AutoML-DAFL : A Drift-Aware Federated Learning Framework for Cyber-Physical Aquaculture Water Quality Monitoring

Hassan Sajjad
UR 3926, Laboratoire de Génie Informatique et d'Automatique de l'Artois (LGI2A), Universite d'Artois
Béthune, France
hassan.sajjad@univ-artois.fr

Adnen El Amraoui
UR 3926, Laboratoire de Génie Informatique et d'Automatique de l'Artois (LGI2A), Universite d'Artois
Béthune, France
adnen.elamraoui@univ-artois.fr

François Delmotte
UR 3926, Laboratoire de Génie Informatique et d'Automatique de l'Artois (LGI2A), Universite d'Artois
Béthune, France
francois.delmotte@univ-artois.fr

***Abstract*—Federated Learning (FL) for Internet of Things (IoT) and cyber-physical sensor networks, such as aquaculture water monitoring, faces critical challenges due to temporal sensor drift, non-IID data distributions, and communication constraints across edge devices, which compromise global model stability and resource efficiency. However, most existing federated approaches lack coherent mechanisms to address these issues, leading to degraded performance in realistic edge deployments. We present AutoML-DAFL, a drift-aware, AutoML-guided federated learning framework with a multi-objective reward controller that jointly optimizes predictive accuracy, model consistency, and communication efficiency. The framework integrates temporal drift detection and mitigation into the federated training loop through MAE-based regularization while ensuring persistent convergence. To evaluate its effectiveness, we benchmark AutoML-DAFL against FedAvg and advanced baselines, including FedNova and SCAFFOLD. Extensive ablation studies and comparative analysis on real aquaculture monitoring data demonstrate the contribution of each reward component: removing the communication-aware term degrades model consistency, while excluding MAE-based smoothing reduces training stability. The full AutoML-DAFL configuration, integrating all reward components through multi-objective optimization, achieves the lowest RMSE (0.0721), highest R^2, and improved fairness across clients, demonstrating strong resilience to non-IID drift and bandwidth constraints. These results highlight the effectiveness of drift-aware AutoML optimization for resource-efficient, stable federated forecasting in cyber-physical monitoring systems.**

***Index Terms*—Federated Learning, AutoML, Drift detection, Water quality forecasting, Aquaculture Monitoring, Cyber-Physical Systems**

Aquaculture has become a primary contributor to the global seafood supply chain, with an increasingly substantial propagation of production originating from controlled aquaculture systems rather than traditional wild fisheries. Ensuring sustainable productivity requires precise regulation of critical environmental parameters, including temperature, pH, dissolved oxygen, and turbidity, that directly influence fish health and farm efficiency. Therefore, continuous monitoring and accurate forecasting of the Water Quality Index (WQI), an integrated indicator of aquatic health, are essential for detecting early anomalies and providing pervasive, real-time decision support across modern aquaculture systems. Recent advances in IoT-based sensing and edge computing have enabled large-scale, real-time environmental monitoring; however, the data produced across distributed aquaculture sites remained siloed due to privacy, ownership, and commercial confidentiality constraints. These constraints limit the feasibility of centralized model training and underscore the need for decentralized, AI-driven privacy-preserving learning frameworks to ensure smart and sustainable automation for decision support. [1].

Federated Learning (FL) offers a promising solution by enabling distributed clients to collaboratively train a shared global model without exchanging raw data [2]. However, standard federated algorithms such as FedAvg and FedProx [3] inherently assume that data distributions across clients remain stationary over time. In practice, real aquaculture systems exhibit pronounced temporal variability driven by seasonal cycles, sensor drift, calibration inconsistencies, and operational interventions. These non-stationary features introduce temporal concept drift, reducing regression accuracy and producing biased global model updates in traditional federated models. While recent studies proposed drift-aware and normalization FL strategies, most have been evaluated primarily on classification benchmarks [4], [5]. Consequently, the challenges of drift-aware regression, particularly for non-stationary, temporally evolving environmental sensor data, remain insufficiently addressed in contemporary literature.

