

Article

A Reputation-Aware Adaptive Incentive Mechanism for Federated Learning-Based Smart Transportation

Abir Raza, Elarbi Badidi *  and Omar El Harrouss 

College of Information Technology, United Arab Emirates University (UAEU), Al Ain P.O. Box 15551, United Arab Emirates; omar.elharrouss@uaeu.ac.ae (O.E.H.)

* Correspondence: ebadidi@uaeu.ac.ae

Highlights

The proposed framework integrates three key innovations: a multi-dimensional reputation evaluation system for reliable client assessment, a dynamic incentive control module that adaptively tunes reward weights in real time to ensure fairness and sustained convergence, and a penalty-based mechanism leveraging anomaly detection to isolate and discourage untrustworthy participants.

What are the main findings?

- A multi-dimensional reputation evaluation system was developed for reliable client selection and incentivization.
- Elimination of malicious clients and penalization with a two-stage mechanism.

What are the implications of the main findings?

- Legitimate clients are selected with a reputation score is greater than or equal to 0.6 and incentivized through adaptive rewards.
- Malicious clients are penalized using reputation slashing and suspended from participating in federated learning rounds.

Abstract

Federated learning (FL) has emerged as a promising paradigm for privacy-preserving distributed intelligence in modern urban transportation systems, where vehicles collaboratively train global models without sharing raw data. However, the dynamic nature of vehicular environments introduces critical challenges, including unstable participation, data heterogeneity, and the potential for malicious behavior. Conventional FL frameworks lack effective trust management and adaptive incentive mechanisms capable of maintaining fairness and reliability under these fluctuating conditions. This paper presents a reputation-aware federated learning framework that integrates multi-dimensional reputation evaluation, dynamic incentive control, and malicious client detection through an adaptive feedback mechanism. Each vehicular client is assessed based on data quality, stability, and behavioral consistency, producing a reputation score that directly influences client selection and reward allocation. The proposed feedback controller self-tunes the incentive weights in real time, ensuring equitable participation and sustained convergence performance. In parallel, a penalty module leverages statistical anomaly detection to identify, isolate, and penalize untrustworthy clients without compromising benign contributors. Extensive simulations conducted on real-world datasets demonstrate that the proposed framework achieves higher model accuracy and greater robustness against poisoning and



Academic Editor: Pierluigi Siano

Received: 19 November 2025

Revised: 7 January 2026

Accepted: 28 January 2026

Published: 4 February 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

gradient manipulation attacks compared to existing baseline methods. The results confirm the potential of our trust-regulated incentive mechanism to enable reliable federated learning in smart cities transportation systems.

Keywords: adaptive rewards; anomaly detection; client selection; dynamic adaptation; federated learning; free-rider mitigation; incentive mechanism; smart transportation; Internet of Vehicles (IoV); reputation system; security; trust management; vehicular networks

1. Introduction

Federated Learning has emerged as a transformative paradigm in machine learning. It enables multiple clients to collaboratively train a shared global model without exchanging raw data [1,2]. With its practical innovation, FL has been deployed in various domains, including smart transportation, smart healthcare, smart construction, and so on [3,4]. The training is performed by a central server that periodically uploads and downloads the current model to a subset of clients [5]. The choice of these clients is crucial because it determines a large portion of the system performance and convergence of the FL process.

Nonetheless, one of the biggest challenges in FL is selecting a subset of clients to participate in each training round. The sampling of the right client nodes is called client selection. The selection of clients in federated learning involves carefully choosing the right subset of available clients that are reliable, have sufficient resources, and can provide high-quality contributions. The aim is to strike the right balance between building an accurate model and keeping the system efficient while minimizing risks from untrustworthy or malicious participants. A well-thought-out client selection process can speed up model training, cut down communication costs, and safeguard the integrity of the global model.

Limiting the number of participating clients helps curtail the exposure of sensitive data to unauthorized parties and the central server as well. Moreover, an efficient selection of clients can minimize inequalities between clients with varied data distributions (heterogeneity of data). In vehicular federated learning, data heterogeneity can be defined as the disparity or inconsistency in the data collected by the participating vehicles. This variation in data occurs as the vehicles operate in discrete physical, behavioral, and environmental settings in a vehicular network.

Adopting a suitable strategy for selecting logical clients can also help reduce communication overhead and computational costs. Furthermore, the selection process mitigates the risk of having compromised, malicious, or corrupted clients within the federated learning process, which would otherwise have affected the completeness and accuracy of the learning model. Conclusively, a well-designed client selection strategy contributes to faster convergence of the global model [6].

Determining the optimal number of clients in federated learning for trust management is a nuanced problem that depends on several factors, including the nature of the application, the diversity and quality of the data, the computational resources available, and the specific trust management strategies employed. The challenge of trust-based client selection becomes particularly important in scenarios where malicious clients aim to compromise the FL system. Such attacks can lead to the poisoning of the global model, reducing its generalization capability and making it vulnerable to adversarial threats. Therefore, enhancing the client selection process with trust evaluation mechanisms such as reputation systems, anomaly detection, or game-theoretic approaches can mitigate these risks and improve the reliability of federated learning.

The lack of intrinsic motivation is evident among participants who incur resource costs (such as computational power and bandwidth) and have concerns related to fairness and data quality. The absence of appropriate incentives in FL systems can cause three issues that damage system performance: insufficient participation combined with free-riders and biased model convergence.

FL incentive mechanisms establish reward structures that stimulate quality input while ensuring efficiency, fairness, and long-term sustainability. Traditional incentive mechanisms often rely on centralized control, and this raises concerns related to fairness and manipulation. Ensuring sustained and high-quality participation is one of the key challenges in federated learning, as contributors incur computational, communication, and energy costs without guaranteed benefits. Research concerning FL continues to face essential gaps in its ability to maintain fairness against efficiency, managing dynamic and heterogeneous participants in decentralized environments. Recent incentive approaches integrate game-theoretic models that use both auction-based and Shapley value techniques along with reputation systems, blockchain-based rewards, and contract theory.

Designing an effective FL incentive mechanism presents two key challenges. The first one is from the participants' perspective. The primary difficulty lies in motivating and retaining a robust client base. This requires ensuring that participants receive competitive rewards while simultaneously safeguarding their data privacy, which is crucial in encouraging sustained engagement. The second major issue lies with the FL platform to accurately evaluate participant contributions. Given the diverse requirements of different FL tasks, the challenge is twofold: (1) maximizing the federation's long-term sustainability and (2) optimizing incentive costs without compromising participation quality. Balancing these competing demands necessitates a carefully calibrated incentive structure.

1.1. Our Contribution

An adaptive incentive mechanism for federated learning that dynamically aligns participant rewards with their genuine contributions is proposed. The framework autonomously balances model accuracy, participant fairness, and system security through reputation-based tokenomics and parameter self-tuning.

The main contributions of this article can be summarized as follows:

1. We design a novel, trust-aware client selection process that leverages real-time multi-dimensional reputation scoring, enabling the system to prioritize reliable clients and effectively filter out malicious vehicles.
2. We propose a dynamic incentive allocation scheme with self-tuning parameters $(\alpha, \beta, \gamma, \lambda)$ that autonomously and adaptively balance data quality, participation stability, and consistency. These parameters are updated in real time via feedback control, ensuring incentives are fair, performance-driven, and robust against network fluctuations.
3. We devise an integrated penalty module that not only identifies and isolates malicious clients using statistical anomaly detection but also administers reputation slashing and temporary suspensions to discourage harmful participation without unfairly punishing trusted nodes.
4. We conduct extensive experiments on real-world datasets, demonstrating the effectiveness of our framework in retaining an overall model performance of above 95% with 93% collaborative fairness. Our method offers a 10% increase in the model accuracy as compared to the traditional Vanilla FL.

1.2. Organization of the Paper

The remainder of this paper is organized as follows: A comprehensive review of related work, encompassing existing incentive models, datasets, performance metrics, and identified limitations, is provided in Section 2. The proposed trust-aware methodology and system architecture are detailed in Section 3. Section 4 presents the experimental setup, evaluation results, and analytical discussion. Finally, Section 5 concludes the paper and outlines potential directions for future research.

2. Related Work

Researchers have concentrated their focus on creating incentive mechanisms for FL to encourage user engagement in the training process. The survey by Zeng et al. [7] identifies and categorizes existing incentive mechanisms in federated learning, covering approaches like Stackelberg games, auctions, contract theory, Shapley value, reinforcement learning, and blockchain-based methods. It offers a taxonomy and comparison of these techniques, discusses incentive problems particular to federated learning, and suggests future research directions. Zhan et al. investigate different incentive mechanism designs in an attempt to evaluate the training data value of each client [8]. The authors discuss the challenges and opportunities in incentive mechanisms for federated learning and propose a general framework with potential solutions to address design challenges [9].

2.1. Auction-Based Incentive Mechanisms

Thi Le et al. introduced an incentive mechanism for federated learning in wireless cellular networks using an auction approach, where the base station is the auctioneer and the mobile users are the sellers [10]. Zeng et al. propose a multi-dimensional auction approach to incentivize FL clients. Here, all the participants submit their bids with promised capacities to the aggregation server [11].

Reverse auction theory has proven effective in enhancing FL economically and structurally, as demonstrated in several studies including [12]. A federated learning incentive mechanism based on reputation and reverse auction theory, where participants bid for tasks and their reputation reflects their reliability and data quality, and the mechanism selects and rewards participants by combining their reputation and bids under a limited budget [13]. Ahmed and Choi proposed an FL incentive scheme based on reverse auction and trust reputation to reward selected clients with limited budgets [14]. Zhang et al. also proposed a reverse auction model that leverages generative AI to enhance the learning performance of FL [15].

Zheng et al. combined local differential privacy and auction mechanism [16]. The privacy budget is used as an auction item and clients with a high privacy budget and low cost are selected to participate in FL training. However, the above works treat the privacy level as the only incentive factor, ignoring the influence of other factors such as data size, model quality, and non-IID degree in actual FL.

2.2. Contract Theory-Based Incentive Mechanisms

The tutorial by Zhang et al. demonstrates the effectiveness of contract theory in designing incentive mechanisms [17]. The study focuses on utilizing contracts in designing effective incentive mechanisms for miscellaneous wireless applications, such as mobile crowdsourcing and spectrum trading. A contract theory-based incentive mechanism is proposed to motivate mobile devices with high-quality data to participate in the training process [18]. This method utilizes the Subjective Logic Method (SLM) to calculate workers' reputations, using the Earth Mover's Distance (EMD) as a metric for assessing the quality of local data. The EMD is also referred to as Wasserstein-1 metric distance. It is a measure

of the distance between two probability distributions. It basically quantifies the minimum work required to transform one distribution into the other. In the vehicular FL context, EMD assesses the strategic quality and diversity of a vehicle's local contribution by determining how different its data is from the global average [19]. Sun et al. proposed a contract-based personalized privacy-preserving incentive mechanism [20]. The contract specifies the level of privacy protection and the corresponding reward, and clients can choose the appropriate contract according to their privacy preferences. Another contract theory-based incentive mechanism was developed by Tian et al. [21]. It was a two-dimensional contract model that comprises data quality and computation effort. Experimental results show that the generalization accuracy of the FL tasks can be improved by the proposed incentive mechanism where contract-based aggregation is applied. The paper proposes a contract-theory-based incentive mechanism to motivate data holders with high-quality data and computing power to participate in federated learning for health crowdsensing, which can make the baseline algorithm converge faster and resist malicious behaviors [22].

2.3. Deep Reinforcement Learning-Based Incentive Mechanisms

A deep reinforcement learning-based incentive mechanism based on the Proximal Policy Optimization (PPO) is designed to motivate edge nodes to contribute to model training [23]. The mechanism uses an actor–critic model to determine the optimal pricing strategy for the parameter server and the optimal training strategies for edge nodes. Xu et al. formulate a multi-agent DRL incentive mechanism for secure federated learning [24]. Wu et al. formulate vehicle selection as a DRL problem [25]. In each FL round, vehicles are selected on a unified score consisting of communication delay, training delay, and data size. High-performing vehicles with good data and low delay are incentivized implicitly using the DRL policy. W. Lin and Y. Zhou designed an explainable incentive mechanism by combining Directed Acyclic Graph (DAG) and DRL [26]. The framework uses a GNN-based trust assessment for reliable node selection. The vehicles with higher trust are selected and rewarded, while malicious nodes are filtered out.

