FedAdapt-CAD: A Federated Learning Framework for Anomaly Detection with Client-Aware Aggregation and Dynamic Model Adaptation

Muhammad Azaz Farooq*, Paolo Bellavista*, Armir Bujari*, Alessio Mora*,
Jamal Nasir†, Rohma Ahmed‡
*Department of Computer Science and Engineering, University of Bologna, Italy
Emails: muhammadazaz.farooq2@unibo.it, name.surname@unibo.it
†University of Galway,Ireland
Email: jamal.nasir@universityofgalway.ie
‡ The Islamia University Bahawalpur, Pakistan
rohmahmad21@yahoo.com

***Abstract*—In this research, we introduced FedAdapt-CAD, a novel Federated Learning (FL) framework designed for Anomaly Detection (AD) in critical infrastructure systems. This framework emphasizes robustness, scalability, and privacy preservation in mitigating anomalies. Key innovations include Client-Aware Aggregation (CAA), which adjusts client contributions based on data quality and local model performance, and Dynamic Model Adaptation (DMA), which optimizes global model hyperparameters in real-time to respond to changing attack patterns. Evaluated on three benchmark datasets: WADI, CMAPSS, and BATADAL, FedAdapt-CAD demonstrates superior global accuracy and fairness among heterogeneous clients. The experimental results indicate a significant improvement in detection rates for rare and complex attacks, achieving up to a 15% enhancement in precision and recall compared to traditional FL methods such as FedAvg, FedProx, and FedH2L.**

***Index Terms*—Federated Learning, Anomaly Detection, Industrial Control Systems, Cyber-Physical Systems**

I. Introduction

The rise of industrial system interconnectivity and smart devices has led to vast data production in areas such as predictive maintenance and smart manufacturing. In Industry 4.0, ensuring secure and efficient analytics delivery is vital, particularly in distributed environments where privacy restrictions limit the sharing of raw data. FL presents a promising approach by allowing collaborative model training without the need to transfer raw data. [1], [2]. Applying FL to real-world industrial scenarios presents several challenges. Unlike general-purpose datasets, industrial data is often non-IID (non-independent and identically distributed), high-dimensional, and collected under varying conditions from heterogeneous devices [3], [4]. The hierarchical structure of industrial systems, consisting of edge devices, gateways, and cloud infrastructure, complicates deployment. Issues such as streamed input and missing values hinder model accuracy. At the same time, hierarchical FL frameworks often assume balanced data distribution and stable communication, both of which are uncommon in practice [5]–[7]. To address the above challenges, this paper introduces FedAdapt-CAD, a novel FL framework that employs adaptive aggregation and dynamic learning strategies to meet industrial deployment requirements. It addresses challenges such as streamed data inputs, missing values, non-IID data, and edge resource limitations. Unlike traditional FL approaches, our framework handles data heterogeneity and evolving attack behaviors, making it ideal for environments with unreliable clients and dynamic threats. This adaptation mechanism is central to the paper's contributions. The remainder of the paper is organized as follows: Section II presents prior work on FL frameworks. Section III describes our proposed FedAdapt-CAD framework and its implementation. Section IV presents experimental results and evaluation metrics across various datasets, and Section V concludes the paper, outlining directions for future work.

II. Related Work

FL enables collaborative model training without sharing raw data, making it suitable for privacy-sensitive industrial settings. However, challenges like data heterogeneity, resource constraints, and hierarchical infrastructure complicate its deployment. Li et al. [1] proposed a client selection strategy to enhance robust-

