

Article

Game-Theoretic Optimized Federated Learning for Heterogeneous IoT Object Detection

Zepeng Wang ^{1,*} and Jie Chen ^{2,*} ¹ School of Cyberspace Security, Gansu University of Political Science and Law, Lanzhou 730070, China² Department of Automation, University of Science and Technology of China, Hefei 230027, China

* Correspondence: wzp7293@gsupl.edu.cn (Z.W.); chenjie2025@ustc.edu.cn (J.C.)

Abstract

Federated object detection allows distributed IoT cameras to learn a shared detector without exposing raw images. Its performance is limited by non-IID scenes, unequal device resources, intermittent links, and selfish participation. We propose GO-FedDet, a game-theoretic optimized federated detection framework for heterogeneous IoT application. Client participation is formulated as a Stackelberg game, selection is formulated as an exact-potential resource game, and the equilibrium is embedded into a proximal detection objective. A utility-aligned aggregation balances detection contribution, communication/energy cost, and long-term fairness. We prove the existence of a Stackelberg equilibrium and the finite-improvement convergence of the selection game, and we derive a non-convex convergence bound under client drift and partial participation. Experiments on heterogeneous edge object detection show that GO-FedDet improves accuracy, lowers communication cost, and stabilizes fairness compared with representative federated baselines.

Keywords: federated learning; game theory; heterogeneous IoT; object detection; complex networks

1. Introduction

Object detection over IoT vision networks underpins applications such as traffic perception, industrial inspection, smart agriculture, and edge surveillance [1–3]. Cameras and embedded sensors collect local scenes, but raw images are often private, bandwidth-intensive, and subject to site-specific policies [4]. Federated learning (FL) offers a natural paradigm where each client trains locally and a server aggregates model updates [5,6]. However, object detection is harder than image classification in FL: it couples localization and classification losses, relies on dense anchors or queries, and is highly sensitive to rare objects and background shift [7,8].

Recent studies identify FL as a core mechanism for privacy-preserving edge learning but also emphasize statistical heterogeneity, device heterogeneity, communication cost, and reliability as persistent obstacles [9–11]. Specialized studies have evaluated federated detectors for autonomous driving, UAV monitoring, maritime vision, and agricultural scenes [12]. These studies confirm the feasibility of collaborative detection, yet they usually treat client participation as a sampling problem rather than as a strategic interaction among resource-limited IoT nodes.

Client selection, incentives, and aggregation have been extensively studied [13–15]. Hierarchical aggregation, asynchronous coordination, interior-point client selection, and energy-aware FL improve scalability for heterogeneous networks [16,17]. Game-theoretic



Academic Editors: Kiho Lim and Weihua Liu

Received: 10 June 2026

Revised: 29 June 2026

Accepted: 10 July 2026

Published: 12 July 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.

and blockchain-assisted incentive mechanisms indicate that rational clients may reduce participation without proper rewards, reputation, or contracts [18]. However, most approaches optimize a scalar training objective and rarely couple game-theoretic participation with the structured loss of object detection, and they also do not provide a unified convergence analysis for detection-oriented FL under selfish behavior [19].

Game-theoretic and resource-aware studies further show that FL over the IoT is shaped by incentives, congestion, energy budgets, and hierarchical edge coordination [20,21]. Voting-based and clustered selection mechanisms enable scalable coordination for massive networks of devices; federated few-shot detection, privacy-preserving object detection, and secure visual processing have emerged as recent research directions of interest [22]. Blockchain-assisted aggregation and reputation frameworks enhance reliability and trust [23]. From an optimization perspective, personalized FL, model-contrastive learning, matched averaging, and flat-minima regularization provide tools for handling statistical heterogeneity [24]. Classical game theory and potential games supply the equilibrium and convergence foundations used in this work.

The key novelty of this is not the independent use of a Stackelberg game, a potential game, or proximal FL but the closed-loop coupling of these components for federated object detection. Rather than select clients solely by data size or loss, GO-FedDet links four decisions that are usually handled separately: the server reward, the client's willingness to participate, the congestion-aware channel-selection equilibrium, and the detection update used for aggregation. The Stackelberg game determines participation thresholds, the potential game converts feasible participants into a Nash-stable selected subgraph, and the resulting selection bias appears explicitly in the non-convex convergence bound. The proximal detection objective then limits local object-detection drift, while the fairness queue feeds long-term participation history back into both selection and aggregation [25–27]. This coupling provides a methodological contribution beyond a direct combination of existing modules.

This paper proposes GO-FedDet, a game-theoretic optimized federated learning framework for heterogeneous IoT object detection. In each round, the server posts a reward vector and a resource budget. Clients evaluate their marginal detection contribution, energy cost, communication delay, and historical participation. Selected clients train a detector locally with a proximal loss, and the server aggregates updates via a utility-aligned rule that combines validation gain, cost efficiency, and fairness (Figure 1).

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

- (1) We formulate federated IoT object detection as a coupled Stackelberg-potential game, where incentive-compatible participation and congestion-aware resource selection are jointly used to generate an equilibrium-guided client set.
- (2) We design a detection-oriented proximal training and utility-aligned aggregation algorithm that couples reward, drift, energy, delay, validation gain, and fairness within one executable workflow.
- (3) We provide theoretical analysis, including equilibrium existence, finite-improvement convergence of the selection game, bounded participation deficit, and a non-convex convergence bound that explicitly contains the game-induced selection-bias term.

The remainder of this paper is organized as follows. Section 2 introduces the problem description and preliminaries. Section 3 presents the model and game-theoretic analysis. Section 4 gives the algorithm design. Section 5 analyzes convergence and complexity. Section 6 reports simulation results. Section 7 concludes the paper and discusses future work. The main notation and definitions are shown in Table 1.




Figure 1. GO-FedDet framework for heterogeneous IoT detection.

Table 1. Main notation and definitions.

Symbol	Definition
$\mathcal{N}$	Set of IoT clients, $ \mathcal{N} = N$
$\mathcal{S}_t$	Selected client set at round t
θ_t	Global detector parameters
$F_i(\theta)$	Local detection risk of client i
$F(\theta)$	Weighted global detection risk
n_i	Number of local detection samples
Δ_i^t	Model update returned by client i
c_i^t	Composite communication-energy cost
d_i^t	Local drift indicator
q_i^t	Estimated detection contribution
r_i^t	Reward allocated to client i
u_i^t	Utility of client i
B_t	Server-side budget at round t
ρ_i^t	Long-term fairness deficit
μ	Proximal regularization coefficient
η_l	Local learning rate
E	Number of local epochs
L	Smoothness constant of F_i
σ^2	Stochastic gradient variance upper bound
ζ^2	Client heterogeneity upper bound
α_i^t	Aggregation weight for client i
$\beta_1, \beta_2, \beta_3$	Weights for validation gain, data volume, and drift in the contribution score
$\varphi_1, \varphi_2, \varphi_3, \varphi_4$	Weights for contribution, cost, drift, and fairness in aggregation
Φ	Potential function of the selection game

2. Problem Description and Preliminaries

2.1. Heterogeneous IoT Object Detection

Consider an IoT vision network with N clients. Client i owns a private dataset $\mathcal{D}_i = \{(x_{ij}, y_{ij})\}_{j=1}^{n_i}$, where x_{ij} is an image and y_{ij} contains object classes and bounding boxes. Let $\theta \in \mathbb{R}^d$ denote the detector parameters. For both anchor-based and query-based detectors, the local risk can be expressed as

$$F_i(\theta) = \mathbb{E}_{(x,y) \sim \mathcal{D}_i} [\ell_{cls}(\theta; x, y) + \lambda_b \ell_{box}(\theta; x, y) + \lambda_o \ell_{obj}(\theta; x, y)]. \quad (1)$$

The global objective is

$$\min_{\theta} F(\theta) = \sum_{i=1}^N p_i F_i(\theta), \quad p_i = \frac{n_i}{\sum_{j=1}^N n_j}. \quad (2)$$

Unlike centralized detection, the server never accesses $\mathcal{D}_i$; it receives only selected model updates. The training process is further constrained by energy, bandwidth, latency, and strategic client behavior.

