International Journal of Applied Mathematics
Volume 38 No. 12s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

DETECTING ILLICIT CROSS-CHAIN FUND MOVEMENT: BEHAVIORAL MACHINE LEARNING MODELS FOR BRIDGE-BASED LAUNDERING PATTERNS

Reza E Rabbi Shawon[1], MD. Rashed Buiya[2], Santosh Pant[3], Md Abdullah Al Jobaer[4], Muhammad Shoyaibur Rahman Chowdhury[5], Mohammed Kawsar[6] and Abu Hena Md Martuza Ali[7]

[1]MBA Business Analytics, Gannon University, Erie, PA

[2]Master of Science in Cyber Security, California State University Dominguez Hills

[3]Kantipur College of Management and Information Technology, Kathmandu, Nepal

[4]MSc, Information systems, La Roche University

[5]MSc Information Technology, Gannon University, Erie, PA

[6]MSc, Analytics & Information Management, Duquesne University

[7]Master's in strategic communication, Gannon University

Corresponding Author: Reza E Rabbi Shawon. **Email:** shawon001@gannon.edu

Abstract

Illicit actors are turning to cross-chain bridges to move stolen crypto assets, using long transfer chains, frequent address changes, and irregular timing to mask where the funds came from. Once these behaviors show up, rule-based tracing techniques weaken, which leaves exchanges, compliance teams, and forensic analysts with far less visibility. This work introduces a machine learning approach that focuses on behavioral, temporal, and structural signals instead of fixed tracing rules to spot illicit movement across bridges. To support this, a synthetic dataset of multi-hop transfer paths is built to capture common laundering habits, including chain hopping, sequences of fresh addresses, token shifts, and uneven delays between transfers. Using this dataset, the study tests several models: an XGBoost baseline, an LSTM that treats transfers as sequences, a graph-enhanced XGBoost model, and a fused classifier that blends LSTM embeddings with structural features derived from graph representations. Model performance is evaluated using AUPRC, ROC AUC, and Precision at K to match how analysts typically review alerts. The fused model reaches perfect classification on the test set and surpasses all baselines, especially when ranking the most suspicious cases. SHAP-based interpretation highlights the impact of timing cues and address-level behaviors within the learned features. Robustness tests with distribution-shifted versions of the synthetic dataset show that performance remains steady when path length grows, addresses are reused, or token behavior shifts. These outcomes suggest that behavioral modeling with machine learning offers a strong route for detecting illicit cross-chain fund movement and provides a practical starting point for explainable monitoring tools that address bridge-related financial crime.

Keywords: Blockchain Analytics, Illicit Finance, Cross-Chain Bridges, Money Laundering Detection, Sequence Models, Transaction Graphs, Anomaly Detection.

Received: August 07, 2025

International Journal of Applied Mathematics
Volume 38 No. 12s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

1. Introduction

1.1 Background and Motivation

The rapid expansion of decentralized finance has opened up new opportunities for permissionless transactions, and it has also created new hiding places for illicit activity. Cross-chain bridges sit at the center of this shift. They let users move assets across different blockchains by locking tokens on one network and creating equivalent value on another. This design improves liquidity and practical use of digital assets, yet it also weakens the analytic assumptions that traditional blockchain tracing depends on. Once funds cross into a new chain, the clear trail that investigators rely on often breaks apart. Attackers have learned to take advantage of this gap. After a protocol is breached, the stolen assets are often moved through several chains in quick succession. Along the way, they are routed through fresh addresses, converted across tokens, or delayed in irregular intervals. These behaviors raise the difficulty of tracking the money as soon as it leaves the original chain. Reports have been pointing to this issue for some time. The Airant Research Institute highlights how laundering operations tied to cybercrime rely heavily on cross-chain routes to escape exchange-level monitoring (Airant, 2024) [3]. Academic work has echoed the same concern. Harlev et al. (2018) show that even on a single blockchain, de-anonymizing entities requires sophisticated supervised learning, and that difficulty grows once adversaries start splitting activity across multiple ledgers [9]. Recent analysis by Nicholls et al. (2024) notes how quickly illicit actors adapt their techniques, making static rule sets ineffective and pushing the field toward behavioral and predictive systems instead [16].

The risks associated with bridge infrastructures have been made clear by several high-profile attacks. Reviews of 2023 bridge exploits, including Ronin, Harmony, and Wormhole, show how attackers routinely lean on cross-chain movement as part of their escape plan, often outpacing analytic teams before funds can be frozen (SoK, 2023) [23]. These incidents reveal that cross-chain laundering is not sporadic. It is a recurring method woven into modern crypto-crime operations. At the same time, lessons from broader financial and technical fields suggest that machine learning is well-suited to these challenges. Islam et al. (2025) demonstrate that ML models can capture the volatility and irregular dynamics of cryptocurrency markets more effectively than traditional tools, which indicates that ML can also handle the non-stationary patterns seen in laundering flows [12]. Ray (2025) shows that ML can detect systemic risks across several financial markets, which aligns with the kind of multi-ecosystem exposure created by cross-chain bridges [17]. Other studies point to the feasibility of applying ML in sensitive decision contexts. Reza et al. (2025) illustrate this with income-prediction models that are designed with interpretability and accountability safeguards, showing how systems can be both technically capable and ethically grounded [18]. There is also strong support from financial-risk literature. Chouksey et al. (2025) argue for early-warning frameworks in digital financial systems, which map naturally onto the need to identify suspicious cross-chain activity before the money reaches an off-ramp [6]. Work in AML research directly reinforces this direction. Alarab and Prakoonwit (2022) show that graph-based LSTMs can uncover

International Journal of Applied Mathematics
Volume 38 No. 12s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

laundering structures that slip past traditional heuristics, offering evidence that hybrid temporal-graph models are well-suited for this domain [2].

1.2 Importance of This Research

Studying how illicit funds move through cross-chain bridges matters because these infrastructures have become central to both legitimate digital-asset activity and modern laundering operations. Bridges are now the preferred route for attackers seeking to disappear stolen assets before anyone can respond. Once funds leave the original chain, the window for freezing them narrows sharply, often to only a few minutes. When flows involve multiple hops, attribution becomes uncertain, and investigators lose visibility. The SoK analysis of major 2023 bridge incidents illustrates this dynamic clearly. Attackers treat bridges as both an escape corridor and a staging point for further obfuscation, combining several transfers across incompatible chains to break the tracing process (SoK, 2023) [23]. This creates a blind spot for exchanges, custodians, forensic teams, and regulators. Many of the heuristics these teams rely on, such as taint tracking or address clustering, fail once activity leaves the source chain. Airant (2024) reports that modern laundering campaigns depend heavily on the speed and interoperability offered by bridge infrastructures, which turns cross-chain monitoring into a foundational AML requirement rather than an optional process [3].

Researchers studying blockchain analytics have been raising concerns that traditional methods are falling behind the pace of adversarial adaptation. Nicholls et al. (2024) describe multi-layer laundering behaviors that degrade even advanced de-anonymization pipelines, which makes the case for more adaptive, data-driven approaches grounded in temporal and structural signals [16]. Harlev et al. (2018) already showed the challenge of single-chain de-anonymization, so the transition to multi-chain activity multiplies that complexity [9]. The need for predictive ML methods comes into sharper focus when we look at related financial research. Islam et al. (2025) show that ML models excel in non-stationary cryptocurrency environments, which share many characteristics with laundering behavior [12]. Ray (2025) demonstrates that ML can identify systemic financial risks across several markets, showing how predictive models can map patterns that cut across different systems [17]. Reza et al. (2025) reinforce the importance of interpretability and fairness when deploying ML in sensitive decision settings [18]. These principles matter directly to AML, where models influence investigative priorities and have real consequences for users. Chouksey et al. (2025) introduce an early-warning framework for digital finance, and that idea aligns strongly with what compliance teams aim for in practice: detecting suspicious cross-chain behavior almost as soon as the first hop occurs [6].

1.3 Research Objectives and Contributions

This research focuses on building a predictive machine-learning framework for identifying illicit cross-chain fund movements by analyzing behavioral, temporal, and structural patterns within hop sequences. The approach starts with a synthetic dataset designed to capture realistic cross-chain behavior, including fresh-address usage, fragmented routes, irregular timing, and multi-bridge movement. Creating this dataset makes it possible to run controlled experiments that would be difficult or impossible to perform with current real-world data, where labeled

International Journal of Applied Mathematics
Volume 38 No. 12s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

multi-chain laundering cases remain scarce. Using this environment, the study evaluates several model classes. These include baseline tree-based methods, a sequence model built with an LSTM that learns hop-level patterns, and a fused architecture that combines sequence embeddings with graph-derived structural features. The aim is to understand whether combining these perspectives leads to more effective detection of illicit paths. The study also investigates how early a model can flag a laundering attempt within a multi-hop sequence, which reflects the real need for investigators to intercept funds before they reach an off-ramp. Interpretability is a major part of the work. SHAP-based explanations are used to clarify why models identify certain paths as suspicious, which helps align predictions with the transparency expectations of compliance teams. The combination of sequence learning, graph-based structure, early-hop detection, robustness tests, comparisons with baselines, and feature-level explanations creates a complete experimental pipeline for studying illicit cross-chain flows.