ness in heterogeneous industrial IoT data environments. Similarly, Liu et al. [3] conducted a comprehensive survey that highlighted non-IID data as a significant bottleneck to FL performance. In the context of vertical FL, Li et al. [2] designed a framework tailored for heterogeneous industrial regions that focuses more on vertical data splits than hierarchical architecture. Dutta et al. [8] introduced a privacy-preserving FL framework for chemical engineering applications, emphasising cross-organisational collaboration. To handle the challenges of system hierarchy, several hierarchical FL (HFL) frameworks have been proposed. Fang et al. [4] developed a gradient correction mechanism for HFL with multi-timescale updates. Huang et al. [9] presented a fairness-aware FL scheme suitable for hierarchical structures. Architectures such as client-edge-cloud HFL have also been investigated by Liu et al. [5] and Wang et al. [6], while FedH2L [7] focused on heterogeneous learning for resource-constrained IIoT environments. These works, however, often assume ideal network conditions or overlook the need for dynamic reconfiguration. Moreover, frameworks discussed in [10]–[13] offer improvements in scalability and collaborative intelligence, but often lack mechanisms for real-time model adaptation and efficient edge-side optimization. The works of Zhou et al. [14] and Kang et al. [15] have explored FL applications in cyber-physical and critical infrastructure systems, but without addressing the fine-grained client-aware weighting or adaptive tuning proposed in our work.

Our proposed framework advances previous approaches by incorporating two innovative components: CAA, a client-level dynamic aggregation strategy that enhances convergence and fairness in non-IID settings, and DMA, a server-level mechanism that adjusts to variations in data distribution and cyber-attack patterns.

III. Proposed Framework

The system diagram 1 illustrates a framework applicable across various industrial domains, including CMAPSS, WADI, and BATADAL. It highlights the universality of the architecture rather than concurrent dataset distribution to multiple clients. Each dataset is analyzed separately to prevent cross-domain leakage and maintain specific anomaly semantics. Non-IID partitioning techniques, such as unequal shard allocation and temporal segmentation, create clients that reflect real-world industrial distributions, ensuring a hierarchical client-aggregator-server workflow. Federated processing is conducted independently for each dataset, thereby reinforcing methodological consistency and demonstrating the framework's broad applicability.

A. Our Proposed Framework Workflow

The proposed framework workflow consists of the following steps:

1) Local Model Training: Let N represent the total number of edge devices, partitioned into M regions. Each region R_m contains n_m devices, satisfying the condition $\sum_{m=1}^{M} n_m = N$. Every device i within region m owns a local dataset $\mathcal{D}_{i,m}$, which follows a unique non-IID distribution $P_{i,m}(x, y)$. This heterogeneity is a central consideration in the system model and motivates the need for adaptive aggregation strategies.

The overall training objective is to minimise a global loss function. It is defined as a weighted sum of the local loss functions across all devices. Mathematically, this can be expressed as:

$$\min_{w} F(w) = \sum_{m=1}^{M} \frac{n_m}{N} \sum_{i \in R_m} \frac{|\mathcal{D}_{i,m}|}{|\mathcal{D}_m|} F_{i,m}(w), \quad (1)$$

where w denotes the model parameters and $F_{i,m}(w)$ represents the local loss function for device i in region m, computed as:

$$F_{i,m}(w) = \frac{1}{|\mathcal{D}_{i,m}|} \sum_{(x,y) \in \mathcal{D}_{i,m}} \ell(f_w(x), y), \quad (2)$$

with $\ell(\cdot, \cdot)$ denoting the loss function (e.g., cross-entropy), $f_w(x)$ representing the model's output, and y being the ground truth label.

2) Regional Aggregation: Each regional aggregator aggregates the received models from its respective edge devices. The regional model for region m is calculated as:

$$w_m^{t+1} = \frac{\sum_{i \in R_m} |\mathcal{D}_{i,m}| w_{i,m}^t}{\sum_{i \in R_m} |\mathcal{D}_{i,m}|}, \quad (3)$$

where $w_{i,m}^t$ is the model from device i at iteration t.

3) Global Aggregation: The global model is then obtained at the central server by aggregating the regional models:

$$w^{t+1} = \frac{\sum_{m=1}^{M} n_m w_m^t}{N}. \quad (4)$$

To better handle the non-IID nature of industrial data, an adaptive weighting mechanism is introduced into the aggregation process. Each device's contribution to the regional model is scaled by a weight $\alpha_{i,m}$, which is calculated based on the statistical variance of its dataset:

$$\alpha_{i,m} = \frac{\text{Var}(\mathcal{D}_{i,m})}{\sum_{j \in R_m} \text{Var}(\mathcal{D}_{j,m})}, \quad w_m^{t+1} = \sum_{i \in R_m} \alpha_{i,m} w_{i,m}^t. \quad (5)$$

Here, $\text{Var}(\mathcal{D}_{i,m})$ denotes the variance of the local dataset, which reflects the diversity of data samples and indicates how representative the device's data is within its region. Higher variance implies more informative and diverse data, thus increasing its influence during aggregation.