2.2. Detection Loss Decomposition

To establish the notation, we detail the detection loss used in the framework. Let $\mathcal{A}$ denote the set of anchors or decoder queries. For an image x , the detector outputs class probability p_a , objectness score o_a , and bounding box b_a for each $a \in \mathcal{A}$. The local empirical loss reads as

$$\hat{F}_i(\theta) = \frac{1}{n_i} \sum_{j=1}^{n_i} \sum_{a \in \mathcal{A}} \left[\ell_{ce}(p_{a,j}, y_{a,j}) + \lambda_o \ell_{bce}(o_{a,j}, \hat{o}_{a,j}) + \lambda_b (1 - \text{GIoU}(b_{a,j}, \hat{b}_{a,j})) \right]. \quad (3)$$

The three components respond differently to non-IID data: classification is affected by object-class skew, objectness by background ratio and camera placement, and box regression by viewpoint and scale. Consequently, a client may remain useful even when its local loss is large. GO-FedDet therefore estimates contribution via validation improvement and update drift instead of relying solely on the local training loss.

2.3. Rare-Object Weighting

To prevent frequent categories from dominating the aggregation, the server can maintain a public class-prior vector $\bar{\mathbf{p}}^y$. Client i reports a privacy-preserving histogram $\mathbf{h}_i$, optionally with noise. The rare-object score is defined as

$$R_i = \sum_{c=1}^C \frac{h_{ic}}{\sqrt{\bar{p}_c^y + \epsilon}}. \quad (4)$$

The contribution score may then be augmented by $q_i^t \leftarrow q_i^t + \beta_4 R_i$. No image sharing is required; the term simply encourages selection of clients that possess underrepresented classes. This is particularly relevant in industrial inspection and traffic monitoring, where safety-critical defects or rare road objects are inherently sparse.

2.4. Resource and Behavior Model

At round t , client i has available energy e_i^t , channel rate b_i^t , computation frequency f_i^t , and local batch complexity κ_i . The normalized cost is defined as

$$c_i^t = \omega_e \frac{\kappa_i E}{e_i^t} + \omega_b \frac{|\Delta_i^t|}{b_i^t} + \omega_\tau \tau_i^t, \quad (5)$$

where $\omega_e, \omega_b, \omega_\tau \geq 0$ and $\omega_e + \omega_b + \omega_\tau = 1$. The server estimates detection contribution as

$$q_i^t = \beta_1 \max\{0, V(\theta_t) - V(\theta_t + \Delta_i^t)\} + \beta_2 \frac{n_i}{n_{\max}} - \beta_3 d_i^t, \quad (6)$$

where θ_t is the current global detector; Δ_i^t is the update estimated or returned by client i ; $V(\cdot)$ denotes a small public validation (or proxy) risk; and β_1, β_2 , and β_3 are nonnegative coefficients for validation improvement, normalized data volume, and update drift, respectively. The drift indicator is $d_i^t = \|\Delta_i^t\| / (\|\theta_t\| + \epsilon)$. Historical participation fairness is captured by

$$\rho_i^t = \bar{s}^t - s_i^t, \quad \bar{s}^t = \frac{1}{N} \sum_{j=1}^N s_j^t, \quad (7)$$

where s_i^t is the cumulative selection frequency of client i . A positive ρ_i^t indicates that client i has been underselected relative to the average.

2.5. Federated Optimization Preliminaries

Classical FedAvg aggregates local updates by data proportion, and FedProx adds a proximal term to mitigate client drift, while SCAFFOLD uses control variates for drift correction [28]. Adaptive and personalized FL methods further improve optimization under non-IID data [29]. Objective correction, matched averaging, and flat-minima regularization provide additional tools for handling statistical heterogeneity and improving generalization [30]. For object detection, however, the loss landscape is more structured, and the server must also account for the higher cost of visual models such as YOLO, DETR, and their variants [31].

2.6. Heterogeneity Metrics

To quantify the difficulty of federated detection, we define three heterogeneity indicators. The label-skew index is

$$H_y = \frac{1}{N} \sum_{i=1}^N \left\| \mathbf{p}_i^y - \bar{\mathbf{p}}^y \right\|_1, \quad (8)$$

where $\mathbf{p}_i^y$ is the empirical object-class distribution of client i . The scene-skew index is

$$H_x = \frac{1}{N} \sum_{i=1}^N \text{MMD}^2(\mathcal{D}_i^x, \bar{\mathcal{D}}^x), \quad (9)$$

where maximum mean discrepancy captures feature-distribution mismatch. The resource-skew index is

$$H_r = \frac{\sqrt{\frac{1}{N} \sum_i (c_i^t - \bar{c}^t)^2}}{\bar{c}^t + \epsilon}. \quad (10)$$

Those metrics are estimated from local statistics or update metadata and do not reveal extra private information. They allow the server to decide whether the current round should emphasize accuracy, cost, or fairness.

2.7. Complex Network Interpretation

The IoT system can be viewed as a time-varying graph $\mathcal{G}_t = (\mathcal{V}, \mathcal{E}_t)$, whose vertices include clients, edge servers, and the cloud. An edge exists when a client can deliver an update within the delay bound. The weighted adjacency matrix $W_t = [w_{ij}^t]$ depends on the link rate, interference, and route congestion. Client selection thus acts as a control input that changes the active subgraph, and GO-FedDet optimizes learning over this subgraph. This interpretation aligns with the Special Issue theme of dynamics, control, and optimization in complex networks.

Remark 1. This section illustrates that heterogeneous IoT detection cannot be reduced to a direct reuse of standard FL. The decision variables include model parameters, participation actions, rewards, and resource allocation. The resulting system is a coupled learning-control problem over a complex network, motivating a game-theoretic optimization framework rather than a purely statistical aggregation rule.

3. Model and Analysis

3.1. Stackelberg Participation Game

In the Stackelberg formulation, the server acts as the leader and publishes a reward vector $\mathbf{r}^t = (r_1^t, \dots, r_N^t)$ subject to $\sum_i r_i^t \leq B_t$. Each client i , as a follower, chooses a binary action $a_i^t \in \{0, 1\}$ (participate or not). Client utility is

$$u_i^t(a_i^t; \mathbf{r}^t) = a_i^t(r_i^t q_i^t - \chi c_i^t + \psi \rho_i^t) - \frac{\gamma}{2}(a_i^t - s_i^t)^2, \quad (11)$$

where χ converts the normalized resource cost into client utility, ψ controls the compensation for underselected clients, and γ penalizes abrupt deviation from the historical participation level. The server utility is

$$U_s^t(\mathbf{a}^t, \mathbf{r}^t) = \sum_{i=1}^N a_i^t (q_i^t - \lambda_c c_i^t + \lambda_f \rho_i^t) - \lambda_r \sum_{i=1}^N r_i^t \quad (12)$$

Client i finds participation individually rational when

$$r_i^t q_i^t - \chi c_i^t + \psi \rho_i^t - \frac{\gamma}{2} (1 - s_i^t)^2 + \frac{\gamma}{2} (s_i^t)^2 \geq 0. \quad (13)$$

This condition yields a minimum reward threshold

$$r_{i,\min}^t = \frac{\chi c_i^t - \psi \rho_i^t + \frac{\gamma}{2} (1 - 2s_i^t)}{q_i^t + \epsilon}. \quad (14)$$

Thus, a rational server compensates a client only when its expected detection gain outweighs the combined cost and fairness deficit.

3.2. Potential Game for Client Selection

After reward screening, the set of feasible clients is $\mathcal{C}_t = \{i : r_i^t \geq r_{i,\min}^t\}$. Because simultaneous participation is constrained by bandwidth and delay, each feasible client selects a resource channel $k \in \mathcal{K} \cup \{0\}$, where 0 means no participation. The corresponding payoff is

$$\pi_i^t(k_i, \mathbf{k}_{-i}) = g_i^t(k_i) - h_i^t(k_i) n_{k_i}(\mathbf{k}) - \xi \mathbb{I}\{\tau_i^t(k_i) > \tau_{\max}\}, \quad (15)$$

where $g_i^t = q_i^t - \lambda_c c_i^t + \lambda_f \rho_i^t$, h_i^t is a congestion sensitivity parameter, while n_k the number of clients on channel k . The potential function

$$\Phi(\mathbf{k}) = \sum_{i:k_i \neq 0} g_i^t(k_i) - \sum_{k \in \mathcal{K}} \sum_{m=1}^{n_k(\mathbf{k})} \bar{h}_{k,m}^t - \xi \sum_i \mathbb{I}\{\tau_i^t(k_i) > \tau_{\max}\} \quad (16)$$

which makes the game an exact potential game whenever clients sharing a channel face identical congestion prices $\bar{h}_{k,m}^t$. Under this condition, any unilateral payoff change equals the corresponding change in Φ ; therefore, the distributed best-response dynamics converge to a pure Nash equilibrium in finitely many steps.