Beyond predictive accuracy, maintaining performance across distributed sites is critical in environmental monitoring applications. Fairness-oriented FL methods such as q-Fedavg, FairFed [6], [7], aim to mitigate inter-client disparities assuming that clients data remain static over time. However, temporal drift unfolds progress at uneven rates across the clients, causing static fairness mechanisms to miscalibrate. Thus, achieving fairness in federated regression under temporal shifts remains a largely unexplored research problem.

To address these challenges, we proposed a Drift-Aware Federated Learning (DAFL) framework designed for temporally evolving aquaculture systems. Our key contributions are summarized as follows:

- We measure temporal drift across aquaculture stations using a feature-wise Wasserstein distance with exponential moving-average smoothing, providing a stable, geometric-aware estimate of evolving sensor distributions.

- We introduce a β controlled server-side aggregation rule that dynamically adjusts client contributions based on drift severity, while keeping client-side computation overhead unchanged.
- A ϵ greedy controller automatically adjusts the drift-sensitivity β during training by optimizing a light-weight, multi-objective reward function that balances model accuracy, drift stability, and communication efficiency. It eliminates the need for manual hyperparameter tuning and enables responsive adaptation to drift patterns.
- We demonstrate that the proposed approach consistently improves accuracy and stability compared to widely adopted federated optimization methods, i.e., FedNova, SCAFFOLD, and FedAvg under realistic drift settings, while an ablation analysis confirms meaningful contributions of each component of the adaptive β mechanism to overall robustness.

I. Related Work

Federated learning (FL) enables privacy-preserving model training across distributed clients without sharing sensitive data. Various foundational methods, such as FedAvg algorithm [2], FedProx [3], and SCAFFOLD [8], improve robustness under data heterogeneity by utilizing control variates to correct client drift. Communication-oriented variants like FedNova [9] optimize communication efficiency by normalizing local update steps. However, these approaches presume stable client data distributions, which rarely hold in continuous, non-stationary IoT sensor-driven systems.

Aquaculture water-quality monitoring generates time-series streams that are inherently non-stationary in nature. Prior work on water quality prediction, including LSTM-based forecasting [10], dissolved oxygen prediction [11], and IoT-enabled aquaculture monitoring systems [12], demonstrates that environmental variables exhibit strong temporal variability. These studies highlight the importance of modeling temporal patterns in aquaculture but are centrally trained and do not address federated settings.

Concept drift is a primary challenge in sensor-driven environments, where seasonal variations, operational challenges, and sensor degradation alter underlying data distribution over time. Prior studies on concept drift [13] and sensor-focused approaches [14] demonstrate that temporal shifts in data distribution can significantly reduce predictive performance unless the model explicitly compensates for these distributional shifts. Existing drift detection techniques span statistical tests and distributional measures [4], [15], yet most adaptation strategies, including sliding window updates and ensemble weighting, are predominantly designed for centralized or classification-centric settings, providing limited insights for regression tasks.

Recent federated learning efforts have explored normalization-based, multi-modal drift handling [5], yet drift-aware FL for temporally continuous regression tasks, i.e., water quality, remains largely unexplored. This gap is critical for aquaculture environments, where uninterrupted temporal coherence is essential for early anomaly detection and reliable WQI forecasting.

Fairness in FL aims to achieve consistent performance across heterogeneous clients with varying data. Approaches such as FairFed and q-FedAvg [7], [16] incorporate fairness using the Jain index and the Gini coefficient to regulate client influence during training. These methods are widely explored in classification scenarios, offering limited insights for regression tasks affected by temporal variability, i.e., water quality prediction.

Automated machine learning (AutoML) has largely focused on centralized pipelines for architecture search and hyperparameter optimization. In contrast, Federated AutoML remains limited with few efforts addressing, i.e., neural architecture search, i.e., FedNAS [17], and restricted tuning of client-side configurations. Classical Bandit algorithms, including ϵ, UCB, Thompson sampling, provide principled exploration-exploitation strategies for online decision making [18]. Our framework leverages ϵ greedy for dynamic β-selection during training. While recent FL systems employ bandits for client selection and resource scheduling [19], [20], they don't use bandits for aggregation parameter tuning,drift-aware adaptation, and fairness-oriented optimization. Hyperparameter tuning in FL is still limited with β, and other parameters are usually fixed, guided by domain heuristics and grid search, and remain unchanged throughout training rather than evolving to client data distributions. Multi-objective optimization has been explored to balance accuracy and communication efficiency [21], [22], but lacks mechanisms for fairness-aware and round-wise automated hyperparameter optimization.