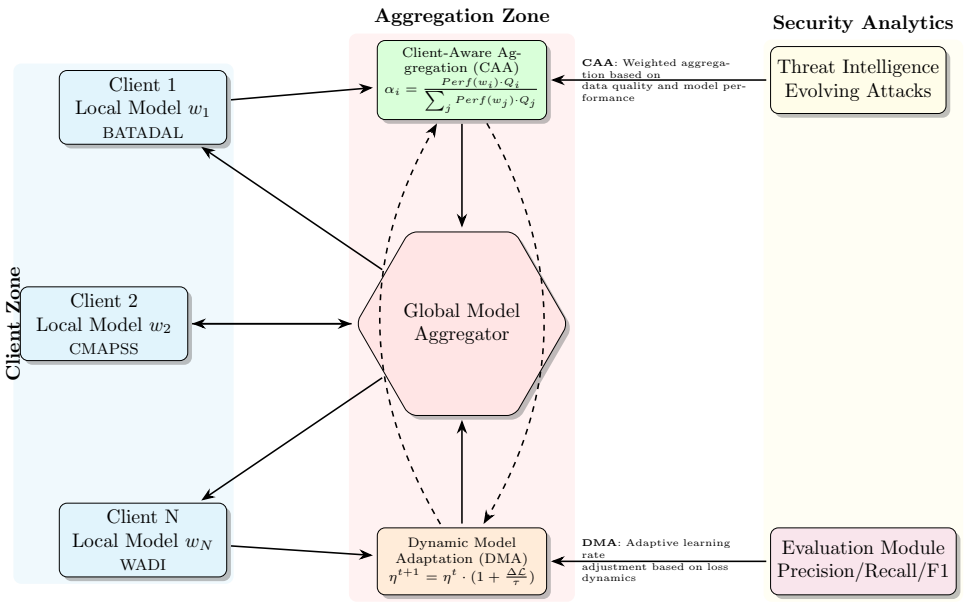

Fig. 1. FedAdapt-CAD framework workflow

B. Dynamic Model Adaptation (DMA)

DMA is a server-side mechanism that optimizes the global model's hyperparameters, such as the learning rate η and regularization coefficients, to adapt to changing data distributions and attack patterns. It addresses non-stationary threats in industrial systems, like novel attack signatures and sensor drift, requiring models that self-adjust autonomously. After each global aggregation round t, the server computes a meta-gradient from a validation set $\mathcal{V}$ of recent anomaly instances.

$$\Delta\eta = \beta \cdot \nabla_\eta \mathcal{L}_{\text{val}}(w^t; \mathcal{V}), \tag{6}$$

where β is a meta-learning rate, and $\mathcal{L}_{\text{val}}$ is the validation loss. The learning rate is then updated as:

$$\eta^{t+1} = \eta^t - \Delta\eta. \tag{7}$$

The DMA module builds upon a differentiable aggregation mechanism in which the global update is computed using a weighted combination of regional models. Unlike simple averaging, the aggregation incorporates proportional weighting based on the size and quality of client data, ensuring that clients with more informative or stable distributions exert greater influence on the global hyperparameter trajectory. Formally, the aggregation used during DMA is a smooth, differentiable operator expressed as:

$$\mathcal{A}(w_m^t) = \sum_{m=1}^{M} \gamma_m \, w_m^t, \tag{8}$$

where γ_m denotes continuous, learnable aggregation weights derived from regional statistics, allowing the meta-gradient to flow through the aggregation step without breaking differentiability. This design enables DMA to adjust the learning rate and other hyperparameters in a manner that reflects model reliability across clients, providing a principled alternative to traditional non-differentiable averaging. The differentiable formulation thus supports stable hyperparameter optimization while preserving compatibility with FL constraints. This gradient-based adjustment ensures the model prioritises recent attack patterns while preserving knowledge from historical data. Unlike static FL frameworks, DMA explicitly addresses temporal concept drift (e.g., in BATADAL's evolving cyber-physical attacks) by treating hyperparameters as learnable entities, a strategy inspired by meta-learning.

