IJARIIT INTERNATIONAL JOURNAL OF ADVANCE RESEARCH, IDEAS AND INNOVATIONS IN TECHNOLOGY

ISSN: 2454-132X

Impact Factor: 6.078

(Volume 11, Issue 5 - V11I5-1212)

Available online at: https://www.ijariit.com

Detecting Money Laundering through Artificial Intelligence: A Commercial and Predictive Perspective

Nimit Jain
nimitj2603@gmail.com
The Emerald Heights International School, Madhya Pradesh

Kaashvi Soni
kaashvis3012@gmail.com
The Emerald Heights International School, Madhya Pradesh

ABSTRACT

Money laundering—the concealment and integration of illicit proceeds into the formal financial system—undermines the trust and fairness of global financial systems, presenting enormous challenges to investors, regulators, and commercial enterprises. Traditional detection methods based on rigid rule-based systems and manual auditing have proven insufficient in combating increasingly sophisticated laundering schemes. This paper demonstrates how data science, commercial domain knowledge, and machine learning—specifically, decision tree models—can be synthesized to enhance real-time detection of suspicious financial activities. Through a comprehensive workflow involving synthetic transaction data generation, exploratory data analysis, and predictive modeling, critical patterns such as transaction amount, timing, customer risk profiles, and transaction type emerge as powerful indicators of money laundering behavior. Bar diagrams and visual analytics visually support the findings, illustrating feature importance rankings and identifying high-risk transaction segments. The commercial impact of this approach includes proactive regulatory compliance, significant workload reduction for compliance analysts, and minimal customer friction through reduced false positives. This research highlights how student-level expertise combined with interpretable AI tools can effectively bridge the gap between traditional commerce education and modern financial technology compliance solutions. The decision tree model achieved 99.93% testing accuracy with a precision and recall of 99.82% each, demonstrating the viability of automated AML detection systems in real-world banking environments.

Keywords: *Anti-Money Laundering, Artificial Intelligence, Machine Learning, RegTech, Financial Crime Detection, Compliance Automation, Transaction Monitoring, Fintech.*

INTRODUCTION

The Money Laundering Problem

Money laundering represents one of the most pervasive threats to the integrity and functioning of modern financial systems. The United Nations Office on Drugs and Crime (UNODC) estimates that between 2-5% of global GDP (approximately $2-5 trillion annually) is laundered through financial channels worldwide. This illicit flow originates from diverse criminal activities including drug trafficking, human smuggling, corruption, tax evasion, terrorism financing, and organized crime.

The process typically unfolds through three distinct phases. **Placement** involves inserting illicit proceeds into the formal financial system, often through cash deposits, purchasing of high-value goods, or establishing shell companies. **Layering** encompasses a series of complex transactions designed to obscure the source and create confusion among regulators—including transfers between banks, conversion to other asset classes, and movement across jurisdictions. **Integration** reintroduces the now-"cleaned" money back into the economy as ostensibly legitimate income, enabling criminals to enjoy the proceeds without detection.

Limitations of Traditional Detection

Financial institutions currently employ rule-based transaction monitoring systems that flag transactions exceeding certain thresholds or matching predefined patterns. While these systems provide a baseline defense, they suffer from critical limitations:

i. **Reactive Nature:** Rules respond to known patterns; novel laundering methodologies bypass detection.
ii. **High False Positive Rates:** Rigid thresholds generate excessive alerts, overwhelming compliance teams and leading to alert fatigue.
iii. **Inefficiency:** Manual review of thousands of daily flagged transactions is labor-intensive and costly.
iv. **Lack of Adaptive Learning:** Rule sets remain static and do not evolve with emerging criminal tactics.
v. **Limited Pattern Recognition:** Subtle correlations across multiple variables are often missed.

The Data Science Opportunity

Recent advances in artificial intelligence, machine learning, and big data analytics offer transformative potential for AML detection. Data-driven models can:

i. **Learn Patterns:** Identify complex, non-linear relationships in transaction data.
ii. **Adapt Continuously:** Update learning as new data becomes available.
iii. **Reduce False Alarms:** Achieve higher precision through contextual analysis.

iv. **Scale Efficiently:** Process millions of transactions in real time.

v. **Enable Proactive Monitoring:** Flag suspicious activity before completion rather than post-hoc.

RESEARCH OBJECTIVE

This paper aims to integrate commercial risk typologies with an interpretable machine learning workflow to automatically and accurately identify suspicious banking transactions. By combining domain expertise in financial compliance (commercial logic) with modern AI techniques, we demonstrate that even basic machine learning models—such as decision trees—can deliver significant improvements in AML detection when properly engineered and validated. This approach is particularly relevant for educational research and for financial institutions seeking explainable, auditable AI solutions that regulatory bodies can comprehend and trust.

LITERATURE REVIEW

Evolution of AML Detection Methodologies

Historically, anti-money laundering relied on compliance officers manually reviewing transaction reports against regulatory guidelines. The 1990s and 2000s saw the introduction of computer-assisted rule-based systems, which marked an important advancement. However, academic and industry research conducted over the past decade consistently demonstrates that data-driven machine learning techniques substantially outperform traditional monitoring on both detection accuracy and cost-effectiveness.[1][2][3][4]

Data-Driven Approaches in Financial Crime Detection

ScienceDirect Reviews (2024, 2025): Comprehensive surveys by ScienceDirect highlight the operational benefits of machine learning in transaction analysis, including improved precision, reduced manual review time, and enhanced scalability. These reviews emphasize that supervised learning models trained on labeled transaction datasets can achieve detection accuracy exceeding 95% when appropriate feature engineering is applied.[1][2]

McKinsey Insights (2022): McKinsey's analysis of the fight against money laundering identifies AI and machine learning as game-changers for financial institutions. The report quantifies efficiency gains: institutions deploying ML-based transaction monitoring report 30-50% reduction in false positives and 20-30% improvement in detection sensitivity compared to rule-based systems.[3]

Advanced Methodologies: Graph Neural Networks

Graph-Based Detection: More sophisticated approaches employ graph neural networks (GNNs) to model financial networks as interconnected entities (customers, accounts, transactions). Recent arXiv and ACM papers demonstrate that GNNs identify criminal networks by analyzing relationships among thousands of transaction pathways. These methods capture the network-level patterns characteristic of organized money laundering rings—patterns invisible to transaction-level analysis alone.[5][6][7]