AutoML-DAFL addresses limitations in existing federated AutoML approaches by combining drift-aware aggregation with ϵ greedy bandit-based online β selection.Its multi-objective reward formulation jointly accounts for predictive performance, fairness, and drift sensitivity, enabling automated,round-wise hyperparameter optimization without relying on manually tuned settings. Overall, it provides a unified AutoML drift-aware federated learning framework suitable for non-stationary temporal regression tasks.

II. Drift-Aware Federated Learning (DAFL)

A. Problem Formulation

We consider N aquaculture stations, each contributing a local dataset $\mathcal{D}_i = \{(\mathbf{x}_{i,j}, y_{i,j})\}_{j=1}^{n_i}$, where $\mathbf{x}_{i,j} \in \mathbb{R}^d$ denotes sensor variables and $y_{i,j}$ represents the computed Water Quality Index (WQI). Standard Federated Averaging (FedAvg) aggregates client updates using $\alpha_i = n_i / \sum_j n_j$ [2]. This approach fails in non-stationary aquaculture environments where temporal data drift arises from sensor aging, fouling asynchronous calibrations, leading to biased global model updates [13]. The AutoML-DAFL approach addresses these limitations by explicitly incorporating drift magnitude into the aggregation mechanism.

B. Drift Quantification and Smoothing

We quantify the distributional shift of these drifted clients using feature-wise Wasserstein-1 distance (W_1) [23], a stable metric comparing the initial marginal $P_{i,k}^{\text{init}}$ to the current distribution $P_{i,k}^{\text{val}}$ for each feature t:

$$W_i^{(t)} = \frac{1}{d} \sum_{k=1}^{d} W_1\big(P_{i,k}^{\text{init}}, P_{i,k}^{\text{val}}\big)\,. \tag{1}$$

To reduce computational cost, if a distribution contains more than 1000 samples, uniform sampling is applied.For addressing fluctuating raw drift values across rounds of training, we apply an exponential moving average (EMA):

$$\bar{W}_i^{(t)} = \lambda\, \bar{W}_i^{(t-1)} + (1-\lambda) W_i^{(t)}. \tag{2}$$

C. Drift Normalization and Adaptive Aggregation

The smoothed drift estimates are max-normalized to allow fair comparison across clients:

$$\tilde{W}_i^{(t)} = \frac{\bar{W}_i^{(t)}}{\max_j \, \bar{W}_j^{(t)}}. \tag{3}$$

This ensures that drift values lie within a comparable range, preventing clients with large absolute drift from disproportionately influencing the aggregation.

The drift-aware aggregation weight $\alpha_i^{(t)}$ for client i at round t is defined by scaling the size-based weight with a dynamic drift factor:

$$\alpha_i^{(t)} = \frac{n_i\big(1 + \beta^{(t)}\, \tilde{W}_i^{(t)}\big)}{\sum_j n_j\big(1 + \beta^{(t)}\, \tilde{W}_j^{(t)}\big)}. \tag{4}$$

Here $\beta^{(t)} \in \mathcal{B}$ is a dynamically selected drift factor at round t that controls the intensity of aggregation. This adaptation ensures that the global model adapts effectively to temporal dynamics. The resulting global update is computed as:

$$\mathbf{w}^{(t+1)} = \sum_{i=1}^{N} \alpha_i^{(t)} \mathbf{w}_i^{(t)}. \tag{5}$$

By scaling each client contribution using drift-aware weights, this aggregation approach produces global updates that remain stable during distributional shifts, ensuring the global model adapts to real-world, evolving conditions.

D. Adaptive β selection via multi-objective bandit

To overcome the need for static hyperparameter β tuning, We introduce an AutoML component via a server-side ϵ greedy bandit controller. The controller adaptively selects the optimal aggregation $\beta^{(t)}$ from a predefined set of candidates $B = \{0.5, 0.8, 1.0, 1.5\}$.

