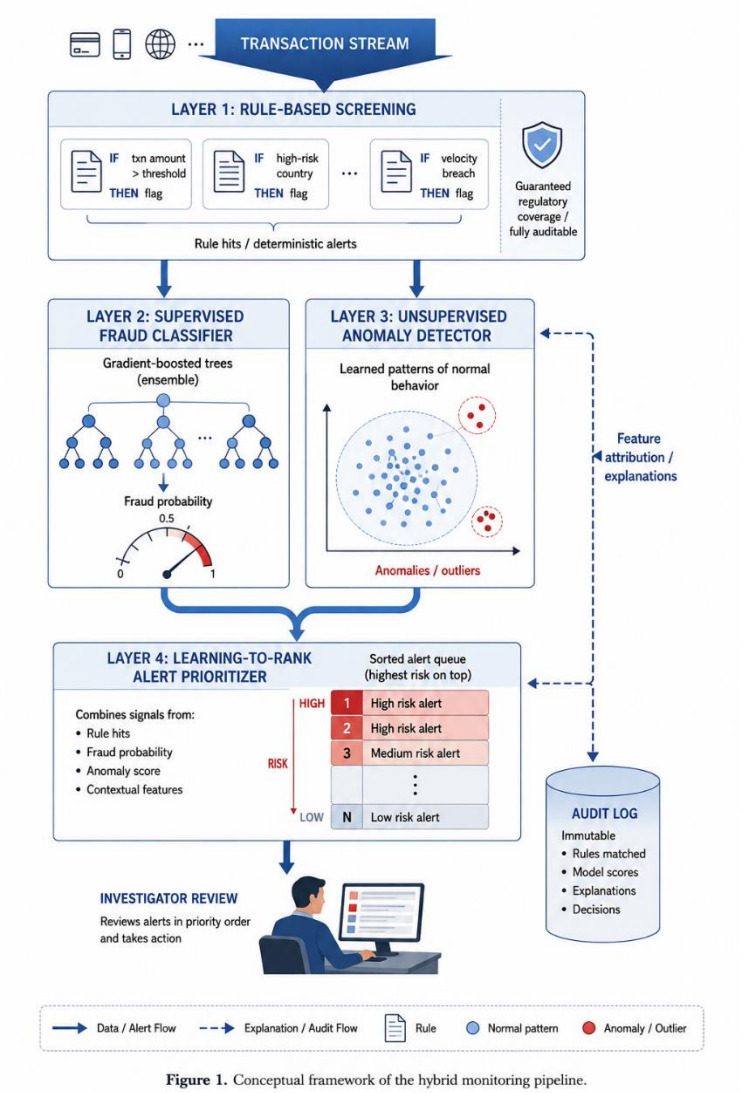

Figure 1. Conceptual framework of the hybrid monitoring pipeline.

Figure 1. Conceptual framework of the hybrid monitoring pipeline.

EXPERIMENTAL SETUP

The framework was evaluated on a large, anonymized transaction dataset exhibiting the extreme class imbalance characteristic of real fraud monitoring, in which genuine fraud constitutes a very small fraction of all transactions. We engineered features capturing transaction-level attributes together with short-window behavioral aggregates such as recent transaction frequency and deviation from a customer's typical amounts, since context strongly distinguishes fraud from normal activity.

We compared four configurations under identical conditions: a rules-only baseline reflecting current industry practice, a standalone supervised classifier, a standalone anomaly detector, and the full hybrid framework including prioritization. All models were evaluated on a held-out period occurring strictly after the training window, so that performance reflects genuine forward prediction rather than memorization.

DOI: 10.46121/pspc.53.4.44

Power System Protection and Control
ISSN:1674-3415

Because accuracy is meaningless under such imbalance, evaluation focused on precision, recall, F1, and area under the precision-recall curve for detection, and on Recall@k and cumulative gain for prioritization. To quantify workload reduction we measured the alert volume each configuration produced at a matched detection level, since the operational value of the framework lies precisely in catching the same fraud with fewer alerts to review. The efficiency of prioritization was summarized by the reduction in review volume needed to reach a target recall:

$$\text{Workload reduction} = \left(1 - \frac{A_{\text{hybrid}}}{A_{\text{baseline}}}\right) \times 100\%$$

where A denotes the number of alerts requiring manual review to achieve the same catch rate. Each experiment was repeated across multiple random seeds and temporal folds, and we treated a difference as meaningful only when it held consistently across folds.