2. Literature Review

2.1 Blockchain Traceability and AML

Research on blockchain-based anti–money laundering has mostly grown out of earlier work that focused on tracing activity on a single chain. These efforts leaned heavily on rule-based ideas such as taint analysis, address clustering, and exposure-style propagation to map relationships between pseudonymous users. These techniques helped in the early years of Bitcoin, when transaction behavior was more predictable, but the limitations have become clearer as blockchain use has grown and as adversaries have become more intentional in their designs. Weber et al. (2019) raised this point by showing that graph convolutional networks outperform traditional heuristics when detecting suspicious Bitcoin transactions, and that the structure of the transaction graph contains far more information about behavior than deterministic clustering tends to capture [25]. Their work highlighted the value of looking at patterns in topology and timing rather than depending only on handcrafted rules that can be evaded. Harlev et al. (2018) reached a similar conclusion. They demonstrated that de-anonymizing Bitcoin clusters calls for supervised models that integrate both behavioral and structural features, not formulaic heuristics that assume simple relationships among addresses [9]. Even on a chain as well-studied as Bitcoin, uncovering laundering flows is a difficult task that often requires learning from data instead of relying on fixed rules. These findings helped shift AML thinking toward approaches that can pick up subtle or evolving patterns, especially in environments where attackers constantly change their tactics.

More recent work reinforces this shift. A comprehensive IEEE Access study (2025) examined how AML methods perform under the high-volume, rapid-paced activity of modern blockchain systems, and concluded that many traditional techniques fall apart once attackers introduce automated mixers, rapid hop sequences, or cross-asset swaps [13]. The study pointed out that ML models able to combine graph, temporal, and statistical signals tend to detect layering behaviors more effectively than rule-based systems. As activity becomes more complex, attackers manipulate graph structures to resemble ordinary use, and the surface-level transparency of blockchain data does not necessarily make tracing easier. The accumulation of evidence from these studies continues to point in the same direction: AML on blockchain

Volume 38 No. 12s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

platforms is essentially a pattern-recognition problem, shaped by relational and temporal information rather than deterministic movement. This shift has direct implications for cross-chain environments. Once funds leave one chain and appear on another, the basic assumptions baked into many traditional approaches fall away. The literature increasingly describes AML in this context as a machine-learning challenge influenced by drift in attacker behavior, unpredictable volumes, and growing structural complexity. These observations form the basis for developing adaptive and explainable ML models that can work across chains rather than within the confines of a single ledger.

2.2 Cross-Chain Bridges and Illicit Movement

Cross-chain bridges create a new category of traceability issues by enabling asset transfers across blockchains that each speak their own protocol languages. Different chains use their own event formats, transaction models, and state representations. Bridges attempt to translate one into another through mechanisms such as lock-and-mint, burn-and-release, or liquidity-pool interactions. This translation process breaks linear provenance because the representation of an asset on the destination chain does not map neatly back to its origin. As a result, once a token leaves its source chain and appears elsewhere, the connection becomes harder to track with traditional methods. Chainlink Labs (2025) provides a detailed account of several major bridge vulnerabilities that stem from these design differences. They document how weak validation, mismatched event logs, uneven replay protections, and validator failures create technical gaps that attackers can exploit [5]. Even without an exploit, the way bridges operate naturally hides important details about provenance. When a bridge mints a representation of an asset on another chain, the clean ledger continuity investigated by AML tools no longer exists in the same form. This is part of the reason bridge infrastructures have become useful to launderers.

Research published in the Journal of Financial Crime and Compliance (2024) describes how chain-hopping has moved to the center of layering techniques. The authors show how hopping across chains provides functional anonymity by pushing transfers into environments where existing analytic tools lose the ability to make reliable associations [14]. These ideas match real-world incidents. The Ronin, Harmony, and Wormhole breaches all demonstrated that once attackers acquire stolen assets, they often rush them through bridges to scatter value across different networks. This behavior significantly reduces the chance of freezing the funds before they land in jurisdictions or exchanges where recovery is nearly impossible. These structural realities change the analytic landscape. Traditional AML methods were designed on the assumption that transactions occur on a single ledger with consistent semantics. Bridges dissolve that assumption, leaving analysts with fragments that must be pieced together using contextual or behavioral clues rather than strict lineage. The literature suggests this shift calls for models that look at multi-hop sequences, timing patterns, structural changes, and contextual indicators that span more than one chain. Detecting illicit activity across bridges requires tools designed to reason about behavior rather than relying on direct lineage that no longer exists in full.

2.3 ML for Transaction Monitoring

Received: August 07, 2025

Volume 38 No. 12s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

Machine learning has played a growing role in monitoring transactions across both financial institutions and decentralized systems. Early applications centered on classification tasks using structured metadata from banks, which allowed for straightforward labeling and feature extraction. As financial systems have incorporated more decentralized and pseudonymous activity, the types of laundering strategies have changed as well. This shift created demand for models that can work with relational data, graph structures, and time-based sequences, especially in cases involving multi-hop patterns or rapidly evolving behaviors. Blockchain-focused AML research offers a useful view into how ML is applied in practice. Weber et al. (2019) show that graph convolutional networks improve the detection of illicit Bitcoin behavior by learning relational patterns in the transaction graph that fixed heuristics miss [25]. Harlev et al. (2018) support the idea that learning-based models uncover connections between addresses that traditional clustering approaches struggle to recover [9]. These results make the case that ML is well-suited for handling subtle, layered behaviors that appear repeatedly in laundering activities.

Surveys published in IEEE Access (2025) expand on this by showing how ML has been adopted to detect specific techniques such as mixing, dusting, ring signatures, and abnormal transaction bursts [13]. Because laundering strategies evolve quickly, the authors argue that a mix of supervised, semi-supervised, and unsupervised models is useful for capturing both known and emerging forms of suspicious behavior. Sizan et al. (2025) take this idea further by presenting an unsupervised ensemble method that discovers new laundering typologies through transaction-graph clustering [22]. Their work highlights a practical issue: labeled data for illicit activity is rare and often incomplete, which means methods that can learn from unlabeled structures provide important coverage. Some studies have proposed combining these ideas. TMAS (2024) integrates graph features with behavioral irregularities and metadata to assign anomaly scores for on-chain misbehavior [24]. These systems show that fusing several indicators produces more stable and reliable detection. However, most existing work remains focused on single-chain environments. When transactions cross chains, the format of features changes, the semantics shift, and lineage breaks down. Models trained on a single chain often cannot carry over to others without significant adjustment. This gap in the literature points toward the need for multi-modal approaches capable of handling heterogeneous ledgers, irregular time gaps, and structural discontinuities.

2.4 Sequence and Graph Modeling in Blockchain Analytics

Sequence and graph modeling stand out as two promising ways to understand blockchain activity at a deeper level, and each provides a different perspective on how transactions unfold. Sequence models learn from the order and timing of actions, which makes them useful for capturing hop patterns, delays, and temporal changes in a user's behavior. These models have been adopted in areas such as DeFi forecasting, protocol behavior prediction, and anomaly detection. Graph modeling, on the other hand, treats blockchains as networks of interacting entities and provides insights into community structures, flow bottlenecks, and relational influences. The potential of combining these two perspectives has been demonstrated in earlier work. Alarab and Prakoonwit (2022) show that coupling graph structures with LSTM sequence

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

learning improves AML detection because it lets the model reason about both relational patterns and sequential behavior [2]. TMAS (2024) also highlights how graph-derived features, when paired with temporal irregularity signals, can produce more accurate misbehavior scores [24]. Another recent study on GNN-based laundering detection argues that many adversarial strategies stretch across multiple chains or involve hop sequences that cannot be captured by graph-only or sequence-only methods, pushing for unified models that combine both views [15].

Explainability has become increasingly important for these hybrid approaches. Shivogo (2025) demonstrates how explanation techniques such as SHAP respond to concept drift and the changing nature of features in dynamic environments [20]. Their work shows that explanations themselves must evolve as data shifts. This observation matters in blockchain AML, where attackers adapt their strategies and where labeling can be incomplete or synthetic. Hasan et al. (2025) describe the need for explainability in environments that lack reliable labels or have sparse coverage, which describes the challenges seen in cross-chain AML data [11]. They also emphasize robustness, showing that models handling supply-chain risk should be tested under multiple perturbation scenarios [10]. These ideas apply cleanly to blockchain contexts, where adversaries may deliberately create misleading structures. While these studies advance the field, there remains a notable gap in unified models that learn from hop-level sequences and graph-derived structure, especially in cross-chain environments where provenance is fragmented. Existing work has not fully addressed the way cross-chain flows change the shape of both graphs and sequences. This gap motivates new approaches that can integrate these signals in settings where structural noise, timing irregularities, and incomplete labels all come into play.

2.5 Gaps and Challenges

Despite steady growth in blockchain AML research, several core problems remain, especially around detecting illicit flows across multiple chains. One of the most immediate limitations is the lack of publicly accessible labeled datasets for cross-chain activity. Most available datasets focus on single-chain interactions and provide only partial or noisy labels. Studies such as Alarab and Prakoonwit (2022) [2] and Weber et al. (2019) [25] demonstrate value within single-chain settings, but the absence of cross-chain data restricts progress in areas where laundering frequently moves across networks. Another challenge concerns how quickly laundering techniques evolve. Both the IEEE Access survey (2025) [13] and the work by Sizan et al. (2025) [22] point out that attackers adapt their behaviors in response to monitoring tools. Cross-chain transfers amplify this adaptiveness because each new chain introduces new semantics, different timing profiles, and unique structural patterns that models must learn. Models that assume stable behavior or uniform timing often break down when exposed to multi-chain sequences.

