International Journal of Advanced Technology & Engineering Research (IJATER)
www.ijater.com

FEDPROX-KQE: TRUST-WEIGHTED FEDERATED LEARNING FOR MULTI-CLASS INTRUSION DETECTION IN INDUSTRIAL IOT NETWORKS

Pushpraj Shrivastava[1], Shweta Chouksey[2], Priyanka Verma[3]
Research Scholar[1], Associate Professor [2], Assistant Professor[3]
Department of Computer Science, LNCT University, Bhopal, MP, India[1,2]
University of Galway, Ireland[3]

***Abstract*—Securing Industrial Internet of Things (IIoT) networks against intrusion is fundamentally difficult because edge devices are resource-constrained, sensitive operational data can-not be shared across factory boundaries, and different factories encounter different attack types in practice. Centralised detection architectures fail on all three counts, while standard Federated Learning (FL) approaches treat detection as a binary problem, use aggregation rules that ignore client contribution quality, and do not account for the non-independent and identically distributed (non-IID) data distributions that arise across hetero-geneous deployments. This paper proposes FedProx-KQE, which addresses these limitations within a single framework. A Vari-ational Autoencoder–Long Short-Term Memory (VAE-LSTM) local model is extended with a multi-class classification head to identify individual attack categories rather than collapsing all threats into one class. Training data is partitioned across four factory clients with non-overlapping attack profiles to reflect realistic heterogeneity, and FedProx regularisation is applied to prevent local updates from drifting too far from the global model. Aggregation weights are replaced by a composite trust score that accounts for each client's local accuracy, data volume, and class coverage. Experiments on the X-IIoTID benchmark under both binary and multi-class settings show that FedProx reduces the impact of data heterogeneity compared to standard Federated Averaging (FedAvg), and a Convolutional Neural Network (CNN) ablation study confirms that model capacity is not the binding constraint in the federated multi-class setting.**

***Index Terms*—Federated Learning, IIoT Intrusion Detection, Non-IID, FedProx, Trust-Weighted Aggregation**

I. INTRODUCTION

Networked sensors, actuators, and edge computing nodes embedded in manufacturing plants generate continuous streams of operational data that are used to control physical processes in real time. These Industrial Internet of Things (IIoT) deployments underpin smart manufacturing, energy dis-tribution, and critical infrastructure, where Industrial Control Systems (ICS) issue safety-critical commands that directly govern physical equipment [1], [2]. A cyberattack against such a system is not merely a data breach; it can produce physical damage, halt production, or disrupt national-level infrastructure [3], [4].

Modern ICS are connected to the Internet and therefore exposed to adversaries who exploit both the limited computa-tional resources of constrained edge devices and the connec-tivity between Operational Technology (OT) and Information Technology (IT) domains [5], [6]. Attacks on Supervisory Control and Data Acquisition (SCADA) systems and pipeline control infrastructure have demonstrated that IIoT networks attract sophisticated adversaries whose objectives extend to physical destruction [7], [8].

Classical intrusion detection approaches such as firewalls, signature-based Intrusion Detection Systems (IDS), and cen-tralised machine learning are inadequate for IIoT environments for three reasons. First, centralised approaches require raw sensor data to be transmitted to a cloud server, introduc-ing latency for real-time detection, consuming bandwidth, and concentrating sensitive operational data [9]. Second, the volume of data generated by large-scale IIoT deployments makes purely cloud-based training architecturally infeasible without significant communication overhead [10], [11]. Third, sharing raw industrial data across organisational boundaries raises privacy and intellectual-property concerns that preclude collaborative centralised learning [12], [13].

Federated Learning (FL) addresses these constraints by training local models on private device data and sharing only model parameters with a central aggregator [14]. Each IIoT edge device trains on its own traffic data and transmits only weight updates, not raw data, preserving privacy and reducing bandwidth while enabling collaborative improvement across geographically distributed factories [3], [6].