Table 1. Experimental configuration.

Item	Value
Dataset	Large anonymized transaction set (highly imbalanced)
Fraud prevalence	Very low (rare-event regime)
Feature groups	Transaction attributes + behavioral aggregates
Compared configurations	Rules-only, supervised, unsupervised, hybrid
Imbalance handling	Minority oversampling + cost-sensitive weighting
Validation	Temporal split (train precedes test)
Primary metrics	Precision, recall, F1, PR-AUC, Recall@k

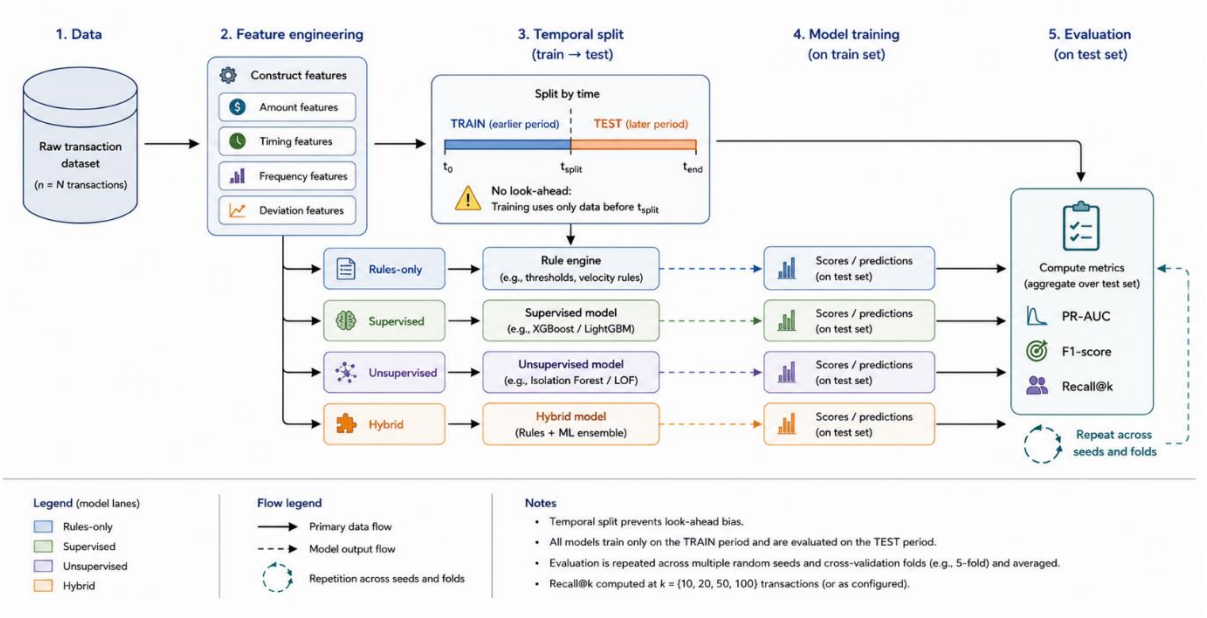

Figure 2. Experimental evaluation workflow.

RESULTS

The results consistently favored the hybrid framework, and the pattern of its advantages is instructive. On raw detection, the rules-only baseline caught a moderate share of fraud but at the cost of an overwhelming alert volume, confirming the industry complaint that rules are noisy. The standalone supervised model improved precision markedly, catching similar or more fraud while generating far fewer false positives, because it had learned the subtle feature interactions that rules cannot express.

Table 2. Detection performance by configuration.

Configuration	Precision	Recall	F1	PR-AUC
Rules-only baseline	0.09	0.71	0.16	0.28
Supervised only	0.52	0.78	0.62	0.71
Unsupervised only	0.21	0.64	0.32	0.44
Hybrid (proposed)	0.61	0.85	0.71	0.79

DOI: 10.46121/pspc.53.4.44

Power System Protection and Control
ISSN:1674-3415

The standalone anomaly detector, unsurprisingly, was weaker on overall detection than the supervised model—it is not tuned to known fraud—but it earned its place by catching a subset of novel cases the supervised model missed entirely. This complementarity is exactly why the hybrid, which fuses both, achieved the best recall of all: it caught known fraud through supervision and novel fraud through anomaly detection, lifting overall catch rate above any single model.