The controller uses a Q-learning approach to learn the long-term value of each β candidate.In each round, it employs an ϵ-greedy strategy: it explores by sampling a random β with probability ϵ, and exploits the best current β ($\arg\max_\beta Q(\beta)$) otherwise.The controller holds β fixed for the first three warm-up rounds, and then adapts it every two rounds. It begins with ϵ=0.2, decays this value by a factor of 0.95 after each selection with a floor of 0.005.

Algorithm 1 AutoML-DAFL with Adaptive β Controller

Require: Rounds T, epochs E, smoothing factor λ, β candidates $\mathcal{B}$, bandit factors ϵ, γ, α

Initialize global model $\mathbf{w}^{(0)}$ and drift buffers $\bar{W}_i^{(0)} = 0$
Initialize bandit Q-values for all $\beta \in \mathcal{B}$
for $t = 0$ to $T-1$ **do**
  Broadcast $\mathbf{w}^{(t)}$ to all clients
  Select $\beta^{(t)}$ using ϵ-greedy bandit
  for each client i in parallel **do**
    Train locally to obtain $\mathbf{w}_i^{(t)}$
    Compute drift $W_i^{(t)}$ using (1)
    Return $(\mathbf{w}_i^{(t)}, W_i^{(t)}, C_i^{(t)})$ and metrics
  end for
  Update smoothed drift $\bar{W}_i^{(t)}$ using (2)
  Normalize drift $\tilde{W}_i^{(t)}$ using (3)
  Compute weights $\alpha_i^{(t)}$ using (4)
  Update global model $\mathbf{w}^{(t+1)}$ using (5)
  Compute reward $R^{(t)}$ using (6)
  Update bandit Q-values using (7)
end for
return $\mathbf{w}^{(T)}$

After each round, the controller calculates a reward function $R^{(t)}$ that simultaneously balances three critical objectives: convergence improvement (ΔMAE), drift stability cost (D_{cost}), and communication overhead (C_{cost}).

To ensure numerical stability and predictable weighting across objectives, each term is normalized using a sigmoid-like function $f(x)$, which maps all non-negative inputs to a bounded interval $[0, 1)$ $f(x) = \frac{x}{1+x}$. The multi-reward function $R^{(t)}$ is defined as :

$$R^{(t)} = w_{\text{MAE}}\, f(\Delta\text{MAE}) - w_{\text{Drift}}\, f(D_{\text{cost}}^{(t)}) - w_{\text{Comm}}\, f(C_{\text{cost}}^{(t)}) \tag{6}$$

We normalize $w_{\text{MAE}} = 1.0$, $w_{\text{Drift}} = 0.3$, and $w_{\text{Comm}} = 0.2$ tuned for the aquaculture domain, to prioritize model convergence.

By bounding the components,$f(x)$ prevents any single metric, e.g., large communication cost, from dominating or saturating the reward signal. This allows the bandit controller to consistently learn the optimal $\beta^{(t)}$ that maximizes the weighted trade-off (defined by w) between performance gain and resource penalties.

Finally, the Q-values are updated using an EMA-based Q-learning rule as follows:

$$Q^{(t)}(\beta_t) \leftarrow (1-\alpha) Q^{(t-1)}(\beta_t) + \alpha[R^{(t)} + \gamma Q^{(t)}(\beta_t)], \tag{7}$$

Where α = 0.5 learning rate, γ = 0.95 discount factor, while drift estimates are smoothed using α = 0.6.

Algorithm 1 outlines the training workflow of the proposed drift-aware federated framework, integrating drift estimation, smoothing, and adaptive aggregation in the standard FL loop.

The next section describes the data and experimental setup utilized to evaluate the DAFL approach.

III. Experimental Setup

We evaluate our experiments on a real-world aquaculture IoT sensor dataset (Bangladesh aquaculture farms) [24], comprising real-time sensor readings collected from both extensive and intensive aquaculture environments in Bangladesh across five modalities: temperature, pH, dissolved oxygen, turbidity, and total dissolved solids. The target variable is the Water Quality Index (WQI) derived from these parameters following standard aquaculture assessment controls.

A. Federated Partitioning

We partition the dataset into five clients via a controlled temporal strategy inducing *gradual feature-space concept drift* [13], where sensor distributions shift monotonically across clients while target relationships remain locally stable. Instead of allocating disjoint segments to clients, we extract overlapping rolling windows of length 12 timestamps with a stride of 1 to maximize training data density while preserving short-term temporal dependencies. This approach follows standard practice in time-series drift modeling to ensure a smooth progression of environmental conditions across clients, while using a synthetic federation topology and real-world sensor measurements.