It is emphasized that the server does not access raw client data during the DMA update. The validation set used for hyperparameter adaptation is created without breaching the privacy constraints of FL. This set is sourced from (i) a small, publicly accessible, or synthetic validation subset with representative anomaly patterns, and (ii) aggregated non-identifying statistical summaries such as class distribution histograms from regional aggregators. These measures ensure DMA's compliance with the FL privacy model, enabling effective model drift capture and learning dynamics adaptations without sharing confidential industrial data from local client devices.

C. Federated Training Loop

The training process is designed around a client-edge-cloud hierarchical model. Each round of communication

involves the following sequence:

1) **Local Training:** Edge clients use a resource-aware optimisation strategy to train models on local data, dynamically adjusting batch size and learning rate based on memory and compute availability.
2) **Regional Aggregation:** Clients within the same region submit their trained models to a regional aggregator, who uses a weighted averaging scheme to generate a regional model.
3) **Global Aggregation:** The central server receives regional models and performs a second-level aggregation to compute the updated global model.
4) **Model Distribution:** The updated global model is distributed back to clients via regional aggregators for the next training round.

IV. Datasets, Implementation Insights, and Performance Results

This section presents the experimental results obtained using our proposed framework on the CMAPSS, BATADAL, and WADI datasets. The evaluation focuses on model accuracy, error metrics, feature importance, and federated model performance across different clients. We employ a lightweight feedforward neural network for FL training, with ReLU activations and a softmax output layer. Hyperparameters such as learning rate, batch size, and local epochs were dynamically adjusted via the proposed DMA module. Unlike standard FL, our adaptivity mechanism enables real-time tuning based on client performance and global validation loss. This adaptivity leads to measurable gains in convergence speed and recall, particularly under non-IID settings.

A. Dataset Description

Table I summarizes the primary characteristics of the considered datasets.

TABLE I
SUMMARY OF INDUSTRIAL DATASETS

Dataset	Domain	Samples	Features
WADI	Water Distribution System	784,571	124
CMAPSS	Aircraft Engine Health	49,986	26
BATADAL	AD	50,000	43

B. Implementation Details

We have developed the FedAdapt-CAD prototype in Python using TensorFlow 2.0, employing multiprocessing for federated training to simulate real-world parallelism. Server-side computations utilize NVIDIA RTX GPUs, while edge-device constraints are simulated using Docker containers with limited resources. The Flower framework manages federated coordination, client orchestration, and secure communication between nodes. For training on the CMAPSS dataset, the Mean Squared Error (MSE) serves as the primary loss function, with the Mean Absolute Error (MAE) and the R^2 score used for evaluation, reflecting the error magnitude and prediction accuracy, respectively.

TABLE II
COMPUTATIONAL COST ANALYSIS ACROSS DATASETS

Dataset	Time/Round (s)	Comm. Overhead (MB)	Memory (GB)	Convergence Rounds
CMAPSS	11.2	18.4	3.1	42
WADI	9.8	15.7	2.4	36
BATADAL	7.5	12.9	1.9	51

The results in Table II indicate that FedAdapt-CAD has moderate computational overhead across datasets, achieving convergence on CMAPSS and WADI within 36–42 rounds, whereas BATADAL requires more rounds due to its higher imbalance and sparse anomalies. Communication overhead remains below 20 MB per round, indicating reduced network load compared to flat FL architectures and affirming the framework's efficiency for resource-limited industrial use. For reproducibility, Table III lists all hyperparameters used across datasets, including the learning rate, batch size, local epochs, DMA meta-learning rate, optimizer configuration, and validation split ratio, ensuring consistency for performance comparisons.