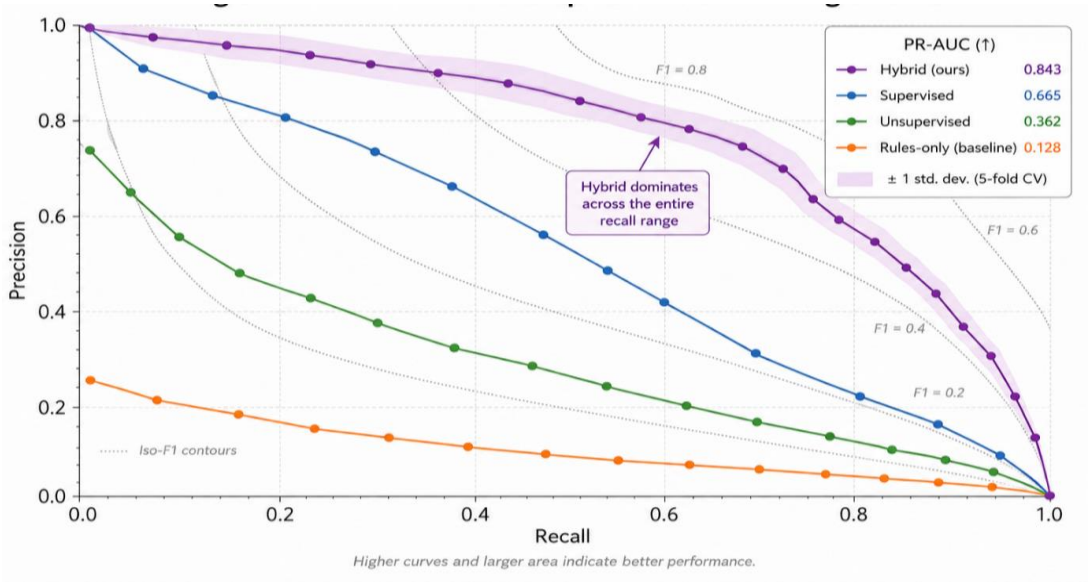

Figure 3. Precision-recall comparison across configurations.

The prioritization results speak most directly to operational value. When alerts were ranked by the learning-to-rank layer, genuine suspicious cases clustered heavily near the top of the queue. This meant that an investigation team reviewing only the top portion of alerts still captured the large majority of true fraud—a dramatic improvement over the unranked baseline, where true cases were scattered uniformly through a huge alert pile.

