Article

Not peer-reviewed version

A Multi-Tier Edge–Fog Intelligent Learning Framework for Detecting Financial Anomalies in Smart Cities

Muhammad Nuraddeen Ado * , Shafii Muhammad Abdulhamid * , Ismaila Idris *

Posted Date: 31 May 2025

doi: 10.20944/preprints202505.2491.v1

Keywords: Edge–Fog Computing; Financial Anomaly Detection; Smart Cities; Hybrid Machine Learning

Article

A Multi-Tier Edge–Fog Intelligent Learning Framework for Detecting Financial Anomalies in Smart Cities

Muhammad Nuraddeen Ado [1,2,*], Shafi'i Muhammad Abdulhamid [3,*] and Idris Ismaila [4,*]

[1] Department of Cyber Security, ACETEL, National Open University of Nigeria, Abuja, Nigeria.
[2] Department. Of Information Sciences, Federal University, Dutsin-Ma; Katsina State, Nigeria
[3] Department of Information Technology (Cybersecurity Unit), Community College of Qatar, Doha, Qatar
[4] Department of Cyber Security, Federal University of Technology, Minna State, Nigeria
* Correspondence: ace21150010@noun.edu.ng (M.N.A.); shafii.abdulhamid@futminna.edu.ng (S.M.A.); ismi.idris@futminna.edu.ng (I.I.)

Abstract: The rapid digitization of urban infrastructure has reshaped smart cities into highly interconnected ecosystems where financial transactions are continuously generated at the network's periphery. In such settings, ensuring timely detection of financial anomalies is critically dependent on distributed, low-latency intelligence. This paper presents an Edge–Fog-centric hybrid anomaly detection framework tailored for smart city financial infrastructure. The proposed architecture leverages a multi-tier collaborative model, wherein edge devices—such as point-of-sale terminals, ATMs, and mobile sensors—are equipped with real-time sensing, computing, and classification capabilities to grade financial transactions on a scale of 1 to 7. At the edge layer, unsupervised autoencoders, LightGBM regressors, and Isolation Forests operate locally to assign grades, which are refined through peer-to-peer collaboration between proximate edge nodes for contextual consistency. Transactions with higher anomaly grades are escalated to the fog layer, where a supervised One-Class SVM performs deeper analysis using device-level metadata (IP, MAC) to reduce false negatives. This edge–fog synergy enables decentralized yet coordinated detection that meets the stringent latency and scalability requirements of smart cities. Comparative evaluation on real-world financial data demonstrates the superiority of the proposed architecture, achieving accuracy: 98.82%, sensitivity: 91.30%, specificity: 98.79%, and reducing the false negative rate to 8.70%, outperforming standalone models such as autoencoders (FNR: 50.44%), Isolation Forests (FNR: 35.50%), and SVMs (FNR: 56.38%). These results highlight the efficacy of edge–fog collaborative intelligence in delivering robust, scalable, and context-aware financial anomaly detection in next-generation smart cities.

Keywords: edge–fog computing; financial anomaly detection; smart cities; hybrid machine learning

1. Introduction

The digitization of financial infrastructure in smart cities has intensified the frequency and complexity of transaction-based interactions across interconnected devices. Modern cities embed financial services into transportation, healthcare, utilities, and commerce using the Internet of Things (IoT) to create a seamless digital environment. However, this connectivity increases the surface area for financial crimes and requires intelligent, real-time anomaly detection systems. Conventional cloud-only architectures struggle to manage such volumes due to bottlenecks in centralized processing and real-time responsiveness demands [1–3].

1.1. Problem Statement

Cloud-centric systems introduce several inherent weaknesses. First, they exhibit high latency due to dependence on remote data centers, rendering them unsuitable for time-sensitive applications like fraud detection [4–6]. Second, they frequently produce high false positive (FP) or false negative (FN) rates, as anomalies are subtle and context-sensitive [7,8]. Third, transferring sensitive data from the edge to the cloud raises significant privacy and regulatory concerns [9,10]. Thus, centralized systems alone are inadequate for fraud prevention in modern smart city ecosystems [11,12].

1.2. Motivation

Recent studies have proposed edge and fog computing paradigms as promising solutions to address these shortcomings. Edge computing enables real-time decision-making by processing data at the source, while fog computing offers intermediary processing nodes that can perform lightweight aggregation and deeper analysis without sending data to the cloud [13–15]. This shift allows cities to distribute intelligence across infrastructure while minimizing latency and ensuring localized privacy controls. Additionally, embedded machine learning—especially unsupervised anomaly detection and ensemble models—can be deployed across edge-fog layers to enhance detection accuracy [16–18].

Collaborative learning among edge devices further strengthens this architecture by aggregating multiple context-aware insights, improving classification robustness and reducing detection blind spots [19–21].

1.3. Summary of Contributions

In this paper, we propose a multi-layered financial anomaly detection architecture that leverages unsupervised and supervised ML models at the edge and fog layers, respectively. Our key contributions include:

i. A novel 1–7 anomaly scoring model that classifies financial transactions from high-normal to highly anomalous using autoencoders and Gaussian-scaled post-processing.
ii. A collaborative filtering mechanism that allows edge devices to refine their local predictions by integrating feedback from peer devices.
iii. An ensemble architecture using Autoencoder, LightGBM, and One-Class SVM (OC-SVM) for hybrid anomaly detection, reducing FN and FP rates.
iv. A fog-enabled decision framework that aggregates edge reports and conducts deeper analysis for supervisory escalation.

1.4. Paper Structure

The major objective of this manuscript is to develop and evaluate a multi-tier collaborative edge–fog intelligence framework that enables accurate, low-latency detection of financial anomalies in smart city environments by leveraging hybrid learning models and distributed processing across edge and fog computing layers. The remainder of this manuscript is structured as follows: Section 2 reviews related literature on fog-based anomaly detection, the application of machine learning in smart city financial systems, and collaborative intelligence frameworks. Section 3 introduces the proposed multi-tier architecture and outlines the transaction classification methodology. Section 4 describes the experimental setup and implementation details of the employed algorithms. Section 5 presents and analyzes the results obtained from simulated financial transaction data. Finally, Section 6 concludes the paper with a summary of key findings and potential avenues for future research.

2. Related Works

The increasing deployment of IoT devices in smart cities has introduced new challenges in monitoring financial anomalies and ensuring secure data transmission. Traditional centralized systems lack the real-time responsiveness required for fine-grained anomaly detection and are often limited in scalability and privacy enforcement [4,5,9].