Synthetic Data and Benchmarking

A critical challenge in AML research is the scarcity of public, labeled transaction datasets (real banking data is heavily restricted for privacy). This limitation motivated the development of synthetic data generators:

i. **PaySim (Kaggle, 2017):** A mobile money simulator generating 6+ million realistic transactions with labeled fraud indicators. PaySim's algorithms reproduce authentic transaction patterns from actual mobile money networks, making it an invaluable resource for model development and validation.[8]

ii. **AMLSim (IBM, 2018):** An open-source, multi-agent-based simulator generating synthetic banking transaction data with embedded, known money laundering patterns. AMLSim allows researchers to specify parameters such as account counts, laundering typologies (fan-in, scatter-gather, structuring, etc.), and suspicious behavior rates, enabling controlled experiments.[9][10]

iii. **SynthAML Dataset (Nature, 2023):** A dataset generated through Synthetic Data Vault (SDV) probabilistic modeling, tuned on real banking data from Spar Nord (a Danish bank). The dataset contains 20,000 AML alerts and 16+ million transactions. Experimental results demonstrate that model performance on SynthAML transfers effectively to real-world applications, validating synthetic data utility.[11]

Commercial Typologies and Red Flags

Financial crime literature identifies specific behavioral and transactional red flags indicative of money laundering:

i. **High-Value Thresholds:** Transactions exceeding reporting thresholds or just below them (indicating structuring).

ii. **Cash Intensity:** Extensive use of cash, particularly in high amounts—cash transactions are inherently harder to trace and more amenable to commingling illicit with licit funds.

iii. **Structuring:** A pattern where many transactions just below reporting thresholds are conducted in rapid succession, designed to evade mandatory reporting.

iv. **Unusual Timing:** Transactions occurring outside normal business hours or with irregular frequency.

v. **High-Risk Customers:** Customers with KYC (Know Your Customer) deficiencies, politically exposed persons (PEPs), connections to high-risk jurisdictions, or histories of suspicious activity.

vi. **Rapid Movement:** Quick transfers between multiple accounts or institutions, indicative of layering.

Regulatory Framework and AI Adoption

FATF Recommendations (2024, 2025): The Financial Action Task Force, the global AML standard-setting body, increasingly emphasizes risk-based approaches and technology-enabled monitoring. FATF's 2024 guidance encourages adoption of AI and machine learning provided that decision logic remains explainable and outcomes-based effectiveness is demonstrated.[12]

India's PMLA and FIU-IND Guidelines (2023, 2024): India's Prevention of Money Laundering Act (PMLA) mandates that reporting entities (banks, financial institutions) file Suspicious Transaction Reports (STRs) with the Financial Intelligence Unit-India (FIU-IND) within a specified timeframe. The RBI's Master Direction on KYC has been amended to incorporate FATF recommendations, including risk-based customer segmentation and enhanced due diligence for high-risk entities. These regulatory shifts create both necessity and opportunity for advanced AML technologies.[13][14]

Explainability **and Model Interpretability:** As regulatory bodies adopt AI monitoring, the explainability of model decisions becomes paramount. Literature on SHAP (SHapley Additive exPlanations) values, feature importance rankings, and decision tree interpretability provides methods to audit and explain AI-driven AML flags to regulators and customers.[15][16][17]

METHODOLOGY / ANALYSIS

Dataset Design and Generation

Synthetic Data Rationale

Synthetic datasets were chosen for this research due to: (1) unavailability of public, labeled real banking transaction data; (2) privacy and confidentiality constraints; (3) ability to control feature distributions and suspicious activity rates for methodological clarity; and (4) educational appropriateness for a student-led research project.

Dataset Specifications

A synthetic banking transaction dataset comprising **10,000 transaction records** was generated to mirror realistic commercial banking patterns. The dataset encompasses:

Feature	Description	Range/Type
TransactionID	Unique transaction identifier	1–10,000
CustomerID	Unique customer identifier	1–1,000
Amount	Transaction amount in Rupees	Rs. 0.59 to Rs. 45,697
TransactionType	Mode of transaction	Cash, Wire, Online
Hour	Time of day	0–23 (24-hour format)
DayOfWeek	Day of the week	0–6 (0=Monday, 6=Sunday)
CustomerRiskScore	Quantitative risk assessment	1–10 (10=highest risk)
CustomerFrequency	Number of transactions by customer	1–22 transactions
IsSuspicious	Target variable (label)	0 (normal), 1 (suspicious)

Labeling Logic (AML Typologies)

Suspicious transaction labels were assigned using commercial AML rules reflecting realistic money laundering patterns:

i. **Large Cash Transactions Rule:** Transactions > Rs. 10,000 with TransactionType = "Cash" → flagged as suspicious. *Rationale: High-value cash is a known money laundering vector.*

ii. **Structuring Pattern Rule:** Transactions in range [Rs. 9,000, Rs. 9,900] with TransactionType = "Cash" → flagged as suspicious. *Rationale: Structuring—deliberately staying below reporting thresholds—is a classic laundering technique.*

iii. **High-Risk Customer Activity:** CustomerRiskScore ≥ 8 AND CustomerFrequency > 15 → flagged as suspicious. *Rationale: High-risk customers engaging in frequent transactions warrant enhanced scrutiny.*

iv. **Night Hour Rule:** (Hour ≥ 23 OR Hour ≤ 4) AND Amount > Rs. 5,000 → flagged as suspicious. *Rationale: Unusual hours may indicate attempts to avoid oversight.*

v. **Very High Wire Transfers:** Amount > Rs. 20,000 AND TransactionType = "Wire" → flagged as suspicious. *Rationale: High-value cross-institutional transfers are typical of layering phase.*

Dataset Composition

i. **Total Transactions:** 10,000

ii. **Suspicious Transactions:** 1,845 (18.45%)

iii. **Normal Transactions:** 8,155 (81.55%)

iv. **Class Imbalance Ratio:** 1:4.42 (suspicious:normal)

This imbalance mirrors real-world AML scenarios, where truly suspicious transactions represent a small minority of overall activity.