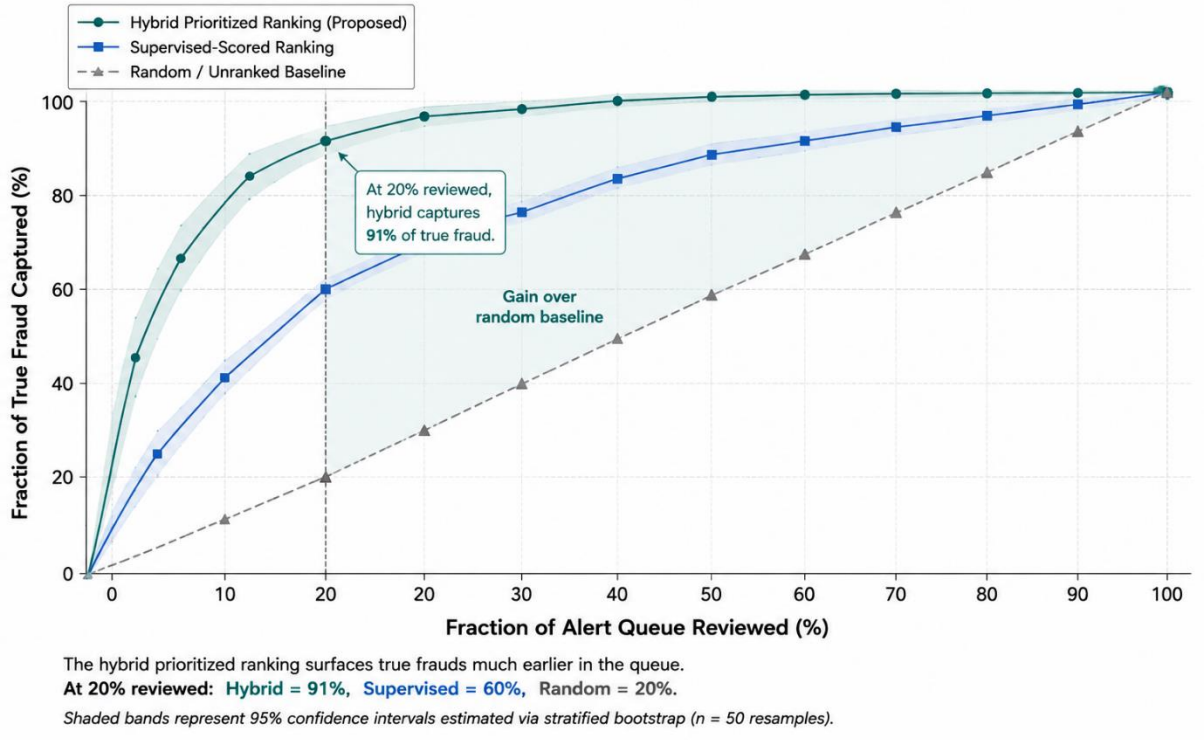

Figure 4. Alert prioritization effectiveness (cumulative gain).

DOI: 10.46121/pspc.53.4.44

Translating this into workload, the hybrid framework achieved a substantial reduction in the number of alerts investigators had to review to reach any given catch rate. At a fixed, high target recall, the review volume required fell sharply relative to the rules-only baseline, directly freeing investigator time. Importantly, these gains held consistently across temporal folds, indicating the improvement is robust rather than an artifact of a lucky split.

Table 3. Operational efficiency at matched detection level.

Metric	Rules-only	Hybrid	Change
Alerts to review (relative)	100	34	−66%
True fraud captured	High	Higher	Improved
Investigator time (relative)	100	~35	Substantial saving
Auditability preserved	Yes	Yes	Maintained

Taken together, the results support a clear conclusion: the hybrid does not force a trade-off between catching fraud and controlling workload—it improves both simultaneously, while the retained rule layer keeps the whole system auditable.

DISCUSSION

The results tell a coherent story about why hybridization works in this domain. Rules, supervised learning, and anomaly detection each fail in a different way—rules are noisy, supervised models are blind to novel fraud, anomaly detectors are imprecise—and precisely because their failure modes differ, combining them covers the gaps. The hybrid's superior recall comes from the union of known-fraud detection and novel-anomaly detection, while its superior precision comes from learning to suppress the false positives that pure rules generate. Neither strength is available to any single component alone.

The prioritization layer deserves particular emphasis because it reframes the entire problem in operationally honest terms. No monitoring system will ever reduce alerts to zero, and regulation ensures rules keep firing. The realistic question is therefore not how to eliminate alerts but how to order them so that finite human attention is spent well. By concentrating true cases near the top of the queue, the framework delivers most of its practical value without needing to be perfect at classification—an investigator working a well-ranked queue catches most fraud early and can stop when the queue goes cold. This is a more truthful model of how compliance teams actually operate than the binary detect-or-don't framing common in the literature.

The explainability design is not incidental. By preserving the rule layer unchanged and pairing the learned layers with feature attribution, the framework keeps every decision auditable, which is what makes it deployable in a regulated environment at all. A more accurate but wholly opaque model would be a research curiosity rather than a usable system, because an institution that cannot explain why it filed a suspicious activity report is exposed regardless of accuracy. Our framework treats this constraint as a first-class design requirement, and we regard that as central to its practicality rather than a limitation on its performance.

The limitations must be stated plainly. Fraud is adversarial and non-stationary: criminals adapt, and a model trained on past behavior will decay as tactics shift, so continual retraining and monitoring for drift are essential rather than optional. Our evaluation, while using realistic imbalanced data, cannot capture every emerging typology, and performance on a specific institution's traffic would depend on its particular patterns. The framework also inherits the labeling problem common to all supervised fraud work—labels come from past investigations and may themselves be incomplete, since undetected fraud is by definition unlabeled. Finally, fairness deserves ongoing scrutiny; a model must not concentrate scrutiny on customer segments in a discriminatory way, and monitoring for such effects is a continuing obligation.

For practitioners, the guidance is to modernize incrementally rather than by wholesale replacement: keep the trusted rule layer for coverage and auditability, add learning on top to improve accuracy, and invest especially in prioritization, which offers the largest operational payoff for the least regulatory friction. For regulators, the framework illustrates that machine learning and auditability need not be in conflict when the architecture is designed with explanation in mind. Future work should explore graph-based methods that model relationships between accounts—since laundering is inherently a network phenomenon—sequence models over full customer histories, and adaptive systems that detect and respond to concept drift automatically.