Despite this progress, three gaps limit existing FL-based IIoT security systems. First, most FL intrusion detection works treat detection as a binary task, normal versus attack [3], [6], [15]. Knowing that an attack occurred provides no actionable information about its nature. In production IIoT environments, the response to a Denial of Service (DoS) flood differs entirely from the response to a Man-in-the-Middle (MITM) attack, ransomware, or a brute-force credential attack. Second, different factories face different attack distributions due to their distinct network topologies, device profiles, and operational contexts. Standard FedAvg [14] weights client contributions by sample count, which performs poorly when local distributions diverge significantly [16], [17]. When client data distributions are non-overlapping, as is common in heterogeneous factory environments, FedAvg produces a global model that oscillates between locally-biased solutions across rounds [18]. Third, Fe-dAvg is insensitive to whether a client's local model is actually useful. A client with a large but narrowly distributed dataset may receive a higher

aggregation weight than a smaller client whose model generalises better across attack categories [17], [19].

This paper proposes FedProx-KQE, a federated intrusion detection framework that addresses all three gaps. The contributions of this work are as follows:

1) A supervised multi-class classification head is integrated into a VAE-LSTM architecture, enabling 19-class attack classification on the X-IIoTID dataset alongside the standard reconstruction objective, advancing beyond the binary detection paradigm dominant in existing FL-based IDS literature.

2) Training data is partitioned across four simulated IIoT factory clients with heterogeneous attack-class distributions, and FedProx regularisation is applied to constrain local update divergence under these non-IID conditions, addressing the client-drift problem that undermines Fe-dAvg in realistic heterogeneous deployments.

3) An aggregation scheme is proposed that weights each client's contribution by a composite Knowledge Quality Evaluation (KQE) score derived from local accuracy, data volume, and class diversity, replacing standard sample-count weighting with a signal sensitive to both quantitative and qualitative dimensions of client knowledge.

4) All four configurations, Centralised, FedAvg, FedProx, and FedProx+KQE, are evaluated on both binary and multi-class tasks, and a CNN ablation study is conducted that isolates data heterogeneity rather than model capacity as the primary performance bottleneck in the federated multi-class IDS setting.

II. RELATED WORK

A. Anomaly and Intrusion Detection for ICS and IIoT

The problem of detecting cyberattacks in ICS has been studied extensively, with early work establishing process-mining and one-class classification baselines for anomaly scoring on constrained devices [7], [8]. Myers et al. [8] proposed a process-mining method that exploits log data from ICS devices to improve anomaly detection efficiency, while Nader et al. [7] investigated Support Vector Data Description (SVDD) and Kernel Principal Component Analysis (KPCA) for intrusion detection in SCADA systems. Subsequent work applied dimension-reduction techniques in combination with deep learning to improve efficiency on resource-constrained edge hardware [20].

Long Short-Term Memory (LSTM) networks have been widely applied to time-series anomaly detection in IIoT, exploiting their capacity to model long-range temporal dependencies in sensor data streams [10], [11]. Variational Autoen-coders (VAEs) have been adopted for unsupervised anomaly scoring, using a probabilistic latent space to distinguish normal from anomalous traffic without labelled attack data [21], [22]. The VAE-LSTM hybrid model was proposed in the FL context by Huong et al. [3], who demonstrated detection performance on SCADA data while satisfying the compute constraints of Raspberry Pi edge hardware. The architecture uses fully connected layers in the encoder and decoder to minimise model complexity, confirming deployment feasibility in real IIoT settings. However, this and all related architectures in the literature address only binary anomaly detection.

Liao et al. [6] demonstrated that neglecting cyber anomalies introduced by data loss from imperfect wireless channels pro-duces significant degradation in detection accuracy, motivating robust detection frameworks that account for wireless channel impairments alongside classical anomaly detection.

B. Federated Learning for Industrial Security

The application of FL to ICS and IIoT security has expanded rapidly since McMahan et al. [14] introduced FedAvg for communication-efficient collaborative learning from decentralised data. Huong et al. [3] demonstrated the first complete edge-FL deployment for ICS anomaly detection, using the Message Queuing Telemetry Transport (MQTT) protocol for model weight exchange between Raspberry Pi edge devices and a cloud aggregator, establishing a practical reference architecture for resource-constrained IIoT deployments.

Liao et al. [6] advanced FL for smart factories by combining Generative Adversarial Network (GAN)-based data imputation with FL-based collaborative training across multiple factories. Their impact-factor-assisted factory selection method balances local model quality and training execution time, demonstrating that selective client participation improves training efficiency at scale. However, their approach addresses only binary clas-sification and does not model non-IID data distributions ex-plicitly.

