International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 10 Issue: 11 | Nov - 2026 SJIF Rating: 8.659 ISSN: 2582-3930

Real-Time Payment Fraud Scoring and Agentic Anti-Money-Laundering Investigation Platform

Anshuman Sinha[1], Pradeepkumar[2], Sachin[3],Sivabakkiyan[4],Sri Harini[5]

Sri Shakthi Institue of Engineering and Technology,Coimbatore

---

Abstract

Real time payment fraud scoring angd agentic anti money laundering investigation platform is a proposed financial-crime analytics platform combining real-time payment fraud scoring with an agent-assisted anti-money-laundering (AML) investigation workflow. Its fraud-scoring path is designed to use machine-learning models, engineered transaction features, online feature retrieval, and explainability outputs to identify potentially suspicious payments. An entity-graph layer is intended to reveal relationships among accounts, devices, IP addresses, and transactions that may indicate coordinated activity. A multi-agent AML workflow is designed to support link analysis, transaction-pattern review, adverse-media retrieval, investigation narrative drafting, and quality checks. Human review remains central to consequential decisions and suspicious activity report submission. This paper summarises the system concept, architecture, methodology, reported implementation progress, evaluation approach, limitations, and planned next steps. No numerical performance results are claimed because verified experimental metrics were not provided.

Key Words: payment fraud detection, anti-money laundering, machine learning, graph analytics, explainable AI, agentic workflow.

1. INTRODUCTION

Digital payment systems enable fast and convenient transfers, but their scale and speed can also be exploited for fraud and money laundering. Detection systems must assess transactions quickly while helping analysts understand why an event was flagged. Separate tools for scoring, relationship analysis, and investigation can fragment workflows and increase manual effort. Real time payment fraud scoring angd agentic anti money laundering investigation platform proposes a connected pipeline for transaction risk scoring, explainable decisions, graph-based context, and AI-assisted investigation.

2. PROBLEM STATEMENT

The project addresses five practical challenges: detecting suspicious payments under low-latency constraints; reducing avoidable false positives while retaining useful detection coverage; identifying linked entities and potential fraud rings; producing understandable reason codes for model outputs; and helping investigators organise evidence into reviewable case narratives. The platform is intended to assist investigation rather than autonomously determine guilt or submit regulatory reports without human approval.

3. OBJECTIVES

- Develop a machine-learning-based payment risk-scoring workflow for near-real-time assessment.
- Support consistent offline feature construction and online feature retrieval, with parity checks between both paths.
- Represent relationships among accounts, devices, IP addresses, and transactions through entity-graph features.
- Use explainability methods such as SHAP to provide reason codes for reviewers.
- Design an agentic AML workflow for evidence gathering, link analysis, narrative drafting, and quality review.
- Establish evaluation, model tracking, auditability, and human-review checkpoints.

4. PROPOSED SYSTEM ARCHITECTURE

Data and features: transaction records are cleaned and transformed into consistent feature definitions. An offline builder supports model development and evaluation, while an online path is designed to retrieve current features through Redis. PostgreSQL is identified as the offline storage layer in the project specification.

Fraud scoring: gradient-boosted tree models such as XGBoost or LightGBM are considered for risk scoring. Calibration and experiment tracking with MLflow support model comparison and monitoring. The design includes a scoring API and a target p95 response latency below 100 ms; this is a design target, not a verified benchmark.

Explainability and graph context: SHAP-based reason codes are intended to communicate influential model features. An entity graph links relevant accounts, devices, IP addresses, and transactions; graph-derived features can add relationship context to transaction-level scores.

Agentic AML investigation: a multi-agent workflow, proposed using LangGraph, is designed to coordinate link analysis, transaction-pattern review, adverse-media retrieval through retrieval-augmented generation (RAG), case narrative drafting, and quality assurance. Human approval is required for final consequential actions, including suspicious activity report (SAR) submission.