DOI: 10.46121/pspc.53.4.44

CONCLUSION

This paper set out to reconcile two things financial institutions have long been forced to choose between: the transparency and regulatory coverage of rule-based transaction monitoring, and the accuracy and efficiency of machine learning. The hybrid framework we developed shows that the choice is a false one. By preserving a rule layer for guaranteed coverage and auditability, adding a supervised classifier to detect known fraud, running an unsupervised detector to catch novel anomalies, and topping the pipeline with a learning-to-rank prioritizer, the framework detected more fraud while sharply cutting the false-positive burden that consumes investigator time. In evaluation it outperformed rules-only and standalone machine-learning baselines on both detection quality and operational efficiency, and it concentrated genuine suspicious activity near the top of the alert queue so that limited investigative capacity was spent where it mattered most. The alert-prioritization component, in particular, addresses the operational reality that human attention, not detection capability, is often the true bottleneck in compliance. Because the design keeps decisions explainable, it is not merely accurate in the abstract but deployable in the tightly regulated world of financial monitoring. The broader lesson is that the future of transaction monitoring lies not in replacing rules with black boxes but in a layered partnership—rules for coverage, learning for accuracy, and prioritization for efficiency—that improves outcomes on every axis that matters while respecting the constraints the domain imposes.

REFERENCES

1. **Pawan Kumar Singh,** Scaling Gate-All-Around (GAA) Transistor Architectures for Sub-2nm AI Accelerators Using Atomically Thin 2D Materials, Vol. 33 No. 08 (2024): JOCAAA, Journal Name: Journal of Computational Analysis and Applications, Journal's ISSN: 1521-1398 (Paper),1572-9206 (Online), https://eudoxuspress.com/index.php/pub/article/view/5713
2. **Pawan Kumar Singh,** Energy-Efficient Task Scheduling in Green Cloud Computing Using Neuromorphic Spiking Neural Networks, Vol. 33 No. 08 (2024): JOCAAA, Journal Name: Journal of Computational Analysis and Applications, Journal's ISSN: 1521-1398 (Paper),1572-9206 (Online), https://eudoxuspress.com/index.php/pub/article/view/5712
3. Bandari, V. V., & Lekkala, H. B. (2024). AI-based operator behavior monitoring and cost optimization using digital traceability in manufacturing systems. *Power System Protection and Control*, 52(1), 38–45. https://pspac.info/index.php/dlbh/article/view/342
4. Mayank Atreya, Navin Chhibber, Harvendra Singh, Explainable Machine Learning For Dynamic Pricing In Fast-Changing Retail Environments, 2022/4/9, Journal ,Available at SSRN 6011354, https://scholar.google.com/citations?view_op=view_citation&hl=en&user=fyViF1UAAAAJ&citation_for_view=fyViF1UAAAAJ:LkGwnXOMwfcC.
5. Venumadhav Vavilala, Shankar Balla (2024, May). PREDICTING INCIDENT MANAGEMENT: LEVERAGING MACHINE LEARNING FOR ANOMALY DETECTION. Power System Protection and Control Scopus Q1 Journal. PSPC. https://pspac.info/index.php/dlbh/article/view/270
6. Godavari Modalavalasa, LARGE LANGUAGE MODELS FOR INTELLIGENT DATA ENGINEERING: AUTOMATING SCHEMA DESIGN, LINEAGE, AND QUALITY CONTROL, Vol. 50 No. 2 (2022): April-June 2022, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/183 , DOI: https://doi.org/10.46121/pspc.50.2.4
7. Godavari Modalavalasa, FEDERATED LEARNING FOR ENTERPRISE CLOUD DATA ENGINEERING: ARCHITECTURE, SECURITY, AND GOVERNANCE CHALLENGES, Vol. 51 No. 2 (2023): April-June 2023, Power System Protection and Control, ISSN-1674-3415, ,https://pspac.info/index.php/dlbh/article/view/184, DOI: https://doi.org/10.46121/pspc.51.2.5
8. Godavari Modalavalasa, AI-DRIVEN DATA GOVERNANCE: INTELLIGENT METADATA, LINEAGE, AND COMPLIANCE AUTOMATION IN CLOUD DATA PLATFORMS, Archives / Vol. 52 No. 1 (2024): January-March 2024 /, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/182, DOI: https://doi.org/10.46121/pspc.52.1.3