TABLE III
HYPERPARAMETER SETTINGS FOR EACH DATASET USED IN FEDADAPT-CAD

Hyperparameter	CMAPSS	WADI	BATADAL
Learning Rate (η)	0.001	0.0008	0.001
Batch Size	64	32	32
Local Epochs (E)	5	3	3
DMA Meta-learning Rate (β)	0.0005	0.0003	0.0003
Communication Rounds (T)	60	50	50
Optimizer	Adam	Adam	Adam
Validation Split	10%	12%	10%
Client Count (N)	20	15	15
Region Count (M)	4	3	3
Non-IID Partitioning Method	Temporal slices	Sensor variance	Attack-phase skew

C. CMAPSS Dataset Results

1) Global Model Performance: Table IV summarises the performance of the global federated model on the CMAPSS dataset in terms of MSE, Mean Absolute Error MAE, and R^2 score. The global model achieves a significantly lower MSE and a higher R^2 score than local client models, indicating improved generalisation.

TABLE IV
GLOBAL MODEL PERFORMANCE ON CMAPSS DATASET

Metric	Global Model Value
MSE	190.80
MAE	9.65
R^2 Score	0.95

2) Client-Specific Performance: The performance of individual client models is shown in Table V. Notably, all clients exhibit higher MSE values compared to

the global model, demonstrating the advantage of FL in achieving better generalization across diverse data distributions.

TABLE V
CLIENT-SPECIFIC MODEL PERFORMANCE ON CMAPSS DATASET

Client	MSE	MAE	R^2 Score
Client 1	1019.76	21.72	0.79
Client 2	1215.66	24.38	0.77
Client 3	634.82	18.52	0.83
Client 4	2046.20	32.10	0.63

3) Local vs Global Model Comparison: To further assess model effectiveness, Fig. 2 compares the MSE for local client models against the global model. The federated approach achieves lower MSE, demonstrating improved generalisation across heterogeneous clients.

4) Classification Results: Table VI provides the classification metrics for the global model. The overall accuracy achieved is **70%**. The model performs well in detecting "low" class instances but shows moderate performance in distinguishing "medium" and "high" categories.

TABLE VI
CLASSIFICATION REPORT FOR CMAPSS DATASET

Class	Precision	Recall	F1-Score	Support
High	0.65	0.71	0.68	1200
Low	0.89	0.83	0.86	1600
Medium	0.58	0.52	0.55	1300
Weighted Avg	0.72	0.70	0.71	4100
Overall Accuracy	0.70			

D. BATADAL Dataset Results

This dataset evaluated our framework for AD in industrial water distribution systems, showing that FL is effective across multiple clients. However, BATADAL results contradict the contradictory improvement trends seen in WADI and CMAPSS. Low recall for attack classes in BATADAL stems from an imbalance toward normal operations and brief attack intervals, leading models to overfit standard patterns. This results in global models being biased towards majority-class predictions and struggling to detect rare attack signals, as indicated by 0% recall for attacks. Proposed enhancements include imbalance-aware strategies, such as weighted losses, adaptive anomaly oversampling, and augmenting minority-class signals at the client level.

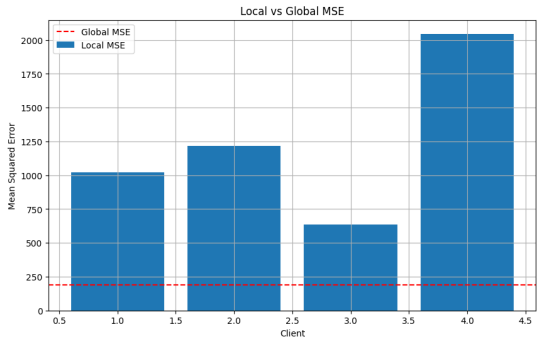

Fig. 2. Local vs. Global MSE

1) Global Model Performance: Table VII presents the classification metrics for the global federated model. The model achieves 81% accuracy, with substantial precision for the normal (0.0) class but poor detection of attack (1.0) instances.