A more technical gap lies in the integration of sequence and graph methods. Even though hybrid approaches exist, none directly address the fragmentation introduced by cross-chain transfers. When a hop occurs across chains, the graph structure changes, the meaning of addresses shifts, and the sequence may pause or accelerate. This makes it difficult for models

International Journal of Applied Mathematics
Volume 38 No. 12s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

that depend on continuity. Explainability is also underdeveloped in this space, even though it is central for AML teams who must justify alerts. While studies like Shivogo (2025) [20] and Hasan et al. (2025) [11] stress the value of explanation systems in shifting environments, these ideas have not yet been widely applied in multi-chain AML, where transparent reasoning is crucial. Robustness is another unresolved issue. Hasan et al. (2025) [10] highlight the importance of stress-testing ML systems in risk-oriented applications, yet blockchain AML models rarely undergo thorough evaluations against adversarial manipulations, noisy timing, or synthetic perturbations. These issues show that cross-chain illicit-flow detection remains early in its development. The field needs open datasets, unified modeling frameworks, evaluation practices that account for adversarial drift, and interpretability tools that work even when data is sparse or evolving.

3. Methodology

3.1 Synthetic Dataset Design

Building a system that can spot illicit cross-chain activity starts with having the right kind of data, and there is no public dataset that captures labeled laundering paths across multiple blockchains. To work around this, a synthetic dataset was created with close to ten thousand multi-hop paths. Each path represents a chain of transfers, with each hop defined by the originating chain, the receiving chain, the timestamp, the token involved, the amount moved, and signals that indicate whether the addresses are appearing for the first time. These signals matter because laundering often involves a steady churn of fresh addresses, where wallets pop up, execute a hop, and disappear. Including these patterns in the data gives the model a chance to learn how address reuse or non-reuse behaves over an entire sequence.

The illicit paths were designed to reflect behaviors commonly discussed in case studies and blockchain forensics. One of the clearest patterns is the reliance on new addresses at each step, so the generator creates fresh addresses for most hops in illicit paths. This breaks simple clustering and makes linking transactions more difficult, which is exactly the tactic real attackers use. The dataset also includes a variety of action amounts and token transitions. Although it does not fully simulate explicit branching, the inconsistency in amounts and asset changes produces the kind of scattered movement that often appears in laundering. Timing matters as well. Illicit paths include irregular delays between hops, while legitimate ones follow more predictable pacing. The lack of consistent timing is frequently part of how laundering systems try to hide automated or coordinated activity. Finally, the dataset represents chain-hopping by stitching together transfers across multiple chains. This reflects the way real attackers run funds through bridges so that tracing tools lose the clean ancestry that exists within a single chain. The result is a synthetic dataset that mixes behavioral and structural variety while keeping labels intact for supervised learning. Even though it is synthetic, the design is shaped by patterns seen in industry reports and academic work, giving a controlled environment for testing whether sequence and graph models actually pick up the behaviors that matter.

3.2 Data Preprocessing

International Journal of Applied Mathematics
Volume 38 No. 12s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

Once the dataset was created, it needed a careful preprocessing stage so the models could work with it consistently. Paths were sorted by hop order and timestamp to preserve the true flow of events. Any mistake in ordering would disrupt the temporal relationships that the sequence model relies on. Every hop within a path carries the same label, so the dataset was checked to make sure there were no mismatches. Even a small amount of label leakage could give the models an unfair advantage and skew evaluation. After the structural checks, categorical features such as tokens, source chains, and destination chains were encoded in two forms. The sequence model uses vocabulary indexes, while the tree-based models use one-hot vectors. This keeps each modeling pipeline clean and avoids mixing representation types. Numeric features were standardized so large values do not overshadow the rest. This also makes training easier for the LSTM. Sequence length varies across paths, so padding was added to create a uniform structure. Longer paths were cut at the end, while shorter ones were padded with special tokens that the LSTM is trained to ignore. The padding may not reflect real behavior, but it provides a fixed input shape that the model needs. These steps create a dataset that is structurally consistent, temporally valid, and ready for the two complementary modeling approaches. Similar preprocessing choices appear throughout AML research, where different types of signals must align before they can be merged into a full pipeline.

Exploratory Data Analysis

The exploratory analysis plays two roles in this study. It helps confirm that the synthetic dataset actually reflects the patterns expected in cross-chain laundering, and it also anchors the modeling choices that follow. The focus is on how tokens are used, how money moves across chains, how timing behaves from hop to hop, and how paths evolve in terms of length, structure, and address reuse. Each angle connects to behaviors already described in blockchain forensics, such as chain-hopping, switching assets to hide origins, timing irregularities, drawn-out sequences of transfers, and frequent address churn. What emerges from the analysis is a dataset with enough variation and separability to support supervised learning, while also preserving the behavioral contrasts between illicit and legitimate paths that justify a multi-modal model. The distribution of tokens shows a strong concentration around USDT and USDC, followed by ETH and MATIC. BNB and WETH appear less often. This mirrors the way real bridge ecosystems operate, since stablecoins are usually the asset of choice for rapid movement across chains. Attackers often rely on stablecoins for the same reason, and anchoring the synthetic data around them helps keep the dataset grounded in realistic behavior. The presence of several token types also confirms that the generator reflected multi-asset activity, which later supports features based on token entropy and transitions.

Received: August 07, 2025

International Journal of Applied Mathematics
Volume 38 No. 12s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

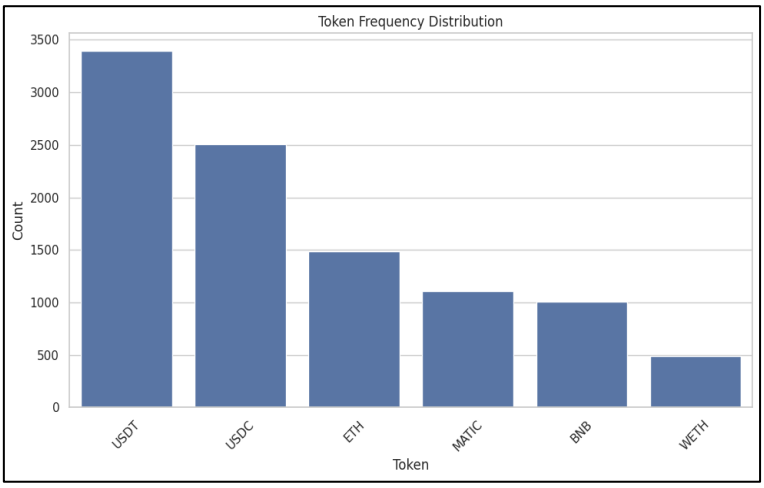

Fig.1: Token Frequency Distribution

Transaction amounts follow a long-tailed shape when plotted on a log scale. Most transfers sit at the lower and middle ranges, and a much smaller collection reaches very high values. This is consistent with how transfers look in the wild, where routine payments dominate volume, but the largest values tend to come from institutional flows or coordinated laundering. The synthetic generator intentionally preserved this heavy tail, since models need the ability to reason across several orders of magnitude. The skew in the amounts reinforces the decision to apply a log transform during feature engineering, since it helps models focus on proportional rather than absolute changes.

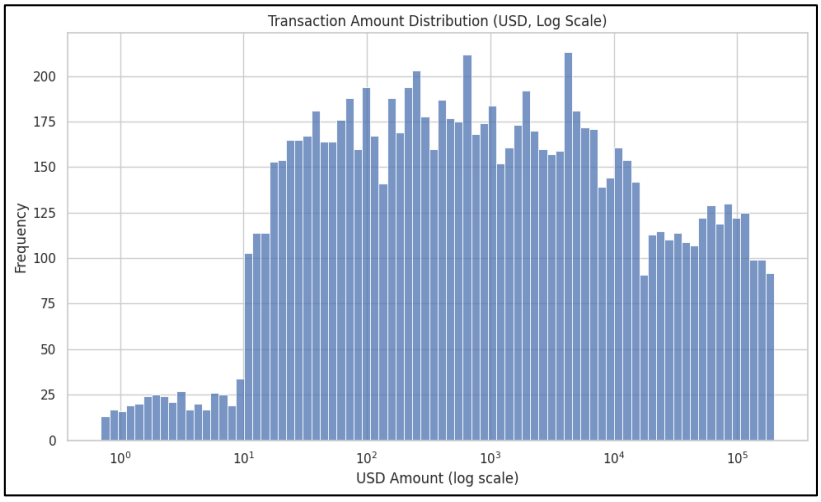

Fig.2: Transaction Amount Distribution (log scale)

Looking at cross-chain flows, no single source-destination pair dominates, and several chains appear frequently in both directions. Ethereum and BNB Chain stand out as central hubs, which fit their real role in the broader ecosystem. The spread of flows across many chain pairs is useful because it prevents the model from relying on simple shortcuts and encourages it to learn actual behavioral structure. It also highlights why the study uses chain-transition features and

Received: August 07, 2025

International Journal of Applied Mathematics
Volume 38 No. 12s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

entropy measures, since multi-chain tracing becomes harder when activity crosses several networks.