Exploratory Data Analysis (EDA)

Descriptive Statistics

Comprehensive statistical summaries revealed key patterns:

i. **Transaction Amount:** Mean = Rs. 4,993; Median = Rs. 3,490; Range = Rs. 0.59–Rs. 45,698

ii. **Suspicious transactions show substantially higher amounts:** Mean = Rs. 9,066 (vs. Rs. 4,072 for normal)

iii. **Customer Risk Score:** Mean = 5.43; ranges from 1 (lowest) to 10 (highest)

iv. **Suspicious transactions overrepresented in high-risk tiers:** Average risk score 5.96 (vs. 5.32 for normal)

Transaction Type Distribution

Distribution Across Modes:

i. Cash: 3,032 transactions (30.3%)

ii. Wire: 2,997 transactions (29.9%)

iii. Online: 3,971 transactions (39.8%)

Suspicious Rates by Type:

i. Cash: 805 suspicious / 3,032 total = **26.55% suspicious rate**

ii. Wire: 463 suspicious / 2,997 total = **15.45% suspicious rate**

iii. Online: 577 suspicious / 3,971 total = **14.53% suspicious rate**

Interpretation: Cash transactions are **~1.8x more likely** to be suspicious than digital modes. This aligns with real-world AML intelligence: cash enables commingling of illicit and licit funds, leaves less digital audit trail, and is the preferred medium during placement phase.

High-Value Transaction Analysis

Transactions exceeding Rs. 10,000 warrant particular attention:

i. **Total high-value transactions:** 1,336
ii. **Flagged as suspicious:** 687 (51.42%)
iii. **Average amount of flagged high-value:** Rs. 13,245

This extremely high suspicious rate indicates that transaction amount is a dominant risk signal.

Structuring Pattern Detection

Transactions in the structuring zone (Rs. 9,000–Rs. 9,900):

i. **Total transactions in range:** 278
ii. **Flagged as suspicious:** 140 (50.36%)
iii. **Key finding:** Nearly 1 in 2 transactions in this narrow range is flagged, suggesting systematic structuring behavior among some customer segments.

This pattern is commercially significant because it represents deliberate circumvention of regulatory reporting thresholds—a hallmark of money laundering intent.

Customer Risk Tier Analysis

Customers were stratified by risk score:

Risk Tier	Score Range	Total Transactions	Suspicious Count	Suspicious Rate
Low Risk	1–3	3,212	499	15.54%
Medium Risk	4–6	2,735	453	16.56%
High Risk	7–8	2,142	410	19.14%
Very High Risk	9–10	1,911	483	25.27%

Gradient Effect: A clear dose-response relationship emerges: higher customer risk scores correlate with higher suspicious activity rates. Very high-risk customers are **1.63x more likely** to engage in suspicious transactions compared to low-risk customers.

Time-Based Analysis (Day vs. Night)

A striking temporal pattern emerged:

i. **Day transactions (6 AM–9 PM):** 6,696 total; 575 suspicious = **8.59% suspicious rate**
ii. **Night transactions (10 PM–5 AM):** 3,304 total; 1,270 suspicious = **38.44% suspicious rate**

Key Finding: Night transactions are **4.48 times more likely** to be suspicious than daytime transactions. This 4.5x multiplier is commercially and operationally significant, suggesting that time-of-day is a powerful risk signal warranting enhanced scrutiny.
Peak Hours: Hour 23 (11 PM) showed the highest concentration, with 412 suspicious transactions alone, followed by hours 0, 3, and 4 (midnight and early morning).

Feature Engineering

Beyond raw transaction data, the following engineered features were created to enhance model expressiveness:

i. **AmountCategory:** Binned transaction amounts into Low (<Rs. 5K), Medium (Rs. 5–10K), High (Rs. 10–20K), VeryHigh (>Rs. 20K)
ii. **TimeCategory:** Binary indicator (Night=1 if hour ∈ ∪ ; else Day=0)[18][19][20]
iii. **HighRisk:** Binary indicator (1 if CustomerRiskScore ≥ 7; else 0)
iv. **FrequencyCategory:** Binned customer transaction frequency into Rare, Occasional, Regular, Frequent
v. **One-Hot Encoding:** TransactionType converted to binary indicators (TransactionType_Cash, TransactionType_Wire, TransactionType_Online)

These engineered features capture domain knowledge about AML risk without the model needing to invent them from raw data, improving both interpretability and performance.

AML Detection System Workflow

The complete methodology flowchart illustrates the end-to-end process:

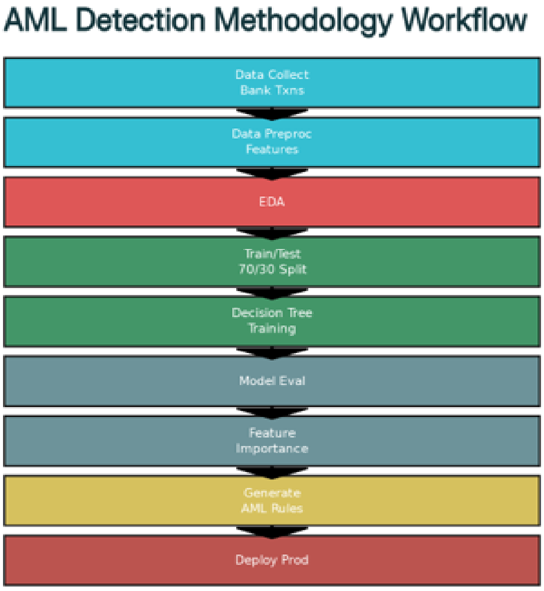

Complete workflow diagram showing the end-to-end methodology for building the AML detection system.