2.4. Game Theory-Based Incentive Mechanisms

The paper discusses various game-theoretical models to study the economic behaviors of participants in federated learning [27]. Another research proposes a game-theoretic framework for designing incentive mechanisms in federated learning to encourage data providers to truthfully report their private information and contribute all their data to the federation [28]. The paper proposes an incentive mechanism for on-device federated learning that selects users most likely to provide reliable data and compensates them for privacy leakage, formulated as a two-stage Stackelberg game [29]. Khan et al. model the incentive-based interaction between a global server and participating edge devices using a Stackelberg game [30]. In this game, FL users can strategically set the number of local iterations to maximize their utility. A two-stage Stackelberg game approach is used to optimize utility for both the server and workers by designing an incentive mechanism for federated learning [31]. The paper proposes an incentive mechanism for federated learning based on random client sampling and modeling the interaction between the server and clients as a Stackelberg game [32]. The paper proposes a game-theoretic incentive mechanism for federated learning with randomized client participation to achieve an unbiased and high-performance model [33].

2.5. Miscellaneous Incentive Mechanisms

Zhan et al. proposed an incentive mechanism based on the bargaining approach for mobile crowd sensing [34]. The incentive mechanism by Yu et al. dynamically divides a given budget in a context-aware manner among data owners by maximizing the collective

utility [35]. The study adopts a quality-conscious approach to FL, incorporating quality estimation, a quality-based incentive mechanism, and model synthesis [36]. It measures the contribution of clients based on their data quality by quantifying the data quality of clients using the loss reduction and then creates a reverse auction in conjunction with historical quality records, thus attracting more high-quality clients. Gao et al. worked out a fair incentive mechanism based on a dynamic real-time worker assessment mechanism [37]. The method measures the trustworthiness of clients by using the Subjective Logic Model (SLM). Zhang et al. combined participant reputation and bidding within a limited budget, using reputation to indirectly gauge participant reliability and data quality [13]. Yang et al. propose an incentive mechanism based on Portfolio theory [38]. The authors designed a resource investment incentive mechanism based on a weighted potential game [39]. Four algorithms, namely the GBD-based algorithm, the weighted potential game best-response algorithm, the distributed stochastic learning algorithm, and the policy gradient-based algorithm, are used to solve resource input strategies considering resource dynamics.

Table 1 provides an overview of the literature on federated learning incentive mechanisms.

2.6. Comparative Analysis and Design Insights

Although the preceding discussion has provided a detailed taxonomy of incentive mechanisms in federated learning, it is equally important to critically examine their suitability for highly dynamic vehicular environments. Table 2 summarizes the relative strengths and limitations of the major classes of approaches. Moreover, an analytical comparison of these incentive categories is carried out in Table 3.

Auction- and contract-theoretic schemes offer efficient and mathematically tractable reward allocation models, yet they generally assume static client sets and stable connectivity. Such assumptions do not hold in Internet-of-Vehicles scenarios characterized by high mobility and intermittent participation. Blockchain-based mechanisms introduce decentralized trust and transparent auditing; however, they impose considerable communication and computational overheads that are incompatible with the latency and resource constraints of edge-assisted vehicular networks. Game-theoretic and reinforcement-learning-based designs provide strategic adaptability, but they require long convergence times and often neglect trust dynamics and fairness across heterogeneous nodes. Reputation-driven approaches can successfully filter unreliable clients, yet most existing models rely on one-dimensional scoring and lack the ability to quantify multi-faceted contribution quality in terms of data reliability, stability, and behavioral consistency.

To the best of our knowledge, the existing body of work demonstrates valuable progress but falls short of meeting the combined requirements of adaptivity, trust awareness, fairness, and scalability demanded by next-generation Vehicular Federated Learning (VFL).

2.7. Mapping of Research Gaps to Proposed Modules

The limitations of the past literature directly motivate the design of the proposed Reputation-aware Dynamic Adaptive Incentive Mechanism presented in Section 3, which integrates multi-dimensional reputation scoring, self-tuning incentives, and an adaptive penalty framework to ensure robust and equitable collaboration among heterogeneous, mobile participants. Table 4 establishes a direct correspondence between the observed research gaps in the literature and the design modules of our proposed reputation-aware dynamic adaptive incentive mechanism. This mapping serves as a conceptual bridge to Section 3, clarifying how each methodological component addresses a specific limitation in existing work.

Table 1. Summary of literature review on incentive mechanisms in federated learning.

Study	Year	Method	Dataset	Performance Metric(s)	Limitation(s)
Kang et al. [18]	2019	Multiweight SLM, EMD, Contract theory	MNIST	EMD, attack strength, local data accuracy, task publisher profit	Lack of empirical implementation and evaluation of the proposed reputation-based worker selection and blockchain reputation management framework.
Le et al. [10]	2021	Primal–dual greedy auction	Mobile local data	Social welfare	The auction mechanism introduces significant communication overhead and latency during participant coordination.
Zhang et al. [13]	2021	Reputation and reverse auction theory	MNIST, Fashion-MNIST, IMDB	Model accuracy, economic efficiency, participant utility, social welfare, convergence rate	Computational complexity increases rapidly with the number of participating clients, limiting scalability.
Ahmed et al. [14]	2023	Reverse auction and trust reputation	MNIST, Fashion-MNIST, CIFAR10	Model convergence accuracy, economic benefit (%)	Trust mechanism remains vulnerable to misleading or adversarial gradients.
Li et al. [22]	2022	Contract theory	MNIST, UCI-HAR	Model accuracy, incentive cost, participant utility, convergence rate	Assumes fixed client attributes, reducing applicability in dynamic environments.
Zhan et al. [23]	2020	DRL with PPO, Stackelberg game	Edge node local data	Training cost, server utility	The DRL-driven incentive mechanism may reduce edge node utility under certain operating conditions.
Hu et al. [29]	2020	Two-stage Stackelberg game	Synthetic data	Server utility, user utility, model accuracy, convergence rate	Scalability is limited due to increased communication and computation overhead.
Khan et al. [30]	2020	Stackelberg game, CoCoA	MNIST	Communication cost, local relative accuracy	Primarily focuses on resource optimization and does not fully address heterogeneous communication constraints.
Xiao et al. [31]	2020	Two-stage Stackelberg game	MNIST	Server utility, client utility, equilibrium analysis	Limited empirical validation beyond small-scale experiments.
Wu et al. [32]	2020	Stackelberg game	MNIST	Server utility, client utility, model accuracy, convergence rate	Assumes fixed client data quality and cost factors, limiting adaptability to dynamic client behavior.
Luo et al. [33]	2020	Game theory	Synthetic, MNIST, EMNIST	Model performance, client profit, convergence rate	Relies on static client characteristics, reducing applicability in real-world FL deployments.
Yu et al. [35]	2020	Polynomial-time algorithm	Real-world data	Payoff received, waiting time for payoffs	Possible mismatch between user contributions and rewards due to FL system constraints.
Gao et al. [37]	2021	Reputation based on SLM	Randomly generated data	System revenue (%)	Susceptible to free-rider behavior and performs poorly under highly non-IID data settings.

Table 2. Strengths and weaknesses of federated learning incentive mechanisms.

Mechanism	Strengths	Weaknesses
Auction Theory	- Enables cost-efficient task and resource allocation through competitive bidding. - Provides flexible pricing and can adapt to dynamic client participation. - Encourages truthful cost reporting under certain auction designs.	- Communication overhead increases with the number of bidders. - Vulnerable to collusion or strategic manipulation among participants. - Winner determination and payment rules become complex at scale.
Blockchain	- Ensures decentralized, transparent, and tamper-proof reward management. - Smart contracts automate incentive distribution and verification. - Provides auditability independent of the FL server.	- High computational and energy cost limits use on resource-constrained devices. - Scalability is restricted by block generation rate and transaction throughput. - Consensus protocols introduce latency and communication burden.
Contract Theory	- Addresses asymmetric information between server and clients. - Encourages truthful reporting and aligns incentives with client capabilities. - Flexible in modeling heterogeneous client attributes.	- Requires accurate knowledge or estimation of client types. - Designing scalable contracts becomes challenging as client count grows. - Frequent contract updates may be needed to track dynamic environments.
DRL-Based	- Automatically learns long-term optimal incentive policies. - Adapts to dynamic and uncertain FL environments. - Can model complex and multi-dimensional reward structures.	- Requires large training datasets, increasing computation cost. - High training complexity makes real-time deployment difficult. - Produces black-box policies with limited interpretability.
Game Theory	- Models strategic interactions among rational FL participants. - Provides equilibrium-based stability for incentive outcomes. - Supports both cooperative and competitive FL settings.	- Assumes rational clients, which may not reflect real-world behavior. - Scalability decreases as the number of players increases. - Computing Nash or Stackelberg equilibria can be expensive.
Reputation Mechanisms	- Encourages honest participation and discourages free-riding. - Allows filtering of unreliable, low-quality, or inconsistent clients. - Simple to integrate with existing FL pipelines.	- May be susceptible to manipulation if not securely designed. - Requires reliable and timely updates to maintain accuracy. - New clients face cold-start difficulty due to initial low reputation.
Shapley Value	- Provides rigorous and fair contribution evaluation based on marginal utility. - Ensures cooperative fairness grounded in game-theoretic principles. - Useful for data-dependent incentive design.	- Exact computation has exponential complexity ($O(2^n)$). - Assumes equal accessibility to contribution evaluation, limiting practicality. - Sensitive to non-IID client data distributions.

Table 3. Analytical comparison of existing incentive mechanisms in vehicular federated learning.

Approach Category	Typical Strengths	Limitations in Dynamic IoV Context
Auction/Contract Theory	- Enables cost-efficient resource allocation. - Incentivizes truthful reporting via formal economic models. - Flexible reward structures for heterogeneous clients.	- Assumes stable client availability and accurate type information. - High communication overhead during bidding/contract negotiation. - Limited capacity to track rapidly changing mobility, channel quality, and handoff frequency.
Blockchain-based	- Provides tamper-proof auditability of incentives. - Eliminates reliance on a central authority. - Smart contracts support automated, transparent reward settlement.	- High computation and energy cost unsuitable for vehicular edge nodes. - Consensus delays degrade timeliness in fast-moving IoV environments. - Scalability issues arise under dense RSU deployments or frequent connectivity disruptions.
Game Theory	- Captures strategic and competitive behavior among vehicles. - Supports equilibrium-based incentive guarantees. - Applicable to both cooperative and competitive participation.	- Convergence to equilibrium is slow in large, dynamic fleets. - Assumes rational, fully informed agents—often unrealistic in IoV. - Sensitive to intermittent connectivity and fluctuating participation.
Deep RL-based	- Learns adaptive incentive policies directly from interactions. - Handles high-dimensional and non-linear reward structures. - Adapts to evolving client behavior over time.	- Requires extensive training data and computational resources. - Prone to instability when vehicle availability and channel conditions change rapidly. - Difficult to interpret and validate in safety-critical IoV settings.
Reputation-based	- Lightweight and easy to integrate into FL pipelines. - Enables filtering of unreliable or low-quality contributors. - Supports long-term behavioral assessment.	- Often relies on single-dimensional or heuristic scoring. - Reputation updates may lag behind fast mobility-induced behavior changes. - Susceptible to collusion attacks and cold-start bias in newly joining vehicles.

Table 4. Mapping of research gaps to design objectives.

Identified Research Gaps	Corresponding Design Objective/Solution (Section 3)
Static incentive mechanisms fail to adapt to heterogeneous, mobile IoV environments.	Introduce a dynamic adaptive incentive mechanism with self-tuning parameters $(\alpha, \beta, \gamma, \lambda)$ that adjust rewards in real time according to data quality, stability, and consistency.
Lack of robust reputation metrics for distinguishing low-quality and high-quality updates.	Employ a multi-dimensional reputation model integrating data quality (cosine similarity), stability, and contribution consistency to prioritize trustworthy clients.
Inadequate filtering and penalization of malicious or free-riding participants.	Incorporate multi-stage anomaly detection (gradient-based Z-score, behavioral consistency) and a penalty module with reputation slashing and suspension.
Absence of fairness-preserving reward allocation under dynamic participation.	Develop a trust-aware incentive disbursement policy proportional to reputation scores and energy levels, ensuring fairness and sustainability.
Limited empirical validation on dynamic, heterogeneous networks.	Conduct extensive experiments on benchmark datasets (MNIST, Fashion-MNIST) to evaluate performance, fairness, and robustness across varying network conditions.

3. Proposed Framework

In this paper, we consider a universal cross-device federated learning context containing multiple smart vehicles (SV) with roadside infrastructure. Our proposed pipeline comprises four main elements: client selection using a multi-dimensional reputation score, identification and filtering of malicious clients, penalising nodes with malicious behavior, and rewarding the legitimate contributing nodes. Figure 1 illustrates a simplified representation of our proposed system's architecture.