3.3. Server Reward Optimization

The server can adjust its reward vector through a projected primal-dual procedure. Define the Lagrangian

$$\mathcal{L}_s(\mathbf{r}, \lambda) = -U_s^t(\mathbf{a}(\mathbf{r}), \mathbf{r}) + \lambda \left(\sum_i r_i - B_t \right). \quad (17)$$

A practical update employs projected gradients:

$$r_i^{t+1} = \Pi_{[0, r_{\max}]} \left[r_i^t - \eta_r \left(\frac{\partial \mathcal{L}_s}{\partial r_i} + \lambda^t \right) \right], \quad (18)$$

$$\lambda^{t+1} = \left[\lambda^t + \eta_\lambda \left(\sum_i r_i^{t+1} - B_t \right) \right]_+. \quad (19)$$

Since followers' best responses are discontinuous at participation thresholds, the server uses a smoothed response $\hat{a}_i^t = \sigma((r_i^t - r_{i,\min}^t)/\nu)$ during reward adaptation. As $\nu \rightarrow 0$, the smoothed response recovers the binary decision. This design avoids oscillations when clients lie near their thresholds.

3.4. Fairness-Controlled Aggregation

The fairness term in the aggregation weight (24) can be viewed as a Lyapunov control mechanism. Introduce a virtual queue $Q_i^t = [Q_i^{t-1} + \bar{s} - a_i^t]_+$, where $\bar{s}$ denotes a target participation frequency. Setting $\rho_i^t = Q_i^t / (1 + \sum_j Q_j^t)$ gives higher priority to historically underselected clients, but only when their updates still contribute positively. Fairness therefore avoids blindly including poor updates; it acts as a stabilizer that prevents the long-term exclusion of useful minority clients.

3.5. Dynamic Control View of Participation

The participation frequency evolves as

$$s_i^{t+1} = (1 - \omega_s)s_i^t + \omega_s a_i^t, \quad (20)$$

where the forgetting factor $\omega_s \in (0, 1]$. The fairness deficit ρ_i^t thus provides feedback from the participation state to the game utility. The overall closed-loop dynamics can be written as

$$(\theta_{t+1}, \mathbf{s}^{t+1}) = \mathcal{T}(\theta_t, \mathbf{s}^t, \mathbf{c}^t, \mathbf{q}^t, \mathbf{b}^t), \quad (21)$$

where $\mathcal{T}$ encapsulates reward thresholding, potential-game selection, local optimization, and aggregation. This perspective shows that learning and participation co-evolve. Removing fairness feedback may cause the fixed point to concentrate on a small subset of clients; excessive reward feedback may induce threshold oscillations. GO-FedDet balances these two effects through smoothed reward adaptation and virtual fairness queues.

3.6. Connection to Social Welfare

Round- t social welfare is

$$W_t = U_s^t + \sum_{i=1}^N u_i^t. \quad (22)$$

After internal reward transfers between server and clients are canceled, welfare depends primarily on detection gain, resource cost, and fairness. Hence, the game does not merely maximize the server's short-term utility; it approximates a welfare-aware learning policy. This property is essential in multi-stakeholder IoT systems, where long-term collaboration requires that participating clients receive reasonable benefit.

3.7. Detection-Oriented Proximal Objective

Each selected client i solves a regularized local problem

$$\min_{\theta} \tilde{F}_i^t(\theta) = F_i(\theta) + \frac{\mu}{2} \|\theta - \theta_t\|^2 + \lambda_d \mathcal{D}_{KD}(z_i(\theta), z_g(\theta_t)), \quad (23)$$

where $\mathcal{D}_{KD}$ is a lightweight logit-distillation term computed on local unlabeled images or a public proxy set. This term stabilizes objectness and classification logits across clients. The aggregation weight is

$$\alpha_i^t = \frac{\exp\{\varphi_1 q_i^t - \varphi_2 c_i^t - \varphi_3 d_i^t + \varphi_4 \rho_i^t\}}{\sum_{j \in \mathcal{S}_t} \exp\{\varphi_1 q_j^t - \varphi_2 c_j^t - \varphi_3 d_j^t + \varphi_4 \rho_j^t\}}, \quad (24)$$

where φ_1 , φ_2 , φ_3 , and φ_4 determine the relative importance of validation contribution, communication-energy cost, update drift, and long-term fairness, respectively. The server updates the global detector as $\theta_{t+1} = \theta_t + \sum_{i \in \mathcal{S}_t} \alpha_i^t \Delta_i^t$.

Remark 2. The model separates *who* participates (via game theory) from *how* each participant updates (via proximal detection learning). The game layer handles rational behavior and resource

coupling, and the learning layer addresses non-IID detection drift. Their interaction forms a closed-loop optimization mechanism.

4. Algorithm Design

4.1. Design Motivation

A plain federated detector is likely to repeatedly select clients with strong hardware or large datasets. While this can reduce the loss in the short term, it tends to ignore rare scenes, harms fairness, and introduces instability when low-resource clients occasionally upload large, drifting updates. GO-FedDet counteracts these tendencies through three intertwined mechanisms: reward thresholding, potential-game selection, and utility-aligned aggregation.

4.2. Overall Workflow

The procedure can be summarized as “screen–compete–train–aggregate”. The server first screens clients using the Stackelberg reward thresholds. Feasible clients then compete for communication resources via a finite potential game. The selected clients perform proximal detection training, after which the server aggregates their updates by trading off detection contribution, cost, drift, and fairness. Table 2 provides the compact pseudo-code.

Table 2. Pseudo-code of GO-FedDet.

Step	Operation
Input	Initial detector θ_0 , client set $\mathcal{N}$, total rounds T , budget B_t , local epochs E , learning rate η_l , proximal coefficient μ .
1	For $t = 0, 1, \dots, T - 1$, the server broadcasts θ_t , B_t , a validation proxy, and the reward rule.
2	Each client computes its local cost c_i^t , drift proxy d_i^t , contribution q_i^t , and fairness deficit ρ_i^t .
3	The server calculates $r_{i,\min}^t$ via (14) and forms the feasible set $\mathcal{C}_t$.
4	Feasible clients play the potential resource game; each updates its channel action by best response until no unilateral payoff improvement remains.
5	The server obtains the selected set $\mathcal{S}_t$ from the Nash-stable action profile and sends local training instructions.
6	Each selected client minimizes (23) for E epochs and returns the compressed update Δ_i^t .
7	The server computes weights α_i^t by (24) and updates $\theta_{t+1} = \theta_t + \sum_{i \in \mathcal{S}_t} \alpha_i^t \Delta_i^t$.
8	The server updates cumulative participation s_i^{t+1} , fairness deficits ρ_i^{t+1} , and the reward budget for the next round.
Output	Final detector θ_T for IoT object detection.

4.3. Implementation Notes

The framework is detector-agnostic and can accommodate YOLO-style one-stage architectures as well as transformer-based detectors. For resource-constrained IoT deployment, the backbone may be frozen during early rounds and gradually unfrozen later. The communication payload can consist of full model weights, low-rank adapters, or only selected detection heads [32]. The reward variables need not represent monetary value; they can be interpreted as access priority, battery credits, bandwidth tokens, or edge-service quotas.

4.4. Parameter Tuning Advice

The reward coefficient should be sufficiently large to incentivize clients whose marginal detection contribution is high despite elevated resource cost. A practical starting point is to set χ to the median of the ratios q_i^t / c_i^t . The fairness coefficient ψ is best kept small during early rounds, when the detector requires rapid convergence, and increased once the validation loss stabilizes. The proximal coefficient μ governs the trade-off between local adaptation and global consistency. A useful adaptive rule sets $\mu_t = \mu_0(1 + \text{median}_i d_i^t)$, so that stronger regularization is applied when drift is large.