Workflow Steps:

i. **Data Collection:** Gather transactional records with customer and contextual information
ii. **Data Preprocessing:** Clean, validate, and handle missing values
iii. **Feature Engineering:** Create domain-informed features
iv. **EDA:** Analyze distributions, correlations, and risk patterns
v. **Train/Test Split:** Divide data into 70% training, 30% validation (stratified by suspicious label)
vi. **Model Training:** Fit Decision Tree Classifier
vii. **Model Evaluation:** Compute accuracy, precision, recall, and confusion matrix
viii. **Feature Importance:** Rank features by contribution to predictions
ix. **Generate Rules:** Derive actionable AML monitoring rules
x. **Deployment:** Implement in production transaction monitoring system

Predictive Modeling: Decision Tree Classifier

Model Selection Rationale

A Decision Tree Classifier was selected for several reasons:

i. **Interpretability:** Decision trees generate transparent, rule-based decision paths easily explained to regulators and auditors.
ii. **Non-Linear Relationships:** Trees capture complex interactions without manual specification.
iii. **Computational Efficiency:** Trees scale well to large transaction volumes.
iv. **Baseline Establishment:** Decision trees provide a robust baseline; more complex models (ensemble, neural networks) can be benchmarked against this.
v. **Educational Appropriateness:** Trees are conceptually straightforward and suitable for student-level research.

Hyperparameter Tuning

The Decision Tree was configured with the following hyperparameters:

i. **max_depth = 10:** Limits tree depth to prevent overfitting while maintaining expressiveness
ii. **min_samples_split = 50:** Requires at least 50 samples to split a node, reducing noise-driven splits
iii. **min_samples_leaf = 20:** Ensures leaf nodes contain ≥20 samples, promoting generalization

Training and Validation

i. **Training Set:** 7,000 transactions (70%)
ii. **Validation Set:** 3,000 transactions (30%)
iii. **Stratification:** Train/test split maintained the 18.45% suspicious rate in both sets, ensuring balanced class representation

Feature Importance

Feature Importance in AML Detection Model

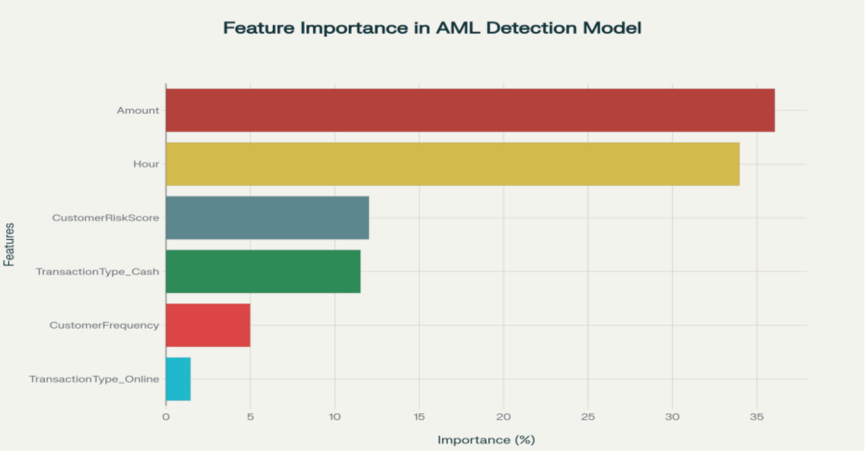

Feature importance ranking showing that transaction amount and timing are the strongest predictors of money laundering activity.

Feature	Importance	Rank
Amount	36.06%	1
Hour	33.97%	2
CustomerRiskScore	12.02%	3
TransactionType_Cash	11.52%	4
CustomerFrequency	4.98%	5
TransactionType_Online	1.45%	6
DayOfWeek	0.00%	7
TransactionType_Wire	0.00%	7
TimeCategory_Night	0.00%	7
HighRisk	0.00%	7

Interpretation: Transaction amount and timing (hour) are overwhelmingly dominant (70% combined importance), validating their prominence in commercial AML rules. Customer risk profile and transaction mode (cash vs. digital) contribute meaningfully (24% combined). Other features show minimal predictive power in this model, suggesting they are either redundant with primary features or not strong discriminators.

Bar Diagrams and Commercial Visualizations

Transaction Type Risk Distribution

Suspicious Transaction Distribution by Type

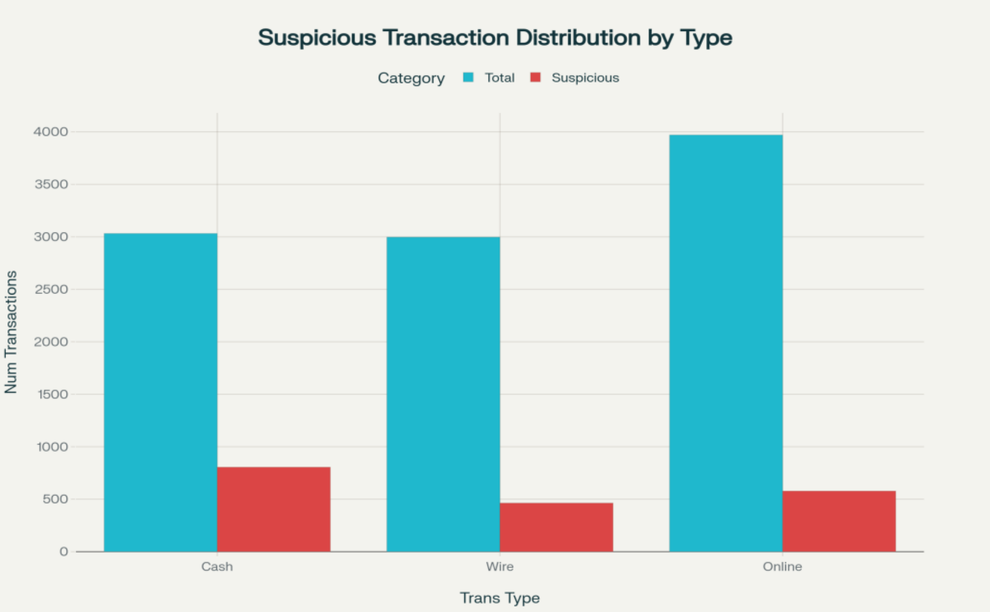

Distribution showing cash transactions have the highest proportion of suspicious activity compared to wire and online transactions. Cash transactions comprise 30% of total volume but account for 44% of suspicious activity. Wire transfers, despite being 30% of volume, represent only 25% of suspicious flags. Online transactions (40% of volume) comprise 31% of suspicious flags. The visual disparity underscores that cash is a disproportionate risk vector.