Figure 1. System design of the dynamic adaptive incentive mechanism. The Red crosses imply the rejected clients (with a reputation less than 0.6), while the subset of selected clients are encircled in Green (with a reputation value of 0.6 or greater).

3.1. Clients Selection with Multi-Dimensional Reputation

As multiple distributed devices, e.g., IoT devices, smartphones, or edge nodes, collaborate to train a shared global model without sharing raw data in an FL environment, selecting the right clients is crucial. In FL, client selection is a process of choosing a subset of participating devices (clients) to contribute towards model training in each round to achieve efficiency, fairness, and fast convergence.

Our proposed adaptive incentive mechanism (AIM) is based on the principles of Zero Trust Architecture (ZTA), which requires that every network node, user, or application be evaluated for trust and continuously monitored [40]. In our system, each FL client is carefully assessed for trustworthiness and monitored for erratic behavior, ensuring that only reliable participants are selected for training.

In the proposed pipeline, the SVs are evaluated using a multi-dimensional reputation criterion that is influenced by data quality, stability, consistency, their respective weights, and reputation decay. The three weight factors, α , β , and γ , are self-tuning parameters that manage the significance of data quality, stability, and consistency in the calculation of the overall reputation score. These weights are updated automatically in real time using

feedback control. Each SV v_i has a reputation $R_i(t) \in [0, 1]$ updated at time t , and the vehicles are prioritized using Equation (1).

$$R_i(t) = \underbrace{\lambda(t)}_{\text{Forgetting Factor}} R_i(t-1) + (1 - \lambda(t)) \left(\underbrace{\alpha(t)Q_i(t)}_{\text{Data Quality}} + \underbrace{\beta(t)S_i(t)}_{\text{Stability}} + \underbrace{\gamma(t)C_i(t)}_{\text{Consistency}} \right) \quad (1)$$

where:

$Q_i(t)$	: Cosine similarity of gradients
$S_i(t)$	: Uptime ratio (sliding window)
$C_i(t)$	: $1 - \text{std}(Q_i^{t-k:t})$ (Contribution consistency)
$\alpha(t)$	: Data quality weight (tuned by accuracy trends)
$\beta(t)$	: Stability weight (adapted to participation rates)
$\gamma(t)$	: Consistency weight
$\lambda(t)$	: Forgetting factor (updated via network churn)
$\alpha(t) + \beta(t) + \gamma(t)$	= 1 (Normalization)

It greedily selects the clients with the highest reputation values. Greedy-based approaches are widely used in recent research for implementing quality-aware budgeted incentive mechanisms [41,42].

3.1.1. Data Quality (Q_i)

Traditional incentive mechanisms quantify clients' contributions based on data size with the assumption that larger datasets enhance model performance. However, evaluating contributions based solely on data size is not adequate. To overcome this limitation, data quality has been incorporated as a metric to evaluate clients' contributions [16,20,36,37,43–45]. Adaptive incentive mechanisms use data quality metrics, e.g., consistency checks or outlier detection to filter out untrustworthy participants as well [46].

In the context of FL, the data distributions submitted by clients are the key factor for contribution evaluation. It is possible that the clients may submit fake data distributions to increase their contribution. To address this problem, we design a free-rider detection mechanism leveraging cosine similarity to reward clients with aligned data (high similarity). In VFL, vehicles can behave erratic sometimes due to mobility or connection loss and send poisoned model updates. The Similarity _{i} metric gauges a vehicle's utility in the FL process.

Data quality Q_i defines the usefulness and trustworthiness of a vehicle i through Equation (2). The model alignment value is calculated using Cosine Similarity between a vehicle's gradients (∇W_i) and the global gradients (∇W_{global}) using Equation (3). Meanwhile, noise or miscellaneous attacks are anticipated using Gradient Variance with Equation (4).

$$Q_i = \underbrace{\text{Model Alignment}}_{\text{Cosine Similarity}} - \underbrace{\text{Noise/Attack Risk}}_{\text{Penalty Terms}} \quad (2)$$

$$\text{Model Alignment} = \max \left(0, \frac{\langle \nabla W_i, \nabla W_{\text{global}} \rangle}{\|\nabla W_i\| \cdot \|\nabla W_{\text{global}}\|} \right) \quad (3)$$

$$\text{Noise Penalty} = \lambda \cdot \text{Var}(\nabla W_i) \quad (4)$$

Finally data quality (Q_i) of each vehicle i is calculated using Equation (5).

$$Q_i^{(t)} = \max \left(0, \frac{\langle \nabla W_i^{(t)}, \nabla W_{\text{global}}^{(t-1)} \rangle}{\| \nabla W_i^{(t)} \| \cdot \| \nabla W_{\text{global}}^{(t-1)} \|} \right) - \underbrace{\lambda_{\text{noise}} \cdot \text{Var}(\nabla W_i^{(t)})}_{\text{Noise Penalty}} \quad (5)$$

where:

$\nabla W_i^{(t)}$: Local model gradients from vehicle i

$\nabla W_{\text{global}}^{(t-1)}$: Previous global model gradients

λ_{noise} : Noise sensitivity parameter (default: 0.2)

$\text{Var}(\cdot)$: Gradient variance over last k updates

Data quality is bounded in $[0, 1]$: $Q_i^{\text{cos}} \in [0, 1]$. High-quality data leads to improved model updates and better overall performance, while low-quality data can degrade the model.

3.1.2. Data Quality Weight (α)

This scales the influence of gradient similarity on reputation and directly controls model convergence speed. $\alpha(t) \in [0.5, 0.9]$. High $\alpha \rightarrow$ Prioritizes model accuracy and low $\alpha \rightarrow$ Reduces sensitivity to noisy data.

3.1.3. Stability $S_i(t)$

Frequent fluctuations in mobility patterns and vehicular network density lead to communication interruptions. This carries massive communication overheads and compromises the efficiency and reliability of vehicular systems [47–49]. The stability term is incorporated into the framework to address the key issues of high mobility and varying density of vehicles of the IoV network. It primarily ensures that our incentive mechanism steers towards vehicles that are persistent and offer reliable contributions.

Stability refers to a vehicle's reliability in participation and performance over time in the FL process. It quantifies how predictably a client contributes to the training process without dropouts or sudden performance shifts. It can be considered as the predicted uptime from the mobility patterns—how long a vehicle v_i stays connected. If a client's stability $S_i(t)$ drops (e.g., starts missing rounds), its reputation decays via the forgetting factor $\lambda(t)$.

Stability in reputation calculation ensures that the FL system reliably incentivizes desired behavior, maintains long-term participant engagement, and resists manipulation using Equation (6). If there is no stability, a client's reputation might change suddenly because of reputation oscillations [50]. Using the stability factor prevents clients from unfair punishments for short-term issues, encouraging sustained participation.

A smart vehicle with high stability ($S_i(t) \approx 1$) receives a higher reputation and rewards, so it joins every round, uploads data and maintains steady battery levels. On the contrary, a car with low stability ($S_i(t) \approx 0$) frequently drops out due to poor connectivity or erratic behaviour and is penalised via temporary suspension or reputation decay ($\lambda(t) \cdot R_i(t-1)$) dominates.

$$S_i(t) = \underbrace{\frac{1}{k} \sum_{j=t-k}^{t-1} \mathbb{I}_{\text{participated}}^{(j)}}_{\text{Uptime Ratio}} \times \underbrace{\left(1 - \frac{\text{Handoff Count}}{\text{Max Allowable}} \right)}_{\text{Mobility Penalty}} \quad (6)$$

where:

k : Sliding window size (e.g., 5 rounds)

$\mathbb{I}_{\text{participated}}$: Indicator function (1 if participated, 0 otherwise)

Handoff Count: RSU transitions in window k

Max Allowable: Network-specific threshold (e.g., 3 handoffs)

3.1.4. Stability Weight (β)

This determines the importance of uptime/participation and maintains sufficient participant diversity. $\beta(t) \in [0.1, 0.5]$. High $\beta \rightarrow$ Favors vehicles with reliable connectivity and low $\beta \rightarrow$ Tolerates intermittent participation.

3.1.5. Consistency (C_i)

Unlike stability $S_i(t)$, which measures participation reliability, consistency focuses on the semantic coherence of a client's contributions to the model's training progress. It refers to a vehicle's ability to produce updates that align reliably with the global model's learning objectives over time.

In a dynamic adaptive incentive mechanism, Consistency ensures that clients are evaluated fairly and predictably over time. This prevents manipulation, encourages long-term participation, and improves model performance. Without consistency checks, the system may be victimized by free-rider attacks; clients may submit high-quality updates to gain rewards, then disappear. If clients join or leave so frequently or randomly, the aggregation becomes unstable.

The server calculates the consistency using Equations (7) and (8). High $C_i(t)$ increases $R_i(t)$, while low values trigger penalties.

$$C_i^{(t)} = \text{cosine\_similarity}(\nabla_{\text{local}}^{(t)}, \nabla_{\text{global}}^{(t)}) \quad (7)$$

OR

$$C_i^{(t)} = \frac{\nabla_{\text{local}}^{(t)} \cdot \nabla_{\text{global}}^{(t)}}{\|\nabla_{\text{local}}^{(t)}\| \|\nabla_{\text{global}}^{(t)}\|} \quad (8)$$

3.1.6. Consistency Weight (γ)

Consistency Weight scales the importance of behavioral consistency in reputation scores. It basically controls how much the incentive mechanism prioritizes immediate rewards vs. long-term contributions. A high gamma ($\gamma \rightarrow 1$) means the system focuses on future contributions mostly (e.g., rewarding long-term participation), whereas a low gamma ($\gamma \rightarrow 0$) indicates that the IM values short-term performance (e.g., immediate data quality). If $C_i(t)$ drops (e.g., due to adversarial behavior), the term $\gamma(t)C_i(t)$ reduces the reputation.

3.1.7. Forgetting Factor (λ)

The reputation score must decay over time if a smart vehicle stops contributing or provides substandard data. The reputation decay here determines how fast a client's historic reputation declines, so it does not ignore their earlier actions while paying attention to their recent behaviors. The forgetting factor reduces the weight of past reputation over time to adapt to the changes in VANETs using Equation (9).

It balances stability vs. responsiveness in reputation tracking and dynamically adjusts incentives based on participation rates. $\lambda(t) \in [0.7, 0.95]$ (time-varying). High $\lambda \rightarrow$ Slow forgetting (prioritizes long-term reliability), and low $\lambda \rightarrow$ Fast forgetting (responsive to recent behavior).

$$\lambda(t) = \lambda_{\min} + (\lambda_{\max} - \lambda_{\min}) \cdot e^{-\tau \cdot \text{ChurnRate}(t)} \quad (9)$$

where:

$\lambda_{\min}$: Lower bound (e.g., 0.7)
 $\lambda_{\max}$: Upper bound (e.g., 0.95)
 τ : Decay rate (e.g., 2.0)
 $ChurnRate(t)$: $\frac{\text{Disconnected vehicles}}{\text{Total vehicles}} \in [0, 1]$

Rationale for Trust Metric Selection: The choice of cosine similarity and gradient variance as trust estimators was motivated by three key factors: (1) computational efficiency for edge devices, (2) interpretability in capturing behavioral alignment, and (3) empirical robustness under vehicular mobility and dynamic participation.

Cosine similarity has been widely adopted in trust-based federated learning to measure the directional coherence between local and global gradients [51]. It provides a lightweight and parameter-free measure of alignment that reflects whether a client's update is consistent with the global optimization trajectory, without requiring distributional reconstruction or access to raw data. This property is critical in IoV settings, where communication latency, energy constraints, and privacy preservation dominate system design. Alternative metrics such as the Wasserstein distance or KL divergence, while theoretically expressive, entail high computational complexity and sensitivity to stochastic mini-batch sampling, making them less suitable for real-time vehicular learning scenarios.

Gradient variance, on the other hand, serves as a direct statistical indicator of update reliability and noise resilience. Following the insights from Byzantine-resilient aggregation methods [52] and trust bootstrapping techniques [43], variance quantification enables the detection of inconsistent or adversarial updates that deviate significantly from the mean gradient trend.

The combined use of alignment (cosine similarity) and dispersion (gradient variance) therefore yields a dual perspective of both the directional trustworthiness and statistical reliability of participating clients. This integration ensures that the reputation assessment balances sensitivity to malicious gradients with robustness to benign variability, while maintaining low computational overhead and compatibility with large-scale VFL deployments [23].