5. METHODOLOGY

- Step 1 — Feature definitions: document transaction features so offline training and online inference use consistent meanings.
- Step 2 — Offline processing: build reproducible feature datasets for model training and evaluation.
- Step 3 — Online feature path: retrieve features through Redis and run parity tests against the offline feature path.
- Step 4 — Entity graph: construct relationships among entities and calculate graph-based features.
- Step 5 — Model development: train and compare fraud models, calibrate scores, and track experiments with MLflow.
- Step 6 — Evaluation and governance: evaluate model quality and latency, and define champion/challenger model registry policies.

- Step 7 — Investigation workflow: integrate agent-assisted evidence gathering, draft generation, quality checks, and human review as subsequent implementation work.

6. IMPLEMENTATION PROGRESS

Based on the project progress information supplied for this report, the following components have been implemented or worked on:

- Feature definitions module and offline feature builder.
- Online Redis feature path and offline/online parity tests.
- Entity graph and graph-derived features.
- Fraud models, calibration, and MLflow experiment tracking.
- Evaluation suite and champion/challenger registry policy.

Status boundary: this list reflects implementation progress reported by the project team. The scoring API, complete AML agent workflow, React investigation console, security/audit integration, and end-to-end validation are listed as planned next steps unless separately verified as complete.

7. EVALUATION PLAN

Evaluation should measure fraud-detection quality and operational behaviour together. Useful measures include precision, recall, F1-score, precision-recall area under the curve (PR-AUC), false-positive rate, calibration quality, and performance across relevant transaction segments. Operational tests should measure p50/p95/p99 latency, throughput, error rate, and feature freshness. Offline/online parity tests should detect differences in feature calculations. Graph and agent outputs should be reviewed for relevance, traceability, unsupported claims, and consistency with the evidence available to the investigator.

Results should be reported with the dataset description, data-split strategy, baseline comparisons, thresholds, and limitations. No numeric model-performance result is claimed because verified experimental metrics were not provided.

8. RESPONSIBLE USE, SECURITY, AND LIMITATIONS

- A fraud score is a risk signal, not proof of wrongdoing; decisions should involve authorised human review.
- Agent-generated summaries and adverse-media matches require source verification and should distinguish evidence from inference.
- Access control, data minimisation, secure handling of financial information, audit logs, and retention policies are important deployment requirements.
- Model performance can be affected by class imbalance, changing fraud patterns, incomplete labels, data leakage, and distribution shifts.
- The report describes the design and progress provided by the project team; production readiness and target latency require independent validation.

9. PLANNED NEXT STEPS

- Complete and validate the real-time scoring service and its API contract.
- Implement the AML multi-agent workflow and source-grounded evidence retrieval.
- Develop the React investigation console for risk scores, reason codes, linked entities, and case review.
- Add security controls, audit trails, and human-approval checkpoints.
- Run end-to-end tests and publish reproducible evaluation results with documented limitations.

10. CONCLUSION

AegisFlow proposes an integrated approach to payment fraud scoring and AML investigation by combining machine learning, online/offline feature processing, graph analytics, explainability, and agent-assisted case workflows. The implementation work reported so far establishes foundational feature, graph, modelling, and evaluation components. Further integration, security review, end-to-end testing, and evidence-based performance reporting are necessary before the platform can be described as production-ready.

ACKNOWLEDGEMENT

The authors may add acknowledgements for individuals or organisations that supported the work.

REFERENCES

1. Chen, T. and Guestrin, C. (2016). XGBoost: A Scalable Tree Boosting System. Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. https://doi.org/10.1145/2939672.2939785
2. Lundberg, S. M. and Lee, S.-I. (2017). A Unified Approach to Interpreting Model Predictions. Advances in Neural Information Processing Systems (NeurIPS). https://arxiv.org/abs/1705.07874
3. IEEE-CIS Fraud Detection dataset, Kaggle. https://www.kaggle.com/competitions/ieee-fraud-detection (dataset reference; use only if this dataset is actually used in the project).
4. LangGraph documentation. https://langchain-ai.github.io/langgraph/ (technical resource for the proposed agent workflow).
5. MLflow documentation. https://mlflow.org/docs/latest/index.html
6. FastAPI documentation. https://fastapi.tiangolo.com/

AUTHOR BIOGRAPHIES (OPTIONAL)

Add a 5–6 line biography for each author and an author photograph if required by the journal. Replace all author and affiliation placeholders before submission.