In the final experiments, the hyperparameters were selected on the validation proxy and then fixed for all test comparisons. Unless otherwise stated, we used $\mu = 0.01$, $\lambda_d = 0.1$,

$\chi = 0.5$, $\psi = 0.5$, $\gamma = 0.1$, $\lambda_c = 0.3$, $\lambda_f = 0.5$, $\lambda_r = 0.01$, $\omega_e = 0.4$, $\omega_b = 0.4$, $\omega_\tau = 0.2$, $(\beta_1, \beta_2, \beta_3, \beta_4) = (1.0, 0.3, 0.2, 0.1)$, and $(\varphi_1, \varphi_2, \varphi_3, \varphi_4) = (1.0, 0.5, 0.5, 0.3)$. These values were obtained by a coarse grid search around the tuning rules above; the test set was not used for parameter selection.

4.5. Privacy and Security Aspects

GO-FedDet never transmits raw images or bounding-box annotations; however, model updates can still leak information through gradient inversion attacks. The framework is compatible with standard safeguards such as secure aggregation, differential privacy, and update clipping [33]. From a game-theoretic standpoint, privacy protection can be introduced as an additional cost term. A client with stringent privacy requirements incurs a higher effective cost, and the reward threshold automatically compensates for this cost or excludes the client if participation is no longer socially efficient.

Remark 3. *The algorithm is designed for practical deployment rather than solely for theoretical completeness. Every decision rests on quantities readily available in FL systems: validation loss, update norm, channel state, energy level, and participation history. Consequently, GO-FedDet can be implemented on edge servers without ever collecting raw images from IoT clients.*

5. Algorithm Analysis

5.1. Equilibrium Existence

Theorem 1 (Stackelberg Equilibrium). *For each round t , suppose $q_i^t > 0$, $c_i^t < \infty$, and B_t is finite. Then the Stackelberg participation game possesses at least one pure-strategy equilibrium. Moreover, the server admits an optimal reward allocation that does not overpay: for every selected client, $r_i^t = r_{i,\min}^t$ holds at an optimum.*

Proof. Fix a reward vector $\mathbf{r}^t$. Each client has a binary action set $\{0, 1\}$ and a real-valued utility (11); hence, the best-response set is nonempty. The participation condition reduces to a threshold rule, and so the follower response correspondence maps any reward to at least one binary profile. The server's feasible set $\mathcal{R}_t = \{\mathbf{r} : r_i \geq 0, \sum_i r_i \leq B_t\}$ is compact. The server utility is piecewise continuous over $\mathcal{R}_t$ because follower actions change only at finitely many thresholds. On each region where the response profile stays fixed, the server utility is linear in $\mathbf{r}$. Therefore, a maximum exists on the closure of each region, and the global maximum over finitely many regions exists as well. If a selected client receives $r_i > r_{i,\min}$, reducing its reward to $r_{i,\min}$ leaves the follower action unchanged while weakly increasing server utility (because $\lambda_r \geq 0$). Hence, an optimal allocation avoids overpaying marginal participants. Pairing the optimal leader action with the induced best responses yields a Stackelberg equilibrium. Thus, the proof of Theorem 1 is completed. $\square$

5.2. Finite-Improvement Convergence

Theorem 2 (Potential-Game Convergence). *When clients sharing the same channel face identical congestion prices, the client selection subgame is an exact-potential game with potential Φ in (16). Any asynchronous best-response sequence converges to a pure Nash equilibrium in finitely many steps.*

Proof. Consider a unilateral deviation of client i from k_i to k'_i , holding other actions fixed. The change in the first summation of Φ exactly matches the change in g_i^t . For the channel being left, the congestion sum drops the last marginal congestion term; for the channel being joined, it adds the new marginal term. Since clients in the same channel share the same congestion price sequence, the net congestion change equals the congestion component of $\pi_i^t(k'_i, \mathbf{k}_{-i}) - \pi_i^t(k_i, \mathbf{k}_{-i})$. The delay penalty is individual and changes identically in

payoff and potential. Thus, the unilateral payoff difference coincides with the potential difference. The action space is finite. Every strict best response strictly increases Φ , which is bounded above on a finite domain. Consequently the process terminates at a profile with no profitable unilateral deviation—a pure Nash equilibrium. Thus, the proof of Theorem 2 is completed. $\square$

5.3. Non-Convex Convergence Bound

Assumption 1. Each F_i is L -smooth and bounded below. The stochastic gradient $g_i(\theta)$ satisfies $\mathbb{E}[g_i(\theta)] = \nabla F_i(\theta)$ and $\mathbb{E}\|g_i(\theta) - \nabla F_i(\theta)\|^2 \leq \sigma^2$. Client heterogeneity satisfies $\sum_i p_i \|\nabla F_i(\theta) - \nabla F(\theta)\|^2 \leq \zeta^2$.

Theorem 3 (Convergence of GO-FedDet). Let $\eta_l \leq 1/(4LE)$ and $\mu \geq 0$. Suppose the aggregation weights satisfy $\sum_{i \in \mathcal{S}_t} \alpha_i^t = 1$ and $\alpha_i^t \geq 0$. Then after T rounds,

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \|\nabla F(\theta_t)\|^2 \leq \frac{2(F(\theta_0) - F^*)}{\eta_l ET} + C_1 \eta_l L \sigma^2 + C_2 \eta_l^2 E^2 L^2 \zeta^2 + C_3 \bar{\delta}, \quad (25)$$

where $\bar{\delta} = T^{-1} \sum_t \mathbb{E} \|\sum_{i \in \mathcal{S}_t} (\alpha_i^t - p_i) \nabla F_i(\theta_t)\|^2$, and C_1, C_2, C_3 are positive constants.

Proof. By L -smoothness,

$$F(\theta_{t+1}) \leq F(\theta_t) + \langle \nabla F(\theta_t), \theta_{t+1} - \theta_t \rangle + \frac{L}{2} \|\theta_{t+1} - \theta_t\|^2. \quad (26)$$

Define $G_t = \sum_{i \in \mathcal{S}_t} \alpha_i^t G_i^t$, where G_i^t accumulates the local stochastic directions over E steps. Then $\theta_{t+1} = \theta_t - \eta_l G_t$ up to the proximal correction; the correction only shrinks the distance from θ_t and can be absorbed into the constants when $\eta_l \mu E$ is bounded. Taking conditional expectation,

$$\mathbb{E}_t F(\theta_{t+1}) \leq F(\theta_t) - \eta_l \langle \nabla F(\theta_t), \mathbb{E}_t G_t \rangle + \frac{L \eta_l^2}{2} \mathbb{E}_t \|G_t\|^2. \quad (27)$$

Now, subtract $\nabla F(\theta_t)$ in the inner product:

$$\langle \nabla F, \mathbb{E}_t G_t \rangle = \|\nabla F\|^2 + \langle \nabla F, \mathbb{E}_t G_t - \nabla F \rangle. \quad (28)$$

Using Young's inequality, the second term has the lower bound $-\frac{1}{4} \|\nabla F\|^2 - \|\mathbb{E}_t G_t - \nabla F\|^2$. The bias $\mathbb{E}_t G_t - \nabla F$ contains two sources: local-drift bias from E local steps, bounded by $O(\eta_l^2 E^2 L^2 \zeta^2)$, and selection bias $\delta_t = \|\sum_{i \in \mathcal{S}_t} (\alpha_i^t - p_i) \nabla F_i(\theta_t)\|^2$ due to game-based partial participation. The variance term is bounded by $C\sigma^2$ because the weights are nonnegative and sum to one. Consequently,

$$\mathbb{E}_t F(\theta_{t+1}) \leq F(\theta_t) - \frac{\eta_l E}{2} \|\nabla F(\theta_t)\|^2 + C \eta_l^2 E L \sigma^2 + C \eta_l^3 E^3 L^2 \zeta^2 + C \eta_l E \delta_t. \quad (29)$$

Summing over $t = 0, \dots, T-1$, telescoping the left side, dividing by $\eta_l ET/2$, and using $F(\theta_T) \geq F^*$ yields (25). Thus, the proof of Theorem 3 is completed. $\square$

Corollary 1. If $\eta_l = \mathcal{O}(1/\sqrt{T})$, E is bounded and the game-induced bias $\bar{\delta}$ is controlled by fairness-aware selection, then GO-FedDet reaches an $\mathcal{O}(1/\sqrt{T})$ stationary-point rate, up to heterogeneity and selection-bias terms.