Feature-wise Wasserstein-1 analysis confirms structured non-IID heterogeneity, validating a structured monotonic drift across the five clients. Early clients exhibit warmer, higher oxygen levels, while later clients shift towards lower oxygen, cooler conditions. Figure 1 illustrates the resulting Wasserstein heatmap across all partitioning clients, confirming this distributional shift. This provides a time-series benchmark necessary for evaluating drift-aware aggregation.

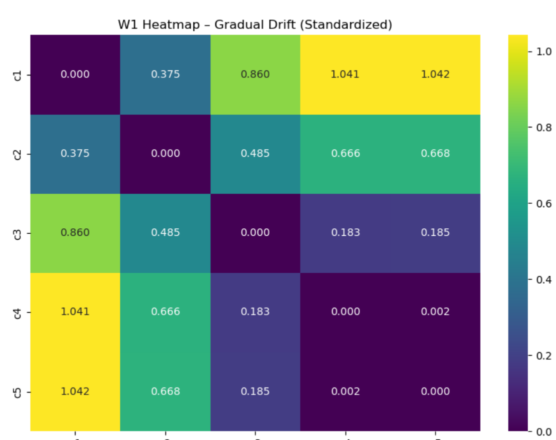

Fig. 1. Wasserstein-1 Distance Heatmap across client partitions validating the non-IID partitioning strategy

B. Data Pre-processing

During preprocessing, feature standardization is performed using a global standard scaler across the entire dataset, ensuring a consistent feature space for Wasserstein drift computation. The global mean and standard deviation are computed once; each client applies this pre-computed global scaler locally, eliminating the need for data sharing. After normalization, temporal sequences are constructed using a rolling window encoder (window=12, stride 1) to capture short-term transitions in water quality measurements.

C. Implementation Details

We conduct four experiments using identical model architectures, data partitions, and training configurations:

- **Experiment 1 - FedAvg (Baseline)** Standard FedAvg with data-proportional weighting .
- **Experiment 2 - FedNova** Normalized aggregation to handle heterogeneous local updates .
- **Experiment 3 - SCAFFOLD** Control variates correction to handle concept drift .
- **Experiment 4 - AutoML-DAFL (Proposed)** Drift-aware aggregation with adaptive β selection via ϵ-greedy bandit controller.Controller selects β from candidates $\{0.5, 0.8, 1.0, 1.5\}$ based on multi-objective reward.Drift smoothing parameter is λ= 0.6.

1) Model Architecture: All four approaches (FedAvg, FedNova, SCAFFOLD, AutoML-DAFL) employ the identical hybrid CNN-LSTM architecture with 457,217 parameters to ensure fair comparison across aggregation strategies. Table I summarizes the architecture of the proposed model. The model utilized 1-D convolution for extracting short-term temporal patterns, followed by Batch and Layer normalization to improve training consistency. These normalized features are fed to two Bidirectional LSTM layers for capturing long-term dependencies. Dropout and L2 regularization are applied for generalization.

TABLE I
Hybrid CNN–BiLSTM Model Architecture

Layer	Output Shape	Params
Input	(3, 5)	0
Conv1D (128 filters)	(3, 128)	2,048
BatchNorm	(3, 128)	512
LayerNorm	(3, 128)	256
Bi-LSTM 1 (128 units, return seq.)	(3, 256)	263,168
Bi-LSTM 2 (64 units)	(128)	164,352
Dense (128 units)	(128)	16,512
Dropout ($p = 0.3$)	(128)	–
Dense (64 units)	(64)	8,256
Dropout ($p = 0.3$)	(64)	–
Dense (32 units)	(32)	2,080
Output (1 unit, linear)	(1)	33
	Total Parameters	**457,217**

2) Training Configuration: Federated learning is implemented using the Flower framework with simulation mode. We ran 15 communication rounds with 3 local epochs per round, batch size 64, Adam optimizer (learning rate 1×10^{-5}). All clients participate in every training round.