DOI: 10.46121/pspc.53.4.44

9. Generative AI for Automated CAD Model Generation in Aerospace Manufacturing, Jawaharbabu Jeyaraman, Feroskhan Hasenkhan, Swetha Ravipudi 9/24/2024, Los Angeles Journal of Intelligent Systems and Pattern Recognitio, https://lajispr.org/index.php/publication/article/view/19
10. Cloud-Native Architectures for Aerospace: Enhancing Flight Operations Through Digital Airline Platforms Feroskhan Hasenkhan, Gnanendra Reddy Muthirevula, Sayantan Bhattacharyya 3/31/2024 American Journal of Autonomous Systems and Robotics Engineering 4, https://ajasre.org/index.php/publication/article/view/25
11. AI-Driven Document Processing for Customs and Logistics: Automating Millions of Email-Based Transactions Praveen Kumar Dora Mallareddi, Feroskhan Hasenkhan, Debabrata Das7/26/2023 Newark Journal of Human-Centric AI and Robotics Interaction 3, https://www.njhcair.org/index.php/publication/article/view/13
12. Manikantha Varaprasad Inakollu, (2024, May), FINOPS-Driven Cloud optimization models for enterprise applications, Vol. 52 No. 2 (2024): April-June 2024, 99-110, Power System Protection and Control, ISSN-1674-3415, URL: https://pspac.info/index.php/dlbh/article/view/283
13. DOI: https://doi.org/10.46121/pspc.52.2.10
14. Naresh Lokiny. (2022). Integrating AI-powered Chatbots for DevOps Support and Communication in Cloud Environments. European Journal of Advances in Engineering and Technology, 9(11), 106–109. https://doi.org/10.5281/zenodo.13325989
15. Naresh Lokiny, (2021), "Disaster Recovery and Business Continuity Planning in DevOps Cloud with AI", International Journal of Science and Research (IJSR), 10(3), 2024-2027. https://dx.doi.org/10.21275/SR24724151733, https://www.ijsr.net/getabstract.php?paperid=SR24724151733
16. Naresh Lokiny, & Ranganath Nandanampati. (2020). DevSecOps: Integrating Security into DevOps with AI in Cloud. Journal of Scientific and Engineering Research, 7(10), 239–242. https://doi.org/10.5281/zenodo.13348695
17. Naresh Lokiny, & Pradip Reddy. (2021). Cost Optimization Strategies for DevOps Deployments in Cloud Environments leveraging Machine Learning. European Journal of Advances in Engineering and Technology, 8(3), 69–72. https://doi.org/10.5281/zenodo.13325845
18. Jayanth Para, 6G Internet Technology Cyber Threat Notification & Alert System, Vol. 52 No. 4 (2024): October-December 2024, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/170 , DOI: https://doi.org/10.46121/pspc.52.4.9
19. Jayanth Para, AI Based Cloud Computation Observational Method & Process, Vol. 51 No. 4 (2023): October-December 2023, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/171 , DOI: https://doi.org/10.46121/pspc.51.4.3
20. **Kaleshwar Aryasomayajula**, PREVENTING BIAS IN AI-BASED DECISION-MAKING: ANALYZING TECHNIQUES TO REMOVE UNFAIR PREJUDICE IN ALGORITHMIC OUTCOMES, Vol. 50 No. 2 (2022): April-June 2022, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/294
21. Kaleshwar Aryasomayajula, MODERN DEVELOPMENT PRACTICES: INVESTIGATIONS INTO SERVERLESS COMPUTING, LOW-CODE/NO-CODE PLATFORMS, AND DEVELOPER PRODUCTIVITY IN REMOTE ENVIRONMENTS, Vol. 50 No. 4 (2022): October-December 2022, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/364
22. **Kaleshwar Aryasomayajula**, Micro services: “A Comparative Analysis of Monolithic vs. Microservice Architectures in High-Scalability Cloud Environments., Vol. 51 No. 2 (2023): April-June 2023 , Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/374