3.2. Multi-Stage Malicious Clients Detection and Filtering

Our dynamic AIM filters out malicious clients in federated learning-based IoV through a multi-stage process designed to exclude adversarial nodes and optimize FL performance.

3.2.1. Stage 1: Gradient Anomaly Detection

Gradient anomaly detection helps detect gradient inversion and label-flipping attacks. Anomaly detection using Z-score on cosine similarity is a statistical method to identify malicious or malfunctioning devices by analyzing how much their gradient updates deviate from the norm.

Using this technique, abnormal client gradient updates are detected by comparing each gradient to the z-score of recent historical values. For instance, clients with $|Z_i| > \tau$ (e.g., $\tau = 2.5$) are flagged as outliers. Low Z_i indicates adversarial updates (e.g., poisoning attacks), while high Z_i may suggest rare cases such as over-aligned Sybil attacks.

The Z-score alignment between a client's gradient (∇_{client}) and the global gradient (∇_{global}) is measured using Equation (10).

$$Z_i = \frac{\text{Similarity}_i - \mu_{\text{similarity}}}{\sigma_{\text{similarity}}} \quad (10)$$

where:

$\mu_{\text{similarity}}$: Mean similarity across clients or $\mu_Q(t)$

$\sigma_{\text{similarity}}$: Standard deviation of similarities or $\sigma_Q(t)$

3.2.2. Stage 2: Behavioral Consistency Check

In this proposed dynamic adaptive incentive mechanism, security and reliability are measured through a behavioral consistency check that evaluates whether a client's behavior (e.g., gradient updates, participation patterns) aligns with expected behavior or not over time. The behavioral consistency check verifies that a client's contributions are aligned with the global model's learning trajectory, follow predictable patterns, and match historical performance. It is defined as the Variance of recent reputation scores to detect Sybil attacks and colluding vehicles using Equation (11). To summarize, a low $C_i(t)$ triggers investigation. For example, flag if $\text{TrustScore}_i < 0.6$.

$$\text{TrustScore}_i = 1 - \text{std}(R_i^{t-4:t}) \quad (11)$$

3.2.3. Stage 3: Adaptive Threshold Update

To ensure that our dynamic incentive mechanism remains responsive to the fluctuating IoV conditions, an adaptive threshold update is introduced. The adaptive threshold update is a dynamic rule that is defined to automatically adjust the sensitivity of the malicious data filter over time.

The decision threshold θ is dynamically updated based on the relative imbalance between false positives (FP) and false negatives (FN) in the incentive system and is expressed using Equation (12). Consequently, the incentive mechanism is continuously calibrated based on recent detection performance. This makes our incentive system more effective at identifying malicious clients adaptively and robust to changing data landscapes.

$$\theta^{\text{new}} = \theta^{\text{prev}} + \eta \cdot \left(\frac{\text{FP} - \text{FN}}{\text{Total vehicles}} \right) \quad (12)$$

where:

θ : Anomaly detection threshold

η : Learning rate (e.g., 0.01)

FP/FN: False positives/negatives

3.3. Malicious Clients Penalizing

In our proposed method, malicious clients are penalized by employing two punitive actions:

1. Reputation Slashing: The reputation score of malicious clients is reduced using Equation (13).
2. Temporary Suspension: Clients with erratic behavior are prevented from participating in some specific number of communication rounds using Equation (14).

$$R_i^{\text{new}} = R_i^{\text{prev}} \cdot e^{-\kappa \cdot \text{severity}} \quad (13)$$

where:

κ : Slashing coefficient (default: 1.5)

severity: Normalized anomaly score $\in [0, 1]$

$$T_{\text{suspend}} = \min(T_{\text{max}}, \tau \cdot \lfloor 2^{n-1} \rfloor) \quad (14)$$

where:

n : Number of prior violations

τ : Base suspension (e.g., two FL rounds)

T_{max} : Cap (e.g., 20 FL rounds)

3.4. Incentivizing Legitimate Clients

The disbursement of rewards is based on the differentiated reward principle, i.e., the clients with higher quality and reputation should receive better rewards, suggesting higher utility. Here, rewards are calculated dynamically using the current reputation score, remaining battery percentage, and disbursement rules. A higher reputation acquires bigger rewards. More battery % suggests the devices with sufficient battery levels are rewarded. And there exist transparent policies for distributing rewards.

3.4.1. Incentivization

The server collects local gradients from the selected clients to train the global model. The dynamic Incentivizing module calculates and awards incentives to the participants based on their reputation and contribution. The dynamic incentive mechanism here ensures the fair distribution of rewards that adapts to the number of active clients in the round. The server pays each client i a reward R_i . Thus, the incentive mechanism is represented using Equation (15).

$$P = P_1 + P_2 + \cdots + P_k \quad (15)$$

Finally, the updated global model and the awarded incentives are sent back to the clients, commencing a new training cycle. The above steps will be repeated until the optimization problem of the global loss function converges or achieves the desired model accuracy.

3.4.2. Disbursement Rules

As a general rule, vehicles with higher ratings of Q_i receive more rewards. The “per vehicle reward” calculation rewards clients with high data quality ($Q_i(t)$). Rationally, the rewards or incentives $P_i(t)$ are directly proportional to $R_i(t)$. For a vehicle i at round t , the base reward is calculated using Equation (16):

$$P_i(t) = \underbrace{\frac{R_i(t)}{\sum_{j=1}^N R_j(t)}}_{\text{Reputation Share}} \cdot B(t) \quad (16)$$

where:

$P_i(t)$: Payment allocated to client i at time t

$R_i(t)$: Reputation score of client i at time t

$\sum_{j=1}^N R_j(t)$: Sum of reputation scores of all N participating clients at time t

$B(t)$: The total budget (in tokens or monetary units) available for distribution in each FL round is expressed using Equation (17).

$$B(t) = B_0 \cdot \left(0.5 + \frac{N_{\text{active}}(t)}{2N_{\text{average}}} \right) \quad (17)$$

where:

$B(t)$: Total budget at time t

B_0 : Base budget (e.g., \$1000). Scales with current participation to maintain motivation.

$N_{\text{active}}(t)$: The actual number of vehicles meeting the participation criteria (active) in round t is calculated using Equation (18)

N_{average} : The average number of long-term active vehicles in the network over a period is calculated using Equation (19)

$$N_{\text{active}}(t) = \sum_{i=1}^N \mathbb{I}_{\{R_i(t) \geq R_{\min}\}} \quad (18)$$

where I is an indicator function that excludes low-reputation or malicious nodes from the active set. It returns 1 if the reputation $R_i(t)$ of node i at time t is greater than or equal to the minimum reputation threshold $R_{\min}$, and 0 otherwise.

$$N_{\text{average}} = \frac{1}{T} \sum_{t=1}^T N_{\text{active}}(t) \quad (19)$$

where T is the total FL rounds observed (e.g., $T = 100$). This value is updated periodically (e.g., every 50 rounds).

Legitimate clients with a normal (Z_i) score receive rewards, and these are aligned with contribution and constraints. The system automatically adjusts rewards based on real-time performance and is very useful in smart vehicles where battery and data stability are critical.

3.4.3. Dynamic Incentive Allocation

The dynamic budget scales with network density; there are higher per-vehicle rewards in sparse rural networks, suggesting a larger pool of active vehicles. For instance, a reward of \$10/vehicle in an urban area and \$50/vehicle in a rural area is given. The final incentive combines reputation-based fairness, budget availability, and energy-aware fairness. An energy term in the reward allocation mechanism favours vehicles with higher remaining battery (E_i). The incentive payout P_i for each participant i at time t is computed using Equation (20).

$$P_i(t) = \underbrace{\frac{R_i(t)}{\sum_{j=1}^N R_j(t)}}_{\text{Reputation Share}} \cdot B(t) \cdot \underbrace{\left(1 - \frac{E_{\max} - E_i}{E_{\max}}\right)}_{\text{Energy Fairness}} \quad (20)$$

where:

$R_i(t)$: Current Reputation score of participant i

$\sum_{j=1}^N R_j(t)$: Total system reputation, the sum of all participants' reputations at time t

$\frac{R_i(t)}{\sum_{j=1}^N R_j(t)}$: This fraction is the reputation share, representing client i 's proportional entitlement to the incentive pool based on its relative reputation

$B(t)$: Total budget at time t

E_i : Remaining vehicle battery (%)

$E_{\max}$: Full battery level

4. Experimental Setup and Results

In this section, experimental results are presented that demonstrate the effectiveness of our proposed dynamic adaptive incentive mechanism within a federated learning framework.

4.1. Testbed Setup and Simulation Tools

All the experiments in this adaptive incentive mechanism were implemented in Python 3.9 using the PyTorch 2.6.0 deep learning framework. And the proposed pipeline was executed in the Google Colab cloud environment using NVIDIA Tesla T4 Tensor Core GPU. The testbed emulates a vehicular federated learning scenario with 100 smart vehicles interacting over 50 synchronous communication rounds.

4.2. Experimental Setup

A federated learning process is simulated to evaluate the effectiveness of our proposed incentive framework through a pattern recognition task on two benchmark datasets. The vehicles locally train the model for 5 epochs with a batch size of 100 and a learning rate of 0.01. Table 5 lists the key parameters that provide context for our results:

Table 5. Simulation parameters.

Symbol	Description	Value
N	Number of Vehicles	100
T	Number of Federated Learning Rounds	50
α	Data Quality Weight	0.4
β	Stability Weight	0.3
γ	Consistency Weight	0.3
λ	Forgetting Factor	0.5
λ_{noise}	Noise Penalty Factor	0.1
R	Minimum Reputation Score for Selection	0.6
Z	Initial Z-score Threshold for Malicious Filtering	−2.0
I	Initial Trust Score Threshold for Malicious	0.9
κ	Penalty Slashing Coefficient	1.5
P	Penalty Threshold	0.3
η	Adaptive Threshold Learning Rate	0.01
τ	Base Suspension Rounds	1
τ_{max}	Maximum Suspension Rounds	5
B_0	Total Budget	10,000
E_{max}	Maximum Energy Level	100
E_{min}	Minimum Energy Level	10

4.3. Simulation Results

This subsection explains the internal dynamics of the proposed incentive mechanism. It reports the preliminary experimentation results that illustrate how reputation scoring, filtering logic, and incentive allocation evolved over FL rounds.

As the first stage, clients are scrutinized based on their reputation scores, and those fulfilling the minimum reputation criteria are selected for further screening. A higher reputation threshold results in a much more selective process and guarantees that clients must consistently demonstrate good conduct to qualify for training and rewards. With a higher reputation threshold, fewer clients are selected, and ultimately, the model's accuracy is consequently reduced. In our case, the minimum reputation score is set to 0.6. Based on this criterion, the clients selected in each training round, along with their reputation scores, are listed in Table 6.

It can be observed that fewer clients are chosen in the early rounds, but as the simulation progresses, more clients meet the criteria. In our proposed AIM, reputations evolve based on clients' participation, quality updates, and consistent behavior. The reputation scores of the first ten vehicles across five communication rounds are presented in Table 7. The reputation evolution data highlights the dynamic nature of our reputation mechanism and its ability to adapt to changing client behavior.

Table 6. Reputation scores per FL round.