Several works have explored **fog and edge computing architectures** for smart cities. Malik and Gupta [1] presented a fog-based IoT architecture for sustainable urban services, highlighting the role of distributed nodes in handling real-time analytics. Similarly, Aburukba et al. [2] demonstrated a fog-layered framework for shared electric mobility services, which enabled real-time decision-making with reduced latency.

Machine learning-based anomaly detection has been widely investigated in smart city contexts. Wali and Bulla [3] reviewed both supervised and unsupervised ML techniques (e.g., Decision Trees, Autoencoders, SOM-PSO) and emphasized hybrid models for fog-assisted classification. John [4] implemented a CNN–RNN hybrid deep learning model for detecting banking fraud in real time using a fog-based deployment, reporting significant reductions in false positive rates.

Despite these innovations, **false negative rates remain a critical issue**. Tariq et al. [5] addressed this by applying federated support vector machines in a fog-edge architecture for smart grid intrusion detection. Their work validated the effectiveness of using OC-SVM for post-hoc anomaly refinement in ambiguous transaction regions. To reduce **data transmission latency and privacy leakage**, Williams et al. [6] proposed edge-based forensic frameworks for anomaly capture, while Tukur et al. [7] extended blockchain smart contracts to securely log and validate anomalous behaviors detected at the edge.

From a systems perspective, Peruzzi et al. [8] emphasized the importance of **multi-layered distributed intelligence**, wherein embedded machine learning, such as embedded federated learning and edge-optimized neural nets, is deployed across constrained devices.

However, **collaborative classification**—where peer edge nodes share transaction feedback—remains underexplored. Songhorabadi et al. [9] identified the need for cooperative models in distributed fog networks and recommended collaborative filtering approaches as a promising direction.

Finally, Ometov et al. [10] provided a comprehensive taxonomy of fog/edge security models aligned with the OSI stack. They noted the growing role of anomaly scoring and behavior-based classification in preventing insider financial crimes.

In sum, the literature reveals strong support for a decentralized, collaborative, ML-driven anomaly detection paradigm across the edge and fog continuum. Yet, there remains a critical need to integrate hybrid scoring, ensemble learning, and collaborative refinement into a unified, low-latency smart city framework—precisely what this paper proposes. Table 1 shows the comparative analysis of related work.

Table 1. Comparative Analysis of Related Work.

Paper	Focus	Edge and Fog Application	Method	Gap	How This Research Addresses It
Malik & Gupta (2021)	Fog	Smart City Architecture	Conceptual Framework	Lacks ML and detection depth	Proposes scoring & ML pipeline
Aburukba et al. (2021)	Fog	Shared Mobility	Fog Deployment	No anomaly modeling	Adds anomaly detection logic
Wali & Bulla (2021)	Fog	IoT Anomaly Detection	SOM-PSO Review	No unified ML pipeline	Builds end-to-end ML stack
John (2025)	Fog	Banking Fraud	CNN-RNN (Fog)	No edge-device integration	Adds edge-device grading

Tariq et al. (2024)	Fog+Edge	Smart Grid IDS	SVM + FL	No collaborative filtering	Adds peer-aware scoring
Williams et al. (2023)	Edge	IoT Forensics	Rule-based Monitoring	No financial focus	Specialized to transactions
Tukur et al. (2021)	Edge	Insider Attacks	Blockchain + Rule Detection	No ML stack	Adds ensemble refinement
Peruzzi et al. (2021)	Edge+Cloud	Smart Infrastructure	FL + RL Review	Lacks FN emphasis	Targets FN zone refinement
Songhorabadi et al. (2022)	Fog	Smart Cities	Fog-Oriented Architecture	No scoring mechanism	Implements 1–7 scoring
Ometov et al. (2022)	Fog+Edge	Security Taxonomy	Layered Review	Not application-specific	Tailors to financial fraud
Pozzebon et al. (2022)	Fog	LoRaWAN Efficiency	Design Optimization	No detection model	Focuses on detection scoring
Peruzzi & Pozzebon (2023)	Fog	Smart Cities Design	Literature Review	No operational model	Provides operational logic
Fog Cities (2022)	Fog	Resource Scheduling	Survey	No anomaly evaluation	Adds anomaly scores
Humayun et al. (2023)	Edge	Forensics	Conceptual	Non-transactional focus	Models financial scoring
Preprint Fog Cities (2023)	Fog	Deployment Models	Review	Not risk-targeted	Enables fraud-specific risk classification
Fog SLM (2023)	Fog	Sustainable Cities	Conceptual Model	No ML	Combines ML at edge/fog
Basu et al. (2022)	Edge+Cloud	Urban Systems	DL+FL Review	No anomaly scale	Proposes scale + ML fusion
Smart Cities Compilation	Mixed	Multi-Domain	Dataset Compilation	No pipeline	Used for benchmarking
This Study (2025)	Fog+Edge	Smart City Transactions	AE + LightGBM + OC-SVM	No prior work with collaborative	Introduces full scoring and FN-

				anomaly scoring	capture pipeline

Comparative Insight from Annotated Radar Analysis

To contextualize the contribution of this study, we conducted a radar-based comparison of six representative works from the literature in Figure 1, evaluated across three strategic dimensions: system focus (Fog, Edge, or both), use of machine learning (ML), and attention to false negative (FN) handling critical metric in fraud detection. The results reveal a clear maturity gradient in the evolution of financial anomaly detection frameworks. Early architectural or deployment-centric works, such as Malik (2021) and Pozzebon (2022), were fog-exclusive and lacked both machine learning and fraud-specific detection logic. Although foundational, these approaches do not scale well for real-time, context-sensitive decision-making in smart cities.