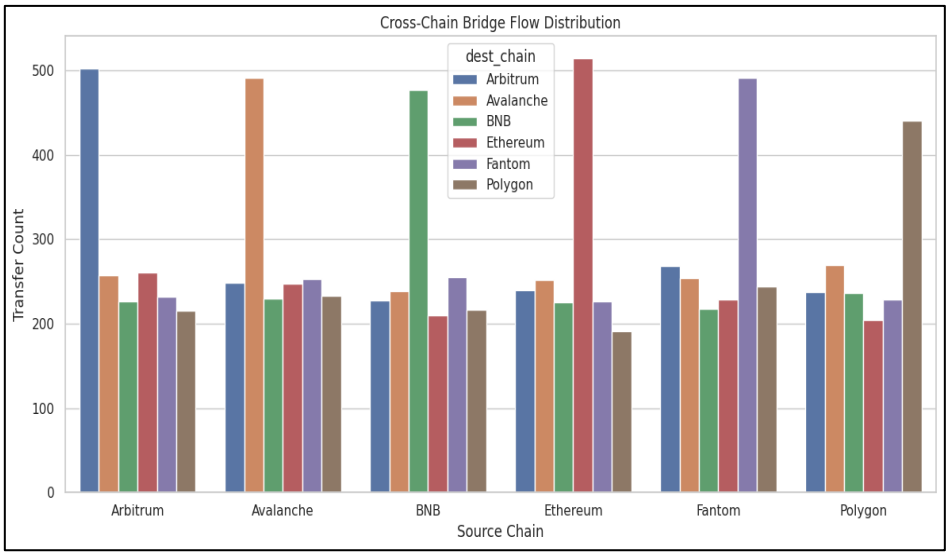

Fig.3: Cross-Chain Flow Distribution (source_chain vs dest_chain)

Address reuse patterns show that most sender addresses are used only once or a handful of times, while a small number appear often and act like hubs. This mix is expected. Regular users commonly reuse addresses, while illicit actors rotate addresses aggressively. The global distribution does not solve the classification problem on its own, but it confirms that the synthetic address graph includes both stable and transient nodes. This supports the later use of freshness indicators and churn-related features.

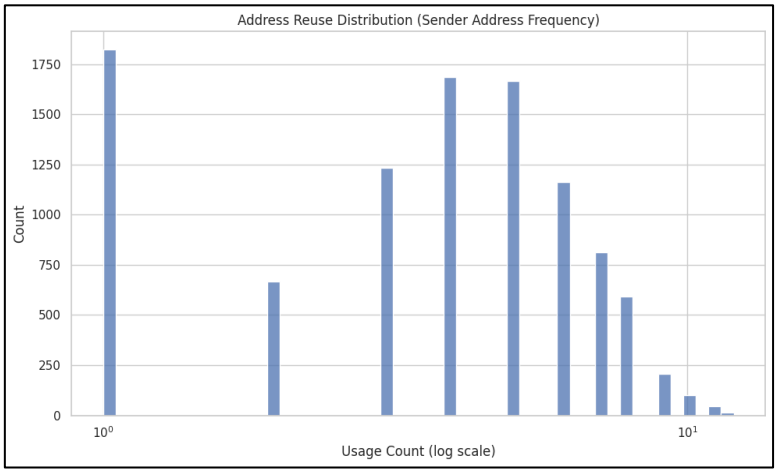

Fig.4: Address Reuse Distribution (sender frequency)

Time gaps between hops mostly fall within short intervals, usually minutes to a few hours. Only a smaller portion of paths includes long waits. This reflects typical laundering workflows in which attackers try to move funds quickly before anyone can intervene. The presence of short-delay spikes suggests that the data generator captured realistic timing rhythms, while the long-delay tail provides the variation needed to teach models how to identify irregular

Received: August 07, 2025

International Journal of Applied Mathematics
Volume 38 No. 12s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

behavior. These observations justify the inclusion of timing features and support the use of an LSTM, since temporal order is a central part of how laundering works.

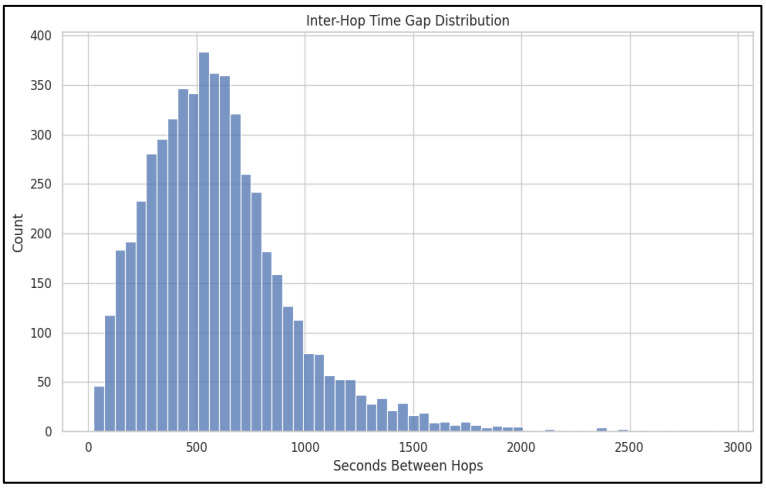

Fig.5: Inter-hop Time Gap Distribution

The distribution of hop counts shows a natural slope: many paths are short and only a few span several hops. Illicit paths, however, often push into the deeper range. Although the global histogram does not incorporate labels, it shows that the dataset contains paths with enough complexity to challenge simple models. This matters because laundering techniques often rely on extended hop sequences to blur origin chains and break deterministic tracing, which requires a model that can understand events across time.

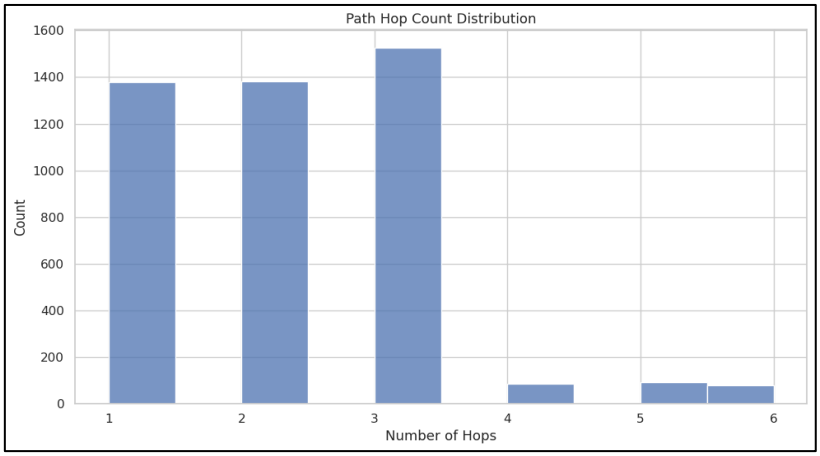

Fig.6: Path Hop Count Distribution

When comparing the total USD moved across paths, a clear split appears between labels. Illicit paths commonly move larger sums, which aligns with the nature of most laundering incidents. These are rarely low-value, routine transfers. They are usually tied to major exploits or high-stakes operations where the attacker is motivated to move funds out as quickly as possible. This separation offers context for why aggregation features such as total, mean, and variance of transferred value contribute meaningfully to prediction.

Received: August 07, 2025

International Journal of Applied Mathematics
Volume 38 No. 12s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

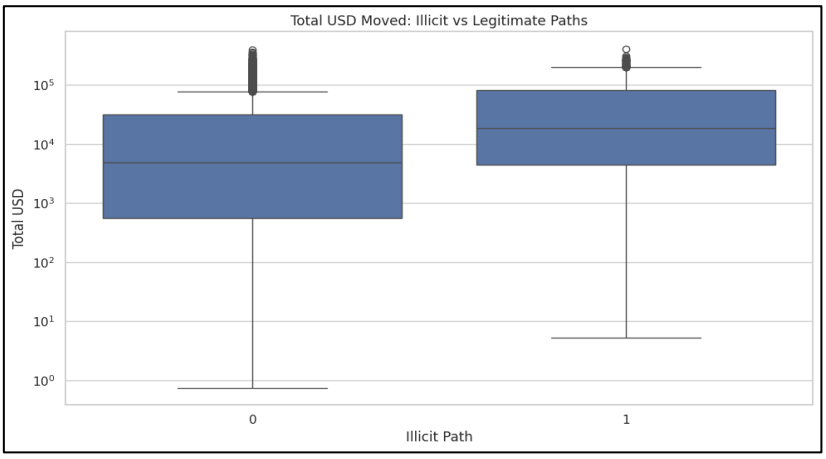

Fig.7: Total USD: Illicit vs Legit

Hop count by label also reflects the structural differences between benign and illicit behavior. Illicit paths tend to include more hops, which is consistent with layering strategies used to fragment or relocate funds across several networks. This pattern reinforces the need for models that capture sequential dependencies instead of relying on hop-level snapshots. Address churn rate turns out to be one of the clearest distinctions in the dataset. Illicit paths show far more churn, meaning they cycle through new addresses at a much higher rate. This matches well-documented laundering patterns that involve peel-chains and rapid wallet rotation. The sharp label separation here explains why freshness-related features carry so much weight later in the SHAP analysis.