Verma et al. [15] proposed the Federated Learning-based Deep Intrusion Detection (FLDID) framework using a hybrid CNN+LSTM+Multi-Layer Perceptron (MLP) model with Pail-lier encryption for secure gradient communication, achieving high accuracy on binary classification of the X-IIoTID dataset. Liu et al. [11] proposed a communication-efficient on-device FL approach for deep anomaly detection in industrial IIoT using gradient compression to reduce uplink cost. Wang et al. [10] introduced a hierarchical FL architecture for anomaly detection in industrial IoT, with intermediate edge servers performing local aggregation before transmitting to the cloud. Ma et al. [23] proposed an integrated FL and transformer anomaly detection framework for cloud manufacturing. A federated fault diagnosis framework, FedCAE, was proposed in [24], showing better detection accuracy than training on a single factory particularly in non-IID scenarios.

Cui et al. [25] proposed privacy-preserving anomaly detection in IoT infrastructures using security-enhanced FL, addressing both detection accuracy and gradient privacy simultaneously. Li et al. [4] conducted a comprehensive

survey on security and privacy of FL in industrial IoT, cataloguing attacks against FL systems including Byzantine faults, model poisoning, and inference attacks, providing context for the trust-weighting approach proposed in this paper.

C. Non-IID FL and Aggregation Strategies

The non-IID problem in federated learning, where differ-ent clients hold data from different distributions, is widely recognised as the primary challenge limiting FL effectiveness in heterogeneous real-world deployments [18]. Li et al. [16] introduced FedProx, which augments local training with aproximal term penalising the squared distance between local and global model weights, constraining client drift under heterogeneous data distributions and providing convergence guarantees that FedAvg lacks in the non-IID setting.

Zhao et al. [18] provided theoretical and empirical evidence that data heterogeneity is the dominant factor limiting FedAvg performance, motivating subsequent work on distribution-aware aggregation. Fu et al. [17] conducted a comprehensive survey of client selection in FL, identifying local model loss, data volume, communication cost, and data quality as the four dominant selection criteria, and showed that naive sample-count weighting systematically disadvantages clients with small but high-quality local datasets. Duan et al. [19] proposed a self-balancing FL mechanism to prevent training bias caused by imbalanced data distributions, directly address-ing the scenario encountered in the heterogeneous factory experiments of this work. The proposed KQE mechanism extends this line of work by incorporating class diversity as an explicit quality signal, providing a principled alternative to sample-count weighting that is specifically designed for the multi-class IDS setting.

III. SYSTEM MODEL

The system model used in this paper is illustrated in Fig. 1. The proposed FedProx-KQE framework consists of three ma-jor components: edge devices representing IIoT factories, a cloud server, and the underlying smart industry process layer at each factory site.

1) Edge Devices: Each edge device represents an IIoT factory client participating in collaborative federated learning. Every device stores local network traffic col-lected from its associated factory and trains a local VAE-LSTM model using the FedProx optimization algorithm with proximal regularization. After local training, the device transmits its updated model parameters to the cloud server and receives the aggregated global model for the next communication round. This iterative process continues until convergence. In this work, a VAE-LSTM model with a multi-class classification head is employed for intrusion detection, while a CNN-augmented classi-fier head is additionally evaluated in the ablation study.

2) Cloud Server: The cloud server coordinates the feder-ated learning process among all participating clients. Its primary responsibilities are: (i) initializing and main-taining the global model parameters, (ii) computing a Knowledge Quality Evaluation (KQE) trust score for each client based on local classification accuracy, train-ing data volume, and class diversity, and (iii) performing trust-weighted aggregation of client model updates to generate a new global model, which is subsequently broadcast to all edge devices for the next federated learning round.

3) Smart Industry Layer: Each edge device receives operational network traffic from its corresponding fac-tory's physical process layer. This layer comprises a Programmable Logic Controller (PLC) responsible for process control, a local Supervisory Control and Data Acquisition (SCADA) system for monitoring industrial operations, and interconnected sensors and actuators that continuously generate network traffic. The captured traffic forms the X-IIoTID dataset samples used for local model training.