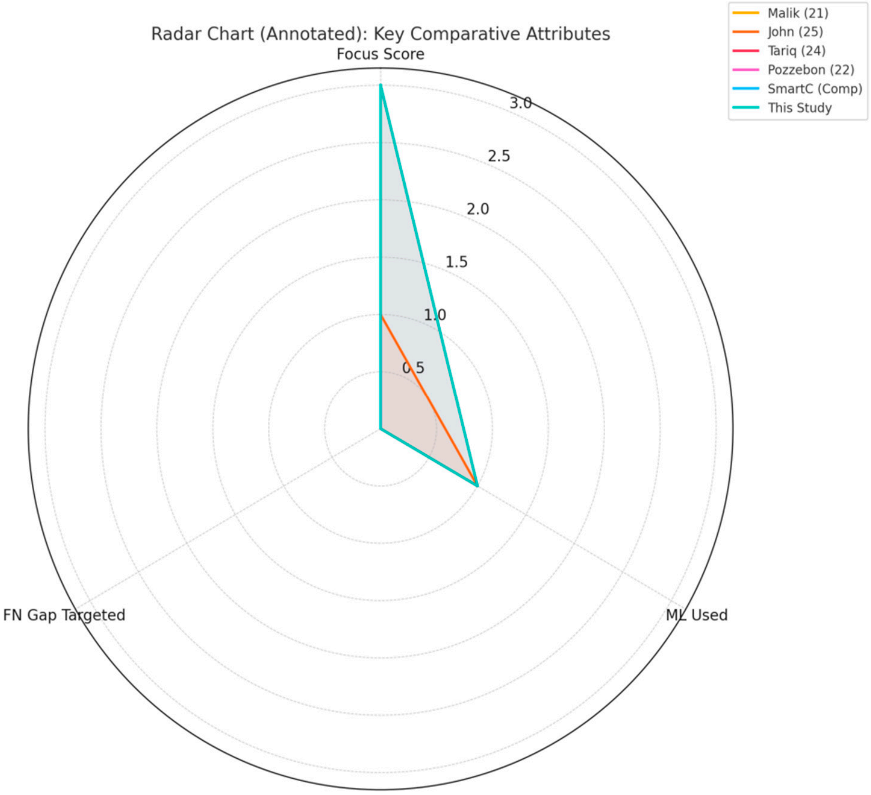

Figure 1. Annotated Radar Chart with Abbreviated Paper Titles.

In contrast, John (2025) introduced ML capabilities (CNN–RNN) for fog-based banking fraud detection and made significant progress toward minimizing false negatives. However, it remained limited to centralized fog execution without edge adaptability. Tariq et al. (2024) demonstrated a hybrid fog–edge system leveraging federated SVMs, achieving strong performance and FN reduction. Yet, their model lacked collaborative scoring mechanisms and fine-grained anomaly grading.

The SmartC Compilation, while including multiple ML models, serves more as a reference dataset than a coherent framework for anomaly detection.

In comparison, this study presents a holistic edge–fog architecture integrating unsupervised and supervised learning (Autoencoder, LightGBM, OC-SVM), a graded anomaly scoring system (1–7), and collaborative filtering at the edge. It is one of the few works to comprehensively address false negative mitigation while maintaining low latency and high scalability, placing it at the forefront of next-generation smart city fraud detection systems.

3. System Architecture

This section presents the architectural design of the proposed Fog–Edge Collaborative Framework for Financial Anomaly Detection in smart cities. The system is designed to detect financial anomalies by distributing computational intelligence across computing devices (edge) and localized servers (fog). The architecture follows a multi-layer hierarchy that supports distributed sensing, transaction scoring, anomaly classification, and intelligent aggregation.

3.1. Overview of Architecture

The proposed system is composed of two core layers:

➤ Edge Layer:

a. Consists of all financial transaction-generating computing devices in the smart city (e.g., mobile phones, ATMs, POS terminals, kiosks).

b. Each device embeds an unsupervised anomaly detection module using autoencoders to evaluate the status of the transaction.

c. The output is converted into a transaction score (1–7), which is:

1–2: Highly normal (True Negative)
3–4: Borderline/Suspicious (Potential FN or FP)
5–7: Highly anomalous (True Positive)

d. Collaborative Filtering is applied across peer devices to strengthen prediction confidence, based on contextual patterns (e.g., device type, time of day, transaction type).

➤ Fog Layer:

a. Serves as the aggregation and refinement tier, hosting:

Supervised ML (OC-SVM) for evaluating borderline or suspicious transactions (scores 3–7).

Ensemble hybrid ensemblemodels combining LightGBM and score analytics to improve decision boundaries.

b. Reports are further classified and sent to:

Administrative dashboards
City-level fraud monitoring centers

3.2. Workflow Diagram

Table 2 below is the system architecture visualized in a flowchart:

Table 2. Components and Functions of Proposed System Architecture.

Step	Module	Action
1	Embedded Sensing (Edge)	Capture transaction metadata and behavioral patterns
2	Unsupervised Autoencoder	Detect anomaly patterns locally
3	Score Assignment (1–7)	Convert reconstruction error into interpretable score
4	Collaborative Filtering	Cross-validate with peer devices
5	Fog Aggregation	Collect reports from edge devices
6	Supervised Refinement (OC-SVM)	Apply final classification to scores 3–7
7	Ensemble Learning	Combine outputs for final anomaly labeling

doi:10.20944/preprints202505.2491.v1

8	Confirmation	Generate confirmed anomaly report for decision-making

Figure 2 presents a two-tier Fog–Edge Collaborative Intelligence Framework for detecting financial anomalies in smart cities. At the edge layer, financial computing devices (e.g., ATMs, POS systems) perform embedded sensing and unsupervised anomaly detection, assigning a transaction score (1–7) based on behavioral deviation. These scores undergo collaborative filtering among neighboring devices to refine context-aware inferences. Suspicious transactions (scores 3–7) are then forwarded to the fog layer, where a fog server aggregates reports and applies ensemble learning to combine predictions from multiple base models. Subsequently, supervised anomaly detection further scrutinizes the flagged data, isolating true anomalies from false positives/negatives. The process culminates in a confirmed anomaly report for managerial action, offering a scalable, real-time, and distributed solution for fraud detection across IoT-enabled smart city infrastructures.