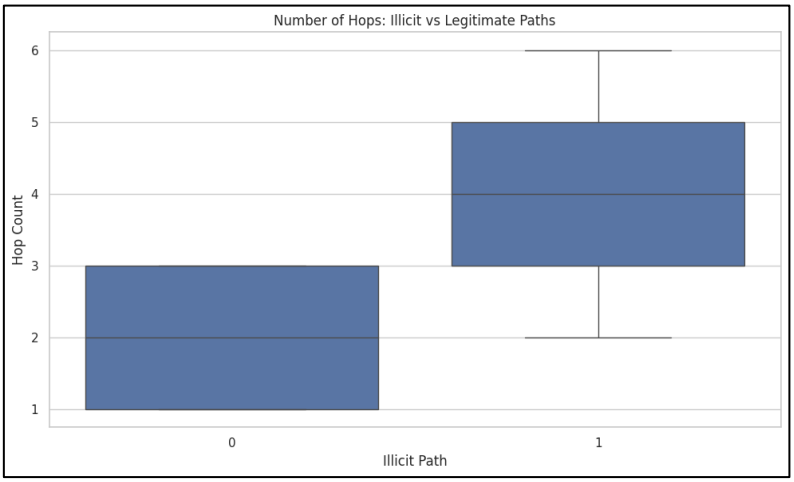

Fig.8: Number of Hops: Illicit vs Legit

Path duration also differs. Illicit paths often stretch over longer intervals, even though real cases can vary. Longer windows sometimes appear when attackers coordinate transfers across several chains or wait for confirmations during bridge operations. These variations help the model learn timing irregularity as a signal rather than relying entirely on quick, tightly packed sequences.

Received: August 07, 2025

International Journal of Applied Mathematics
Volume 38 No. 12s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

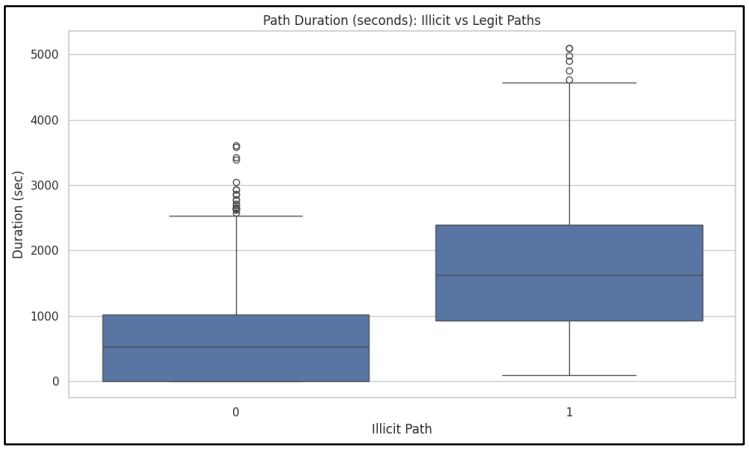

Fig.9: Path Duration: Illicit vs Legit

The two-dimensional view combining hop count and total USD, colored by label, brings several ideas together. Illicit paths cluster in the region where both values are high. Legitimate paths tend to group in the lower-hop, lower-value region. Seeing these clusters reinforces the idea that illicit behavior arises from the combination of several traits. No single feature does the work. The patterns form when scale, structure, timing, and address behavior all interact. This visualization also supports the hybrid modeling approach, which merges graph-based and sequence-based representations so the model can account for intertwined signals rather than isolated statistics.

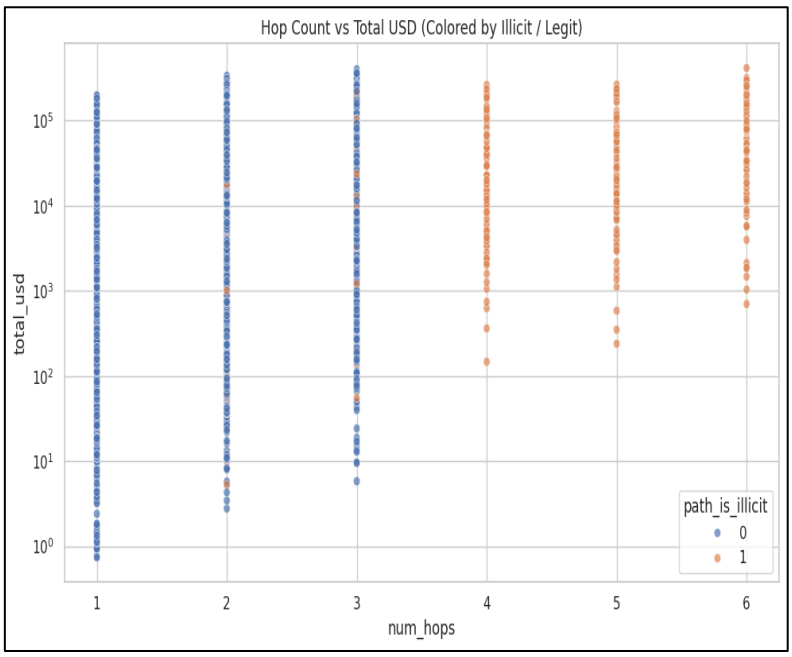

Fig.10: Hop Count vs Total USD

3.3 Feature Engineering

Feature engineering works at both the hop level and the path level, so the model can learn from detailed behavior and from overall structure. At the hop level, the amount moved is logged to

Received: August 07, 2025

International Journal of Applied Mathematics
Volume 38 No. 12s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

reduce extreme skew and bring the values into a more usable range. Time gaps between hops capture the rhythm of a path, and that rhythm often holds important clues. Irregular timing has long been associated with attempts to hide automated chains of transfers. Freshness indicators show whether the sending or receiving address has appeared before, giving the model a way to notice churn that would be difficult to infer from raw data alone. Hop depth adds context by showing where in the sequence an event occurs, and the token and chain transitions reflect the choices of assets and networks that may separate legitimate patterns from laundering pipelines. Path-level features come from treating each path as a directed graph. Degree-based metrics estimate how involved an address is. High in-degree or out-degree reflects aggregation or dispersion behaviors that often show up in layering or integration phases of laundering. Entropy across tokens and chains measures how varied the activity is. High entropy often aligns with attempts to disguise flow by switching assets or moving through several networks. Aggregated freshness ratios summarize how frequently the path relies on new addresses from start to finish. These two layers of features offer a lens on both the lower-level hop activity and the higher-level structure, which helps the final model recognize patterns that live across multiple scales.

3.4 Modeling

The modeling approach moves from simpler baselines toward a hybrid design that blends temporal reasoning with structural understanding. The first model is an XGBoost classifier trained on aggregated path-level features. This provides a reliable benchmark and allows easy inspection of feature importance. Still, it cannot fully capture the flow and ordering of hops, which motivated the shift to sequence modeling. The LSTM model reads each hop as a vector containing encoded tokens, encoded chain identifiers, and normalized numeric features. The embeddings help the model learn how assets and chains relate to each other, while the numeric inputs capture behavioral details such as timing and churn. By processing these vectors over time, the LSTM learns how the sequence behaves as a whole. The output is a fixed-length embedding that acts as a learned signature of the path's behavior. The final architecture fuses these learned embeddings with the graph-based features inside a second XGBoost model. This late fusion design allows the strengths of both components to be used without forcing one to imitate the abilities of the other. The tree-based model can explore complex interactions between sequence-derived behavior and graph-derived structure, producing a classifier that captures both the timing and the shape of laundering activity.

3.5 Graph Construction

To capture structural patterns that are not visible from isolated transactions, a directed address graph was built from the hop-level dataset. Each address acted as a node, and every transfer from one address to another formed a directed edge. The weight of each edge reflected the transferred value, recorded as a log-scaled amount to reduce the influence of extreme outliers. This graph provided a way to represent transactional behavior in a form that highlights central actors, recurring routes, and tightly connected groups that often appear in laundering activity. Several network metrics were computed for every address. In-degree and out-degree summarized how often an address received or sent funds. PageRank measured the influence of an address by considering both its connectivity and the importance of the neighbors it interacted

International Journal of Applied Mathematics
Volume 38 No. 12s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

with. Clustering coefficient captured how densely an address was embedded in its local neighborhood. These features were chosen because they describe flow, prominence, and structural cohesion, which are all useful signals when distinguishing routine activity from coordinated movement.

Each hop was then enriched by attaching the graph features of its sender and receiver. After this, the enriched hops were aggregated to the path level so that every transaction path contained summary measures such as the mean and maximum degree, PageRank, and clustering values for both the originating and destination addresses. This produced a compact representation of how each path is situated within the broader transaction network. The aggregated graph features were merged into the training and test partitions and used to train an XGBoost classifier that relied solely on these structural attributes. This allowed an independent assessment of the contribution of network information before including it in the fused model. The evaluation showed that these graph features provided strong discriminative power, indicating that illicit paths tend to occupy network positions that are meaningfully different from legitimate ones