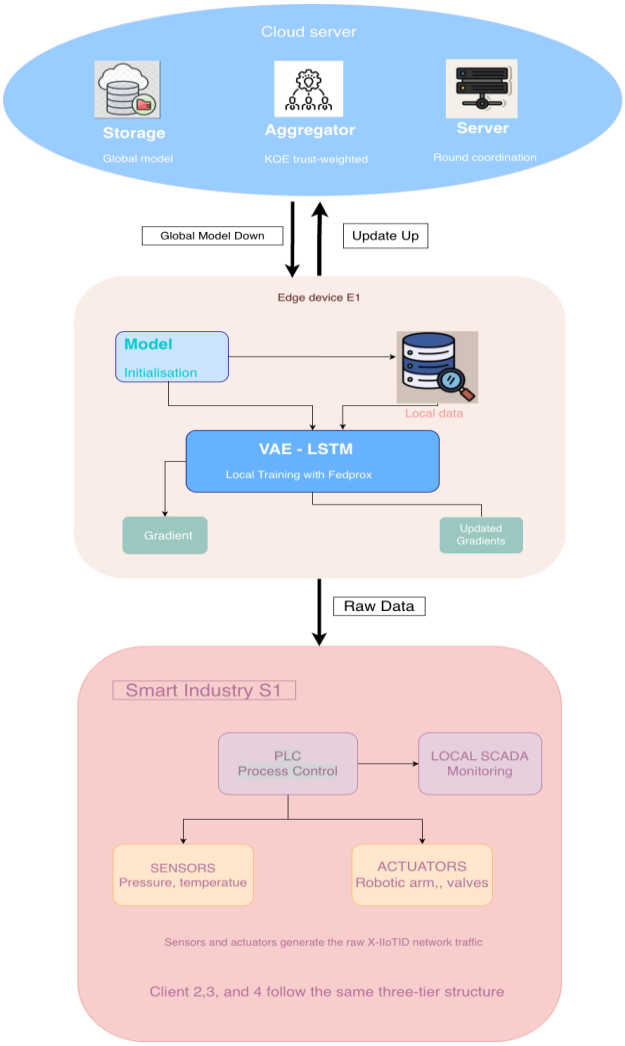

Fig. 1. System architecture of the proposed FedProx-KQE framework, showing the cloud server, edge devices, and the smart industry process layer for a representative IIoT client.

IV. PROPOSED METHODOLOGY

A. System Overview

The proposed FedProx-KQE framework adopts a federated learning paradigm in which four geographically distributed IIoT clients collaborate to train a shared intrusion detection model without disclosing local network traffic data to any centralised entity. In each communication round, clients re-ceive the current global model from the server, perform local training on private traffic data, and return only updated model parameters. The server aggregates these parameters using the proposed KQE mechanism and broadcasts the improved global model for the next round.

Unlike standard FedAvg [14], which weights contributions by data volume alone, KQE dynamically adjusts each client's influence based on three quality dimensions: local model accu-racy, dataset size, and class diversity. This design is motivated by the observation that in heterogeneous IIoT environments, a large local dataset does not necessarily imply a high-quality or representative local model, particularly when class distributions are severely imbalanced.

B. Dataset: X-IIoTID

The X-IIoTID dataset [26] is a connectivity-agnostic, protocol-independent benchmark for IIoT intrusion detection, generated from real IoT device traffic spanning heterogeneous network protocols. It contains samples with 68 features per flow, labelled across 19 classes: one Normal class and 18 distinct attack categories encompassing DoS variants, MITM attacks, network and OS scanning, brute-force attacks, ran somware, backdoor, Cross-Site Scripting (XSS), SQL injec-tion, and password theft.

For federated experiments, the dataset is partitioned across four clients to reflect a realistic non-IID distribution, with each client assigned a distinct subset of attack categories to simulate separate factory environments with non-overlapping threat profiles. For binary experiments, all attack labels are collapsed to a single Attack class and the dataset is balanced by undersampling the majority class.

C. VAE-LSTM Model Architecture

The local model deployed at each client is a VAE with an LSTM encoder [3]. The architecture jointly optimises re-construction, regularisation, and classification objectives. The VAE encoder learns a compact latent representation of network traffic while filtering redundant features, a property that is especially valuable on resource-constrained IIoT edge devices.