IV. Results and Discussion

The proposed AutoML-DAFL framework was benchmarked against FedAvg, FedNova, and SCAFFOLD under temporal data shifts. These baselines serve as state-of-the-art benchmarks for mitigating client drift and update heterogeneity, representing two primary challenges in federated cyber-physical systems.

A. Overall Performance Comparison

Table II summarizes the comparative analysis of evaluated federated learning approaches. Figure 2 illustrates the global convergence of evaluated federated learning methods across communication rounds, illustrating the global convergence of loss (MSE) and MAE under heterogeneous client updates.

The most critical finding is the superior robustness of AutoML-DAFL, which achieved the lowest overall RMSE of 0.0721, while comparative algorithms struggled under continuous temporal drift.FedNova achieved an RMSE of 0.256, while SCAFFOLD experienced performance degraded significantly with an RMSE of 0.7055 and the lowest R^2 of 0.5030. In comparison, AutoML-DAFL achieved a 32% reduction in global RMSE compared to standard Fedavg (0.1062), and other baselines confirm the adaptive aggregation mechanism's effectiveness in controlling peak prediction errors caused by local model divergence.

Simultaneously, the framework minimized global MAE to 0.0304, validating its superiority in average prediction accuracy across the entire sensor network. Furthermore, the global R^2 of 0.9947 confirms that AutoML-DAFL offers more reliable decision support for aquaculture monitoring than any of the evaluated baselines.

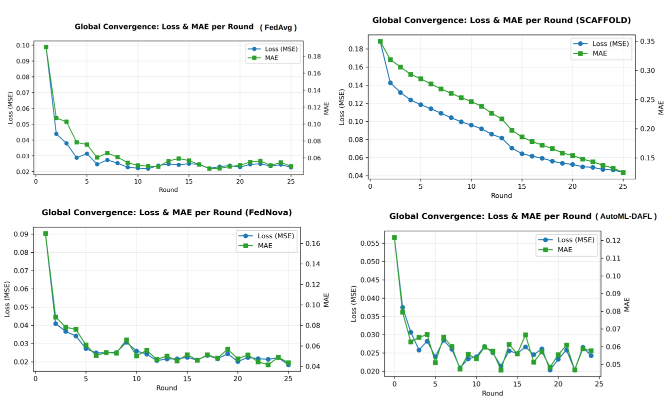

Fig. 2. Global convergence comparison of federated learning algorithms showing loss (MSE) and MAE trends across communication rounds.

B. Efficiency and Fairness analysis

Beyond predictive accuracy, the AutoML-DAFL approach exhibited superior training efficiency. This efficiency gain is critical for resource-constrained pervasive Internet of Things (IoT) systems, where faster convergence reduces computational load and enables more responsive on-device adaptation.

TABLE II
Performance Comparison of Federated Learning Algorithms

Metric	FedAvg	FedNova	SCAFFOLD	AutoML-DAFL
MAE (↓)	0.0331	0.0849	0.5792	**0.0304**
RMSE (↓)	0.1062	0.2526	0.7055	**0.0721**
R^2 (↑)	0.9885	0.9363	0.5030	**0.9947**
Jain (↑)	0.9993	0.9665	0.9991	**0.9997**

Furthermore, Jain's fairness index is consistently high across the federated network, with AutoML-DAFL achieving the highest score of 0.9997. This represents a marginal improvement over FedAvg (0.9993), SCAFFOLD (0.9991), and a distinct advantage over FedNova (0.9665). This stability is crucial in edge-oriented and cyber-physical systems, as it confirms that the accuracy and robustness gains are consistently shared across all clients, ensuring sound, unbiased automated monitoring without overfitting or penalizing any subset of the distributed system.

C. Ablation Study Analysis

The ablation study validates the necessity of each component of the multi-objective DAFL reward, confirming distinct stabilizing effects in heterogeneous, pervasive federated environments (see Table III and Figure 3). The communication-aware term emerges as the primary mechanism for global robustness: its removal (DAFL-no-comm) triggers a severe increase in RMSE of 14.96%, demonstrating that unreliable clients disproportionately influence aggregation and make the federated system highly vulnerable to error amplification. In contrast, the MAE-driven term functions as a local gradient smoothing regularizer that prevents drift following behavior. While its removal (DAFL-No-MAE) achieves the lowest MAE (-8.95%) across clients, this local optimization leads to a loss in global consistency (0.64% RMSE) and noisier gradient updates.