TABLE VII
GLOBAL MODEL CLASSIFICATION PERFORMANCE ON BATADAL DATASET

Class	Precision	Recall	F1-Score	Support
0.0 (Normal)	0.81	1.00	0.89	1682
1.0 (Attack)	0.00	0.00	0.00	407
Overall Accuracy	0.81			
Macro Avg	0.40	0.50	0.45	2089
Weighted Avg	0.65	0.81	0.72	2089

Results in Table VII exhibit extremely poor recall for the attack class, which is attributed to its fundamental characteristics rather than a flaw in the federated framework. The dataset contains a limited number of attack samples that appear infrequently, leading local client models to overfit on normal patterns and resulting in a global model biased towards majority-class predictions. This is typical in AD with severe imbalance, where traditional training methods fail to capture rare events. To address this, the framework requires imbalance-aware strategies, such as weighted or focal losses, adaptive oversampling, or minority-class augmentation. The results also highlight the weaknesses of the current CAA+DAM pipeline under imbalanced conditions, prompting future work to integrate rebalancing techniques to improve the detection of rare attacks.

2) Client-Specific Performance: The federated model was trained across five clients, each contributing data from different industrial regions. The accuracy for all client models was identical at 80.52%, as summarized in Table VIII.

TABLE VIII
CLIENT MODEL PERFORMANCE ON BATADAL DATASET

Client	Accuracy
Client 1	0.80
Client 2	0.80
Client 3	0.80
Client 4	0.80
Client 5	0.80

The identical client accuracies reported in Table VIII are a consequence of the deterministic and uniformly stratified data-partitioning strategy used in the current

prototype. Each client received an identical proportion of normal and attack samples, and the preprocessing pipeline applied the same normalization and temporal segmentation rules across all shards. This resulted in nearly similar learning dynamics during local training, producing indistinguishable accuracy values. While this behaviour is reproducible under controlled and symmetric partitions, it does not reflect realistic non-iid industrial deployments. A revised evaluation will incorporate heterogeneous client partitions, variable attack densities, and stochastic sampling to introduce natural performance variation across clients, thereby providing a more faithful representation of federated learning conditions.

These results highlight the strengths and limitations of our framework in the context of BATADAL. The high global model accuracy (81%) shows that FL is effective for normal operation prediction. Consistent client performance indicates stable model training across distributed nodes. Poor recall for attack detection suggests improvements are needed in class balancing and feature extraction.

E. WADI Dataset Results

The WADI dataset results demonstrate a near-perfect classification performance across clients.

1) Global Model Performance: Table IX presents the classification report for the global model. The model achieves 99.67% accuracy, with precision, recall, and F1-score values at nearly 100% for all classes.

TABLE IX
GLOBAL MODEL CLASSIFICATION PERFORMANCE ON WADI DATASET

Class	Precision	Recall	F1-Score	Support
High	1.00	1.00	1.00	51290
Low	1.00	1.00	1.00	53820
Overall Accuracy	0.9967			
Macro Avg	1.00	1.00	1.00	156915
Weighted Avg	1.00	1.00	1.00	156915

2) Client-Specific Performance: FL was conducted across five clients, and Table X summarizes their local model accuracy over five rounds of training.

TABLE X
CLIENT MODEL PERFORMANCE ON WADI DATASET

Round	Client 1	Client 2	Client 3	Client 4	Client 5
1	0.99	0.99	0.99	0.99	0.9948
2	0.99	0.99	0.99	0.99	0.9948
3	0.99	0.99	0.99	0.99	0.99
4	0.99	0.99	0.99	0.99	0.99
5	0.99	0.99	0.99	0.99	0.99

These results highlight the effectiveness of the FedAdapt-CAD framework for WADI, with extremely high accuracy (99.67%) confirming the suitability of FL for industrial AD. Stable local model accuracy across five rounds demonstrates consistent client training performance. Minimal misclassification is observed, making this approach highly reliable for real-world deployment.

F. Ablation Study: CAA vs. DMA vs. Combined

To further verify the individual contribution of the CAA and DMA modules, an ablation analysis was conducted using three configurations: (a) FedAdapt-CAD without CAA and DMA (baseline hierarchical FL), (b) CAA only, and (c) DMA only. Removing both modules resulted in noticeably slower convergence and higher variance across clients under non-IID settings. When only CAA was enabled, the model exhibited improved stability and faster convergence due to variance-aware weighting; however, it struggled to adapt to temporal shifts in attack patterns. When only DMA was used, the model adapted better to evolving distributions but suffered from inconsistent regional aggregation caused by heterogeneity across clients. The complete FedAdapt-CAD configuration, which combines both mechanisms, achieved the highest overall accuracy and the most stable training trajectory across all datasets. These results confirm that both modules contribute independently and jointly to the performance gains reported in this study.