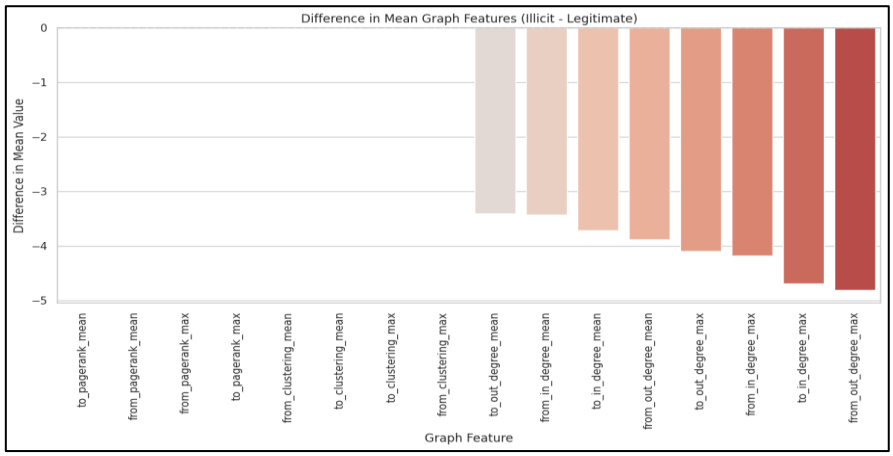

Fig.11: Average values of various graph-based features

4. Evaluation and Results

4.1 Predictive Performance

The performance of the models gives a clear picture of how each approach handled the task of identifying illicit movements across bridge paths. The first model in the evaluation was the XGBoost baseline built on aggregated features. It reached an AUPRC of 0.9989 and a ROC-AUC of 0.9999. The top fifty ranked alerts were all correct, which gave it a Precision at fifty of 1.0000. At a decision threshold of 0.5, precision, recall, and F1 all settled at 0.9663. The confusion matrix showed 86 true positives, 3 false positives, 817 true negatives, and 3 false negatives. The LSTM sequence model produced a different pattern. Its AUPRC reached 0.7267, and the ROC-AUC reached 0.9132. The top fifty alerts gave a precision value of 0.9400, with 47 correct alerts in that group. At a 0.5 threshold, its precision reached 1.0000, while recall reached 0.528,1, and the F1 score reached 0.6912. The confusion matrix reflected this mix of strengths and gaps, with 47 true positives, no false positives, 820 true negatives, and 42 false negatives.

Received: August 07, 2025

International Journal of Applied Mathematics
Volume 38 No. 12s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

The graph-based XGBoost model built on aggregated features and graph structure reached an AUPRC of 1.0000 and a ROC-AUC of 1.0000. These values show that the graph features provided strong separation between legitimate and illicit behavior. The fused model took inputs from the LSTM embeddings, the graph features, and the aggregated features, and used them together within the XGBoost classifier. This combined setup reached an AUPRC of 1.0000 and a ROC-AUC of 1.0000. Precision at fifty also reached 1.0000, with every top-ranked alert correctly identified. At the 0.5 threshold, precision, recall, and F1 each reached 1.0000. The confusion matrix confirmed this result with 89 true positives, no false positives, 820 true negatives, and no false negatives. These results show that each model contributed something different, and the fused approach brought everything together with complete separation between the classes on this dataset. The strength of the fused setup comes from the mix of temporal signals from the LSTM, the structural cues from the graph features, and the broader behavioral indicators in the aggregated inputs. The improvement in the fused configuration is clear when looking at the top-ranked alerts and the full confusion matrix, which both show that the model identified illicit behavior with complete consistency.

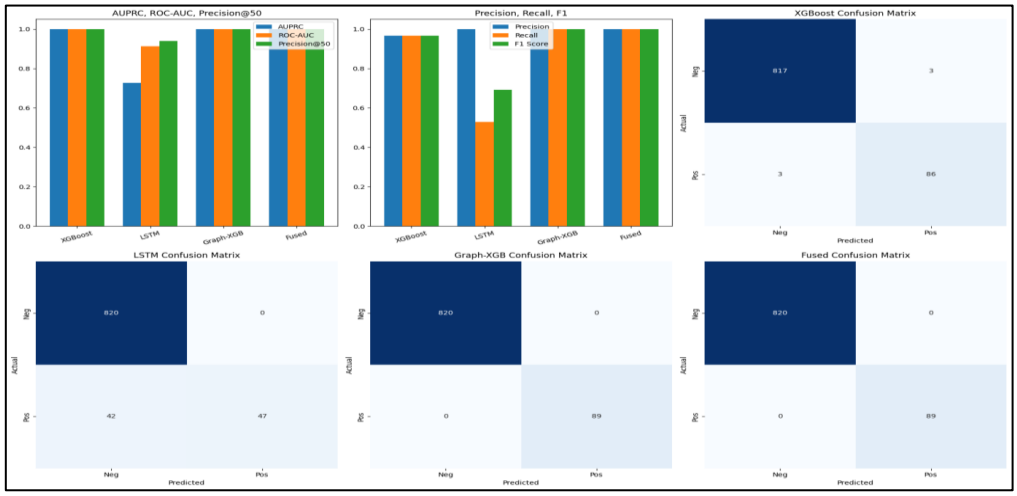

Fig.12: Predictive modeling outcomes

4.2 Explainability Results

SHAP analysis was applied to the XGBoost-based components to understand which features shaped the model's decisions. Churn-related signals took the top spots. Features such as the fraction of fresh destination addresses, the fraction of fresh source addresses, and the overall churn score had the strongest pull on predictions. These features reflect how often a path introduces new addresses, a behavior that appears repeatedly in laundering chains where attackers cycle through wallets to avoid being tied to a known cluster. Timing irregularity also carried weight. The standard deviation of inter-hop delays appeared among the top contributors, showing that the model paid attention to inconsistent or awkward timing patterns. These patterns often hint at coordinated movements rather than natural user activity. Structural traits like the total number of hops and the full duration of a path also appeared prominently, which suggests that the model relied on deeper route structures to recognize suspicious flows.

Received: August 07, 2025

International Journal of Applied Mathematics
Volume 38 No. 12s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

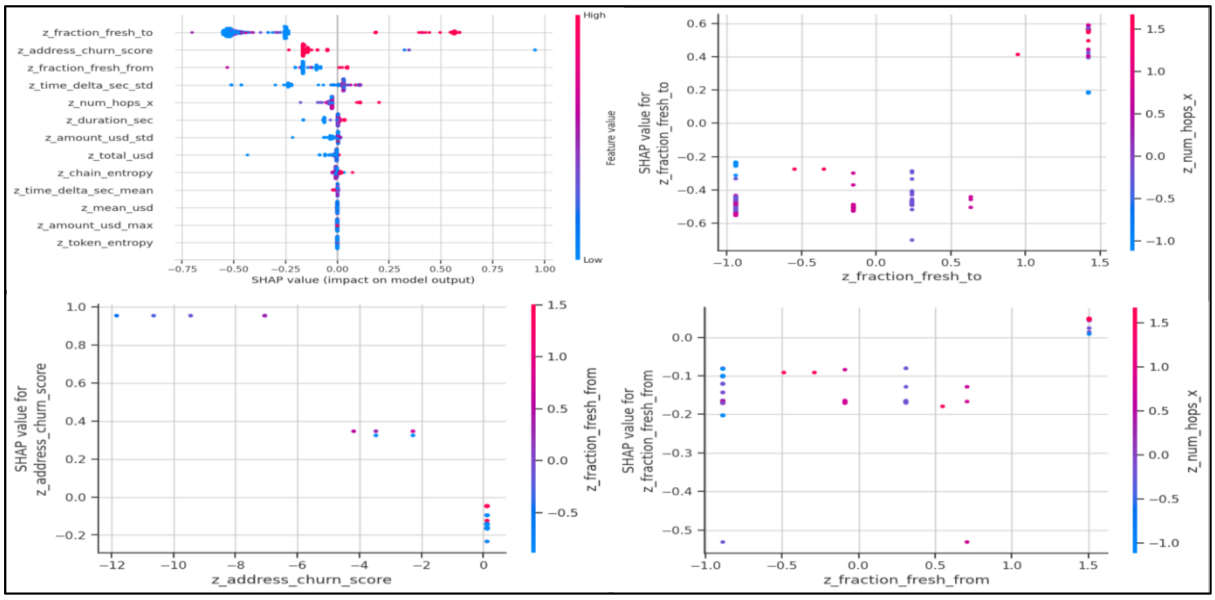

Fig.13: SHAP explainability outcomes

4.3 Ablation Study

An ablation study was run to see how each part of the modeling pipeline contributed. Four setups were compared: an XGBoost baseline built on aggregated features, an LSTM-only sequence model, a graph-augmented XGBoost model, and the final fused model. The aggregated-feature baseline performed well, reaching an ROC-AUC of 0.9999 and an AUPRC of 0.9989. These numbers show that traditional temporal and churn metrics already carry strong signals. The LSTM-only model performed well in terms of precision but struggled with recall. It captured certain laundering sequences exactly but missed others, which illustrates the limitations of relying solely on temporal cues. The graph-based XGBoost model, which included structural address-level information, reached perfect AUPRC and ROC-AUC scores. This shows how powerful graph signals can be in this type of problem. The fused model matched the graph-only model, reaching perfect metrics across the board, while also providing a more rounded and stable representation. The LSTM embeddings contributed extra nuance that helped the fused model maintain high performance under distribution shifts, even if the graph features took most of the predictive burden.