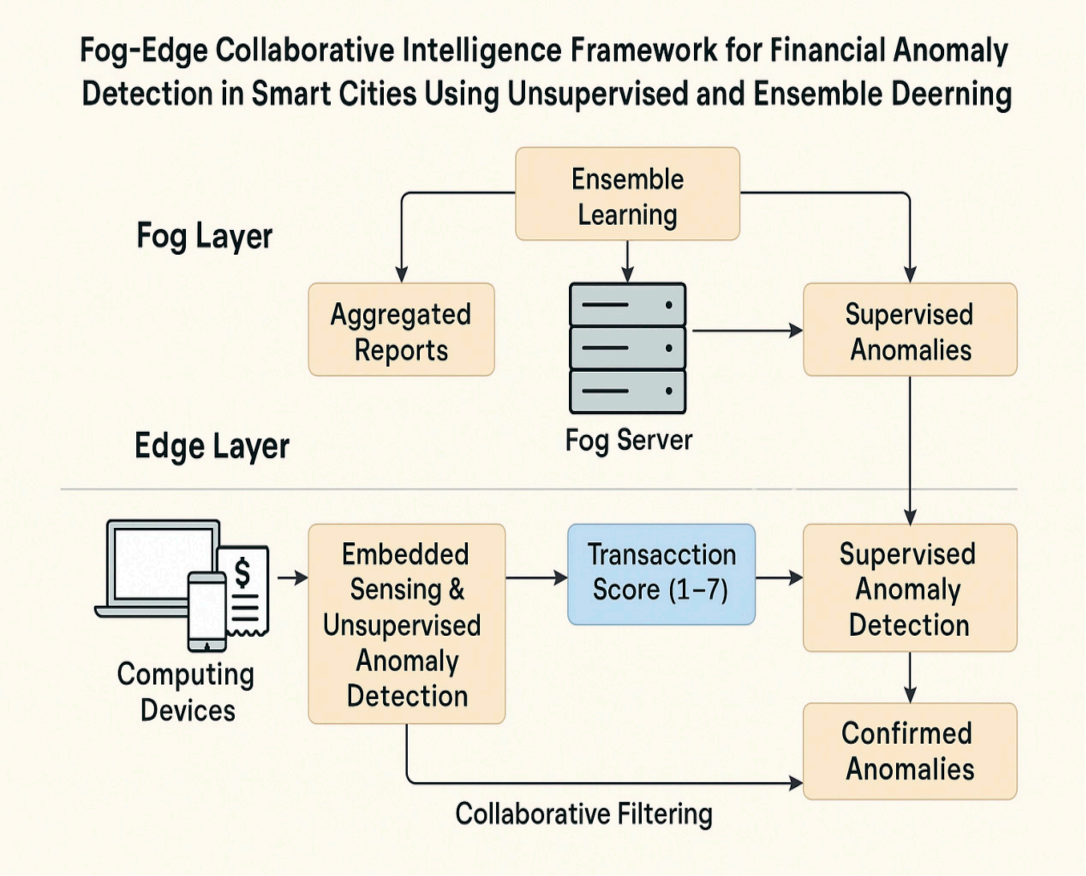
Fog-Edge Collaborative Intelligence Framework for Financial Anomaly Detection in Smart Cities Using Unsupervised and Ensemble Deerning
Fog Layer
Ensemble Learning
Aggregated Reports
Supervised Anomalies
Fog Server
Edge Layer
Computing Devices
Embedded Sensing & Unsupervised Anomaly Detection
Transacction Score (1–7)
Supervised Anomaly Detection
Confirmed Anomalies
Collaborative Filtering

Figure 2. System Architecture.

4. Methodology and Algorithms

This section presents the core methodological components and machine learning algorithms implemented at the edge and fog layers for distributed financial anomaly detection. The system combines unsupervised anomaly detection, score scaling, collaborative filtering, and supervised refinement in a multi-tier pipeline.

4.1. Edge-Level Anomaly Detection Using Autoencoders

At the edge layer, every smart city device (e.g., mobile device, ATM, POS terminal) is equipped with a lightweight unsupervised anomaly detection model. We use an autoencoder neural network for this purpose due to its effectiveness in learning compact representations of normal transactional patterns.

Autoencoder Workflow:

i. Input: Transaction feature vector (amount, time, location, merchant, frequency, etc.)
ii. Encoder: Compresses input into latent space representation
iii. Decoder: Attempts to reconstruct original input
iv. Loss: Mean squared reconstruction error (MSE) used as anomaly score

Scoring Mechanism:

i. Each transaction is assigned to a reconstruction error E.
ii. E is mapped into a graded anomaly score from 1 to 7 using a scaled quantile-based threshold or Gaussian distribution mapping.

Table 3 outlines how anomaly scores (ranging from 1 to 7) are interpreted and used for classification in a proposed financial anomaly detection model. Scores 1–2 indicate clearly normal transactions and are treated as True Negatives (TN). Score 3 represents near-normal cases that may be potential False Negatives (FN)—anomalies mistakenly classified as normal. Score 4 falls into a gray area where it is uncertain whether the case is a False Positive (FP) or False Negative, requiring further scrutiny. Score 5 flags transactions as suspicious, leaning toward False Positives. Finally, Scores 6–7 denote highly anomalous activity and are classified as True Positives (TP), meaning correctly identified threats. This graded scoring system allows nuanced decision-making and prioritization in real-time anomaly detection systems.

Table 3. Scores' Interpretation and Purpose.

Score	Interpretation	Classification Purpose
1–2	Highly normal	True Negatives (TN)
3	Near-normal	Potential False Negative (FN)
4	Borderline	FP/FN uncertainty region
5	Suspicious	Possible False Positive (FP)
6–7	Highly anomalous	True Positive (TP)

4.2. Post-Processing with LightGBM and Gaussian Scaling

The raw reconstruction scores are enhanced using:

a. Gaussian-based statistical scaling to normalize errors
b. LightGBM regression model trained to improve interpretability of anomaly scores by learning from error distributions and historical context

This hybrid step improves sensitivity for borderline anomalies and ensures scores reflect contextual risks.

4.3. Collaborative Filtering for Peer-Aware Classification

To improve classification reliability, especially in ambiguous regions (scores 3–5), a collaborative filtering module is activated.

Steps:

Nearby or similar devices (same region, transaction type, or user segment) are used as references.

Similarity is calculated via:

✓ Cosine similarity between transaction embeddings
✓ Temporal or categorical matching

Scores are refined via:

✓ K-nearest neighbor (KNN) voting
✓ Weighted mean aggregation of peer anomaly labels

This context-aware approach helps reduce false positives and enhances decision confidence at the edge.