1) Encoder: An LSTM with hidden dimension $H = 256$ processes the input sequence $\mathbf{x} \in \mathbb{R}^{B\times 1\times D}$ (where $D = 63$, B is batch size) and produces hidden state $\mathbf{h}_n$. Two linear projections map $\mathbf{h}_n$ to latent mean $\boldsymbol{\mu}$ and log-variance $\log \boldsymbol{\sigma}^2$

2) Decoder: The decoder reconstructs the original input $\hat{\mathbf{x}}$ from the latent vector $\mathbf{z}$ using a second LSTM. Reconstruction serves two purposes: it acts as an unsupervised regulariser preventing encoder collapse, and reconstruction error can serve as an anomaly score in deployment scenarios where attack labels are unavailable.

3) Classifier Head: A two-layer MLP with Batch Normal-isation and Dropout maps $\mathbf{z}$ to class logits:

$$\hat{\mathbf{y}} = \text{Softmax}(W_2\,\text{ReLU}(W_1\mathbf{z} + \mathbf{b}_1) + \mathbf{b}_2)\,. \qquad (3)$$

4) Joint Loss Function:

5) Joint Loss Function: Local training minimises a com-bined objective:

$$\mathcal{L} = \alpha\mathcal{L}_{\text{recon}} + \beta\mathcal{L}_{\text{KL}} + \gamma\mathcal{L}_{\text{class}} \qquad (4)$$

where $L_{recon} = \text{MSE}(\mathbf{x}, \hat{\mathbf{x}})$, $L_{class} = \text{CrossEntropy}(\hat{\mathbf{y}}, \mathbf{y})$, and the KL divergence term is:

$$\mathcal{L}_{\text{KL}} = -\frac{1}{2}\sum_{j}\left(1 + \log\sigma_j^2 - \mu_j^2 - \sigma_j^2\right) \qquad (5)$$

with weights $\alpha = 1.0$, $\beta = 0.01$, $\gamma = 1.0$.

A. Federated Training with FedProx

Federated training follows FedProx [16], which augments the local objective with a proximal regularisation term:

$$L_{local} = L + \frac{\mu}{2}\,\|w - w^t\|^2 \qquad (6)$$

where w are local parameters, w^t are the received global

parameters, and $\mu = 0.001$ is the proximal coefficient. Each client runs E local epochs before returning parameters; $E = 3$ for multi-class and $E = 1$ for binary experiments.

E. Knowledge Quality Evaluation (KQE)

Standard FedAvg weights client contributions by dataset size, which can cause the global model to over-represent clients with large but low-quality local datasets. KQE com-putes a composite trust score for each client using three signals:

- Local Accuracy a_k: classification accuracy on the client's held-out validation set, weighted at $\omega_a = 0.5$.
- Data Size n_k: number of local training samples, weighted at $\omega_n = 0.3$.
- Class Diversity d_k: count of unique attack classes present in the client's data, weighted at $\omega_d = 0.2$.

Each dimension is normalised across all K clients using min-max scaling:

$$\tilde{x}_k = \frac{x_k - \min_j x_j}{\max_j x_j - \min_j x_j}. \quad (7)$$

The composite trust score and aggregation rule are:

$$s_k = \omega_a \tilde{a}_k + \omega_n \tilde{n}_k + \omega_d \tilde{d}_k, \quad \sum_k s_k = 1 \quad (8)$$

$$w^{t+1} = \sum_{k=1}^{K} s_k w_k^t. \quad (9)$$

TABLE I

HYPERPARAMETER CONFIGURATION

Parameter	Value	Applies To
LSTM hidden dim (H)	256	All configs
Latent dim (L)	64	All configs
Recon. weight (α)	1.0	All configs
KL weight (β)	0.01	All configs
Class weight (γ)	1.0	All configs
FedProx coeff. (μ)	0.001	FedProx, KQE
Local epochs, multi-class (E)	3	All FL
Local epochs, binary (E)	1	All FL
FL rounds	10	All FL
Clients (K)	4	All FL
KQE ω_a	0.5	KQE only
KQE ω_n	0.3	KQE only
KQE ω_d	0.2	KQE only

F. CNN Ablation Study Design

To test whether richer feature extraction could improve multi-class accuracy, the MLP classifier head was replaced with a CNN pipeline: (i) $\mathbf{z}$ is reshaped to ($B \times 1 \times L$) and processed by Conv1d (32 filters, kernel 3, padding 1);
(ii) ReLU and MaxPool1d (kernel 2) reduce the sequence length; (iii) pooled features feed a classifier LSTM (hidden 64); (iv) the final hidden state passes through a two-layer dense head with ReLU and Dropout (0.3). The encoder and decoder remained unchanged.