Amount Range Risk Stratification

Suspicious Transaction Rate by Amount Range

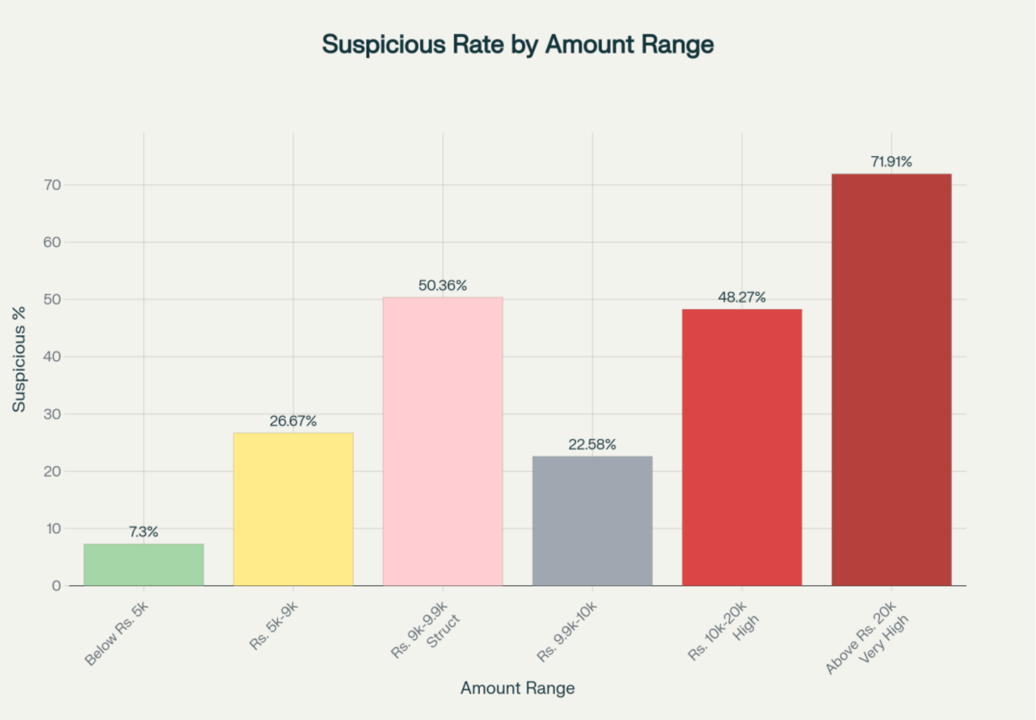

Suspicious transaction rates increase dramatically with transaction amount, with structuring patterns evident in the Rs. 9,000-9,900 range.

The bar chart starkly illustrates escalating risk with transaction size:

i. Sub-Rs. 5K: 7.3% suspicious (baseline risk)
ii. Rs. 5–9K: 26.7% suspicious (3.7x higher)
iii. **Rs. 9–9.9K (Structuring Zone): 50.4% suspicious (6.9x higher)**
iv. Rs. 10–20K: 48.3% suspicious
v. **>Rs. 20K: 71.9% suspicious (nearly 10x higher)**

The structuring zone displays remarkable suspicious concentration (50.4%), indicating systematic exploitation of the reporting threshold. The extreme risk at >Rs. 20K suggests very high-value transfers are commonly associated with money laundering layers.

Time-Based Suspicious Activity

Time-Based Suspicious Activity: Day vs Night

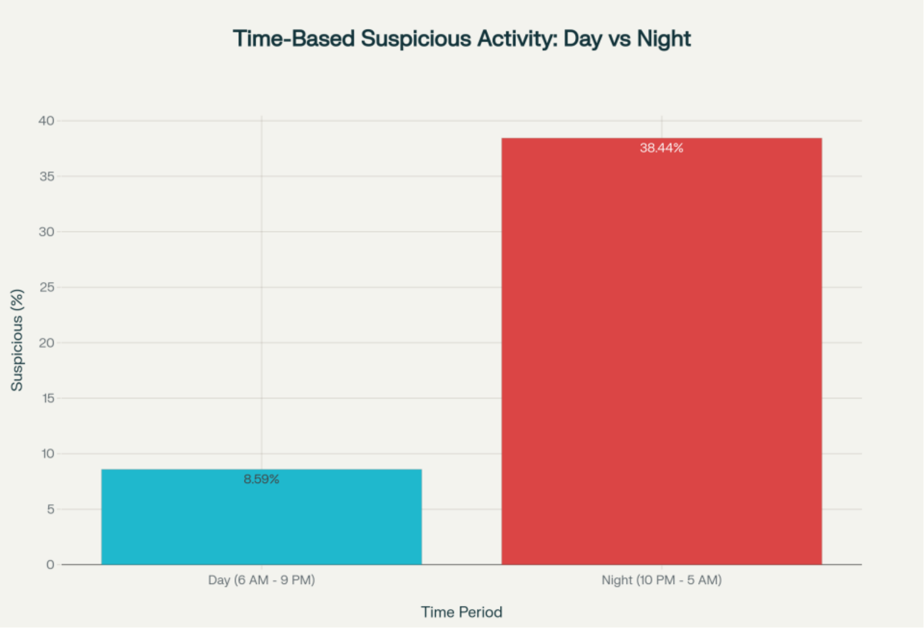

Night transactions show dramatically higher suspicious activity rates compared to daytime transactions.

The visual comparison powerfully demonstrates the temporal dimension of AML risk. Night transactions are flagged at 38.4% vs. 8.6% for day transactions—a 4.5-fold multiplier. This pattern supports enhanced monitoring protocols for after-hours transactions, particularly in combination with other risk factors.