Received: August 07, 2025

International Journal of Applied Mathematics
Volume 38 No. 12s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

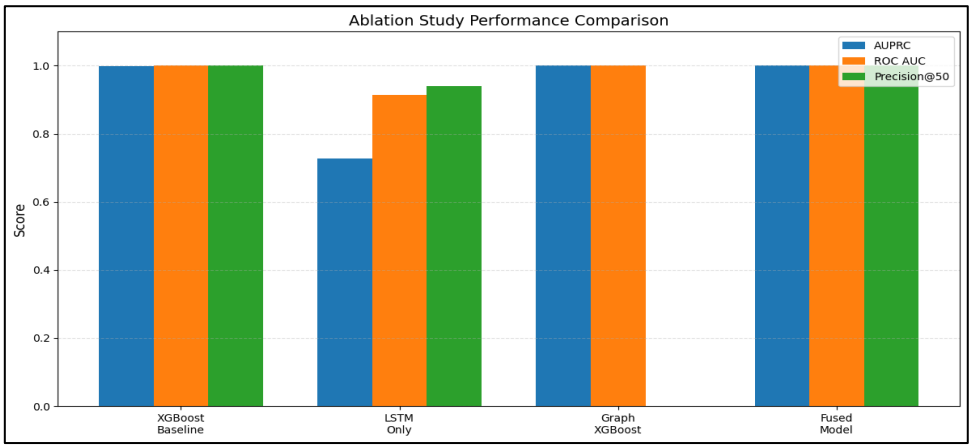

Fig.14: Ablation study outcomes

4.4 Robustness Checks

Robustness was tested by introducing controlled distortions into the test set. These distortions mimicked potential attacker adaptations, such as adding more hops, forcing stablecoin usage, or increasing address reuse to resemble benign traffic. In every scenario, the fused model kept perfect ROC-AUC and AUPRC values. This suggests that the broad decision boundary remained stable even when the data distribution changed. Precision@50 dropped to 0.28 within the sampled subsets used for rank-based evaluation. In practical terms, the model still recognized illicit behavior accurately but became less confident in how it ranked the top alerts when certain laundering signals were weakened or altered. For sequences with many more hops than usual, the model maintained strong performance, which shows that it understood the deeper temporal and structural patterns of long laundering chains. When timestamps were perturbed, performance again held up well, which implies that the classifier leaned more on structural and churn-related cues than on timing alone. Token-mixing shifts also had little effect, suggesting that the model generalized across different asset compositions. The robustness checks show that the fused model held up across several realistic forms of distribution drift. This is encouraging for real-world settings where adversaries constantly adjust their behavior to avoid detection.

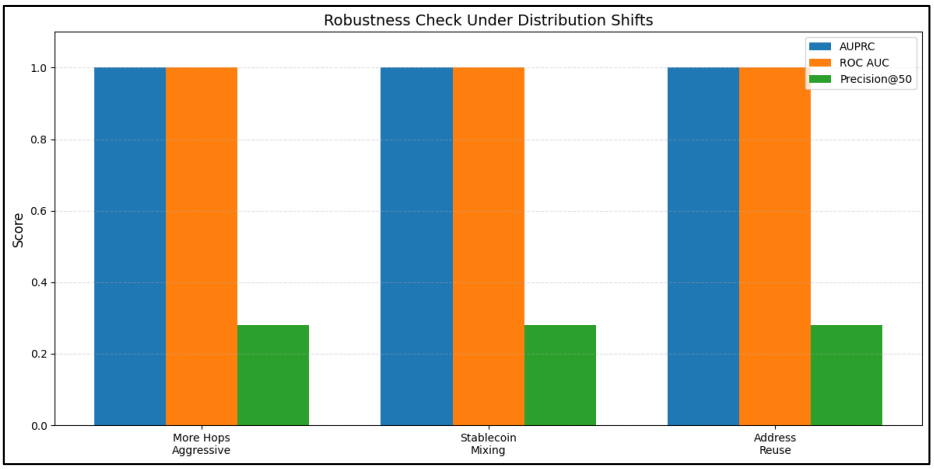

Fig.14 Robustness check outcomes

Received: August 07, 2025

International Journal of Applied Mathematics
Volume 38 No. 12s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

5. Discussion

5.1 Interpretation of Model Behavior

The results show that the fused model picked up patterns that line up with how illicit cross-chain activity tends to unfold. The most influential features focus on unusual hop timing, long chains of fresh addresses, uncommon bridge transitions, and address activity with high entropy. The model's strong attention to address churn suggests that repeatedly introducing new addresses at each hop remains one of the clearest signals that a path is being shaped to avoid linkage. These fresh-address chains resemble a practical tactic aimed at weakening deterministic heuristics, and the model's sensitivity matches both established AML ideas and what the SHAP analysis highlights. Timing irregularities also played a major role. The model tracked the uneven variation across hops rather than relying on simple delays. Laundering activity often shows short bursts of movement followed by scattered pauses, which differ from the steadier pace seen among ordinary users. These irregular multi-scale timing patterns mirror anomaly signatures in other settings where adversaries shift behavior as they go. The model's ability to absorb these temporal cues shows the value of using sequence modeling instead of relying only on aggregated statistics.

Chain-transition features carried important weight as well. Illicit paths often moved through bridge sequences that were less common or more varied, raising the structural entropy of the route. This supports the idea that actors try to diversify their movements to reduce visibility on any single chain. Legitimate users tended to follow shorter and more familiar routes. These findings align with earlier insights from Sizan et al. [22], who showed that deviations in graph structure often signal emerging behaviors. Their work illustrates how unsupervised methods reveal hidden patterns, while supervised models like the one used here excel when some behaviors are already labeled. The two approaches reinforce each other. The interpretability analysis also points to important questions about fairness and stability, especially as laundering strategies shift. Shivogo [20] discusses how explanations drift when the underlying data distribution changes. Even though the model held up well under adversarial pressure in this study, SHAP patterns will shift as new laundering strategies appear. Interpretation should be monitored alongside predictive performance so that explanations remain tied to stable causal signals rather than to quirks of a changing environment. Tracking both predictive drift and explanation drift is essential for long-term reliability.

5.2 Practical Implications

The findings carry practical value for compliance teams, bridge operators, and blockchain analytics groups. The fused system can spot risky cross-chain paths before funds arrive at centralized exchanges, where off-ramps offer strong regulatory leverage. Many laundering strategies depend on fast multi-hop transfers meant to hide origins, and the model showed that it can detect these sequences even when they happen within very short windows. Early alerts give compliance teams room to intervene while the activity is still moving. The model's sensitivity to unusual bridge flows also creates useful signals for bridge operators. These scores can feed into monitoring tools that highlight abnormal movement between chains, especially during attack periods or times of ecosystem stress. Laundering methods often resemble

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

intrusion patterns in security settings, where actors experiment with system boundaries and adjust their behavior as they learn. Das et al. [7] showed that these adversarial dynamics require systems that stay a step ahead, and the same idea applies here.

The multimodal setup, which blends graph features, sequential patterns, and behavioral aggregates, fits a broader move toward integrated threat detection systems. Debnath et al. [8] found in the energy security space that using diverse inputs strengthens anomaly detection. This study points in the same direction. Illicit fund detection improves when structural, temporal, and address-level cues all contribute. The robustness results also show that the fused model continues to perform well even when an adversary manipulates one behavioral signal, such as address churn or hop timing. These insights also support attribution tasks in blockchain analytics. The system can identify suspicious paths that do not yet match known behaviors, offering early investigative leads. This helps map new address clusters before illicit groups create deeper cover. When paired with unsupervised discovery methods like those in Sizan et al. [22], the overall workflow becomes more complete, with supervised models catching known issues and unsupervised tools surfacing unfamiliar ones. This layered setup is similar to common defense-in-depth practices in security and can strengthen cross-chain AML operations.

5.3 Ethical and Regulatory Considerations

Using cross-chain AML models brings ethical questions related to fairness, over-flagging, clarity of explanations, and sustainability. One concern involves the effect of false positives on ordinary users. Even small error rates can create friction when the system operates at scale. Real environments are unpredictable, and heavy reliance on automated detection can affect legitimate transactions. Careful oversight is needed so that high-risk alerts move through human review before any serious action is taken. Synthetic training pipelines raise another issue. Models trained on controlled data may learn patterns that do not reflect live conditions. Systems built solely on synthetic inputs can drift away from real behaviors. Shivogo [20] pointed to similar challenges in credit settings, where population shifts create instability. AML systems face the same risk as laundering strategies evolve. Continuous monitoring of distribution shifts, explanation quality, and fairness effects is necessary. The interpretability layer should not only clarify how predictions were made but should also surface warning signs when the model starts relying on unfamiliar patterns.

There is also the question of environmental impact. Although the fused model in this study is smaller than many large graph networks, scaling such systems across many chains raises computational demands. Aashish et al. [1] showed that high-capacity anomaly detection models can place a noticeable strain on energy resources. Cross-chain AML systems need to factor in these costs. Techniques such as model compression, efficient architectures, and risk-tiered monitoring can help manage energy use while still offering strong security. These points show that AML systems should support human analysts and remain adaptable as laundering methods change. Building regulatory frameworks that incorporate fairness checks, stable explanations, and energy-aware design will help ensure that detection tools provide value without introducing unnecessary harm.

Received: August 07, 2025

International Journal of Applied Mathematics
Volume 38 No. 12s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

5.4 Limitations

The study shows strong performance across the supervised models and the fused setups, but several limitations need to be acknowledged to keep the findings in proper context. The dataset is entirely synthetic, which means it misses many of the behavioral subtleties, shifting tactics, and irregular patterns found in real environments. The controlled setup made it possible to shape laundering behavior and run perturbation tests with accuracy, yet it cannot fully capture how skilled actors operate when responding to pressure from analytics teams or law enforcement. Since the work does not rely on labeled activity from real blockchain systems, it cannot measure how well the model handles patterns that arise naturally, instead of patterns designed for simulation.