The ablation patterns confirm a fundamental cross-metric trade-off across variants. The single variants achieve rapid convergence or lowest local errors, but these gains are achieved at a cost of global stability. Critically, Jain fairness remains consistently high across all variants, indicating that performance degradation comes from increased model variance and noise rather than inequitable distribution of prediction quality. The trade-off between MAE and RMSE demonstrates that DAFL-No-Comm fails under non-IID data distributions and bandwidth constraints, while DAFL-No-MAE becomes fragile under high noise and temporal drift. Both variants fail to balance accuracy, robustness, and resource awareness, demonstrating that single-objective optimization is inadequate for pervasive edge systems.

DAFL-full emerges as the most optimal configuration that maintains accuracy and stability under simulated heterogeneous conditions. By integrating MAE-guided local smoothing with communication-aware aggregation, DAFL-full achieves the lowest RMSE (0.0721), highest R^2, and consistent conver-

TABLE III
ABLATION STUDY ANALYSIS OF AUTOML-DAFL FRAMEWORK FOR ACCURACY AND FAIRNESS METRICS.

Variant	**MAE**	**RMSE**	**Jain**	**ΔMAE / ΔRMSE**
DAFL-Full	0.0304	0.0721	0.9997	0.00 / 0.00
DAFL-No-MAE	0.0276	0.0725	0.9995	-8.95 / +0.64
DAFL-No-Comm	0.0313	0.0828	0.9997	+3.17 / +14.96

gence across varied data distributions, noise, and communication constraints. The ablation results confirm that integrated multi-objective weighting is essential for federated AutoML systems for handling the compounded effects of data drift, noise, and communication variability in resource-constrained, pervasive environments.

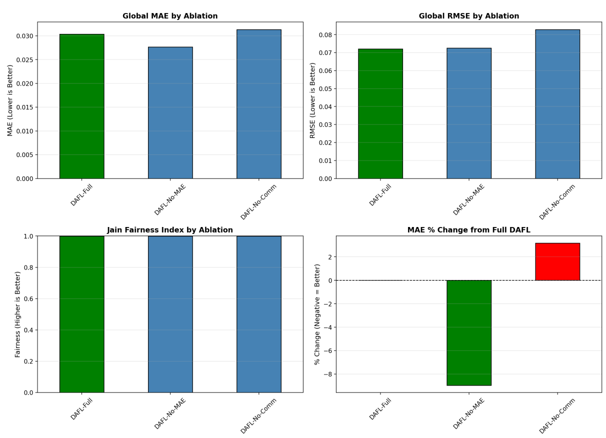

Fig. 3. Ablation study analysis of AutoML-DAFL framework for accuracy and fairness metrics.

V. CONCLUSION AND FUTURE WORK

This study presents an AutoML-guided, drift-aware federated learning framework designed to achieve robust and resource-aligned optimization tailored for cyber-physical and IoT-enabled aquaculture water quality monitoring. The rigorous ablation analysis confirms that the full DAFL configuration is essential, validating the multi-objective reward design by achieving a balanced trade-off between accuracy, stability, and resource efficiency under unstable sensor environments. By outperforming state-of-the-art baselines, our framework provides a superior,drift-resilient mechanism for real-time intelligent automation. It confirms that a communication-aware module is essential for global robustness, mitigating error amplification, while an MAE-aware module mitigates noise-induced variance. To overcome the cross-metric tradeoff of removing any component, we demonstrate that these components together in DAFL-full deliver the lowest RMSE, highest R^2, and fairness, representing a significant advancement in AI-driven decision support systems for sustainable and green automation.

Future work will validate AutoML-DAFL through field-scale edge and aquaculture farm deployments to assess long-term stability, energy efficiency, resource optimization, and controller adaptivity under real-world operational and hardware constraints. To enhance autonomy in industrial processes, we will extend the reward function to mitigate catastrophic forgetting and investigate privacy-enhancing mechanisms, both of which are critical for secure Industry 4.0 production deployment. We will advance the applicability of AutoML-DAFL by exploring scalability to large federations and developing theoretical convergence guarantees for our multi-objective reward formulation.