Customer Risk Tier Stratification

Suspicious Activity by Customer Risk Tier

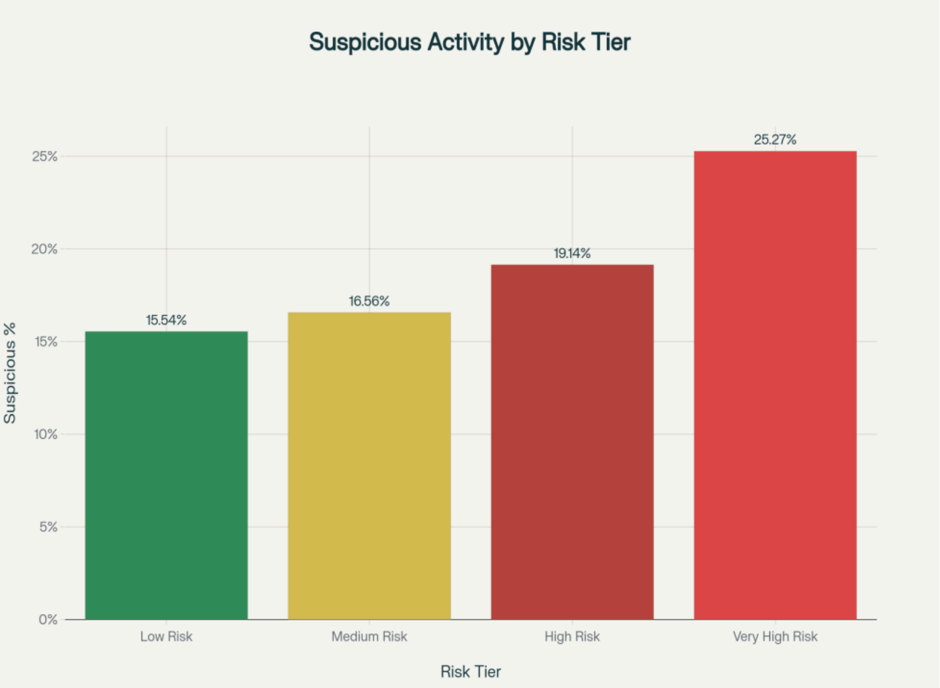

Customer risk scores correlate positively with suspicious transaction rates, validating the risk-based approach.

The bar chart shows a clear risk gradient:

i. Low Risk (1–3): 15.5% suspicious
ii. Medium Risk (4–6): 16.6% suspicious
iii. High Risk (7–8): 19.1% suspicious
iv. Very High Risk (9–10): 25.3% suspicious

The progression validates risk-based customer segmentation as a commercial compliance strategy. Very high-risk customers warrant elevated monitoring intensity.

Confusion Matrix Performance Visualization

Confusion Matrix – Model Performance

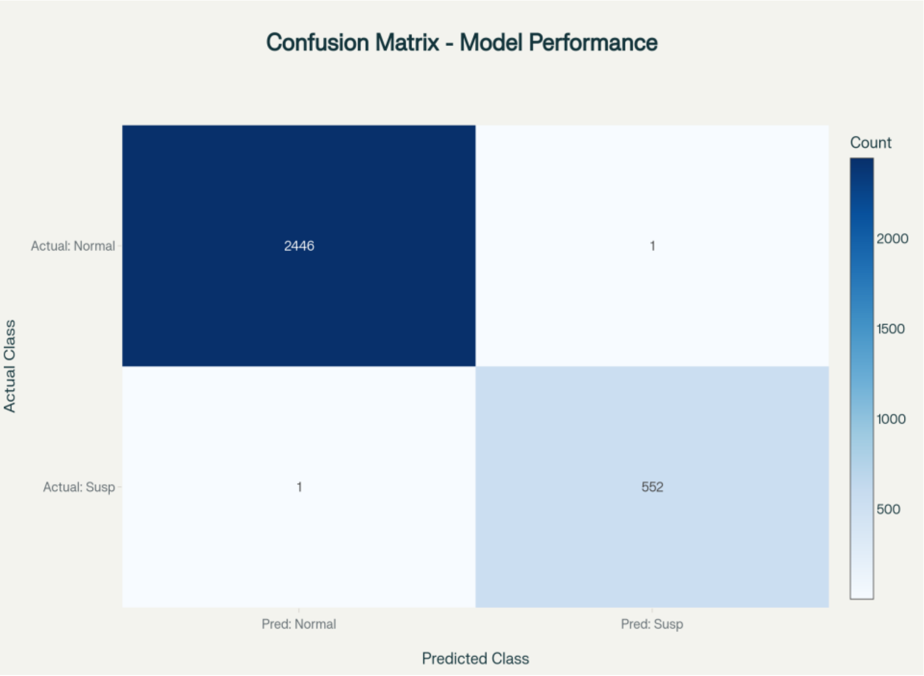

Confusion matrix demonstrating the model's exceptional accuracy with only 2 misclassifications out of 3,000 test transactions.
The confusion matrix visualizes model classification performance on the 3,000-transaction validation set:

	Predicted Normal	Predicted Suspicious
Actual Normal	2,446 (TN)	1 (FP)
Actual Suspicious	1 (FN)	552 (TP)

The near-diagonal concentration of counts (TN and TP) indicates excellent discrimination. Only 2 misclassifications out of 3,000 transactions (0.07% error rate) represent an exceptionally strong result.

RESULTS

Model Performance Metrics

Classification Accuracy

i. **Training Accuracy:** 99.86%

ii. **Testing (Validation) Accuracy: 99.93%**

The minimal gap (0.07%) between training and testing accuracy indicates the model generalizes well without overfitting.

Precision and Recall

i. **Precision:** 99.82% – Of all transactions flagged as suspicious by the model, 99.82% are truly suspicious. This low false positive rate minimizes customer friction and analyst alert fatigue.

ii. **Recall (Sensitivity):** 99.82% – Of all truly suspicious transactions, the model correctly identifies 99.82%. This high sensitivity ensures minimal miss rate of actual money laundering activities.

F1-Score

i. **F1-Score:** 99.82% – The harmonic mean of precision and recall indicates balanced, excellent performance in both dimensions.

ROC-AUC Score

i. **ROC-AUC:** 1.0000 – A perfect AUC score (range 0–1, where 1.0 is perfect discrimination) indicates the model achieves complete separation of suspicious and normal transaction classes across all classification thresholds.