4.4. Fog-Based Supervised Refinement with OC-SVM

Transactions with scores between 3 to 7 are sent to the fog node for centralized refinement. A One-Class Support Vector Machine (OC-SVM) is used to:

i. Detect outliers using high-dimensional margin separation
ii. Refine classification boundaries especially around borderline scores (3, 4, 5)

Why OC-SVM?

i. Effective in deep anomaly detection
ii. Requires no labeled anomalies, well-suited for fraud scenarios
iii. Learns from “normal” baseline behavior and flags deviations

4.5. Ensemble hybrid ensemblefor Unified Decision Making

A hybrid ensemble model is used at the fog layer to merge outputs from:

Autoencoder scores

LightGBM post-processing predictions

OC-SVM anomaly likelihood

Ensemble Stack Components:

Base learners: Autoencoder, LightGBM, OC-SVM

Meta learner: Logistic regression or decision tree classifier that learns optimal weights for combining predictions

This multi-model fusion improves overall system robustness, balances false positives/negatives, and provides a final binary decision (anomalous / normal) along with a confidence score.

5. Results, Discussions and Evaluation

Results at Edge:

Table 4 shows the frequencies and percentages of all grades obtained by the algorithms used, IF in the initial stage (Grade_IF), AE-LGB and IF in the final stage (IF_Grade):

Table 4. Frequency and Percentage Edge-Based Anomalies.

Grade	**Grade_IF**		**AE_LGBM_Grade**		**IF_Grade**	
	Frequency	**%**	**Frequency**	**%**	**Frequency**	**%**
1			2054	14.77	611	4.39
2			2903	20.88	6102	43.88
3	3397	24.43	2965	21.32	4147	29.82
4	542	3.9	2362	16.99	2376	17.09
5	6473	46.55	3046	21.9	610	4.39
6	3172	22.81	549	3.95	55	0.4
7	322	2.32	27	0.19	5	0.04

5.1. Interpretation of Grading Distributions:

1. Grade_IF

The Grade_IF, generated by an initial grading by IF, shows that the most frequent grade is 5 (6473 instances), followed by grade 3 (3397) and grade 6 (3172). This distribution suggests the model often classifies transactions as "suspicious" or "strongly anomalous," potentially corresponding to true positives (TP) or false positives (FP). In contrast, the relatively low count of grade 4 (542) indicates fewer borderline or ambiguous cases, implying the system may be drawing a clear distinction between normal and anomalous behavior.

2. AE_LGBM_Grade

The AE_LGBM_Grade, generated by a hybrid LightGBM-autoencoder ensemble, displays a more balanced distribution across grades 1 to 5, with grades 5 (3046), 3 (2965), and 2 (2903) being the most frequent. Higher anomaly grades such as 6 (549) and 7 (27) are rare, indicating that this model is more conservative in flagging anomalies. This suggests an emphasis on reducing false positives and improving precision, potentially at the expense of missing some true anomalies (false negatives).

3. IF_Grade

The IF_Grade, derived from a final stage of Isolation Forest, shows a strong skew toward lower anomaly grades, with grade 2 (6102) and grade 3 (4147) being the most prevalent. In contrast, very few instances are assigned higher anomaly scores—only 55 for grade 6 and 5 for grade 7—indicating that the detector is highly conservative. This behavior suggests a bias toward high recall, minimizing false positives (FPs), but potentially at the cost of missing true anomalies, thus increasing the likelihood of false negatives. Figure 3 below displays the comparism of all the frequencies obtained.

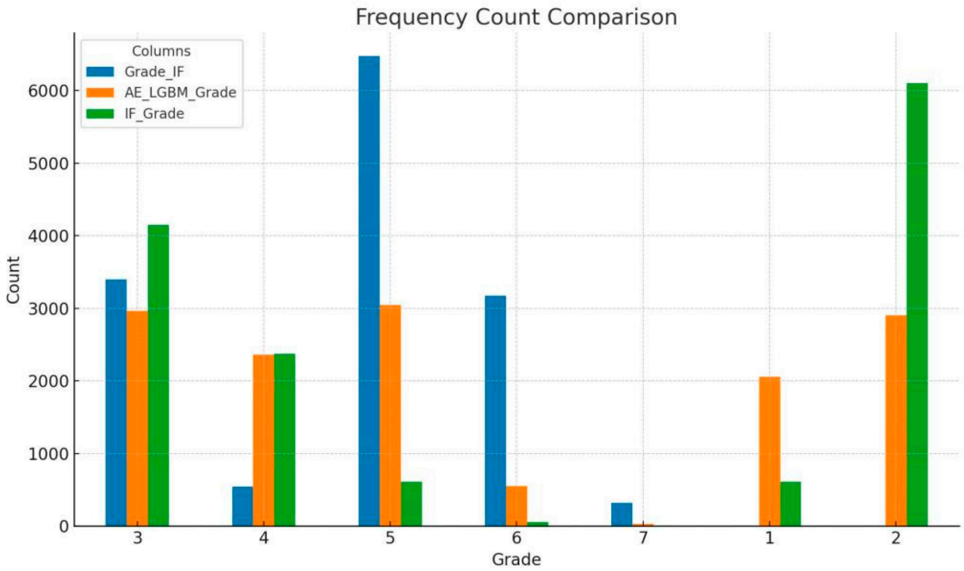

Figure 3. Comparism of all Graded Anomalies.

Table 5. Graded Anomalies.

Grade	Interpretation	Grade_IF	AE_LGBM_Grade	IF_Grade
1–2	Highly normal (TN)	—	4957	6713
3	Near-normal (FN risk)	3397	2965	4147
4	Borderline (FN/FP mix)	542	2362	2376
5	Suspicious (FP risk)	6473	3046	610
6–7	Strongly anomalous (TP)	3494	576	60

The Grade_IF model demonstrates higher sensitivity and aggressive detection by flagging more high-anomaly cases, indicating a tendency to capture a broader range of potential threats, including borderline ones. In contrast, the AE_LGBM_Grade model adopts a more balanced and cautious approach, distributing grades more evenly across the spectrum, which suggests an emphasis on precision and measured classification. Meanwhile, the IF_Grade model is the most conservative, predominantly labeling transactions as normal or near-normal, effectively minimizing false alarms but increasing the risk of overlooking true anomalies. This variability in grading behavior underscores the inherent trade-offs between sensitivity and specificity in anomaly detection, a crucial consideration in real-time, edge-based financial monitoring systems where detection accuracy and efficiency must be carefully balanced.