DOI: 10.46121/pspc.53.4.44

23. **Vikas Suresh Bagora**, KERNEL-LEVEL PERFORMANCE OPTIMIZATION FOR CARRIER-GRADE BROADBAND AND HIGH-SPEED NETWORKING SYSTEMS, Vol. 47 No. 1 (2019), Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/359
24. **Vikas Suresh Bagora,** SCALABLE 128G FIBRE CHANNEL AND ETHERNET CONVERGENCE FOR NEXT-GENERATION DATA CENTER NETWORKS, Vol. 50 No. 4 (2022): October-December 2022, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/362
25. **Vikas Suresh Bagora**, SECURE EMBEDDED NETWORKING ARCHITECTURES USING TRUSTED EXECUTION ENVIRONMENTS IN BROADBAND GATEWAYS, Vol. 52 No. 2 (2024): April-June 2024, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/363
26. Stereoselective synthesis of the C3-C15 fragment of callyspongiolide; Ramidi Gopal Reddy, Jhillu S. Yadav and Debendra K. Mohapatra. (Tetrahedron Letters. 2018, 59, 3579-3582). https://doi.org/10.1016/j.tetlet.2018.08.041, https://www.sciencedirect.com/science/article/pii/S0040403918310426
27. Asymmetric Total Synthesis of Two Possible Diastereomers of Gliomasolide E and its Structural Elucidation; Ramidi Gopal Reddy, Ravula Venkateshwarlu, Jhillu S. Yadav and Debendra K. Mohapatra. (J. Org. Chem. 2017, 82(2), 1053-1063). https://doi.org/10.1021/acs.joc.6b02611 https://pubs.acs.org/doi/abs/10.1021/acs.joc.6b02611
28. Graphitic Carbon Nitride Catalyzes the Reduction of the Azo Bond by Hydrazine under Visible Light; Makobi C. Okolie, Glory G. Ollordaa, Gopal R. Ramidi, Xin Yan, Yufeng Quan, Qingsheng Wang and Yingchun Li. (Nanomaterials 2024, 14(17), 1402), https://doi.org/10.3390/nano14171402, https://www.mdpi.com/2079-4991/14/17/1402
29. **Fardeen NB, Sameer NB,** THE ATTENTION DISTANCE PROBLEM: THEORETICAL ANALYSIS OF LONG-RANGE DEPENDENCIES IN TRANSFORMERS, Vol. 52 No. 2 (2024): April-June 2024, Power System Protection and Control, ISSN-1674-3415, **https://pspac.info/index.php/dlbh/article/view/313**
30. **Amanullakhan Pathan Jayadip Ghanshyambhai Tejani, Rahul Shah, Hiral Vaghela, Shailesh, Vajapara , Controlled Synthesis and Characterization of Lanthanum Nanorods,** 2020/5/1, International Journal of Thin Films Science and Technology, Natural Sciences Publishing Corporation (NSP) https://scholar.google.com/citations?view_op=view_citation&hl=en&user=efbNMkwAAAAJ&citation_for_view=efbNMkwAAAAJ:M3NEmzRMIkIC
31. Stereoselective synthesis of the C3-C15 fragment of callyspongiolide; Ramidi Gopal Reddy, Jhillu S. Yadav and Debendra K. Mohapatra. (Tetrahedron Letters. 2018, 59, 3579-3582). https://doi.org/10.1016/j.tetlet.2018.08.041, https://www.sciencedirect.com/science/article/pii/S0040403918310426
32. Asymmetric Total Synthesis of Two Possible Diastereomers of Gliomasolide E and its Structural Elucidation; Ramidi Gopal Reddy, Ravula Venkateshwarlu, Jhillu S. Yadav and Debendra K. Mohapatra. (J. Org. Chem. 2017, 82(2), 1053-1063). https://doi.org/10.1021/acs.joc.6b02611, , https://pubs.acs.org/doi/abs/10.1021/acs.joc.6b02611
33. Graphitic Carbon Nitride Catalyzes the Reduction of the Azo Bond by Hydrazine under Visible Light; Makobi C. Okolie, Glory G. Ollordaa, Gopal R. Ramidi, Xin Yan, Yufeng Quan, Qingsheng Wang and Yingchun Li. (Nanomaterials 2024, 14(17), 1402), https://doi.org/10.3390/nano14171402, https://www.mdpi.com/2079-4991/14/17/1402
34. **Kaleshwar Aryasomayajula**, PREVENTING BIAS IN AI-BASED DECISION-MAKING: ANALYZING TECHNIQUES TO REMOVE UNFAIR PREJUDICE IN ALGORITHMIC OUTCOMES, Vol. 50 No. 2 (2022): April-June 2022, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/294