5.4. Proof of Bias Decomposition

We further expand the bias term used in Theorem 3. Let

$$B_t = \mathbb{E}_t G_t - \nabla F(\theta_t). \quad (30)$$

It decomposes as

$$B_t = \underbrace{\sum_{i \in \mathcal{S}_t} \alpha_i^t (\nabla F_i(\theta_{i,e}^t) - \nabla F_i(\theta_t))}_{B_t^{\text{loc}}} + \underbrace{\sum_{i \in \mathcal{S}_t} (\alpha_i^t - p_i) \nabla F_i(\theta_t)}_{B_t^{\text{sel}}}, \quad (31)$$

where $\theta_{i,e}^t$ denotes an intermediate local iterate. By smoothness,

$$\|B_t^{\text{loc}}\|^2 \leq L^2 \sum_{i \in \mathcal{S}_t} \alpha_i^t \|\theta_{i,e}^t - \theta_t\|^2. \quad (32)$$

The local distance is bounded by the accumulated stochastic gradients, giving $\mathcal{O}(\eta_l^2 E^2 (\sigma^2 + \zeta^2))$. The second term, B_t^{sel} , is precisely the selection bias introduced by game-based partial participation. The proposed aggregation weights reduce this bias by penalizing high drift and high cost while rewarding contribution and fairness. This decomposition separates the optimization error from the strategic selection error.

5.5. Stability Under Bounded Reward Perturbation

Suppose the reward vector is perturbed by at most ϵ_r due to noisy contribution estimates. If $|r_i^t - r_{i,\min}^t| > \epsilon_r$ for every client, the follower action profile remains unchanged; only clients within the threshold band may switch actions. Hence, the selected set satisfies

$$|\mathcal{S}_t(\mathbf{r} + \Delta \mathbf{r}) \triangle \mathcal{S}_t(\mathbf{r})| \leq \sum_i \mathbb{I}\{|r_i^t - r_{i,\min}^t| \leq \epsilon_r\}. \quad (33)$$

This explains why smoothed response and fairness queues reduce oscillation: they enlarge the margin between stable participants and nonparticipants, making the selection dynamics less sensitive to noisy validation estimates.

5.6. Robustness to Client Dropout

In practical IoT systems, selected clients may fail before uploading their updates. Let $m_i^t \in \{0, 1\}$ indicate successful return. The aggregation rule becomes

$$\theta_{t+1} = \theta_t + \sum_{i \in \mathcal{S}_t} \tilde{\alpha}_i^t m_i^t \Delta_i^t, \quad \tilde{\alpha}_i^t = \frac{m_i^t \alpha_i^t}{\sum_{j \in \mathcal{S}_t} m_j^t \alpha_j^t + \epsilon}. \quad (34)$$

If the dropout probability is bounded by $p_d < 1$ and independent of the stochastic gradients, the variance term in Theorem 3 is inflated by at most a factor proportional to $(1 - p_d)^{-1}$. The selection game can mitigate dropout by incorporating delay and channel reliability into c_i^t , thereby achieving robustness at both the decision level and the aggregation level.

5.7. Communication Compression

GO-FedDet can be combined with update sparsification. Let $\mathcal{Q}_s(\Delta)$ be an unbiased compressor satisfying

$$\mathbb{E} \mathcal{Q}_s(\Delta) = \Delta, \quad \mathbb{E} \|\mathcal{Q}_s(\Delta) - \Delta\|^2 \leq \omega_q \|\Delta\|^2. \quad (35)$$

Then compression merely adds $\omega_q \mathbb{E} \|\Delta\|^2$ to the variance term. The game layer naturally favors compressed updates when bandwidth is scarce, because a smaller upload size reduces c_i^t . This makes the framework compatible with low-bandwidth IoT links.

5.8. Fairness Queue Boundedness

Theorem 4 (Bounded Participation Deficit). *Assume that in every window of W rounds each client has at least one feasible round in which $q_i^t > 0$ and $c_i^t < \infty$. If the fairness coefficient satisfies $\lambda_f > \lambda_c \max_i c_i^t + \xi$, then the virtual fairness queue Q_i^t is uniformly bounded by a constant depending on W , λ_f , and the maximum cost.*

Proof. Consider the Lyapunov function $L_Q(t) = \frac{1}{2} \sum_i (Q_i^t)^2$. From the queue update,

$$(Q_i^{t+1})^2 \leq (Q_i^t)^2 + (\bar{s} - a_i^t)^2 + 2Q_i^t(\bar{s} - a_i^t). \quad (36)$$

Summing over clients gives a one-step drift bound

$$\Delta_Q(t) \leq C_Q + \sum_i Q_i^t \bar{s} - \sum_i Q_i^t a_i^t, \quad (37)$$

where C_Q is finite because $a_i^t \in \{0, 1\}$. When a client's queue becomes large, its fairness deficit increases its payoff through the term $\lambda_f \rho_i^t$. Under the stated coefficient condition, the fairness gain outweighs the maximum resource and congestion penalty in its feasible round. Thus, a best response selects this client whenever the queue exceeds a finite threshold. In any length- W window, an excessively large queue must be served at least once, creating a negative drift outside a compact set. By the Foster–Lyapunov argument for deterministic bounded-increment queues, Q_i^t is uniformly bounded. Hence, participation deficit cannot diverge. Thus, the proof of Theorem 4 is completed. $\square$

5.9. Interpretation of the Convergence Bound

The term $C_1 \eta_l L \sigma^2$ in (25) captures the stochastic gradient noise; it can be reduced by larger local batches or variance-reduction techniques. The term $C_2 \eta_l^2 E^2 L^2 \zeta^2$ represents the local-drift error, which grows with more local epochs and stronger data heterogeneity. The term $C_3 \bar{\delta}$ is the cost of strategic partial participation and is the most distinctive part of the bound. A purely statistical FL analysis often hides this factor inside sampling noise, whereas here it appears explicitly, showing how the game design directly influences learning convergence.

5.10. Complexity Analysis

For each round, reward screening costs $\mathcal{O}(N)$. Potential-game best-response updates cost $\mathcal{O}(IMK)$, where I is the number of improvement iterations, $M = |\mathcal{C}_t|$, and K is the number of channels. Local training dominates computation at $\mathcal{O}(|\mathcal{S}_t| E C_{\text{det}})$, with C_{det} being the per-epoch detector training cost. Server aggregation costs $\mathcal{O}(|\mathcal{S}_t| d)$. Compared with standard FedAvg, the additional decision overhead is negligible when $K \ll d$.

We also recorded the practical runtime overhead in the simulation. On the hardware specified in Section 6.1, FedAvg required about $1.00 \times$ normalized time per round, FedProx required $1.04 \times$ because of the proximal term, SCAFFOLD required $1.11 \times$ because of control-variate communication and update bookkeeping, and GO-FedDet required $1.07 \times$. The extra time needed for GO-FedDet mainly came from reward screening and best-response resource selection, which took less than one second per round for 40 clients and four channels. Since detector training and upload dominate the wall-clock time, the game layer introduces modest overhead while reducing the number of inefficient client uploads.

Remark 4. *The analysis shows why GO-FedDet is both stable and scalable. Equilibrium results guarantee that the selection process terminates. The convergence bound identifies three performance drivers: stochastic noise, local heterogeneity, and game-induced selection bias. This decomposition also provides clear guidance for parameter tuning in simulation and deployment.*

6. Simulation and Analysis

6.1. Experimental Setup

We evaluate GO-FedDet in a simulated heterogeneous IoT object-detection environment built on the PASCAL VOC 2007 + 2012 detection dataset. The dataset contains 20 object categories. Images are resized to 640×640 , and the original training images are partitioned over 40 virtual IoT cameras. The test split is kept centralized only for evaluation. To simulate non-IID IoT scenes, object-class proportions are sampled by a Dirichlet distribution and scene features are clustered before client assignment; the default label-skew factor is 0.3. Image quantities are also heterogeneous, with each client holding 150–800 images.