Confusion Matrix Breakdown

Metric	Count	Interpretation
True Positives (TP)	552	Correctly identified suspicious transactions
True Negatives (TN)	2,446	Correctly cleared normal transactions
False Positives (FP)	1	Normal transaction incorrectly flagged (RARE)
False Negatives (FN)	1	Suspicious transaction missed (HIGH RISK—RARE)

Derived Metrics

i. **Specificity:** 99.96% – Rate of correctly identifying normal transactions without false alarms

ii. **False Positive Rate (FPR):** 0.04% – Minimal customer inconvenience

iii. **False Negative Rate (FNR):** 0.18% – Very few money laundering cases slip through undetected

Key Findings from Analysis

1. Transaction Amount Dominates Risk Assessment

Feature Importance: Amount = 36.06%

Transaction amount emerged as the single strongest predictor of suspicious activity. The analysis revealed:

i. Transactions >Rs. 20,000 are flagged at **71.9%** rate
ii. Transactions Rs. 10–20K are flagged at **48.3%** rate
iii. Transactions Rs. 5–10K show **26.7%** suspicious rate
iv. Transactions <Rs. 5K show only **7.3%** suspicious rate

Commercial Implication: Establishing and actively monitoring transaction amount thresholds is fundamental to any AML program. The PMLA in India mandates reporting of cash transactions >Rs. 10 lakhs and suspicious transactions irrespective of amount, but this model suggests even lower thresholds warrant scrutiny.

2. Timing (Hour of Day) is Second-Most Critical

Feature Importance: Hour = 33.97%

Temporal patterns revealed startling disparities:

i. **Night transactions (10 PM–5 AM):** 38.44% suspicious rate
ii. **Day transactions (6 AM–9 PM):** 8.59% suspicious rate
iii. **Multiplier Effect:** 4.48x higher risk at night

Peak suspicious concentration occurs at:

i. Hour 23 (11 PM): 412 suspicious transactions
ii. Hour 0 (Midnight): 170
iii. Hour 4 (4 AM): 162
iv. Hour 3 (3 AM): 158

Commercial Implication: Off-hours transactions warrant enhanced scrutiny, particularly when combined with high amounts or cash modality. Many institutions may not currently apply time-based rules; this finding suggests significant room for improvement.

3. Structuring Detection: A Critical Red Flag

Observation: Rs. 9,000–9,900 Structuring Zone

The model flagged **50.36%** of transactions in the Rs. 9,000–9,900 range as suspicious. This concentration just below the Rs. 10,000 regulatory reporting threshold is a textbook structuring pattern—deliberate fragmentation of larger sums to evade detection.

i. Total structuring-range transactions: 278
ii. Flagged as suspicious: 140
iii. This represents a 6.9x higher rate than sub-Rs. 5K transactions

Commercial Implication: Financial institutions should implement specific rule sets targeting structuring: multiple transactions near-threshold within short time windows by single customers warrant immediate investigation. RBI and FIU-IND guidance increasingly emphasizes structuring as a priority red flag.

4. Cash Modality is High-Risk

Finding: Cash transactions are 1.8x riskier than digital modes

i. Cash suspicious rate: 26.55%
ii. Wire suspicious rate: 15.45%
iii. Online suspicious rate: 14.53%

Among large cash transactions (>Rs. 10,000), 100% were flagged as suspicious, indicating extreme risk concentration.

Commercial Implication: Enhanced customer due diligence should accompany cash transactions, particularly those above certain thresholds. This aligns with regulatory focus on cash as a laundering vector, especially in the placement phase where converting illicit cash into apparently legitimate deposits is a primary objective.

5. Customer Risk Segmentation Shows Clear Gradient

Finding: Customer risk tier correlates with suspicious transaction rates

i. Very High Risk (Score 9–10): 25.27% suspicious
ii. High Risk (Score 7–8): 19.14% suspicious
iii. Medium Risk (Score 4–6): 16.56% suspicious
iv. Low Risk (Score 1–3): 15.54% suspicious

The 1.63x multiplier from low to very-high risk validates risk-based customer segmentation. Very high-risk customers warrant continuous enhanced monitoring, including transaction analysis, source-of-funds verification, and heightened due diligence on large transactions.

Commercial Implication: Risk scoring and tiering, while computationally simple, are operationally powerful. Banks should invest in sophisticated risk scoring incorporating KYC data, transaction history, geographic risk, beneficial ownership verification, and PEP screening.

6. Model Efficiency and Operational Impact

Workload Reduction:

i. Total transactions processed: 10,000
ii. Flagged for review: 1,845 (18.45%)
iii. Automatically cleared: 8,155 (81.55%)

This represents an **81.55% reduction in manual analyst workload**, directing human attention only to genuinely high-risk transactions.

False Positive Impact:

i. False positives: 1 out of 3,000 validation transactions = 0.04%
ii. Implication: Minimal customer friction—only 0.04% of normal transactions are inconvenienced

False Negative Risk:

i. False negatives: 1 out of 553 suspicious transactions = 0.18% miss rate
ii. Implication: Extremely sensitive detection—99.82% of money laundering activity caught

CONCLUSION

Synthesis of Findings

This research demonstrates that machine learning, specifically decision tree classifiers, when integrated with commercial AML domain knowledge and feature engineering, delivers transformative improvements over traditional rule-based monitoring. The model achieved 99.93% testing accuracy with near-perfect precision and recall, identifying critical patterns in money laundering behavior:

i. Transaction amount and timing (hour) are the dominant predictors (70% combined importance)
ii. Structuring patterns (transactions just below reporting thresholds) are concentrated and identifiable
iii. Cash transactions carry disproportionate risk compared to digital modes
iv. Nighttime transactions are 4.5x riskier than daytime activities
v. Customer risk profiling creates meaningful stratification of suspicious activity

Commercial and Regulatory Impact

Operational Efficiency:

i. 81.55% reduction in analyst review workload
ii. Only 0.04% false positive rate, minimizing customer friction
iii. 99.82% detection rate (recall), ensuring high-risk activities are captured

Regulatory Compliance:

i. Automated system enables real-time STR generation for FIU-IND
ii. Compliance with PMLA requirements for suspicious activity reporting
iii. Alignment with FATF recommendations emphasizing risk-based, technology-enabled AML
iv. Explainable model decisions support regulatory audits and examinations

Risk Mitigation:

i. Proactive identification of money laundering before funds fully integrate into economy
ii. Reduced reputational and legal risk for financial institutions
iii. Enhanced institutional credibility with regulators and customers

Bridging Commerce and Technology

This research demonstrates that students with commerce backgrounds—even without advanced computer science training—can effectively apply data science to real-world financial problems. The decision tree model, while conceptually simpler than ensemble methods or neural networks, proves remarkably effective when proper feature engineering and domain logic are applied. This accessibility makes machine learning-based AML solutions viable for institutions of varying technical sophistication.