Round	Clients Selected for Training	Reputation Score
1	21, 31, 51, 77	0.708, 0.705, 0.737, 0.705 0.725, 0.771, 0.703, 0.763, 0.723, 0.707, 0.707, 0.815, 0.749, 0.710, 0.708, 0.811, 0.738, 0.747, 0.856, 0.710, 0.731, 0.707, 0.813, 0.743, 0.767, 0.727, 0.752, 0.717, 0.729, 0.780, 0.745, 0.718, 0.712, 0.805, 0.757, 0.764 0.769, 0.818, 0.716, 0.745, 0.796, 0.768, 0.741, 0.751, 0.856, 0.786, 0.860, 0.715, 0.782, 0.789, 0.918, 0.751, 0.783, 0.741, 0.704, 0.704, 0.723, 0.868, 0.785, 0.812, 0.780, 0.728, 0.782, 0.747, 0.770, 0.823, 0.785, 0.754, 0.711, 0.746, 0.859, 0.799, 0.823 0.798, 0.832, 0.706, 0.729, 0.770, 0.816, 0.776, 0.761, 0.769, 0.871, 0.717, 0.809, 0.896, 0.723, 0.800, 0.706, 0.817, 0.948, 0.710, 0.772, 0.807, 0.756, 0.734, 0.713, 0.728, 0.890, 0.814, 0.839, 0.804, 0.747, 0.794, 0.746, 0.794, 0.843, 0.804, 0.771, 0.724, 0.764, 0.885, 0.822, 0.860, 0.719 0.817, 0.837, 0.727, 0.732, 0.780, 0.812, 0.785, 0.774, 0.776, 0.883, 0.734, 0.829, 0.915, 0.743, 0.703, 0.809, 0.701, 0.709, 0.832, 0.966, 0.706, 0.782, 0.810, 0.756, 0.757, 0.708, 0.705, 0.728, 0.900, 0.832, 0.845, 0.796, 0.759, 0.796, 0.731, 0.798, 0.861, 0.805, 0.779, 0.725, 0.777, 0.901, 0.831, 0.889, 0.735
2	2, 3, 10, 12, 14, 16, 17, 21, 23, 29, 30, 31, 44, 49, 51, 56, 58, 59, 77, 81, 83, 84, 87, 88, 90, 92, 93, 94, 96, 97, 98, 99	
3	2, 3, 8, 10, 12, 14, 16, 17, 21, 23, 31, 34, 44, 49, 51, 56, 58, 59, 61, 70, 75, 77, 81, 83, 84, 85, 87, 88, 90, 92, 93, 94, 95, 96, 97, 98, 99	
4	2, 3, 4, 8, 10, 12, 14, 16, 17, 21, 22, 23, 31, 34, 44, 46, 49, 51, 55, 56, 58, 59, 61, 70, 75, 77, 81, 83, 84, 85, 87, 88, 90, 92, 93, 94, 95, 96, 97, 98, 99, 100	
5	2, 3, 4, 8, 10, 12, 14, 16, 17, 21, 22, 23, 31, 34, 42, 44, 45, 46, 49, 51, 55, 56, 58, 59, 61, 70, 74, 75, 77, 81, 83, 84, 85, 87, 88, 90, 92, 93, 94, 95, 96, 97, 98, 99, 100	

Table 7. Clients' reputation evolution.

	V1	V2	V3	V4	V5	V6	V7	V8	V9	V10
R1	0.692	0.643	0.716	0.654	0.611	0.651	0.672	0.626	0.697	0.742
R2	0.655	0.725	0.771	0.656	0.627	0.603	0.627	0.903	0.605	0.703
R3	0.691	0.769	0.818	0.705	0.835	0.636	0.871	0.716	0.618	0.745
R4	0.715	0.798	0.832	0.706	0.648	0.656	0.957	0.729	0.625	0.770
R5	0.723	0.817	0.837	0.727	0.672	0.667	0.718	0.732	0.642	0.780

Filtering out clients with potentially lower data quality improves the quality of the aggregated global gradient, as it convenes contributions from high-quality clients only. Additionally, this potentially leads to faster and more accurate model convergence in a real-world FL scenario. To filter out the malicious actors, a two-phased filtering, i.e., Z-score thresholding and consistency behavioral check, is applied in succession to the selected clients. An adaptive threshold update mechanism is also integrated, where thresholds for both Z-score and Trust Score are adjusted based on estimated False Positives (FP) and False Negatives (FN). Table 8 provides a summary of the anomaly detection process. The estimated FP and FN counts provide a proxy evaluation of the detection mechanism's performance against a simplified ground truth based on clients' reputations.

Table 8. Malicious clients' statistics.

	FP	FN	NATST	Filtered Clients
R1	0	0	0.900	None
R2	2	0	0.901	7,13
R3	0	0	0.901	25, 26, 27, 28, 29, 30
R4	0	0	0.901	25, 26, 27, 28, 29, 30
R5	0	0	0.901	25, 26, 27, 28, 29, 30

FP = False Positive, FN = False Negative, NATST = New Adaptive Trust Score Threshold.

To adequately study the impact of our two-stage malicious filtering mechanism, a subset of malicious clients (Clients 25–30) is intentionally introduced from the third round. Their data quality is deliberately reduced to simulate erratic behavior. Table 9 presents the relationship between reputation and state of malicious clients in the FL rounds where a penalty was imposed via reputation slashing and temporary suspension.

Table 9. Impact of reputation slashing and temporary suspension on malicious clients.

Round	Client ID	Is Malicious	Reputation (Before)	Reputation (After)	Status	Violation Count	Suspension End
3	25	True	0.600	0.137	Suspended	1	4
3	26	True	0.600	0.125	Suspended	1	4
3	27	True	0.600	0.136	Suspended	1	4
3	28	True	0.600	0.129	Suspended	1	4
3	29	True	0.600	0.141	Suspended	1	4
3	30	True	0.600	0.122	Suspended	1	4
4	25	True	0.137	0.137	Suspended	1	4
4	26	True	0.125	0.125	Suspended	1	4
4	27	True	0.136	0.136	Suspended	1	4
4	28	True	0.129	0.129	Suspended	1	4
4	29	True	0.141	0.141	Suspended	1	4
4	30	True	0.122	0.122	Suspended	1	4
5	25	True	0.137	0.026	Suspended	2	7
5	26	True	0.125	0.021	Suspended	2	7
5	27	True	0.136	0.025	Suspended	2	7
5	28	True	0.129	0.023	Suspended	2	7
5	29	True	0.141	0.028	Suspended	2	7
5	30	True	0.122	0.021	Suspended	2	7

The deliberate introduction of malicious behavior in round 3 immediately triggered the penalization mechanism. Their reputation scores were significantly slashed, e.g., reputation dropped from 0.6 to 0.137 for Client 25. Concurrently, these clients were placed in a suspended status, indicating their exclusion from subsequent training rounds, and their violation count was incremented. In round 4, these clients remained suspended. By round 5, as the malicious behavior persisted and their suspension expired, they were again identified and penalized, leading to a further drastic reduction in their reputation scores (e.g., Client 25's reputation dropped from 0.137 to 0.026) and an increased suspension period (suspension end round moved from 4 to 7), demonstrating the exponential backoff for repeat offenders.

Based on the detection results from both checks and the adaptive threshold(s), clients identified as malicious are excluded from contributing to the global model aggregation in subsequent rounds.

The total budget set to distribute is 10,000 Bucks, and this suggests that ideally, there are only 100 Bucks for each client. However, the received incentive is different for each client, as observed in Table 10. This is due to the fact that our proposed incentive mechanism is adaptive and dynamic. The distribution of incentives varies between different clients within the same round and the distribution of incentives changes over time for each client across different rounds. This also verifies another notion: that the top four clients throughout the FL process were 31, 51, 77, and 91.

Table 10. Dynamic incentive allocations.

Round	Selected Clients	Reputation Scores	Incentives	Budget Disbursed
1	51, 21, 31, 77	0.737, 0.708, 0.705, 0.705 0.856, 0.815, 0.813, 0.811, 0.805, 0.780, 0.771, 0.767,	587, 552, 463, 471 413, 405, 390, 344, 326, 290, 285, 278, 268, 249,	2073
2	51, 21, 77, 31, 97, 92, 3, 83, 99, 12, 98, 87, 23, 49, 93, 81, 44, 58, 90, 84, 2, 14, 94, 88, 96, 29, 56, 30, 16, 17, 59, 10	0.764, 0.763, 0.757, 0.752, 0.749, 0.747, 0.745, 0.743, 0.738, 0.731, 0.729, 0.727, 0.725, 0.723, 0.718, 0.717, 0.712, 0.710, 0.710, 0.708, 0.707, 0.707, 0.707, 0.703 0.918, 0.868, 0.860, 0.859, 0.856, 0.823, 0.823, 0.818,	242, 236, 231, 230, 212, 211, 209, 201, 199, 195, 190, 181, 106, 93, 87, 79, 85, 53, 52, 51, 50, 44 355, 365, 351, 341, 315, 306, 297, 294, 290, 281, 254, 251, 247, 230, 226, 213, 209, 192, 186, 176, 163, 154, 160, 158, 145, 151, 135, 131, 119, 116, 105, 90, 88, 74, 64, 46, 40	6485
3	51, 77, 31, 97, 21, 92, 99, 3, 83, 98, 12, 49, 23, 81, 93, 58, 44, 87, 84, 90, 2, 14, 94, 17, 56, 88, 96, 10, 16, 59, 85, 75, 8, 34, 95, 61, 70	0.782, 0.782, 0.780, 0.770, 0.769, 0.768, 0.754, 0.751, 0.751, 0.747, 0.746, 0.745, 0.741, 0.741, 0.728, 0.723, 0.716, 0.715, 0.711, 0.704, 0.704 0.948, 0.896, 0.890, 0.885, 0.871, 0.860, 0.843, 0.839, 0.832, 0.822, 0.817, 0.816, 0.814, 0.809, 0.807, 0.804, 0.804, 0.800, 0.798, 0.794, 0.794, 0.776, 0.772, 0.771, 0.770, 0.769, 0.764, 0.761, 0.756, 0.747, 0.746, 0.734, 0.729, 0.728, 0.724, 0.723, 0.719, 0.717, 0.713, 0.710, 0.706, 0.706 0.966, 0.915, 0.901, 0.900, 0.889, 0.883, 0.861, 0.845, 0.837, 0.832, 0.832, 0.831, 0.829, 0.817, 0.812, 0.810, 0.809, 0.805, 0.798, 0.796, 0.796, 0.785, 0.782, 0.780, 0.779, 0.777, 0.776, 0.774, 0.759, 0.757, 0.756, 0.743, 0.735, 0.734, 0.732, 0.731, 0.728, 0.727, 0.725, 0.709, 0.708, 0.706, 0.705, 0.703, 0.701	292, 282, 280, 275, 270, 265, 260, 255, 250, 245, 240, 235, 235, 230, 225, 220, 220, 215, 215, 210, 205, 205, 200, 205, 195, 190, 185, 180, 165, 155, 150, 145, 145, 105, 100, 110, 80, 85, 75, 70, 74, 70 398, 387, 370, 368, 363, 251, 242, 236, 238, 234, 235, 230, 223, 214, 210, 206, 199, 195, 192, 189, 186, 180, 179, 175, 169, 166, 153, 158, 153, 158, 151, 146, 142, 137, 132, 129, 128, 124, 112, 105, 94, 86, 80, 74, 60	7318
4	51, 31, 77, 97, 21, 99, 92, 83, 3, 98, 49, 12, 81, 23, 58, 84, 93, 44, 2, 87, 90, 14, 56, 94, 10, 17, 96, 16, 59, 85, 88, 61, 8, 75, 95, 34, 100, 22, 70, 55, 4, 46	0.706, 0.706 0.966, 0.915, 0.901, 0.900, 0.889, 0.883, 0.861, 0.845, 0.837, 0.832, 0.832, 0.831, 0.829, 0.817, 0.812, 0.810, 0.809, 0.805, 0.798, 0.796, 0.796, 0.785, 0.782, 0.780, 0.779, 0.777, 0.776, 0.774, 0.759, 0.757, 0.756, 0.743, 0.735, 0.734, 0.732, 0.731, 0.728, 0.727, 0.725, 0.709, 0.708, 0.706, 0.705, 0.703, 0.701	205, 205, 200, 205, 195, 190, 185, 180, 165, 155, 150, 145, 145, 105, 100, 110, 80, 85, 75, 70, 74, 70 398, 387, 370, 368, 363, 251, 242, 236, 238, 234, 235, 230, 223, 214, 210, 206, 199, 195, 192, 189, 186, 180, 179, 175, 169, 166, 153, 158, 153, 158, 151, 146, 142, 137, 132, 129, 128, 124, 112, 105, 94, 86, 80, 74, 60	8013
5	51, 31, 97, 77, 99, 21, 92, 83, 3, 49, 81, 98, 23, 2, 12, 58, 44, 93, 90, 84, 87, 14, 56, 10, 94, 96, 17, 16, 85, 61, 59, 34, 100, 22, 8, 88, 75, 4, 95, 46, 70, 55, 74, 42, 45	0.706, 0.706 0.966, 0.915, 0.901, 0.900, 0.889, 0.883, 0.861, 0.845, 0.837, 0.832, 0.832, 0.831, 0.829, 0.817, 0.812, 0.810, 0.809, 0.805, 0.798, 0.796, 0.796, 0.785, 0.782, 0.780, 0.779, 0.777, 0.776, 0.774, 0.759, 0.757, 0.756, 0.743, 0.735, 0.734, 0.732, 0.731, 0.728, 0.727, 0.725, 0.709, 0.708, 0.706, 0.705, 0.703, 0.701	205, 205, 200, 205, 195, 190, 185, 180, 165, 155, 150, 145, 145, 105, 100, 110, 80, 85, 75, 70, 74, 70 398, 387, 370, 368, 363, 251, 242, 236, 238, 234, 235, 230, 223, 214, 210, 206, 199, 195, 192, 189, 186, 180, 179, 175, 169, 166, 153, 158, 153, 158, 151, 146, 142, 137, 132, 129, 128, 124, 112, 105, 94, 86, 80, 74, 60	8557

4.4. Performance Evaluation

This Subsection focuses on system-level effectiveness. Here, we evaluate the proposed method in terms of global model accuracy, fairness, budget balancing, and participation rate. Alongside this, the desirable incentive properties are also validated.