The detector is a lightweight YOLO-style one-stage model with a CSP-like backbone and three detection heads. The model is initialized from a public pretrained checkpoint and then trained federatively for 150 communication rounds. Unless stated otherwise, 25% of the clients are selected in each round, each selected client trains for five local epochs, the local batch size is 16, and the optimizer is AdamW with learning rate 1×10^{-3} and weight decay 5×10^{-4} . The public validation proxy used for contribution estimation consists of 5% class-balanced public images and is not used for testing. All experiments are repeated with five independent random seeds, $\{2026, 2027, 2028, 2029, 2030\}$, which change the client partition, resource sampling, and model initialization. We report the mean and standard deviation unless otherwise specified.

The wireless channel is modeled by $b_i^t = B_i \log_2(1 + P_i g_i^t / (N_0 B_i))$, where B_i is sampled from 2 to 10 MHz, transmit power is 20 dBm, noise density is -174 dBm/Hz, and g_i^t includes distance-dependent path loss with exponent 3.5 and small-scale Rayleigh fading. The transmission delay is $|\Delta_i^t|/b_i^t$, and the energy consumption is computed as the sum of transmission energy and local computation energy, $E_i^t = P_i |\Delta_i^t|/b_i^t + \kappa_i E f_i^2$. The hardware platform is an Ubuntu workstation equipped with an Intel Xeon Silver CPU, 128 GB memory, and one NVIDIA RTX 4090 GPU. Performance is measured by mAP@0.5, mAP@0.5:0.95, normalized communication-energy cost, energy consumption, delay, and Jain's fairness index [34].

In addition to classical FL baselines, we add more directly related comparison methods to examine client selection, incentive, resource awareness, and federated detection. The compared methods are Random-FL, FedAvg [35], FedProx [36], SCAFFOLD [37], FedNova [38], Oort-style utility client selection, Incentive-FL, Resource-FL, and FedDet-style federated object detection. All methods use the same detector, data partitions, training rounds, and client budget.

Figure 2 depicts how participation and fairness evolve under the game layer. The participation ratio climbs steeply in early rounds because the reward threshold eliminates clients whose detection gain cannot offset their cost. After roughly twenty rounds, the curve stabilizes, indicating that the Stackelberg incentive and potential-game selection have reached a balanced regime. The fairness index also improves, since underselected clients accumulate positive fairness deficits and are eventually chosen. This behavior is valuable for IoT detection, where rare scenes often reside on low-resource nodes.




Figure 2. Participation and fairness dynamics.

Table 3 summarizes the overall detection results. GO-FedDet attains the highest mAP values together with the lowest normalized communication-energy cost. Compared with the strongest task-related baseline, FedDet, GO-FedDet improves mAP@0.5 from 51.1% to 52.0% and reduces the normalized cost from 0.92 to 0.74. The improvement does not merely stem from favoring strong clients; the fairness value rises from a range of 0.71–0.82 for the baselines to 0.88 for GO-FedDet, reflecting broader and more balanced client participation. A paired *t*-test over the five runs indicates that the mAP@0.5 improvement over the strongest baseline is statistically significant at the 0.05 level.

Table 3. Overall detection performance over five independent runs.

Method	mAP@0.5	mAP@0.5:0.95	Cost	Fairness
Random-FL	45.6 ± 0.5	27.8 ± 0.4	1.00 ± 0.02	0.71 ± 0.02
FedAvg	47.9 ± 0.5	29.4 ± 0.4	0.96 ± 0.02	0.73 ± 0.02
FedProx	49.1 ± 0.4	30.6 ± 0.3	0.94 ± 0.02	0.76 ± 0.02
SCAFFOLD	49.7 ± 0.4	31.0 ± 0.3	0.98 ± 0.03	0.78 ± 0.02
FedNova	50.2 ± 0.4	31.4 ± 0.3	0.93 ± 0.02	0.79 ± 0.02
Oort-FL	50.5 ± 0.4	31.7 ± 0.3	0.88 ± 0.02	0.77 ± 0.02
Incentive-FL	50.8 ± 0.3	31.9 ± 0.3	0.86 ± 0.02	0.82 ± 0.02
Resource-FL	50.9 ± 0.4	32.0 ± 0.3	0.80 ± 0.02	0.80 ± 0.02
FedDet	51.1 ± 0.3	32.3 ± 0.3	0.92 ± 0.02	0.77 ± 0.02
GO-FedDet	52.0 ± 0.3	33.2 ± 0.2	0.74 ± 0.01	0.88 ± 0.01

Note: Bold values indicate the best result in each comparison.

Figure 3 compares the training loss curves. FedAvg decreases quickly at the beginning but later becomes unstable because non-IID detection gradients cause pronounced local drift. FedProx dampens this instability with a proximal term, yet it still lacks selection incentives. GO-FedDet achieves both the lowest final loss and the smoothest trajectory, corroborating Theorem 3: when selection bias and local drift are jointly controlled, the expected stationarity gap shrinks more steadily. The shaded bands in the revised figure denote the standard deviation over five runs.




Figure 3. Training loss convergence comparison.

Table 4 examines varying non-IID severity. As label skew intensifies, all methods degrade because local clients observe fewer object categories and more biased backgrounds. However, GO-FedDet degrades less than FedAvg and FedProx, indicating that contribution-aware selection and fairness compensation can recover information from diverse clients. This supports the argument that heterogeneous IoT detection should not rely only on data-size-weighted aggregation.

Table 4. Impact of non-IID severity over five independent runs.

Label Skew	FedAvg	FedProx	GO-FedDet
0.1	50.4 ± 0.3	51.0 ± 0.3	52.8 ± 0.2
0.3	48.6 ± 0.4	49.7 ± 0.3	52.0 ± 0.3
0.5	46.2 ± 0.5	47.8 ± 0.4	50.4 ± 0.3
0.7	43.5 ± 0.5	45.1 ± 0.4	48.1 ± 0.4

Note: Bold values indicate the best result in each comparison.

Figure 4 shows the accuracy–cost trade-off across different resource budgets. The mAP increases with a larger communication budget, but the marginal gain gradually saturates, suggesting that blindly adding more clients is inefficient once representative scenes are covered. GO-FedDet uses reward thresholds and congestion-aware selection to locate the useful portion of the curve; hence, the server can reduce cost while preserving most of the detection accuracy.

Table 5 reports resource consumption. SCAFFOLD incurs higher communication cost due to additional control information. FedAvg and FedProx transmit similar model updates and have close upload sizes. GO-FedDet lowers upload volume, energy consumption, and delay because only clients that pass the reward and potential-game filters participate. This result matches the complex-network motivation: learning quality and network resource dynamics need to be optimized jointly.




Figure 4. Accuracy–cost trade-off at different budget levels.

Table 5. Communication and energy consumption.

Method	Upload (MB)	Energy (J)	Delay (s)
FedAvg	1280 ± 21	912 ± 18	184 ± 5
FedProx	1280 ± 20	925 ± 17	189 ± 5
SCAFFOLD	1410 ± 24	960 ± 21	201 ± 6
FedNova	1265 ± 19	904 ± 17	178 ± 5
Resource-FL	1043 ± 18	792 ± 15	154 ± 4
GO-FedDet	946 ± 15	731 ± 13	139 ± 4

Note: Bold values indicate the best result in each comparison.

Figure 5 presents the ablation study. Removing the game layer causes the largest drop in accuracy, confirming that strategic participation is central to the method. Removing the proximal term also hurts performance because non-IID object detection generates substantial local drift. The fairness term and the distillation term bring further gains by protecting rare clients and aligning objectness logits. The complete method performs best because it coordinates incentives, optimization stability, and detection-specific knowledge transfer.




Figure 5. Ablation study of GO-FedDet.

Figure 6 visualizes GO-FedDet under joint statistical and resource heterogeneity. Accuracy declines when label skew becomes stronger or the average resource factor drops, but the decline is gradual rather than abrupt. This indicates that the game layer avoids selecting only high-resource clients, while the proximal term limits harmful drift from low-resource clients. The outcome supports the suitability of GO-FedDet for IoT scenarios that feature dynamic resources and uneven visual distributions.