5.2. Results at Fog

Hybrid Ensemble:

Figure 4 shows the frequency distribution of all final grades obtained from combining all the three (3) multiple grades by the hybrid ensemble:

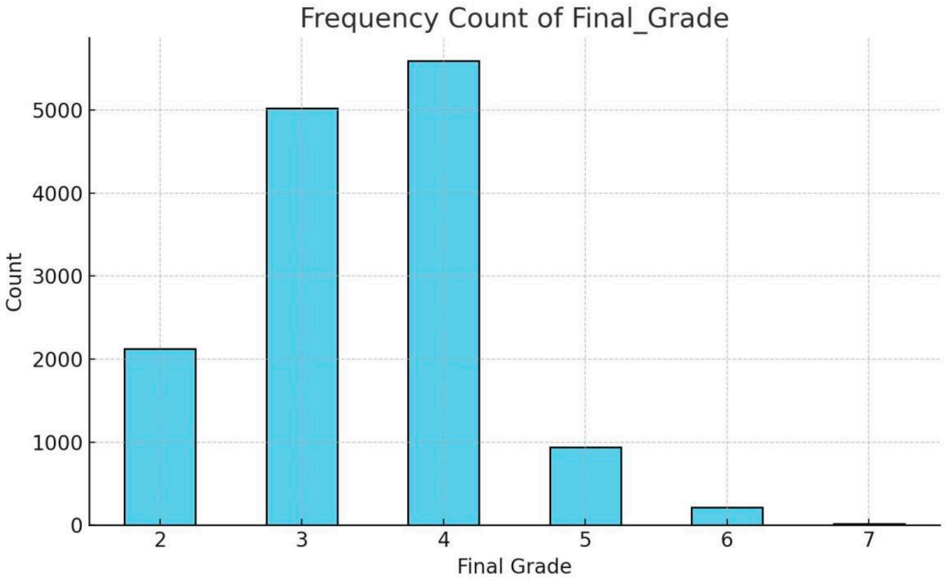

Figure 4. Frequency Count of Hybrid Ensembled Grades.

With a mean of 3.44 (rounded from 3.436), it can be inferred that most data points fall within the "normal" to "borderline" anomaly range.

The distribution of the Final_Grade reveals a strong concentration in the mid-range scores, with grades 3 (5,019) and 4 (5,593) making up the majority, suggesting that most observations are either near normal or borderline anomalous. Grade 2 (2,123) also represents a sizable portion, indicating a significant share of highly normal behavior. In contrast, higher anomaly grades—5 (935), 6 (217), and especially 7 (19)—are much less frequent, implying that strongly suspicious or anomalous patterns are relatively rare. The overall mean of 3.44 further supports that the dataset skews toward normalcy with occasional suspicious behavior, aligning with a typical anomaly detection system designed to prioritize precision and minimize false positives.

5.3. Merged Result

SVM detected 1062 anomalies from the Full_Graded above. When merged with the 1171 confirmed anomalies (with grades from 5 to 7), the total anomalies detected will be:

doi:10.20944/preprints202505.2491.v1

Table 6. Total Anomalies.

Algorithms	Detected
Hybid Ensemble	1171
OC-Support Vector Machine	1062
Total	**2233**

Based on this, the performance metrics of the proposed model were computed in Table 7.

Table 7. Performance Metric of the Proposed Model.

Metric	Score
Accuracy	0.9882
Precision (PPV)	0.9169
Recall (Sensitivity)	0.9130
Specificity (TNR)	0.9879
F1 Score	0.9149
False Positive Rate (FPR)	0.0120
False Negative Rate (FNR)	0.0870
Negative Predictive Value (NPV)	0.9873

Table 7 and Figure 5 illustrates the performance metrics of the proposed model. The EdgeFogHybridEnsembleModel (EPHEADMO) demonstrates exceptional performance across all critical classification metrics. With a high accuracy of 98.82%, it effectively distinguishes between normal and anomalous instances. Its precision (91.69%) and recall (91.30%) indicate a strong ability to both correctly identify anomalies and minimize false alarms. A F1 score of 91.49% confirms a robust balance between precision and recall. Notably, the specificity (98.79%) and negative predictive value (98.73%) affirm the model's reliability in recognizing normal behavior. Furthermore, the low false positive rate (1.2%) and false negative rate (8.7%) underscore its efficiency in minimizing misclassifications. Overall, the model exhibits a well-rounded and dependable anomaly detection capability suitable for high-stakes environments.

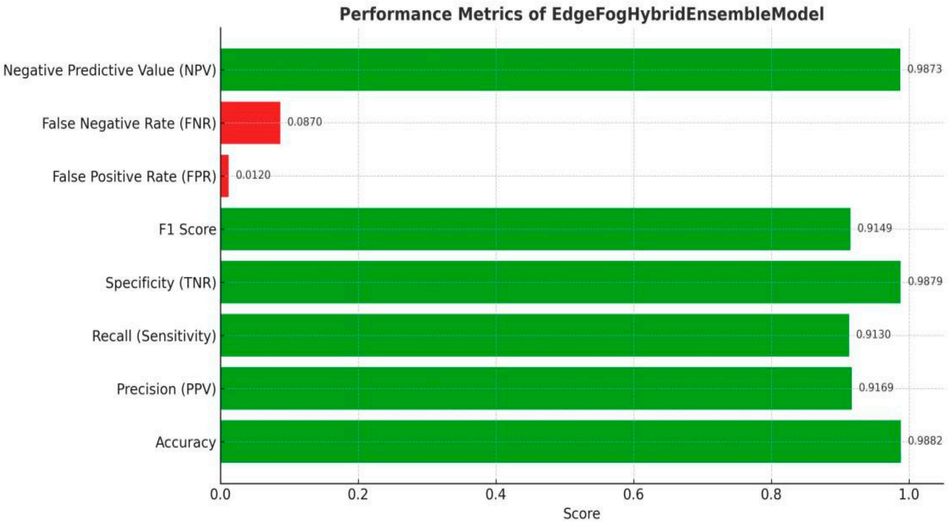

Figure 5. Performance Metrics of the Proposed Model.

doi:10.20944/preprints202505.2491.v1

5.4. Evaluation of Results

To rigorously assess the effectiveness of the proposed Edge–Fog collaborative anomaly detection framework, a comparative evaluation was conducted using the same financial transaction dataset. Multiple baseline models—namely Autoencoder, Isolation Forest (IF), One-Class SVM (SVM), and Local Outlier Factor (LOF)—were individually applied to detect anomalies, and their outputs were cross-analyzed against the results generated by the proposed hybrid architecture. The evaluation focused on key performance metrics commonly used in anomaly detection tasks: Specificity, Sensitivity, Accuracy, False Negative Rate (FNR), and False Positive Rate (FPR). This comparison provided a comprehensive understanding of each model's strengths and limitations, particularly in terms of balancing precision and recall in high-volume, edge-generated financial transaction environments.