G. Hyperparameter Configuration

Table I summarises all hyperparameters used across exper-iments.

V. RESULTS AND DISCUSSION

A. Evaluation Protocol

All models are evaluated using weighted-average Accuracy, Precision, Recall, and F1-score on a held-out centralised server-side test set from X-IIoTID. This protocol follows Verma et al. [15], who used the same dataset and metrics to benchmark federated intrusion detection. Reported results represent the best-performing round over 10 federated rounds.

B. Multi-Class Classification (19 Classes)

1) Overall Performance: Table II presents multi-class re-sults. The centralised VAE-LSTM achieves near-perfect per-formance, confirming the strength of the architecture when all data is available for training. All federated strategies experi-ence a performance drop of approximately 25–30 percentage points relative to this centralised upper bound, reflecting the fundamental challenge of non-IID data distribution across clients.

FedProx outperforms FedAvg by 4.1 percentage points in F1. By constraining local updates to remain close to the global model, FedProx reduces client drift and produces more stable aggregation across rounds [16]. FedProx+KQE achieves 72.05% F1, which is lower than FedProx alone. Analysis of per-round KQE trust scores reveals that in the multi-class setting, local validation accuracy was inconsistent across clients due to severe within-client class imbalance, causing trust scores to default toward data-size-driven weighting and reducing KQE's differentiating power. This motivates per-class accuracy metrics as a future extension to KQE.

2) CNN Ablation Study: Table III reports the results of replacing the MLP classifier head with the CNN+LSTM pipeline.

The CNN-enhanced model produced no meaningful im-provement across all three strategies. This confirms that the performance bottleneck in federated multi-class IDS is data heterogeneity across clients, not the expressiveness of the local classifier, consistent with the observation of Huong et al. [3] that increasing model complexity in federated IIoT settings does not yield commensurate accuracy gains.

A. Binary Classification (Normal vs. Attack)

1) Overall Performance: Table IV presents binary classifi-cation results. All federated methods converge to above 99.2% accuracy, consistent with results achieved by Verma et al. [15] on the same X-IIoTID dataset.

2) KQE Trust Score Analysis: Table V shows representative KQE trust scores from Round 4 of the binary experiment.

TABLE V

KQE TRUST SCORES, BINARY SETTING, ROUND 4

B.

Client	Acc (%)	Data Size	Classes	Trust
Client 1	99.92	21,420	2	0.3011
Client 2	99.79	21,056	2	0.0230
Client 3	100.00	21,227	2	0.3184
Client 4	100.00	21,330	2	0.3575

Clients 3 and 4, achieving 100% local validation accuracy, received the highest trust weights. Client 2, despite holding a comparable dataset size, received the lowest trust score (0.0230) due to lower local accuracy (99.79%). This demon-strates a key property of KQE: it prioritises model quality over raw data volume, directly addressing the limitation identified by Fu et al. [17] in standard FedAvg weighting.

3) Per-Class Detection Analysis: Binary per-class evalua-tion shows near-perfect recall for Normal traffic with strong Attack recall across rounds. The primary error source was Attack flows misclassified as Normal (false negatives), with a False Negative Rate (FNR) below 1.1%. The near-zero False Positive Rate (FPR) is a desirable operational characteristic, as false positives in production IDS deployments trigger unnecessary alerts and carry high operational cost [6].

D. Discussion

1) Impact of Data Heterogeneity: The 25–30% accuracy gap in multi-class classification, which nearly vanishes in binary classification, confirms that this gap is driven by task complexity under non-IID conditions rather than a fundamen-tal limitation of the federated approach. This finding aligns with the analysis of Liao et al. [6], who noted that single-factory data patterns are insufficient for high-quality detection models, and underscores the need for data-level solutions such as personalised FL or cross-client data augmentation, consistent with findings of Duan et al. [19].

2) FedProx vs. FedAvg: FedProx consistently outperforms FedAvg by 4.1% F1 in the multi-class setting. The proximal regularisation mechanism constrains local updates to remain within a bounded distance of the global model, reducing client drift and producing a more stable global model that improves monotonically across rounds. This confirms the theoretical guarantees of Li et al. [16].