Figure 6. Robustness for different heterogeneity factors.

Table 6 studies two important parameters. A moderate proximal coefficient yields the best accuracy because it suppresses drift without stifling local adaptation. An overly large value excessively constrains clients and diminishes their ability to learn local objects. The fairness coefficient improves fairness but can slightly increase cost, as more underselected clients are included. These findings suggest that practical parameter choices should depend on whether the deployment prioritizes pure accuracy, cost efficiency, or long-term client diversity.

Table 6. Sensitivity to key parameters.

Setting	mAP@0.5	Cost	Fairness
$\mu = 0.001$	50.7 ± 0.4	0.73 ± 0.02	0.86 ± 0.02
$\mu = 0.010$	52.0 ± 0.3	0.74 ± 0.01	0.88 ± 0.01
$\mu = 0.100$	50.9 ± 0.4	0.75 ± 0.02	0.87 ± 0.02
$\lambda_f = 0$	50.0 ± 0.4	0.72 ± 0.02	0.76 ± 0.02
$\lambda_f = 0.5$	51.4 ± 0.3	0.73 ± 0.02	0.84 ± 0.02

Note: Bold values indicate the best result in each comparison.

6.2. Discussion

The experimental results should be interpreted from a joint learning and control perspective. Standard FL baselines mainly optimize a statistical loss without explicitly considering whether a client has sufficient energy, whether its channel is congested, or whether participation is fair over time. GO-FedDet translates these deployment factors into game utilities and feeds the game outcome back into detector optimization. This closed loop explains why the method improves both accuracy and cost. It also implies that the best deployment setting is task-dependent: safety-critical applications may increase fairness and validation-gain weights, battery-limited monitoring may emphasize energy-cost weights, and delay-sensitive inspection may increase delay penalties.

6.3. Reproducibility Notes

To improve reproducibility, we specify the main implementation details as follows. The dataset is PASCAL VOC 2007 + 2012 with 20 categories; images are resized to 640×640 ; the number of clients is 40; the default selected-client ratio is 25%; the number of global rounds is 150; the number of local epochs is five; the batch size is 16; the optimizer is AdamW; the learning rate is 10^{-3} ; the weight decay is 5×10^{-4} ; the random seeds are 2026–2030; and the hardware is an RTX 4090 GPU workstation. Detection accuracy is reported as mAP, while cost values are normalized to the cost of FedAvg. The main logs required by the algorithm are local validation loss, update norm, upload size, energy estimate, and round delay, all of which are commonly available in edge-training systems and do not reveal raw visual content.

6.4. Comparison with Alternative Design Choices

A natural alternative would replace the Stackelberg game with a simple heuristic such as the ratio q_i^t / c_i^t . This heuristic is simple but overlooks the fact that clients make participation decisions according to their own utilities: a client with a high score may still decline to participate if energy is low or privacy cost is high. Another option is a multi-armed bandit for selection. Bandits are useful for exploration, yet they typically do not model congestion among simultaneously selected clients. GO-FedDet differs by combining rational participation and resource competition. It can thus be viewed as a middle ground between centralized combinatorial optimization and fully selfish distributed participation.

6.5. Practical Deployment Workflow

A deployment can start from a pretrained lightweight detector. During a warm-up stage, the server uses a small random selection ratio to collect stable estimates of cost, delay, and update drift. In the game-training stage, the reward and potential-game modules are activated. During the final fine-tuning stage, the server increases the fairness weight to include rare clients and avoid overfitting to high-resource nodes. This three-stage workflow keeps early training stable while gradually introducing strategic optimization.

Remark 5. *The simulation results validate the main claims from three perspectives. First, the proposed method improves detection accuracy given non-IID data. Second, it reduces communication-energy cost through strategic participation. Third, it improves fairness across clients. These observations are consistent with the equilibrium and convergence analyses developed in the earlier sections.*

6.6. Limitations

The current study still has several limitations. First, although real public object-detection images are used, the edge resource states, wireless channels, and energy budgets are simulated rather than measured on real IoT camera platforms. Second, the training protocol is synchronous; asynchronous participation and straggler-tolerant aggregation are not yet implemented. Third, the Stackelberg formulation assumes rational clients that truthfully evaluate their cost and contribution, whereas practical clients may be noisy, boundedly rational, or strategically misreporting. Fourth, the contribution estimator relies on a small public or proxy validation set, which may not always be available in privacy-sensitive deployments. Finally, the present experiments focus on benign clients and do not fully investigate poisoning, free-riding, label-flipping, or model-replacement attacks. These limitations motivate future real-platform validation, asynchronous federated detection, robust aggregation, and attack-resistant game mechanisms.

7. Conclusions and Future Research Work

This work introduced GO-FedDet, a game-theoretic optimized federated learning framework for heterogeneous IoT object detection. The design formulates participation as a Stackelberg game, client selection as a potential resource game, and local detector training as a proximal non-convex optimization problem. Reward, detection contribution, resource cost, update drift, and long-term fairness are then combined into a single aggregation rule. On the theoretical side, we proved the existence of a Stackelberg equilibrium and the finite-improvement convergence of the selection dynamics and derived a non-convex convergence bound that explicitly accounts for partial participation and client heterogeneity. Simulation results confirmed that GO-FedDet delivers higher detection accuracy, lower communication-energy overhead, and stronger participation fairness than representative federated baselines.

Future work will proceed along four concrete directions. First, we will deploy GO-FedDet on real edge-camera platforms and practical wireless networks to measure latency, memory footprint, battery consumption, and update loss under hardware constraints. Second, the framework will be extended to asynchronous and event-triggered federated detection, where clients may join training rounds at irregular intervals. Third, we will replace the ideal public-validation assumption with privacy-preserving proxy evaluation or cross-client reputation estimation. Fourth, we will investigate robust game-theoretic aggregation mechanisms that resist poisoning, free-riding, model-replacement, and adversarial object-label attacks.

Author Contributions: Conceptualization, Z.W. and J.C.; methodology, Z.W. and J.C.; software, Z.W.; validation, Z.W. and J.C.; formal analysis, Z.W. and J.C.; investigation, Z.W.; resources, Z.W. and J.C.; data curation, Z.W.; writing—original draft preparation, Z.W.; writing—review and editing, Z.W. and J.C.; visualization, Z.W.; supervision, J.C.; project administration, J.C.; funding acquisition, Z.W. and J.C. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported in part by the Department of Science and Technology of Gansu Province under Grant 24JRRA573, in part by the Gansu University of Political Science and Law under Grant GZF2025XQN15, and in part by the National Natural Science Foundation of China Youth Science Fund under Grant 62103178.

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding author.

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

References

1. Bhatia, S.; Alsuwailam, R.I.; Roy, D.G.; Mashat, A. Improved Multimedia Object Processing for the Internet of Vehicles. *Sensors* **2022**, *22*, 4133.
2. Zhang, J.; Chen, Y.; Zhang, Y.; Guo, B.; Xu, R. Dwha-pcmisp: Salient object detection network in coal mine industrial iot. *IEEE Trans. Ind. Inform.* **2025**, *21*, 5746–5754.
3. Achmadiyah, M.N.; Ahamad, A.; Sun, C.-C.; Kuo, W.-K. Energy-efficient fast object detection on edge devices for iot systems. *IEEE Internet Things J.* **2025**, *12*, 16681–16694.
4. Alves, M.G.; Chen, G.-L.; Kang, X.; Song, G.-H. Reduced CPU Workload for Human Pose Detection with the Aid of a Low-Resolution Infrared Array Sensor on Embedded Systems. *Sensors* **2023**, *23*, 9403.
5. Alshammari, A. Energy-Efficient Cryptographic Protocols for Sustainable IoT Security: A Federated Learning-Enhanced Lightweight Framework with Post-Quantum Resilience. *Sensors* **2026**, *26*, 3656.
6. Yang, W.; Hu, X.; Zhu, X.; Wu, R.; Pedrycz, W.; Liu, X.; Huang, J. Fedmps: Federated learning in a synergy of multi-level prototype-based contrastive learning and soft label generation. *IEEE Trans. Neural Netw. Learn. Syst.* **2026**, *37*, 781–794.
7. Kim, T.; Lin, E.; Lee, J.; Lau, C.; Mugunthan, V. Navigating data heterogeneity in federated learning: A semi-supervised federated object detection. *Adv. Neural Inf. Process. Syst.* **2023**, *36*, 2074–2096.