The efficacy of our method is gauged on the MNIST [53] and Fashion-MNIST [54] datasets. MNIST comprises 70,000 images of handwritten digits (0–9), each of size 28×28 . Fashion-MNIST consists of 70,000 grayscale images of miscellaneous articles related to fashion. Each image size is 28×28 pixels, distributed across 10 mutually exclusive classes on garments, shoes, and accessories.

The standard train–test split is used to partition each dataset; only the training data images are used in the model learning, while the complete test split is dedicated to global model testing. The training data is equally divided into fifty local datasets against fifty vehicles. ResNet-18 and a CNN are employed for MNIST and Fashion-MNIST, respectively.

4.4.1. Performance Measures

A combination of performance metrics is utilized to demonstrate the effectiveness of an incentive mechanism. We are evaluating our dynamic AIM from technical, behavioral, and economic perspectives:

- **Utility:** Utility is a system-level performance measure and it determines the overall usefulness of a client node in a system. In the context of vehicular federated learning, it can be referred to as a vehicle's contribution to the central server and the global model. Utility is often described as a composite score calculated by the federated server. While designing an incentive mechanism, this score is usually considered for client selection and reward allocation.
- **Economic Metrics:** These quantitative metrics determine the viability of an incentive mechanism from an economic-theoretic perspective. In our proposed pipeline, these define how efficiently the incentive budget is spent and how well the rewards are aligned with the client's contributions. Economic metrics also ensure that key economic properties such as budget balance, individual rationality, and payment fairness are satisfied.
- **Behavioral Metrics:** Behavioral metrics assess the effectiveness of an incentive mechanism in terms of participation and cooperation of legitimate clients. In our proposed adaptive incentive framework, these indicators characterize how vehicular nodes behave in response to the adaptive incentive mechanism—i.e., their consistent contributions and participation patterns over time.

4.4.2. Model Performance (Utility)

- **Global Model Accuracy:** We measure the accuracy of the final, aggregated model on MNIST and Fashion MNIST. Our final model's accuracy is high, and this is ideally comparable to a centrally trained model. Figure 2 presents the final accuracy achieved with the two datasets, evaluating the effectiveness of our proposed dynamic adaptive incentive mechanism. The MNIST dataset consistently demonstrates a higher global accuracy, i.e., above 95% in numerous FL rounds. However, it also occasionally shows substantial drops in performance. For instance, in round 20, it reduced to a low of approximately 65.6%. In contrast, the Fashion-MNIST dataset experiences marginally lower accuracy. Its lowest recorded accuracy is approximately 63.8% in round 20, but it reaches almost 99.7% in round 37. The observed fluctuations for both datasets are likely correlated with the dynamic nature of the reputation-based client selection and adaptive reward allocations.

Figure 3 also illustrates the global accuracy achieved in each communication round against the total incentives distributed. For both MNIST and Fashion-MNIST datasets, the scatter plot shows a clear positive correlation between the final model accuracy and the amount of incentives distributed. For both datasets, the majority of high-accuracy (95% or more) is centered around the higher incentive range, typically above 8000 Bucks. This pattern confirms that higher financial investment, driven by the need to reward quality contributions, directly leads to better model performance. Moreover, accuracies often remain in the mid-70s and 80s in rounds with incentives as low as 6500 to 7500 Bucks, the mechanism demonstrates robustness to low incentives as well. This demonstrates the mechanism's ability to operate efficiently within the budget constraints while consistently rewarding quality to avoid model collapse and achieve a reasonable accuracy.




Figure 2. Global accuracy vs. communication rounds.

4.4.3. Economic Metrics

- Fairness:** This evaluates the fair disbursement of rewards or benefits among the contributing vehicles. The contributions are being fairly compensated if the Jain's Index is high or close to 1. The plot in Figure 4 illustrates Jain's Fairness Index of the overall budget distributed among the selected clients in fifty communication rounds, calculated using Equation (21).

$$JFI(\mathbf{x}) = \frac{\left(\sum_{i=1}^N x_i\right)^2}{N \cdot \sum_{i=1}^N x_i^2} \quad (21)$$

where:

N is the total number of participating clients.

$\mathbf{x} = (x_1, x_2, \dots, x_N)$ is the vector of payoffs for each client.

x_i represents the payoff or reward received by client i .

If the fairness index = 1, the rewards are distributed perfectly fairly, with all clients receiving the same payoff. In contrast, if it is approaching zero, the distribution is highly unfair; a client node may be receiving a disproportionately large share of the rewards.

Evidently, the index consistently remains at a very high level (between 0.93 and 0.99) for both MNIST and Fashion-MNIST datasets. Comparatively, the Fashion-MNIST shows a marginally higher Jain's index with a consistently staying closer to the Perfect Fairness benchmark (red dashed line at 1.0). Although both indices show minor fluctuations, they never drop below 0.93, demonstrating a near-perfect distribution of incentives throughout the federated learning process. The sustained high fairness indices confirm a highly equitable distribution of rewards among clients, effectively mitigating the risk of client starvation and ensuring sustained, fair participation. And this satisfies the fairness property of an incentive mechanism.




Figure 3. Global accuracy vs. total incentives distributed.




Figure 4. Jain's fairness index in communication rounds.

- Budget Efficiency:** How smartly the server spends its limited budget to improve the global model performance is referred to as budget efficiency. It is calculated as the ratio of the global model's performance to the total rewards spent by the server. The violin plots in Figure 5 explain the distribution of the total incentives allocated to the selected clients in each communication round. For the MNIST dataset, the total incentives range approximately from 6900 to 9300 Bucks, with an estimated concentration around 7950 Bucks as represented by the central horizontal lines. On the contrary, the Fashion-MNIST dataset exhibits a wider distribution that spans from 6600 to 9600 Bucks. This suggests a higher variability in the total incentives disbursed in each round. Here, the incentives are mostly concentrated around 7900 Bucks (the median value represented by the central horizontal lines). Both the data densities demonstrate that our dynamic AIM efficiently manages the total budget of 10,000 Bucks over the total FL rounds. It is prominently visible that the total incentive value per round consistently remains well below the maximum budget, fulfilling the Budget Balanced property of a good incentive mechanism.




Figure 5. Total incentives distributed in each communication round.

4.4.4. Behavioral Metrics

- Participation Rate:** This measures the percentage of eligible clients who are willing to participate in a given FL round. We can observe a participation rate of at least 40 vehicles in each training round. This suggests that our incentive mechanism has a high and stable participation rate to ensure sufficient data for training. This can be analogous to the high reputation scores of the selected clients. Figure 6 demonstrates the average client reputation across the communication rounds for the MNIST and Fashion-MNIST datasets. Generally, the average reputation for both datasets is maintained at a high level, oscillating mainly between 0.6 and 0.9. This indicates that our proposed method promotes high-quality data contributions from clients throughout the training process. Moreover, both datasets exhibit periodic declines below an average of 0.7, signifying occasional drops in perceived data quality. However, the proposed incentive mechanism consistently drives the average reputation back up, often exceeding 0.8, demonstrating its stability and effectiveness in dynamically

adapting incentives to maintain high client participation and data quality over the entire FL training.




Figure 6. Average clients' reputation in communication rounds.

4.5. Baseline Comparison

To rigorously evaluate the effectiveness of our reputation-based reward allocation scheme, we conducted a comparative study against the following four representative baselines:

1. Vanilla FL: The first baseline is Vanilla FL [55], in which clients with large amounts of data receive higher incentives.
2. The second combines reputation and reverse auction theory [13], i.e., RRAFL. Participants bid for tasks, and their reputation indirectly reflects their reliability and data quality.
3. The third baseline is based on the reverse auction problem, FAIR [36]. It is a quality-aware incentive mechanism that utilizes individual learning quality to facilitate precise user incentives and model aggregation.
4. The fourth is FLORA [56]. It is based on randomized auctions and focuses on social cost minimization.

The results are not compared with a centralized or an all-client model where all clients train a model. In this scenario, one can get a higher accuracy assuming all the clients are legitimate with high-quality updates, but the computational cost and power consumption will be too high.

Table 11 provides a comparative assessment of our proposed adaptive incentive mechanism against the established baselines across behavioral and economic dimensions. Our method offers broader and more consistent behavioral and economic properties. It combines stable client participation, explicit management of unreliable clients, improved budget allocation, and precise contribution-aligned rewards. Collectively, these capabilities result in a reliable and more comprehensive incentive architecture as compared to FAIR, FLORA, Vanilla FL, and RRAFL.

Table 12 presents the computational and communication costs of the proposed method in comparison with other baseline methods.

Table 11. Comparison between the proposed method and the baselines with behavioral and economic perspectives.

Metric Type	Metric	Vanilla FL	RRAFL	FAIR	FLORA	Our Method
Behavioral	Participation Rate	Low (no incentives)	Medium (auction winners)	Medium (quality-based)	Med-High (energy-constrained)	High (reputation-based selection keeps reliable clients active)
	Dropout/Suspension	No modeling	Not modeled	Not modeled	Energy-based dropouts	Explicit suspension + reinstatement based on trust score
	Free-Rider Mitigation	None	Partial (discourages freeloading)	Partial (quality-proportional)	None	Strong (contribution–reward correlation)
	Stability of Selection	Low (random)	Medium	Medium	Medium	High (reputation and trust stabilize participation)
	Malicious/Low-Quality Filtering	None	None	Partial	None	Explicit filtering via consistency + reputation
Economic	Budget Utilization	None	Auction-dependent	Budgeted	Long-term energy budget	Dynamic budget allocation
	Payment Fairness	None	Not guaranteed	Yes (fair incentive distribution)	Not primary goal	Yes (Jain’s Index improves over rounds)
	Incentive Efficiency	No incentives	Moderate	High	High	Highest (reputation-weighted aggregation reduces waste)
	Incentive Precision	N/A	No (auction cost-based)	Yes	No	Yes (reputation + consistency driven)
	Energy Awareness	None	None	None	Yes (energy constraint)	Yes (energy fairness factor)
	Individual Rationality	N/A	Guaranteed	Not explicit	Guaranteed	Guaranteed (non-negative rewards)

Table 12. Computational and communication cost analysis.

Method	Server Computation Cost	Client Computation Cost	Communication Cost
Vanilla FL	$O(P)$ aggregation	$O(EBC)$ local training	Upload/download full model; no extra metadata
RRAFL	$O(N \log N)$ auction + selection	$O(EBC)$ + bidding overhead	Model exchange + bid submission; higher uplink cost
FAIR	$O(N_{\text{sel}}P)$ weighted aggregation	$O(EBC)$	Standard model exchange; no additional messages
FLORA	$O(N) + O(N \log N)$ online randomized auctions	$O(EBC)$ + energy optimization	Model exchange + energy reports; moderate overhead
Our Method	$O(N_{\text{sel}}P + N)$ (reputation, trust, filtering, weighted aggregation)	$O(EBC)$ (no added load)	Same as FedAvg: model upload/download only; no auxiliary communication

Notes: N_{sel} = number of selected clients per round; P = number of model parameters; N = total number of clients; EBC = client-side local training cost per round (E = epochs, B = batches per epoch, C = cost of one forward/backward pass).

For clients or vehicular nodes, our method preserves the same client-side computational and communication footprint as Vanilla FL (employing FedAvg); offering the lowest computational cost. The FL server incurs a modest overhead of $O(N_{\text{sel}}P + N)$ as all adaptive logic, like reputation scoring, malicious filtering, and weighted aggregation, is executed here. This computational cost is slightly higher than pure FedAvg but is significantly lower than the auction-based baselines in RRAFL and FLORA. Summatively, the proposed method is similar to FAIR but more efficient than RRAFL and FLORA in terms of computational costs incurred by the FL server. Regarding communication overheads, the proposed incentive mechanism does not introduce extra communication rounds or metadata exchange, unlike the auction-based techniques used in RRAFL and FLORA.

Overall, the proposed method maintains computational and communication efficiency comparable to the simplest baseline using FedAvg, i.e., Vanilla FL.