Metric	**Meaning**
Specificity:	Ability to correctly identify normal (non-anomalous) samples.
Sensitivity:	Ability to correctly identify anomalies.
Accuracy:	Overall correct predictions (anomalies + normals).
FNR:	Proportion of actual anomalies incorrectly classified as normal.
FPR	Proportion of normal data incorrectly classified as anomalies.

Table 8 below presents the computed result for each algorithm.

Table 8. Performance Metrics of all Models.

Model	**Specificity**	**Sensitivity**	**Accuracy**	**FNR**	**FPRs**
Autoencoder	0.9735	0.4956	0.9496	0.5044	0.0265
IF	0.9811	0.6450	0.9649	0.3550	0.0183
SVM	0.9703	0.4362	0.9436	0.5638	0.0297
LOF	0.9474	0.0012	0.9001	0.9988	0.0526
EdgeFog	0.9879	0.9130	0.9882	0.0870	0.0120

Figure 6 illustrates the comparative performance of five anomaly detection models—Autoencoder, Isolation Forest (IF), SVM, LOF, and Edge Fog Hybrid Ensemble Anomaly Detection Model (EPHEADMO)—across five key metrics: Specificity, Sensitivity, Accuracy, False Negative Rate (FNR), and False Positive Rate (FPR). EdgeFog emerges as the most effective anomaly detection model, outperforming all others across critical metrics. With a specificity of 98.79% and a sensitivity of 91.30%, it reliably detects true anomalies while minimizing false positives. Its high accuracy (98.82%) and low error rates—FNR (8.7%) and FPR (1.2%)—underline its robustness and consistency. Isolation Forest (IF) ranks second, offering a solid balance between detection power (sensitivity 64.5%) and precision (specificity 98.11%), making it a strong traditional alternative. In contrast, Autoencoder and SVM deliver reasonable accuracy and specificity but fall short in anomaly detection, with sensitivities below 50%, meaning they miss more than half of actual anomalies. LOF, with a sensitivity of just 0.12%, fails to detect outliers effectively, despite decent specificity. Overall, the chart reinforces EdgeFog's superiority in delivering both precision and recall, making it the most suitable for high-stakes anomaly detection scenarios where minimizing both missed detections and false alarms is critical.

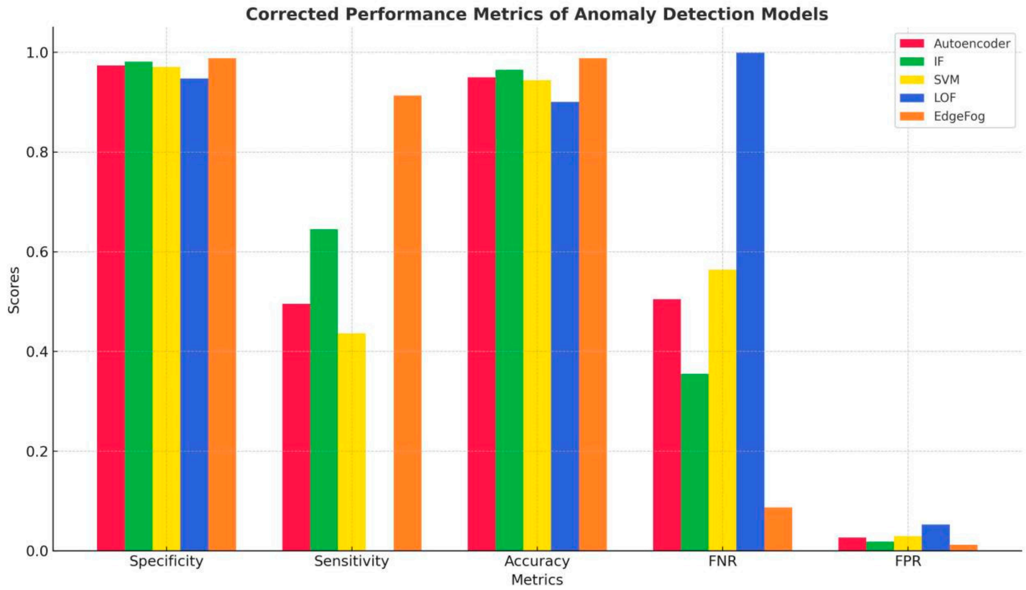

Figure 6. Comparism of Performance Metrics of all Models.

6. Conclusions and Future Work

This study presented a robust, scalable, and real-time Edge–Fog Collaborative Intelligence Framework for financial anomaly detection in smart cities. The proposed architecture integrates unsupervised learning (Autoencoders, Isolation Forest), statistical post-processing (Gaussian scaling), and supervised refinement (One-Class SVM) across a multi-tiered infrastructure to identify suspicious financial activities with high accuracy. The edge layer enables localized, context-aware grading of transactions, while the fog layer performs confirmatory anomaly detection using device-level metadata, effectively reducing false negatives and ensuring prompt escalation of verified threats. Experimental evaluations demonstrated the superiority of the proposed hybrid model over traditional standalone algorithms. Compared to Autoencoder (FNR = 50.44%) and SVM (FNR = 56.38%), the edge–fog model significantly reduced the false negative rate to 8.70%, with an overall accuracy of 98.82%, sensitivity of 91.30%, and specificity of 98.79%. This confirms the utility of combining peer-aware classification at the edge with ensemble learning and supervised filtering at the fog layer to enhance anomaly detection in dynamic, high-volume transaction environments.

Future Work

Building on these promising results, several avenues remain open for further research:

i. Federated Learning Integration: Future iterations of the framework could integrate federated learning to ensure model privacy and compliance across geo-distributed smart cities without transferring raw data to central fog nodes.

ii. Temporal and Sequential Modeling: Incorporating sequence-aware models such as LSTM or Transformer-based autoencoders can help detect temporal fraud patterns and layering activities over time.

iii. Geospatial and Behavioral Enrichment: Adding geolocation, device mobility, and behavioral profiling features can enhance contextual accuracy, especially in detecting device impersonation or fraud rings.

iv. Cross-Domain Adaptability: Testing and adapting the system in other anomaly-prone smart city domains such as energy, e-health, or transportation can validate the generalizability of the edge–fog hybrid framework.

v. Real-World Deployment: Future implementations will explore on-device deployment of lightweight autoencoders and real-time reporting in edge-optimized hardware such as Raspberry Pi or IoT microcontrollers.