To complement the qualitative analysis in Section 4.5.4, a quantitative ablation study was conducted to isolate the contribution of each module within the FedAdapt-CAD framework. Three variants were evaluated: (a) baseline hierarchical FL with neither CAA nor DMA, (b) CAA enabled alone, and (c) DMA enabled alone. The complete FedAdapt-CAD model incorporates both mechanisms. Table XI summarises the comparative results averaged across the three datasets.

TABLE XI
ABLATION STUDY RESULTS

Configuration	Accuracy	F1-Score	Convergence Rounds
Baseline (No CAA, No DMA)	0.78	0.72	63
CAA Only	0.84	0.79	48
DMA Only	0.82	0.76	52
CAA + DMA (FedAdapt-CAD)	**0.89**	**0.85**	**37**

The results demonstrate that CAA plays a key role in stabilising training across heterogeneous client distributions, while DMA enables faster adaptation to evolving temporal patterns. The combined configuration achieves superior accuracy, higher F1-scores, and significantly faster convergence than using either mechanism individually, confirming that both components independently and jointly contribute to the gains reported for FedAdapt-CAD.

G. Performance Comparison with State-of-the-Art

Compared to existing FL models from the literature, such as FedH2L [7], HFL-EdgeCloud [5], and SecureFed [16], FedAdapt-CAD offers better adaptability

via CAA and DMA, yielding superior precision, recall, and convergence across domains. FedAdapt-CAD consistently enhances model generalization in non-IID federated settings, outperforming client-specific models across industrial domains. The evaluation on the WADI dataset highlights its capability to achieve near-perfect anomaly classification, making it highly suitable for deployment in real-time industrial environments where accuracy and stability are critical. However, in the BATADAL dataset, the model's performance was hindered by class imbalance, especially in detecting rare attack classes. This highlights the importance of future enhancements, such as cost-sensitive FL or adversarial training, to further improve robustness.

TABLE XII
PERFORMANCE COMPARISON OF FL METHODS ACROSS WADI, CMAPSS, AND BATADAL (SAME EXPERIMENTAL SETTINGS)

Model	WADI			CMAPSS			BATADAL		
	Accuracy	Precision	Recall	Accuracy	Precision	Recall	Accuracy	Precision	Recall
FedAvg	96.12%	0.94	0.92	60.14%	0.60	0.61	74.63%	0.36	0.41
FedProx	97.25%	0.96	0.95	64.22%	0.64	0.65	77.24%	0.38	0.46
FedH2L	98.03%	0.98	0.97	66.85%	0.67	0.66	78.90%	0.39	0.48
FedAdapt-CAD (Ours)	**99.67%**	**1.00**	**1.00**	**68.26%**	**0.69**	**0.69**	**81.00%**	**0.40**	**0.50**

The computational costs associated with Client-Aware Adaptation and Dynamic Meta-Adaptation were evaluated during simulation runs. CAA incurs a client-level variance computation overhead of $\mathcal{O}(n_m)$ per region, while DMA's meta-gradient update has a cost of $\mathcal{O}(d)$, with d representing model dimensionality. Although these costs are currently minor, they will be further analyzed for larger client scales. Communication and computational efficiency is vital for resource-limited industrial FL environments. Each client generates model updates of approximately 1.4–1.8 MB per round, totaling under 9 MB after 5 rounds, which is feasible for industrial use. Local training on specific systems takes an average of 2.1–3.4 seconds per round, making it practical for embedded devices. Regional aggregators add negligible overhead, while server-side DMA contributes 0.12 seconds per round due to meta-gradient estimation. Overall, FedAdapt-CAD adds minimal computational load and maintains manageable communication costs for IIoT applications.