Our proposed incentive framework consistently achieves the highest global accuracy across the MNIST and the Fashion-MNIST datasets, as illustrated in Figure 7. It maintains a highly competitive accuracy profile on the MNIST dataset, which peaks to nearly 99.7% for Fashion-MNIST. It showcases its ability to drive high-quality training contributions despite budget constraints. Furthermore, our selective reputation-based rewarding process introduces manageable, transient fluctuations in global accuracy, outperforming the baseline methods that may prioritize model-centric contributions or raw data quantity over sustained quality and equitable compensation. This comprehensive trade-off analysis confirms that our dynamic AIM is the most effective approach for federated learning-enhanced IoV networks, which simultaneously optimizes judicious budget utilization, equitable client rewards, and high model accuracy.

4.6. Qualitative Analysis

Our proposed method satisfies most of the desired properties of a well-designed incentive mechanism, namely, Budget Balanced, Collusion Resistance, Fairness, and Individual Rationality.

- **Fairness:** From Figure 4, it can be observed that Jain's Index approaches 1 in nearly all the communication rounds, suggesting the Fairness property of incentive mechanisms is satisfied.
- **Budget Balanced:** As the sum of the budget disbursed among the selected clients upon task completion is less than the initial total budget floated in the FL network

by the central server, as observed in Figure 5. Hence, the Budget Balanced property is satisfied.

- **Individual Rationality:** Figure 5 shows that all the selected clients have positive rewards; it implies the individual Rationality property is met.
- **Collusion Resistant:** As shown in Table 13, the property of Collusion Resistance is satisfied. A higher quality leads to a higher reputation.




Figure 7. Comparison with baseline methods.

Table 13. Final reputations and last timestep's metrics.

	Avg Q	Avg S	Avg C	Avg R	NAC	NSC
R1	0.727	0.808	0.758	0.617	100	96
R2	0.722	0.795	0.765	0.642	98	95
R3	0.673	0.793	0.723	0.693	94	89
R4	0.727	0.833	0.780	0.662	94	89
R5	0.732	0.843	0.779	0.718	94	90

R = Reputation, Q = Data Quality, S = Stability, C = Consistency, NAC = Number of Active Clients, NSC = Number of Selected Clients.

4.7. Key Findings

It is observed that increasing the number of participating clients enhances the global model accuracy. This is largely due to the improved representativeness of the aggregated data distribution. As more clients contribute to the training process, the learning benefits from greater data diversity and reduced sampling bias, which enables the global model to better approximate the underlying population distribution. Several studies report positive correlations between client population size and global model performance under standard FL settings [2,57–59].

Reputation Hyperparameter: We have observed that setting the reputation threshold to 0.5 selects more clients, and the global model accuracy is increased. If this threshold is increased, for instance, to 0.7, fewer clients are selected. And, in return, this decreases the global model accuracy because there is less data available with fewer clients to appropriately train the model. However, the convergence time in this scenario will drastically increase.

Influence of Q, S, and C on Final Reputation: Table 13 provides insights into the relationship between the reputation and quantities of Q, S, and C. Generally, clients with higher final Q, S, and C values tend to have higher final reputations.

Impact of Filtering on Reputation of Remaining Clients: By removing potentially detrimental gradients from the aggregation, the resulting global gradient becomes more aligned with the contributions of legitimate clients. This can positively influence the Q and C scores of legitimate clients in subsequent rounds, potentially leading to higher reputation scores for them and reinforcing the reputation differentiation between client types.

Impact of False Positives on Model Accuracy: A potential concern in malicious client filtering is the impact of false positives. We observed occasional false positives in the malicious-detection stage that had minimal influence on model accuracy or training fairness. Our proposed incentive design ensures that any false penalties are short-lived and have limited influence on participation, fairness, or global model performance. Firstly, our multi-factor filtering logic reduces the probability that a transient variation in client behavior will lead to misclassification. Secondly, the Reputation update rule includes a forgetting factor that smooths short-term volatility in client metrics, ensuring that isolated anomalies have limited influence on long-term reputation. Thirdly, the system includes an explicit reinstatement mechanism that prevents the permanent exclusion of honest participants; it enables temporarily suspended clients to return to active status once their reputation recovers.

Impact of Filtering on Global Model Update: Filtering out clients flagged by the detection mechanisms is expected to improve the quality of the aggregated gradients. While not quantitatively measured in this simulation, in a real-world scenario, this could lead to faster and/or more stable global model convergence by reducing the influence of potentially low-quality or inconsistent contributions.

Global Accuracy Fluctuations: Global accuracy fluxes are inherent in adaptive or reputation-driven federated learning systems. Due to the adaptive nature of the proposed incentive mechanism, the global accuracy exhibited noticeable fluctuations. The participating clients can vary significantly in each FL round due to ever-evolving reputation scores, trust, and suspension and reinstatement of vehicles. These changes in participation patterns lead to corresponding variations in the quality and consistency of aggregated model updates, which give rise to short-term accuracy drops. However, the variability in the global accuracy is diminished as the training progresses, reputations stabilize, and the set of selected vehicles becomes consistent.

Adaptive Threshold Behavior: The adaptive threshold's decrease over time, driven by a higher number of estimated false negatives, highlights the challenge of balancing detection sensitivity. This suggests that the chosen proxy for ground truth and the adaptive mechanism's parameters (like the learning rate η) significantly influence its behavior and would require careful tuning and validation in a real-world application.

To obtain high-quality models in FL, more clients are needed for local training, which inevitably consumes resources such as computation and communication costs. Even in a situation where resources are limited (e.g., the limited payment budgets), our AIM can improve model convergence speed and achieve a better final model accuracy than those solutions that only consider payment costs or ignore edge server personal decision-making.

5. Conclusions

This paper presented a comprehensive trust-aware federated learning framework designed for dynamic vehicular networks, addressing the intertwined challenges of trust evaluation, incentive distribution, and malicious client mitigation. The proposed system introduces a multi-dimensional reputation model that jointly considers data quality, stability, and behavioral consistency, allowing reliable differentiation between trustworthy

and adversarial participants. By integrating these reputation indicators into an adaptive incentive control process, the framework ensures equitable participation while maintaining high model performance under diverse vehicular conditions.

A key contribution of this work lies in the development of a feedback-control mechanism that dynamically adjusts the weighting parameters (α_W , β_W , γ_W , λ_W) based on real-time learning outcomes. This adaptivity enables the incentive mechanism to self-tune its fairness and accuracy trade-offs, effectively reducing the dependence on manual parameter tuning. Furthermore, the integration of a statistical anomaly detection module provides continuous surveillance of client updates, supporting early detection and penalization of malicious behavior through a transparent reputation-slashing process.

Extensive experimental results on vehicular datasets demonstrate that the proposed approach achieves faster convergence, enhanced robustness, and improved fairness compared to static incentive and trust-agnostic baselines. Specifically, the framework reduces the impact of poisoned updates, maintains consistent global accuracy, and improves participation equity across heterogeneous vehicular clients. These findings underscore the viability of incorporating feedback-regulated trust management into large-scale federated learning systems.

Future work will focus on extending this framework toward real-world deployment scenarios involving cross-domain vehicular collaborations and edge-cloud orchestration. Additional research directions include integrating blockchain for tamper-proof reputation recording, leveraging large language models for adaptive trust reasoning, and exploring lightweight encryption mechanisms to further enhance privacy preservation. Ultimately, the proposed architecture establishes a foundation for resilient, fair, and intelligent federated learning in next-generation vehicular and edge computing ecosystems. extend validation to real-world vehicular datasets and multi-agent edge scenarios.

Author Contributions: Conceptualization, A.R. and E.B.; Methodology, A.R.; Software, A.R.; Validation, A.R. and E.B.; Formal Analysis, A.R.; Investigation, A.R., O.E.H. and E.B.; Resources, E.B. and O.E.H.; Writing—Original Draft Preparation, A.R.; Writing—Review & Editing, E.B. and O.E.H.; Visualization, A.R. and O.E.H.; Supervision, E.B.; Project Administration, E.B.; Funding Acquisition, E.B. and O.E.H. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the United Arab Emirates University—UAEU Program for Advanced Research (UPAR) Grant N. G00003443, and the APC was funded by the UAEU Research Startup Grant N. G00004666.

Data Availability Statement: MNIST and Fashion-MNIST datasets are used in this article and are publicly available at [53,54].

Acknowledgments: The authors acknowledge the use of AI-assisted tools for minor editorial tasks, mainly for language and grammatical checks, and assistance with the formatting of LaTeX tables.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Konečný, J. Federated Learning: Strategies for Improving Communication Efficiency. *arXiv* **2016**, arXiv:1610.05492.
2. McMahan, B.; Moore, E.; Ramage, D.; Hampson, S.; y Arcas, B.A. Communication-efficient learning of deep networks from decentralized data. In Proceedings of the Artificial Intelligence and Statistics, PMLR, Lauderdale, FL, USA, 20–22 April 2017; pp. 1273–1282.
3. Li, L.; Fan, Y.; Tse, M.; Lin, K.Y. A review of applications in federated learning. *Comput. Ind. Eng.* **2020**, *149*, 106854. [CrossRef]
4. Banabilah, S.; Aloqaily, M.; Alsayed, E.; Malik, N.; Jararweh, Y. Federated learning review: Fundamentals, enabling technologies, and future applications. *Inf. Process. Manag.* **2022**, *59*, 103061. [CrossRef]
5. Schwanck, F.M.; Leipnitz, M.T.; Carbonera, J.L.; Wickboldt, J.A. A Framework for testing Federated Learning algorithms using an edge-like environment. *Future Gener. Comput. Syst.* **2024**, *166*, 107626. [CrossRef]