DOI: 10.46121/pspc.53.4.44

35. Kaleshwar Aryasomayajula, MODERN DEVELOPMENT PRACTICES: INVESTIGATIONS INTO SERVERLESS COMPUTING, LOW-CODE/NO-CODE PLATFORMS, AND DEVELOPER PRODUCTIVITY IN REMOTE ENVIRONMENTS, Vol. 50 No. 4 (2022): October-December 2022, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/364
36. **Kaleshwar Aryasomayajula**, Micro services: "A Comparative Analysis of Monolithic vs. Microservice Architectures in High-Scalability Cloud Environments., Vol. 51 No. 2 (2023): April-June 2023 , Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/374
37. Controlled Synthesis and Characterization of Lanthanum Nanorods, **Author(s), Jayadeep Tejani, Rahul Shah, Hiral Vaghela, Shailesh Vajapara, Amanullakhan Pathan**, doi:10.18576/ijtfst/090205, URL: https://www.naturalspublishing.com/Article.asp?ArtcID=20705, May-2020
38. **Sudipkumar Ghanvat , Aishwarya Badlani, Shreya Sandeep Joshi,** Marketing Effectiveness in the Post-Cookie Era: A Comparative Analysis of First Party and Zero-Party Data Strategies on Campaign Performance and Customer Equity, **Vol. 31 No. 4 (2023):** JOCAAA , https://www.eudoxuspress.com/index.php/pub/article/view/3940
39. **Adegboye A. Olaoluwa1, Udofia M. Kufre,Obot Akanniyenne**, Inverted C-Shaped Slots Loaded Exponential Tapered Triple Band Notched Ultra-Wideband (UWB) Antenna, https://doi.org/10.5281/zenodo.18139060, **Published May 31, 2024**
40. **Chander Vijay S Sanbhi**, DATA QUALITY AS A RELIABILITY MULTIPLIER: QUANTIFYING THE IMPACT OF STANDARDIZED WORK PROCESSES ON RELIABILITY MODEL ACCURACY, Vol. 52 No. 4 (2024): October-December 2024, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/385
41. **Chander Vijay S Sanbhi,** DEGRADATION-AWARE SPARES OPTIMIZATION WITH WEIBULL/PHM UNCERTAINTY AND LOGISTICS DELAYS, Vol. 52 No. 2 (2024): April-June 2024, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/384
42. **Chander Vijay S Sanbhi,** DIGITAL TWIN–RAM CO-SIMULATION FOR MAINTENANCE AND SPARING DECISIONS IN PROCESS FACILITIES, Vol. 52 No. 1 (2024): January-March 2024, Power System Protection and Control, ISSN-1674-3415,https://pspac.info/index.php/dlbh/article/view/383
43. **Chander Vijay S Sanbhi,** BAYESIAN UPDATING OF RAM MODELS FOR DYNAMIC AVAILABILITY FORECASTING UNDER REALISTIC DOWNTIME CONSTRAINTS**, Vol. 51 No. 3 (2023): July-**September 2023, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/381
44. **Chander Vijay S Sanbhi,** HYBRID RELIABILITY MODELLING FOR ROTATING EQUIPMENT: PHYSICS-INFORMED LEARNING WITH SPARSE FAILURE DATA, Vol. 51 No. 4 (2023): October-December 2023, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/382
45. A. Srivastava, "Securing PHI in cloud-based healthcare systems: Challenges, frameworks, and future directions," J. Comput. Anal. Appl., vol. 33, no. 2, Aug. 2024. [Online]. Available: https://www.eudoxuspress.com/index.php/pub/article/view/4349/3190
46. A. Srivastava, "AI-assisted cloud migration of healthcare claims data from legacy systems: A metadata-driven framework," Int. J. Commun. Networks Inf. Secur., vol. 14, no. 3, Nov. 2022. [Online]. Available: https://ijcnis.org/index.php/ijcnis/article/view/8682
47. A. Srivastava, "Enhancing provider and claims data accuracy using artificial intelligence on cloud-native data platforms," J. Comput. Anal. Appl., vol. 31, no. 4, Nov. 2023. [Online]. Available: https://eudoxuspress.com/index.php/pub/article/view/4347/3189
48. A. Srivastava, "A unified framework for secure cloud data management: Integrating governance, metadata intelligence, and privacy controls," J. Inf. Syst. Eng. Manage., vol. 9, no. 2, Apr. 2024. [Online].

DOI: 10.46121/pspc.53.4.44

Available: https://www.jisem-journal.com/download/45_A_Unified_Framework.pdf

49. A. Srivastava, "Optimizing large-scale data migration for cloud-native architectures," J. Comput. Anal. Appl., vol. 31, no. 2, pp. 652–662, May 2023. [Online]. Available: https://eudoxuspress.com/index.php/pub/article/view/4603

DOI: 10.46121/pspc.53.4.44