REFERENCES

[1] S. Kashem *et al.*, “Extensive and intensive aqua-fisheries dataset for water quality prediction,” *IEEE DataPort*, 2024.

[2] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, “Communication-efficient learning of deep networks from decentralized data,” in *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*. PMLR, 2017, pp. 1273–1282.

[3] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, “Federated optimization in heterogeneous networks,” in *Proceedings of the 3rd Conference on Systems and Machine Learning (SysML/MLSys)*, 2020, pp. 1–16.

[4] E. Jothimurugesan, K. Hsieh, J. Wang, G. Joshi, and P. Gibbons, “Federated learning under distributed concept drift,” in *Proceedings of the NeurIPS 2023 Workshop on Federated Learning*, 2023.

[5] M. Kang, “Fednn: Federated learning on concept drift data using weight normalization and adaptive group normalization,” *Computer Networks*, vol. 234, p. 109961, 2024.

[6] T. Li, M. Sanjabi, A. Beirami, and V. Smith, “Fair resource allocation in federated learning,” *arXiv preprint arXiv:1905.10497*, 2019.

[7] A. Taik *et al.*, “Fairness in federated learning: Fairness for whom?” in *Proceedings of the 2025 AAAI/ACM Conference on AI, Ethics, and Society (AIES)*, 2025.

[8] S. P. Karimireddy, S. Kale, M. Mohri, S. J. Reddi, S. U. Stich, and A. T. Suresh, “Scaffold: Stochastic controlled averaging for federated learning,” in *Proceedings of the 37th International Conference on Machine Learning (ICML)*. PMLR, 2020, pp. 5132–5143.

[9] S. Wang and G. Joshi, “Federated learning with normalized averaging,” *arXiv preprint arXiv:2007.07481*, 2020.

[10] C. Zhang, Y. Liu, and X. Ma, “Deep learning-based water quality prediction using lstm networks,” *Water Research*, 2022.

[11] J. Wang, S. Li, and H. Chen, “Predicting dissolved oxygen in aquaculture using hybrid deep learning models,” *Aquacultural Engineering*, 2021.

[12] M. Ahmed, M. Rahman, and F. Anwar, “Iot-enabled smart aquaculture: Water quality monitoring and predictive analytics using machine learning,” *Computers and Electronics in Agriculture*, 2021.

[13] J. Gama, I. Žliobaitė, A. Bifet, M. Pechenizkiy, and A. Bouchachia, “A survey on concept drift adaptation,” *ACM Computing Surveys*, 2014.

[14] H. Sato and M. Yamamoto, “Drift-aware learning for environmental sensor networks,” *Sensors*, 2020.

[15] M. Arjovsky, S. Chintala, and L. Bottou, “Wasserstein generative adversarial networks,” in *Proceedings of the 34th International Conference on Machine Learning (ICML)*, 2017.

[16] N. Mukhtiar *et al.*, “Federated learning at the forefront of fairness: A survey,” in *Proceedings of IJCAI 2025*, 2025.

[17] C. He, M. Annavaram, and S. Avestimehr, “Fednas: Federated neural architecture search,” in *CVPR*, 2020.

[18] T. Lattimore and C. Szepesvári, *Bandit algorithms*. Cambridge University Press, 2020.

[19] T. Nishio and R. Yonetani, “Fedcs: Federated learning with client selection,” in *ICC*, 2019.

[20] S. Wang, T. Tuor, and T. Salonidis, “Optimizing federated learning on non-iid data with reinforcement learning,” in *INFOCOM*, 2020.

[21] Y. Mansour, M. Mohri, and A. Rostamizadeh, “Three approaches for personalization with applications to federated learning,” in *NeurIPS*, 2020.

[22] P. Wang *et al.*, “Field-guide to federated optimization,” *arXiv:2107.06917*, 2021.

[23] M. Arjovsky, S. Chintala, and L. Bottou, “Wasserstein gan,” in *ICML*, 2017.

[24] M. A. Kashem, N. Ullah, and M. H. Shuvo, “Extensive and intensive aqua-fisheries dataset,” 2024. [Online]. Available: https://dx.doi.org/10.21227/5x62-1189