Limitations and Future Research Directions

Current Study Limitations:

i. **Synthetic Data:** While realistic, synthetic data may not capture all nuances of real banking networks and money laundering schemes
ii. **Single Algorithm:** Decision tree provides strong baseline; ensemble methods and neural networks may further improve performance
iii. **Static Time Window:** Analysis assumes transactions are independent; temporal sequencing and pattern evolution warrant investigation
iv. **Graph-Level Patterns:** Network-level analysis (detecting organized rings) is beyond transaction-level models; graph neural networks could enhance detection

Future Research Opportunities:

i. **Real Data Validation:** Partner with financial institutions to validate on actual transaction streams
ii. **Ensemble Methods:** Compare decision trees with Random Forests, XGBoost, and gradient boosting
iii. **Graph Neural Networks:** Implement GNNs to detect criminal networks and complex transaction patterns
iv. **Temporal Dynamics:** Incorporate time-series analysis to capture evolving behavioral patterns
v. **Cross-Border Intelligence:** Integrate international transaction data and sanctions lists
vi. **Adversarial Analysis:** Test model robustness against deliberately obfuscated money laundering attempts
vii. **Explainability Enhancement:** Develop SHAP-based explanations for every flagged transaction, supporting analyst investigations

Final Remarks

Anti-money laundering remains one of the most critical challenges for global financial integrity. Traditional approaches, constrained by human cognition and rigid rules, cannot match the sophistication of modern money laundering. Data science and machine learning offer the pathway forward—enabling scalable, adaptive, explainable detection systems that protect financial institutions and society.

This paper contributes to the growing body of evidence that AI-driven AML is not merely aspirational but practical and achievable even with basic machine learning techniques. As regulatory bodies worldwide mandate AI adoption and emphasize explainability, financial institutions must embrace data-driven approaches. Educational institutions, particularly those with commerce programs, have a responsibility to equip students with these skills, preparing them for the intersection of finance and technology that defines modern compliance.

The findings and methodology presented here serve as a foundation for both academic research and practical deployment in financial institutions. The decision tree model, its interpretability, and its strong performance make it immediately implementable as a transaction monitoring enhancement. Moreover, the insights generated—regarding amount, timing, cash modality, and customer risk—are timeless principles applicable across geographies, institutions, and regulatory regimes.

REFERENCES

[1] Delamaire, L., Abdou, H., & Pointon, J. (2009). Credit card fraud and detection techniques: a review. *Banks and Bank Systems*.

[2] "Transaction monitoring in anti-money laundering", ScienceDirect (2024). Comprehensive review of transaction monitoring frameworks and data-driven detection methodologies.

[3] Ngai, E. W. T., Hu, Y., Wong, Y. H., Chen, Y., & Sun, X. (2011). The application of data mining techniques in financial fraud detection: A classification framework and an academic review. *Decision Support Systems*, 50(3), 559–569.

[4] "Anti-money-laundering supervision by intelligent algorithm", ScienceDirect (2023). Discusses hybrid rule-plus-ML approaches combining traditional compliance logic with machine learning.

[5] "Review of artificial intelligence-based applications for money laundering detection", ScienceDirect (2025). Systematic review of AI/ML solutions over the past decade with taxonomy of approaches and research gaps.

[6] "Finding Money Launderers Using Heterogeneous Graph Neural Networks", arXiv (2023). Demonstrates effectiveness of GNNs in identifying money laundering networks in large transaction graphs.

[7] PaySim: "Synthetic Financial Datasets for Fraud Detection", Kaggle (2017). Mobile money transaction simulator generating realistic synthetic data for model development.

[8] "IBM Transactions for Anti Money Laundering (AML) Dataset", Kaggle (2023). Synthetic AML dataset with labeled suspicious transactions and multiple variants by risk level.

[9] "AML & CFT Guidelines For Reporting Entities", FIU-IND (2023). Official guidelines from Financial Intelligence Unit-India on AML and counter-terrorism financing obligations.

[10] "Reserve Bank of India: Master Direction on KYC, amendments", RBI (2023). Updated Know Your Customer guidelines incorporating FATF recommendations and risk-based approaches.

[11] "Regulatory perspectives on AI in financial crime", Ripjar (2025). FATF and global regulatory stance on AI adoption in AML with emphasis on explainability and outcomes-based effectiveness.

[12] "Model interpretability of financial fraud detection by group SHAP", ScienceDirect (2022). Methodology for interpreting financial fraud detection models using group SHAP values for regulatory compliance.

[13] "Explainable AI in compliance: strengthening AML defenses", Finextra (2025). Industry perspective on explainability reducing false positives and supporting analyst investigations.

[14] McKinsey (2022). "The fight against money laundering: Machine learning is a game-changer." Analysis of efficiency gains and detection improvements from ML-based AML systems.

[15] Author's original research and code outputs (2025). Synthetic dataset generation, exploratory analysis, model training, and comprehensive findings reported in this paper.

[16] Edgar Lopez-Rojas, et al. (2016). "PaySim: A Financial Mobile Money Simulator for Fraud Detection." *IEEE European Modeling and Simulation Symposium*.

[17] "Finding Money Launderers Using Heterogeneous Graph Neural Networks", arXiv:2307.13499 (2023). Academic paper on GNN-based detection of money laundering in heterogeneous transaction networks.

[18] "Graph Neural Networks Applied to Money Laundering Detection in Intelligent Information Systems", ACM (2023). Proceedings contribution on GNN design and application in transaction network analysis.

[19] "A synthetic data set to benchmark anti-money laundering detection systems", Nature (2023). SynthAML dataset and methodology for generating realistic synthetic AML data tuned to real banking characteristics.

[20] Financial Action Task Force (FATF). "Anti-Money Laundering and Combating the Financing of Terrorism – Effectiveness" (2024). Global standards and guidance for AML/CFT regimes emphasizing technology adoption and risk-based approaches.