3) Comparison with Related Work: Verma et al. [15] achieved high accuracy on binary classification of X-IIoTID using a CNN+LSTM+MLP model in a federated setting. The VAE-LSTM proposed here achieves comparable performance under the same binary task. Unlike [15], this work extends to 19-class identification, explicitly models non-IID distributions, and introduces quality-aware aggregation. The multi-class ex-tension is the primary novel dimension of this work, and the 70–75% F1 range achieved under strict non-IID conditions represents a meaningful first baseline for this harder task on the X-IIoTID dataset.

VI. CONCLUSION

This paper proposed FedProx-KQE, a federated intrusion detection framework that combines a VAE-LSTM local model with FedProx optimisation and a quality-aware aggregation strategy for multi-class and binary intrusion detection in het-erogeneous IIoT environments.

In binary classification, all federated methods exceeded 99.2% F1, confirming that the VAE-LSTM architecture re-liably separates normal from malicious IoT traffic without centralising sensitive data. In multi-class classification across

19 classes, federated methods achieved 70–75% F1 versus near-perfect centralised performance. This gap is attributable to non-IID data heterogeneity across clients and is not re-solved by increasing model complexity, as confirmed by the CNN ablation study. FedProx outperformed FedAvg by 4.1% F1, validating the benefit of proximal regularisation. FedProx+KQE achieved lower F1 than FedProx alone in the multi-class setting due to the unreliability of aggregate local accuracy as a quality signal under severe class imbalance, motivating per-class quality evaluation as the primary direction for future work.

Future directions include personalised federated learning to address non-IID challenges, adaptive KQE weight tuning using per-class accuracy metrics and local confusion matrices, and extension to streaming IIoT traffic detection with concept drift handling. Integrating differential privacy or homomorphic encryption represents a further direction to strengthen gradient privacy guarantees [3], [15].

REFERENCES

[1] E. Sisinni, A. Saifullah, S. Han, U. Jennehag, and M. Gidlund, “Indus-trial internet of things: Challenges, opportunities, and directions,” IEEE transactions on industrial informatics, vol. 14, no. 11, pp. 4724–4734, 2018.

[2] K.-C. Chen, S.-C. Lin, J.-H. Hsiao, C.-H. Liu, A. F. Molisch, and G. P. Fettweis, “Wireless networked multirobot systems in smart factories,” Proceedings of the IEEE, vol. 109, no. 4, pp. 468–494, 2020.

[3] T. T. Huong, T. P. Bac, D. M. Long, T. D. Luong, N. M. Dan, L. A.

Quang, L. T. Cong, B. D. Thang, and K. P. Tran, “Detecting cyberattacks using anomaly detection in industrial control systems: A federated learning approach,” Computers in Industry, vol. 132, p. 103509, 2021.

[4] H. Li, L. Ge, and L. Tian, “Survey: federated learning data security and privacy-preserving in edge-internet of things,” Artificial Intelligence Review, vol. 57, no. 5, p. 130, 2024.

[5] N. Tuptuk and S. Hailes, "Security of smart manufacturing systems," Journal of manufacturing systems, vol. 47, pp. 93–106, 2018.

[6] Y. Liao, Y. Wang, Q. Cui, K.-C. Chen, G. Nan, and X. Tao, "Data-driven cyber-physical anomaly detection with gan in federated smart factories," IEEE Transactions on Industrial Informatics, vol. 21, no. 4,

pp. 3067–3076, 2025.

[7] P. Nader, P. Honeine, and P. Beauseroy, "norms in one-class classifi-cation for intrusion detection in scada systems," IEEE Transactions on Industrial Informatics, vol. 10, no. 4, pp. 2308–2317, 2014.

[8] D. Myers, S. Suriadi, K. Radke, and E. Foo, "Anomaly detection for industrial control systems using process mining," Computers & Security, vol. 78, pp. 103–125, 2018.

[9] A. A. Cook, G. Mısırlı, and Z. Fan, "Anomaly detection for iot time-series data: A survey," IEEE Internet of Things Journal, vol. 7, no. 7,

pp. 6481–6494, 2019.