8. Khan, M.A.; Khalid, O.; Rais, R.N.B. Fed-DTCN: A Federated Disentangled Learning Framework for Unsupervised Zero-Day Anomaly Detection in IoT with Semantic-Aware Augmentation. *Sensors* **2026**, *26*, 1918.
9. Qiao, C.; Li, M.; Liu, Y.; Tian, Z. Transitioning from federated learning to quantum federated learning in internet of things: A comprehensive survey. *IEEE Commun. Surv. Tutor.* **2025**, *27*, 509–545.
10. Tang, X.; Guo, C.; Choo, K.-K.R.; Liu, Y. An efficient and dynamic privacy-preserving federated learning system for edge computing. *IEEE Trans. Inf. Forensics Secur.* **2024**, *19*, 207–220.
11. Su, H.; Zhou, J.; Feng, G.; Niu, X. Trustfed: A reliable federated learning framework with malicious-attack resistance. *IEEE Trans. Cogn. Commun. Netw.* **2025**, *11*, 2411–2425.
12. Huang, Y.; Liu, W.; Lin, Y.; Kang, J.; Zhu, F.; Wang, F.-Y. Flcsdet: Federated learning-driven cross-spatial vessel detection for maritime surveillance with privacy preservation. *IEEE Trans. Intell. Transp.* **2025**, *26*, 1177–1192.
13. Wang, D.; Pang, X.; Hu, J.; Yue, S.; Ren, J. Two-dimensional stackelberg game-based incentive mechanism for differential private federated learning with non-iid data. *IEEE Trans. Mob. Comput.* **2026**, *25*, 12688–12703.
14. Fu, L.; Zhang, H.; Gao, G.; Zhang, M.; Liu, X. Client selection in federated learning: Principles, challenges, and opportunities. *IEEE Internet Things J.* **2023**, *10*, 21811–21819.
15. Seid, A.M.; Erbad, A.; Abishu, H.N.; Boateng, G.O.; Khan, L.U.; Chiasserini, C.F.; Guizani, M. A hierarchical mafdrl-based resource allocation and incentive mechanism for tn-ntn in 6G networks. *IEEE Trans. Mob. Comput.* **2026**, *25*, 2426–2443.
16. Su, Y.; Fan, W.; Meng, Q.; Chen, P.; Liu, Y. Joint adaptive aggregation and resource allocation for hierarchical federated learning systems based on edge-cloud collaboration. *IEEE Trans. Cloud Comput.* **2025**, *13*, 369–382.
17. Li, B.; Gao, W.; Xie, J.; Gong, M.; Wang, L.; Li, H. Prototype-based decentralized federated learning for the heterogeneous time-varying iot systems. *IEEE Internet Things J.* **2024**, *11*, 6916–6927.
18. Tu, X.; Zhu, K.; Luong, N.C.; Niyato, D.; Zhang, Y.; Li, J. Incentive mechanisms for federated learning: From economic and game theoretic perspective. *IEEE Trans. Cogn. Commun. Netw.* **2022**, *8*, 1566–1593.
19. Wang, S.; Xia, W.; Zhao, H.; Wei, K.; Quek, T.Q.S.; Zhu, H. Stackelberg game-based hierarchical incentive mechanism for clustered vehicular federated learning. *IEEE Trans. Commun.* **2025**, *73*, 9071–9086.
20. Ullah, F.; Turab, A.; Ullah, S.; Cacciagrano, D.; Zhao, Y. Enhanced Network Intrusion Detection System for Internet of Things Security Using Multimodal Big Data Representation with Transfer Learning and Game Theory. *Sensors* **2024**, *24*, 4152.
21. Sridhar, V.; Rao, K.; Hussain, S.; Ullah, S.S.; Alroobaea, R.; Abdelhaq, M.; Alsaqour, R. Multivariate aggregated noma for resource aware wireless network communication security. *Comput. Mater. Contin.* **2023**, *74*, 1693–1708.
22. Shin, G.; Albanie, S.; Xie, W. Unsupervised salient object detection with spectral cluster voting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*; IEEE: Piscataway, NJ, USA, 2022; pp. 3971–3980.
23. Li, W.; Cheng, P.; Feng, Y.; Liu, N.; Liu, M.; Li, Y. A blockchain-assisted hierarchical data aggregation framework for iiot with computing first networks. *IEEE Trans. Netw. Serv. Manag.* **2025**, *22*, 3496–3512.
24. Wei, B.; Li, J.; Liu, Y.; Wang, W. Non-iid federated learning with sharper risk bound. *IEEE Trans. Neural Netw. Learn.* **2024**, *35*, 6906–6917.
25. Lu, J.; Zhang, Z.; Wang, J.; Li, R.; Wan, S. A green stackelberg-game incentive mechanism for multi-service exchange in mobile crowdsensing. *ACM Trans. Internet Technol. (TOIT)* **2021**, *22*, 31.
26. Gao, Y.; Ji, Z.; Zhao, K.; de Cola, T.; Li, W. Game-based computation offloading and power allocation for leo constellation networks in distributed and dynamic environment. *IEEE Internet Things J.* **2024**, *11*, 7040–7058.
27. Ju, L.; Zhang, T.; Toor, S.; Hellander, A. Accelerating fair federated learning: Adaptive federated adam. *IEEE Trans. Mach. Commun. Netw.* **2024**, *2*, 1017–1032.
28. An, X.; Shen, L.; Hu, H.; Luo, Y. Federated learning with manifold regularization and normalized update reaggregation. *Adv. Neural Inf. Process. Syst.* **2023**, *36*, 55097–55109.
29. Tan, A.Z.; Yu, H.; Cui, L.; Yang, Q. Towards personalized federated learning. *IEEE Trans. Neural Netw. Learn. Syst.* **2022**, *34*, 9587–9603.
30. Li, B.; Peng, Z.; Li, Y.; Li, Z.; Chen, S.; Shen, C.; Quek, T.Q. Neighborhood and global perturbations supported sharpness-aware minimization in federated learning: From local tweaks to global awareness. *IEEE Trans. Network Sci. Eng.* **2026**, *13*, 5000–5017.
31. Rey, V.; Sánchez, P.M.S.; Celdrán, A.H.; Bovet, G. Federated learning for malware detection in iot devices. *Comput. Netw.* **2022**, *204*, 108693.
32. Hannaan, A.; Shah, Z.; Erbad, A.; Mohamed, A.; Safa, A. Federated learning of low-rank one-shot image detection models in edge devices with scalable accuracy and compute complexity. In *2025 International Wireless Communications and Mobile Computing (IWCMC)*; IEEE: Piscataway, NJ, USA, 2025; pp. 1757–1762.
33. Zhang, B.; Mao, Y.; He, X.; Ping, P.; Huang, H.; Wu, J. Exploring the privacy-accuracy trade-off using adaptive gradient clipping in federated learning. *IEEE Trans. Netw. Sci. Eng.* **2025**, *12*, 2254–2265.
34. Katare, D.; Perino, D.; Nurmi, J.; Warnier, M.; Janssen, M.; Ding, A.Y. A survey on approximate edge AI for energy efficient autonomous driving services. *IEEE Commun. Surv. Tutor.* **2023**, *25*, 2714–2754.

35. Overman, T.; Klabjan, D. Continuous-time analysis of federated averaging. *arXiv*, **2025**, arXiv:2501.18870.
36. Wang, K.; Ding, Z.; So, D.K.C.; Ding, Z. Exploring age-of-information weighting in federated learning under data heterogeneity. *IEEE Trans. Wirel. Commun.* **2025**, *24*, 5023–5036.
37. Luo, R.; Stich, S.U.; Horváth, S.; Takáč, M. Revisiting localsgd and scaffold: Improved rates and missing analysis. *arXiv* **2025**, arXiv:2501.04443.
38. Li, Y.; Lyu, X. Convergence analysis of sequential federated learning on heterogeneous data. *Adv. Neural Inf. Process. Syst.* **2023**, *36*, 56700–56755.

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.