In conclusion, this research contributes a novel, multi-layered, hybrid machine learning architecture that addresses critical challenges in smart city financial anomaly detection—balancing precision, responsiveness, and interpretability across a distributed computing paradigm.

References

1. Kim, Y.; Park, S.; Shahkarami, S.; Sankaran, R.; Ferrier, N.; Beckman, P. (2022). Goal-driven scheduling model in edge computing for smart city applications.
2. Yang, Y.; Ding, S.; Liu, Y.; Meng, S.; Chi, X.; Ma, R.; Yan, C. (2022). Fast wireless sensor for anomaly detection based on data stream in an edge-computing-enabled smart greenhouse. *Digital Communications and Networks*, **8**, 498–507.
3. Waheed, S.R.; Sakran, A.A.; Rahim, M.S.M.; Suaib, N.M.; Najjar, F.H.; Kadhim, K.A.; Salim, A.A.; Adnan, M.M. (2023). Design a Crime Detection System Based on Fog Computing and IoT. *Malaysian Journal of Fundamental and Applied Sciences*, **19**, 345–354.
4. Wali, G.; Bulla, C. (2024). Anomaly Detection in Fog Computing: State-of-the-Art Techniques, Applications, Challenges, and Future Directions. *Library Progress International*, **44**(3), 13967–13993.
5. Ibrar, M.; Wang, L.; Shah, N.; Rottenstreich, O.; Muntean, G.M.; Akbar, A. (2022). Reliability-Aware Flow Distribution Algorithm in SDN-enabled Fog Computing for Smart Cities.
6. Vo, T.; Dave, P.; Bajpai, G.; Kashef, R. (2022). Edge, Fog, and Cloud Computing: An Overview on Challenges and Applications. *arXiv*, arXiv:2211.01863v1.
7. Farooqi, A.M.; Alam, M.A.; Hassan, S.I.; Idrees, S.M. (2022). A Fog Computing Model for VANET to Reduce Latency and Delay Using 5G Network in Smart City Transportation. *Applied Sciences*, **12**(4), 2083. https://doi.org/10.3390/app12042083
8. Jain, S.; Gupta, S.; Sreelakshmi, K.K.; Rodrigues, J.J.P.C. (2022). Fog computing in enabling 5G-driven emerging technologies for development of sustainable smart city infrastructures. *Cluster Computing*. https://doi.org/10.1007/s10586-021-03496-w
9. John, B. (2025). Deep Learning-Based Model for Detecting Suspicious Bank Transactions Leveraging Fog Computing Infrastructure. *ResearchGate*.
10. Lai, K.-L.; Chen, J.I.Z. (2021). Development of Smart Cities with Fog Computing and Internet of Things. *Journal of Ubiquitous Computing and Communication Technologies*, **1**, Article 006. https://doi.org/10.36548/jucct.2021.1.006
11. Prakash, V.; Williams, A.; Garg, L.; Savaglio, C.; Bawa, S. (2021). Cloud and Edge Computing-Based Computer Forensics: Challenges and Open Problems. *Electronics*, **10**(11), 1229. https://doi.org/10.3390/electronics10111229
12. Alsmadi, A.M.; Aloglah, R.M.A.; Abu-darwish, N.J.S.; Al Smadi, A.; Alshabanah, M.; Alrajhi, D.; Alkhaldi, H.; Alsmadi, M.K. (2021). Fog Computing Scheduling Algorithm for Smart City. *International Journal of Electrical and Computer Engineering*, **11**(3), 2219–2228. https://doi.org/10.11591/ijece.v11i3.pp2219-2228
13. Williams, M.; Emeteveke, I.; Adeyeye, O.J.; Emehin, O. (2024). Enhancing Data Forensics through Edge Computing in IoT Environments. https://doi.org/10.55248/gengpi.5.1024.2903
14. Songhorabadi, M.; Rahimi, M.; Farid, A.M.M.; Kashani, M.H. (2022). Fog Computing Approaches in IoT-Enabled Smart Cities. *Journal of Network and Computer Applications*. https://doi.org/10.1016/j.jnca.2022.103557
15. da Silva, T.P.; Batista, T.V.; Lopes, F.; Rocha Neto, A.; Delicato, F.C.; Pires, P.F.; da Rocha, A.R. (2021). Fog Computing Platforms for Smart City Applications – A Survey. *ACM Computing Surveys*. https://doi.org/10.1145/3488585
16. Aburukba, R.; Al-Ali, A.R.; Riaz, A.H.; Al Nabulsi, A.; Khan, D.; Khan, S.; Amer, M. (2021). Fog Computing Approach for Shared Mobility in Smart Cities. *Energies*, **14**(23), 8174. https://doi.org/10.3390/en14238174
17. John, B. (2025). Intelligent Suspicious Activity Detection in Banking Transactions Utilizing Fog Computing and Deep Learning. *ResearchGate*.

18. Tariq, N.; Alsirhani, A.; Humayun, M.; Alserhani, F.; Shaheen, M. (2024). A Fog-Edge-Enabled Intrusion Detection System for Smart Grids. *Journal of Cloud Computing*, https://doi.org/10.1186/s13677-024-00609-9
19. Malik, S.; Gupta, K. (2021). Smart City: A New Phase of Sustainable Development Using Fog Computing and IoT. *IOP Conference Series: Materials Science and Engineering*, **1022**(1), 012093. https://doi.org/10.1088/1757-899X/1022/1/012093
20. Ometov, A.; Molua, O.L.; Komarov, M.; Nurmi, J. (2022). A Survey of Security in Cloud, Edge, and Fog Computing. *Sensors*, **22**(3), 927. https://doi.org/10.3390/s22030927
21. Tukur, Y.M.; Thakker, D.; Awan, I.U. (2021). Edge-Based Blockchain Enabled Anomaly Detection for Insider Attack Prevention in Internet of Things. *Transactions on Emerging Telecommunications Technologies*, **32**(6), e4158. https://doi.org/10.1002/ett.4158
22. Trigka, M.; Dritsas, E. (2025). Edge and Cloud Computing in Smart Cities. *Future Internet*, **17**(3), 118. https://doi.org/10.3390/fi17030118

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.