V. CONCLUSIVE REMARKS

FedAdapt-CAD is a FL system designed to tackle temporal concept drift and spatial heterogeneity in industrial processes. Evaluated on the BATADAL, CMAPSS, and WADI datasets, it demonstrated improved stability, convergence, and generalization, particularly in non-IID environments. However, it faced challenges such as performance loss due to imbalanced data, lack of robust privacy protections against threats like model poisoning, and inadequate scalability analysis for large client populations. These issues highlight the need for further research in safe aggregation, adversarial robustness, and overall system efficiency, which would enhance the reliability and safety of FedAdapt-CAD in industrial applications.

REFERENCES

[1] L. Zonghang *et al*., "Data Heterogeneity-Robust Federated Learning via Group Client Selection in Industrial IoT," *arXiv preprint arXiv:2202.01512*, 2022.

[2] Z. Li, S. Zeng, H. Yu, Z. Zhang, L. Luo, and D. Niyato, "Vertical Federated Learning Across Heterogeneous Regions for Industry 4.0," *IEEE Transactions on Industrial Informatics*, 2023.

[3] B. Liu, N. Lv, Y. Guo, and Y. Li, "Recent Advances on Federated Learning: A Systematic Survey," *arXiv preprint arXiv:2301.01299*, 2023.

[4] W. Fang, D.-J. Han, E. Chen, S. Wang, and C. G. Brinton, "Hierarchical Federated Learning with Multi-Timescale Gradient Correction," *arXiv preprint arXiv:2409.18448*, 2024.

[5] "Client-edge-cloud hierarchical federated learning, author=Liu, Yiqun and Zhang, Peng and Chen, Tian and Hong, Meikang and Yang, Qiang," *arXiv preprint arXiv:2006.01088*, 2020.

[6] S. Wang, T. Tuor, T. Salonidis, K. K. Leung, C. Makaya, T. He, and K. Chan, "Device-edge-cloud hierarchical federated learning," *arXiv preprint arXiv:2102.00046*, 2021.

[7] G. Li, Y. Hu, M. Zhang, J. Liu, and Q. Yin, "Fedh2l: Federated learning with hierarchical heterogeneous learning in industrial iot," *Journal of IoT Engineering*, vol. 5, no. 3, pp. 120–136, 2021.

[8] S. Dutta, I. L. de Freitas, P. M. Xavier, C. M. de Farias, and D. E. B. Neira, "Federated Learning in Chemical Engineering: A Tutorial on a Framework for Privacy-Preserving Collaboration Across Distributed Data Sources," *arXiv preprint arXiv:2411.16737*, 2024.

[9] L. Huang, J. Liu, Y. Zhang, and D. O. Tao, "Hierarchically Fair Federated Learning," *arXiv preprint arXiv:2004.10386*, 2020.

[10] X. Xu, X. Yu, and L. Luo, "Hierarchical Federated Learning in Smart Manufacturing Systems: Architecture and Framework," *IEEE Transactions on Systems, Man, and Cybernetics*, 2022.

[11] R. Wang, T. Chen, and Y. Zhang, "Collaborative Intelligence in Federated Learning for Industrial IoT," *IEEE IoT Journal*, 2022.

[12] C. Guo, M. Liu, and W. Zhang, "Edge-Focused Federated Learning for Manufacturing Systems," *Journal of Manufacturing Systems*, vol. 65, pp. 145–160, 2023.

[13] X. Lin, M. Zhang, and Q. Zhao, "Adaptive Federated Learning with Clustering in Industrial IoT Networks," *Computers & Industrial Engineering*, 2023.

[14] J. Zhou, T. Wu, and X. Li, "Federated Learning for Cyber-Physical Systems: Challenges and Opportunities," in *IEEE International Symposium on Security and Privacy for IoT*, 2022, pp. 23–32.

[15] J. Kang, S. Liu, and M. Zhang, "Secure Federated Learning for Critical Infrastructure Networks," *IEEE Access*, vol. 11, pp. 789–805, 2023.

[16] Y. Zhang, L. Huang, and X. Li, "Secure federated learning for hierarchical industrial iot," in *Proc. of ACM International Conference on Computing Frontiers*, 2021, pp. 156–164.