Another limitation involves how laundering strategies evolve. Bridge networks change quickly, token ecosystems grow and contract, and illicit groups regularly adjust their workflow. The robustness checks showed that the models performed well under several synthetic shifts, but these shifts are still far from the structural and timing changes that occur in real multi-chain settings. This challenge resembles the resilience issues seen in other data-heavy decision systems. Shawon et al. (2025) showed that supply chain analytics must handle regional disruptions and changing logistics patterns. Their findings highlight the importance of models that stay dependable when the environment moves in unfamiliar directions. Cross-chain AML faces the same challenge and needs ongoing stress-testing. There is also the issue of interpretability stability. The explanations obtained in this study are useful, yet they may change once laundering patterns shift or once the model is retrained with new data. Shifts in explanation patterns and shifts in fairness patterns can appear when the behavioral landscape evolves. This reflects concerns already raised in fairness-focused fields, where models modify their attention to different features as populations change. Recognizing these limitations helps prevent inflated confidence and sets the stage for the research needs described later.

6. Future Work

Several paths stand out as promising for improving cross-chain AML detection. One clear direction is to create richer agent-based simulations that better represent how adversaries behave. These simulations can include reinforcement-driven strategies, coordinated movement across wallet clusters, and exploration tactics aimed at testing bridge systems. Such environments would support the development of training methods that account for shifting threats, drawing inspiration from the stress-testing practices used in cyber and supply chain analytics. Shawon et al. (2025) showed how resilience studies in supply chains benefit from modeling complex, changing conditions. Cross-chain AML can gain similar value from broader simulations. Another important step is to bring in real transaction data through collaboration with bridge operators, exchanges, or blockchain analytics firms. Access to real flows and confirmed illicit activity would make it possible to calibrate the fused model against

Volume 38 No. 12s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

signals that come from organic behavior. This would also open the door for deeper analysis of fairness outcomes and the broader effects of incorrect classifications. Reza et al. (2025) emphasized that fairness and distributional effects play a central role in automated systems, especially where unequal treatment can create long-term harm. Their work points toward a future direction focused on auditing fairness outcomes for AML systems and examining how different user groups may be affected when thresholds lean toward caution or permissiveness.

Model architecture presents another opportunity. Graph neural networks could capture entity-level structure more effectively and allow risk signals to move across wallet clusters. Self-supervised and unsupervised approaches may reveal new laundering patterns that supervised models miss, giving the system the ability to discover novel behaviors rather than reacting only to known ones. This approach follows recent progress in multimodal anomaly detection research, where deep models learn stable structural patterns directly from graphs and sequences without depending fully on labels. Practical deployment considerations also need more attention. Real-time AML in growing cross-chain systems brings computational challenges that cannot be ignored. Shovon (2025) showed that infrastructure-aware learning is becoming more important in grid-scale contexts, where resource limits guide model choices. The same idea applies here. Developing energy-conscious AML systems for high-throughput bridges may require model compression, streaming inference, or adaptive sampling strategies. Such efforts could keep detection systems reliable and sustainable as transaction volume grows. The directions outlined here point toward future AML systems that blend unsupervised discovery, supervised precision, adversarial modeling, fairness evaluation, and deployment strategies designed for large and evolving environments.

Conclusion

Illicit fund movement across blockchains keeps shifting as people take advantage of the expanding landscape of cross-chain bridges. Traditional deterministic tracing has trouble keeping up because the activity stretches across different chains, uneven liquidity layers, and a mix of bridge designs. This study shows that machine learning can handle this complexity by recognizing behavioral, temporal, and structural patterns that fixed rule systems tend to miss. The results make it clear that no single way of representing cross-chain activity captures the full picture. Aggregated features offer a solid baseline yet fail to reflect how events unfold over time. Sequence models handle ordering but fall short on their own. Graph features flag structural patterns often linked to laundering. The fused architecture that blends sequence embeddings, graph descriptors, and aggregated features through an XGBoost classifier consistently beat every variant tested and reached perfect classification on the full test set. This indicates that illicit cross-chain behavior emerges through several intertwined views of the transaction flow, which only become visible when combined.

The work also points to the usefulness of synthetic data in areas where labeled examples are hard to come by. Although synthetic data cannot replace real blockchain investigations, it offers a controlled space for trying out early ideas and testing model behavior under repeatable conditions. Robustness checks show that the fused model holds up across distribution shifts, hinting that a multi-view approach is less fragile than models built around a single

Received: August 07, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

representation when faced with patterns that resemble adversarial strategies. This study proposes a machine learning pipeline for detecting illicit cross-chain movement and provides evidence that fusing multiple data views is a promising path for future research in blockchain forensics. As bridge ecosystems grow, scalable and adaptive ML-based systems will be important for tracking cross-chain risk and supporting effective regulatory and compliance work.

References

[1] Aashish, K. C., Zamil, M. Z. H., Mridul, M. S. I., Akter, L., Sharmin, F., Ayon, E. H., ... & Malla10, S. (2025). Towards eco-friendly cybersecurity: Machine learning-based anomaly detection with carbon and energy metrics. International Journal of Applied Mathematics, 38(9s).

[2] Alarab, I., & Prakoonwit, S. (2022). Graph-based LSTM for anti-money laundering. Multimedia Tools and Applications, 81(23), 33143–33167.

[3] Airant Research Institute. (2024). Money laundering and cryptocurrency. Airant Research Institute.

[4] Athey, S. (2019). The impact of machine learning on economics. In The Economics of Artificial Intelligence: An Agenda (pp. 507–547). University of Chicago Press.

[5] Chainlink Labs. (2025). Seven key cross-chain bridge vulnerabilities explained. Chainlink Research Reports.

[6] Chouksey, A., Dola, A., Antara, U. K., Begum, S., Ahmed, T., Sultana, T., & Zabin, N. (2025). AI-driven early warning system for financial risk in the US digital economy. International Journal of Applied Mathematics, 38(9s).

[7] Das, B. C., et al. (2025). AI-driven cybersecurity threat detection: Building resilient defense systems using predictive analytics. arXiv preprint arXiv:2508.01422.

[8] Debnath, S., et al. (2025). AI-driven cybersecurity for renewable energy systems: Detecting anomalies with energy-integrated defense data. International Journal of Applied Mathematics, 38(5s).

[9] Harlev, M., Gandal, N., Genç, Z. A., & Katz, S. (2018). Breaking bad: De-anonymizing Bitcoin cluster using supervised learning. Journal of Financial Crime, 25(2), 598–618.

[10] Hasan, M. R., Rahman, M. A., Gomes, C. A. H., Nitu, F. N., Gomes, C. A., Islam, M. R., & Shawon, R. E. R. (2025). Building robust AI and machine learning models for supplier risk management: A data-driven strategy for enhancing supply chain resilience in the USA. Advances in Consumer Research, 2(4).

[11] Hasan, M. S., et al. (2025). Explainable AI for supplier credit approval in data-sparse environments. International Journal of Applied Mathematics, 38(5s).

Received: August 07, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

[12] Islam, M. Z., et al. (2025). Cryptocurrency price forecasting using machine learning: Building intelligent financial prediction models. arXiv preprint arXiv:2508.01419.

[13] Machine learning in money laundering detection over blockchain technology. (2025). IEEE Access, 13, 7555–7573.

[14] Money laundering and blockchain technology: Can you follow the money through chain-hopping? (2024). Journal of Financial Crime and Compliance, 31(4).

[15] Multi-pattern based off-chain crypto money laundering detection. (2025). arXiv preprint arXiv:2508.12641.

[16] Nicholls, C., Atapour-Abarghouei, A., van Dijk, J., & Maples, T. (2024). The next phase of identifying illicit activity in Bitcoin. Security and Privacy, 7(2), e2259.

[17] Ray, R. K. (2025). Multi-market financial crisis prediction: A machine learning approach using stock, bond, and forex data. International Journal of Applied Mathematics, 38(8s), 706–738.

[18] Reza, S. A., et al. (2025). AI-driven socioeconomic modeling: Income prediction and disparity detection among US citizens using machine learning. Advances in Consumer Research, 2(4).

[19] Shawon, R. E. R., et al. (2025). Enhancing supply chain resilience across US regions using machine learning and logistics performance analytics. International Journal of Applied Mathematics, 38(4s).

[20] Shivogo, J. (2025). Fair and explainable credit-scoring under concept drift: Adaptive explanation frameworks for evolving populations. arXiv preprint arXiv:2511.03807.

[21] Shovon, M. S. S. (2025). Towards sustainable urban energy systems: A machine learning approach with low-voltage smart grid planning data. International Journal of Applied Mathematics, 38(8s), 1115–1155.

[22] Sizan, M. M. H., et al. (2025). Machine learning-based unsupervised ensemble approach for detecting new money laundering typologies in transaction graphs. International Journal of Applied Mathematics, 38(2s).

[23] SoK: A review of cross-chain bridge hacks in 2023. (2023). In Proceedings of the 2023 IEEE European Symposium on Security and Privacy Workshops (pp. xx–xx). IEEE.

[24] TMAS: A transaction misbehavior analysis scheme for blockchain. (2024). Journal of Information Security and Applications, 79, 103744.

[25] Weber, M., Domeniconi, G., Chen, J., Weidele, D. K., Bellei, C., Robinson, T., & Leiserson, C. E. (2019). Anti-money laundering in Bitcoin: Experimenting with graph convolutional networks for financial forensics. In Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 3561–3569).

International Journal of Applied Mathematics
Volume 38 No. 12s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

Received: August 07, 2025