[10] X. Wang, S. Garg, H. Lin, J. Hu, G. Kaddoum, M. J. Piran, and M. S. Hossain, "Toward accurate anomaly detection in industrial internet of things using hierarchical federated learning," IEEE Internet of Things Journal, vol. 9, no. 10, pp. 7110–7119, 2021.

[11] Y. Liu, S. Garg, J. Nie, Y. Zhang, Z. Xiong, J. Kang, and M. S. Hossain, "Deep anomaly detection for time-series data in industrial iot: A communication-efficient on-device federated learning approach," IEEE Internet of Things Journal, vol. 8, no. 8, pp. 6348–6358, 2020.

[12] J. Feng, L. T. Yang, R. Zhang, and B. S. Gavuna, "Privacy-preserving tucker train decomposition over blockchain-based encrypted industrial iot data," IEEE Transactions on Industrial Informatics, vol. 17, no. 7,

pp. 4904–4913, 2020.

[13] X. Su, K. Fan, and W. Shi, "Privacy-preserving distributed data fusion based on attribute protection," IEEE Transactions on Industrial Infor-matics, vol. 15, no. 10, pp. 5765–5777, 2019.

[14] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in Artificial intelligence and statistics. Pmlr, 2017, pp. 1273–1282.

[15] P. Verma, J. G. Breslin, and D. O'Shea, "Fldid: Federated learning enabled deep intrusion detection in smart manufacturing industries," Sensors, vol. 22, no. 22, p. 8974, 2022.

[16] T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith, "Federated optimization in heterogeneous networks," Proceedings of Machine learning and systems, vol. 2, pp. 429–450, 2020.

[17] L. Fu, H. Zhang, G. Gao, M. Zhang, and X. Liu, "Client selection in federated learning: Principles, challenges, and opportunities," IEEE Internet of Things Journal, vol. 10, no. 24, pp. 21 811–21 819, 2023.

[18] Y. Zhao, M. Li, L. Lai, N. Suda, D. Civin, and V. Chandra, "Federated learning with non-iid data," arXiv preprint arXiv:1806.00582, 2018.

[19] M. Duan, D. Liu, X. Chen, R. Liu, Y. Tan, and L. Liang, "Self-balancing federated learning with global imbalanced data in mobile systems," IEEE Transactions on Parallel and Distributed Systems, vol. 32, no. 1, pp. 59–71, 2020.

[20] M. Kravchik and A. Shabtai, "Efficient cyber attack detection in indus-trial control systems using lightweight neural networks and pca," IEEE transactions on dependable and secure computing, vol. 19, no. 4, pp. 2179–2197, 2021.

[21] L. Gjorgiev and S. Gievska, "Time series anomaly detection with variational autoencoder using mahalanobis distance," in International Conference on ICT Innovations. Springer, 2020, pp. 42–55.

[22] H. Xu, W. Chen, N. Zhao, Z. Li, J. Bu, Z. Li, Y. Liu, Y. Zhao, D. Pei,

Y. Feng et al., "Unsupervised anomaly detection via variational auto-encoder for seasonal kpis in web applications," in Proceedings of the 2018 world wide web conference, 2018, pp. 187–196.

[23] S. Ma, J. Nie, J. Kang, L. Lyu, R. W. Liu, R. Zhao, Z. Liu, and D. Niyato, "Privacy-preserving anomaly detection in cloud manufacturing via feder-ated transformer," IEEE Transactions on Industrial Informatics, vol. 18, no. 12, pp. 8977–8987, 2022.

[24] Y. Yu, L. Guo, H. Gao, Y. He, Z. You, and A. Duan, "Fedcae: A new federated learning framework for edge-cloud collaboration based machine fault diagnosis," IEEE Transactions on Industrial Electronics, vol. 71, no. 4, pp. 4108–4119, 2023.

[25] L. Cui, Y. Qu, G. Xie, D. Zeng, R. Li, S. Shen, and S. Yu, "Security and privacy-enhanced federated learning for anomaly detection in iot infrastructures," IEEE Transactions on Industrial Informatics, vol. 18, no. 5, pp. 3492–3500, 2021.

[26] M. Al-Hawawreh, E. Sitnikova, and N. Aboutorab, "X-iiotid: a connectivity-agnostic and device-agnostic intrusion data set for industrial internet of things. ieee internet things j 9(5):3962–3977," 2022.