6. Soltani, B.; Haghighi, V.; Mahmood, A.; Sheng, Q.Z.; Yao, L. A survey on participant selection for federated learning in mobile networks. In Proceedings of the 17th ACM Workshop on Mobility in the Evolving Internet Architecture, Sydney, NSW, Australia, 21 October 2022; pp. 19–24.
7. Zeng, R.; Zeng, C.; Wang, X.; Li, B.; Chu, X. A comprehensive survey of incentive mechanism for federated learning. *arXiv* **2021**, arXiv:2106.15406. [\[CrossRef\]](#)
8. Zhan, Y.; Zhang, J.; Hong, Z.; Wu, L.; Li, P.; Guo, S. A survey of incentive mechanism design for federated learning. *IEEE Trans. Emerg. Top. Comput.* **2021**, *10*, 1035–1044. [\[CrossRef\]](#)
9. Zhan, Y.; Li, P.; Guo, S.; Qu, Z. Incentive mechanism design for federated learning: Challenges and opportunities. *IEEE Netw.* **2021**, *35*, 310–317. [\[CrossRef\]](#)
10. Le, T.H.T.; Tran, N.H.; Tun, Y.K.; Nguyen, M.N.; Pandey, S.R.; Han, Z.; Hong, C.S. An incentive mechanism for federated learning in wireless cellular networks: An auction approach. *IEEE Trans. Wirel. Commun.* **2021**, *20*, 4874–4887. [\[CrossRef\]](#)
11. Zeng, R.; Zhang, S.; Wang, J.; Chu, X. FMore: An incentive scheme of multi-dimensional auction for federated learning in MEC. In Proceedings of the 2020 IEEE 40th International Conference on Distributed Computing Systems (ICDCS), Singapore, 29 November–1 December 2020; IEEE: New York, NY, USA, 2020; pp. 278–288.
12. Chen, Z.; Zhang, H.; Li, X.; Miao, Y.; Zhang, X.; Zhang, M.; Ma, S.; Deng, R.H. FDFL: Fair and discrepancy-aware incentive mechanism for federated learning. *IEEE Trans. Inf. Forensics Secur.* **2024**, *19*, 8140–8154. [\[CrossRef\]](#)
13. Zhang, J.; Wu, Y.; Pan, R. Incentive mechanism for horizontal federated learning based on reputation and reverse auction. In Proceedings of the Web Conference 2021, Ljubljana, Slovenia, 19–23 April 2021; pp. 947–956.
14. Ahmed, A.; Choi, B.J. Frimfl: A fair and reliable incentive mechanism in federated learning. *Electronics* **2023**, *12*, 3259. [\[CrossRef\]](#)
15. Zhang, H.; Li, P.; Dai, M.; Wu, Y.; Qian, L. Efficient federated learning with quality-aware generated models: An incentive mechanism. *IEEE Internet Things J.* **2024**, *12*, 1628–1642. [\[CrossRef\]](#)
16. Zheng, S.; Cao, Y.; Yoshikawa, M.; Li, H.; Yan, Q. FL-market: Trading private models in federated learning. In Proceedings of the 2022 IEEE International Conference on Big Data (Big Data), Osaka, Japan, 24–25 August 2022; IEEE: New York, NY, USA, 2022; pp. 1525–1534.
17. Zhang, Y.; Pan, M.; Song, L.; Dawy, Z.; Han, Z. A survey of contract theory-based incentive mechanism design in wireless networks. *IEEE Wirel. Commun.* **2017**, *24*, 80–85. [\[CrossRef\]](#)
18. Kang, J.; Xiong, Z.; Niyato, D.; Xie, S.; Zhang, J. Incentive mechanism for reliable federated learning: A joint optimization approach to combining reputation and contract theory. *IEEE Internet Things J.* **2019**, *6*, 10700–10714. [\[CrossRef\]](#)
19. Zhang, X.; Wang, J. FedMG: Vehicular Edge Federated Learning for Mobile Scenarios With Geo-Dispersed Data. *IEEE Trans. Veh. Technol.* **2024**, *74*, 1520–1533. [\[CrossRef\]](#)
20. Sun, P.; Che, H.; Wang, Z.; Wang, Y.; Wang, T.; Wu, L.; Shao, H. Pain-FL: Personalized privacy-preserving incentive for federated learning. *IEEE J. Sel. Areas Commun.* **2021**, *39*, 3805–3820. [\[CrossRef\]](#)
21. Tian, M.; Chen, Y.; Liu, Y.; Xiong, Z.; Leung, C.; Miao, C. A contract theory based incentive mechanism for federated learning. *arXiv* **2021**, arXiv:2108.05568. [\[CrossRef\]](#)
22. Li, L.; Yu, X.; Cai, X.; He, X.; Liu, Y. Contract-theory-based incentive mechanism for federated learning in health crowdsensing. *IEEE Internet Things J.* **2022**, *10*, 4475–4489. [\[CrossRef\]](#)
23. Zhan, Y.; Li, P.; Qu, Z.; Zeng, D.; Guo, S. A learning-based incentive mechanism for federated learning. *IEEE Internet Things J.* **2020**, *7*, 6360–6368. [\[CrossRef\]](#)
24. Xu, M.; Peng, J.; Gupta, B.B.; Kang, J.; Xiong, Z.; Li, Z.; Abd El-Latif, A.A. Multiagent federated reinforcement learning for secure incentive mechanism in intelligent cyber-physical systems. *IEEE Internet Things J.* **2021**, *9*, 22095–22108. [\[CrossRef\]](#)
25. Wu, Q.; Wang, S.; Fan, P.; Fan, Q. Deep reinforcement learning based vehicle selection for asynchronous federated learning enabled vehicular edge computing. In Proceedings of the International Congress on Communications, Networking, and Information Systems, Guilin, China, 25–27 March 2023; Springer: Berlin/Heidelberg, Germany, 2023; pp. 3–26.
26. Lin, W.; Zhou, Y. Improved Federated Learning Incentive Mechanism Algorithm Based on Explainable DAG Similarity Evaluation. *Mathematics* **2025**, *13*, 3507. [\[CrossRef\]](#)
27. Zeng, R.; Zeng, C.; Wang, X.; Li, B.; Chu, X. Incentive mechanisms in federated learning and a game-theoretical approach. *IEEE Netw.* **2022**, *36*, 229–235. [\[CrossRef\]](#)
28. Cong, M.; Yu, H.; Weng, X.; Yiu, S.M. A game-theoretic framework for incentive mechanism design in federated learning. In *Federated Learning: Privacy and Incentive*; Springer: Berlin/Heidelberg, Germany, 2020; pp. 205–222.
29. Hu, R.; Gong, Y. Trading data for learning: Incentive mechanism for on-device federated learning. In Proceedings of the GLOBECOM 2020—2020 IEEE Global Communications Conference, Taipei, Taiwan, 7–11 December 2020; IEEE: New York, NY, USA, 2020; pp. 1–6.
30. Khan, L.U.; Pandey, S.R.; Tran, N.H.; Saad, W.; Han, Z.; Nguyen, M.N.; Hong, C.S. Federated learning for edge networks: Resource optimization and incentive mechanism. *IEEE Commun. Mag.* **2020**, *58*, 88–93. [\[CrossRef\]](#)

31. Xiao, G.; Xiao, M.; Gao, G.; Zhang, S.; Zhao, H.; Zou, X. Incentive mechanism design for federated learning: A two-stage stackelberg game approach. In *Proceedings of the 2020 IEEE 26th International Conference on Parallel and Distributed Systems (ICPADS)*, Hong Kong, 2–4 December 2020; IEEE: New York, NY, USA, 2020; pp. 148–155.
32. Wu, H.; Tang, X.; Zhang, Y.J.A.; Gao, L. Incentive mechanism for federated learning based on random client sampling. In *Proceedings of the 2022 IEEE Globecom Workshops (GC Wkshps)*, Rio de Janeiro, Brazil, 4–8 December 2022; IEEE: New York, NY, USA, 2022; pp. 1640–1645.
33. Luo, B.; Feng, Y.; Wang, S.; Huang, J.; Tassiulas, L. Incentive mechanism design for unbiased federated learning with randomized client participation. In *Proceedings of the 2023 IEEE 43rd International Conference on Distributed Computing Systems (ICDCS)*, Hong Kong, 18–21 July 2023; IEEE: New York, NY, USA, 2023; pp. 545–555.
34. Zhan, Y.; Xia, Y.; Zhang, J. Incentive mechanism in platform-centric mobile crowdsensing: A one-to-many bargaining approach. *Comput. Netw.* **2018**, *132*, 40–52. [\[CrossRef\]](#)
35. Yu, H.; Liu, Z.; Liu, Y.; Chen, T.; Cong, M.; Weng, X.; Niyato, D.; Yang, Q. A sustainable incentive scheme for federated learning. *IEEE Intell. Syst.* **2020**, *35*, 58–69. [\[CrossRef\]](#)
36. Deng, Y.; Lyu, F.; Ren, J.; Chen, Y.C.; Yang, P.; Zhou, Y.; Zhang, Y. Fair: Quality-aware federated learning with precise user incentive and model aggregation. In *Proceedings of the IEEE INFOCOM 2021—IEEE Conference on Computer Communications, Virtual*, 10–13 May 2021; IEEE: New York, NY, USA, 2021; pp. 1–10.
37. Gao, L.; Li, L.; Chen, Y.; Zheng, W.; Xu, C.; Xu, M. Fifi: A fair incentive mechanism for federated learning. In *Proceedings of the 50th International Conference on Parallel Processing*, Lemont, IL, USA, 9–12 August 2021; pp. 1–10.
38. Yang, J.; Cao, S.; Zhao, C.; Niu, W.; Tsai, L.C. Portfolio-based incentive mechanism design for cross-device federated learning. In *Proceedings of the 2024 IEEE Wireless Communications and Networking Conference (WCNC)*, Dubai, United Arab Emirates, 21–24 April 2024; IEEE: New York, NY, USA, 2024; pp. 1–6.
39. Yuan, S.; Dong, B.; Li, J.; Guo, S.; Chen, H.; Wu, C.; Wu, J.; Zhao, W. Adaptive Incentivize for Federated Learning with Cloud-Edge Collaboration under Multi-level Information Sharing. *IEEE Trans. Comput.* **2025**, *74*, 2445–2460. [\[CrossRef\]](#)
40. Rose, S.; Borchert, O.; Mitchell, S.; Connelly, S. *Zero Trust Architecture*; Technical Report; National Institute of Standards and Technology, U.S. Department of Commerce: Gaithersburg, MD, USA, 2020. [\[CrossRef\]](#)
41. Wang, H.; Guo, S.; Cao, J.; Guo, M. MeLoDy: A long-term dynamic quality-aware incentive mechanism for crowdsourcing. *IEEE Trans. Parallel Distrib. Syst.* **2017**, *29*, 901–914. [\[CrossRef\]](#)
42. Song, B.; Shah-Mansouri, H.; Wong, V.W. Quality of sensing aware budget feasible mechanism for mobile crowdsensing. *IEEE Trans. Wirel. Commun.* **2017**, *16*, 3619–3631. [\[CrossRef\]](#)
43. Cao, X.; Fang, M.; Liu, J.; Gong, N.Z. Fltrust: Byzantine-robust federated learning via trust bootstrapping. *arXiv* **2020**, arXiv:2012.13995.
44. Zhang, Z.; Dong, D.; Ma, Y.; Ying, Y.; Jiang, D.; Chen, K.; Shou, L.; Chen, G. Refiner: A reliable incentive-driven federated learning system powered by blockchain. *Proc. Vldb Endow.* **2021**, *14*, 2659–2662. [\[CrossRef\]](#)
45. Hui, D.; Zhuo, L.; Xin, C. Quality-aware incentive mechanism design based on matching game for hierarchical federated learning. In *Proceedings of the IEEE INFOCOM 2022—IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, New York, NY, USA, 2–5 May 2022; IEEE: New York, NY, USA, 2022; pp. 1–6.
46. Zheng, Z.; Hong, Y.; Li, K.; Chen, Q. An incentive mechanism based on AoU, data quality, and data quantity for federated learning. In *Proceedings of the 2024 27th International Conference on Computer Supported Cooperative Work in Design (CSCWD)*, Tianjin, China, 8–10 May 2024; IEEE: New York, NY, USA, 2024; pp. 1832–1837.
47. Lansky, J.; Rahmani, A.M.; Hosseinzadeh, M. Reinforcement learning-based routing protocols in vehicular Ad Hoc Networks for Intelligent Transport System (ITS): A survey. *Mathematics* **2022**, *10*, 4673. [\[CrossRef\]](#)
48. Chen, C.; Si, J.; Li, H.; Han, W.; Kumar, N.; Berretti, S.; Wan, S. A high stability clustering scheme for the internet of vehicles. *IEEE Trans. Netw. Serv. Manag.* **2024**, *21*, 4297–4311. [\[CrossRef\]](#)
49. Mukhtaruzzaman, M.; Atiquzzaman, M. Junction-based stable clustering algorithm for vehicular ad hoc network. *Ann. Telecommun.* **2021**, *76*, 777–786. [\[CrossRef\]](#)
50. Fu, Y.; Li, Z.; Liu, S.; Li, C.; Yu, F.R.; Cheng, N. An incentive mechanism for long-term federated learning in autonomous driving. *IEEE Internet Things J.* **2023**, *11*, 15642–15655. [\[CrossRef\]](#)
51. Kang, J.; Xiong, Z.; Niyato, D.; Xie, S.; Zhang, J. Reliable Federated Learning for Mobile Networks. *IEEE Wirel. Commun.* **2022**, *29*, 100–107. [\[CrossRef\]](#)
52. Blanchard, P.; Mhamdi, E.M.E.; Guerraoui, R.; Stainer, J. Machine Learning with Adversaries: Byzantine Tolerant Gradient Descent. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*, Long Beach, CA, USA, 4–9 December 2017; pp. 119–129.
53. LeCun, Y.; Bottou, L.; Bengio, Y.; Haffner, P. Gradient-based learning applied to document recognition. *Proc. IEEE* **1998**, *86*, 2278–2324. [\[CrossRef\]](#)

54. Xiao, H.; Rasul, K.; Vollgraf, R. Fashion-MNIST: A novel image dataset for benchmarking machine learning algorithms. *arXiv* **2017**, arXiv:1708.07747. [[CrossRef](#)]
55. McMahan, H.B.; Moore, E.; Ramage, D.; y Arcas, B.A. Federated learning of deep networks using model averaging. *arXiv* **2016**, arXiv:1602.05629v4.
56. Yuan, Y.; Jiao, L.; Zhu, K.; Zhang, L. Incentivizing federated learning under long-term energy constraint via online randomized auctions. *IEEE Trans. Wirel. Commun.* **2021**, *21*, 5129–5144. [[CrossRef](#)]
57. Li, T.; Sahu, A.K.; Zaheer, M.; Sanjabi, M.; Talwalkar, A.; Smith, V. Federated optimization in heterogeneous networks. *Proc. Mach. Learn. Syst.* **2020**, *2*, 429–450.
58. Kairouz, P.; McMahan, H.B.; Avent, B.; Bellet, A.; Bennis, M.; Bhagoji, A.N.; Bonawitz, K.; Charles, Z.; Cormode, G.; Cummings, R.; et al. Advances and open problems in federated learning. *Found. Trends Mach. Learn.* **2021**, *14*, 1–210. [[CrossRef](#)]
59. Raza, A.; Badidi, E. A Trust-based Client Selection Framework for Federated Learning in the Internet of Vehicles. In *Proceedings of the 2025 International Wireless Communications and Mobile Computing (IWCMC), Abu Dhabi, United Arab Emirates, 12–16 May 2025*; IEEE: New York, NY, USA, 2025; pp. 1180–1185.

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.