Lightweight Secure Federated Learning for Energy-Constrained IoT: A Case Study in Smart Irrigation*

Zohra Dakhia[1,2], Adrián Cánovas-Rodríguez[3], Mariateresa Russo[1], Aurora González-Vidal[3], Massimo Merenda[1], and Antonio F. Skarmeta-Gómez[3†]

[1] University Mediterranea of Reggio Calabria, Reggio Calabria, Italy.
{zohra.dakhia, massimo.merenda, mariateresa}@unirc.it
[2] University Federico II of Naples, Naples, Italy.
[3] University of Murcia (UMU), Murcia, Spain.
{adriancr, aurora.gonzalez2, skarmeta}@um.es

Abstract

Federated learning (FL) is increasingly adopted in Internet of Things (IoT) environments, yet real deployments must address both energy constraints and security threats. This paper presents a lightweight and energy-aware FL framework designed for heterogeneous IoT devices, with smart irrigation used as a representative case study. The framework adapts client participation according to available energy budgets while securing model updates with AES-128 encryption and SHA-256 integrity verification. A real-world dataset from six agricultural patches and a Docker-based testbed with heterogeneous clients are employed to evaluate the approach. Results over 30 training rounds show that the framework sustains predictive accuracy, mitigates risks of tampering and eavesdropping, and introduces less than 1% computational overhead. The study demonstrates that lightweight cryptographic protection combined with energy-aware scheduling provides a practical balance between efficiency and trustworthiness in federated IoT systems.

1 Introduction

Federated learning (FL) has emerged as a key paradigm for distributed intelligence across the IoT–edge–cloud continuum, where multiple devices collaboratively train a model without centralizing raw data [1]. This approach is particularly attractive in environments with heterogeneous IoT devices, since it enables computation at the edge while preserving local data privacy. However, deploying FL in the continuum also introduces 2 major challenges: (i) devices operate under strict energy and computation budgets, and (ii) communications between clients and the server are exposed to security and privacy threats, such as tampering, eavesdropping, or poisoning [2, 3, 4]. To address these challenges, we propose a lightweight and energy-aware FL framework for constrained IoT devices. The solution adapts client participation according to available energy budgets while integrating cryptographic protection to secure communications. Specifically, we employ AES-128 encryption and SHA-256 integrity verification to safeguard model updates with negligible overhead, thereby balancing efficiency with trustworthiness in distributed training. As a representative case study, we apply this framework to smart irrigation in semi-arid regions such as Murcia, Spain, where sustainable water management is critical and where the Segura River Basin Authority allocates irrigation quotas and regulates water distribution. It is essential to preserve the privacy of plot-level water-use records. Smart irrigation systems rely on sensor networks to monitor soil, climate, and irrigation events [5, 6], but

*Proceedings of the 9th International Conference on Mobile Internet Security (MobiSec'25), Article No. S6, December 16-18, 2025, Sapporo , Japan.
†Corresponding author

their deployment is hindered by the use of heterogeneous energy-limited devices. Our approach demonstrates how energy-aware FL with lightweight security can support secure and efficient collaboration in such resource-constrained environments. The central challenge we investigate is whether FL can be effectively deployed on heterogeneous, energy-constrained IoT devices while maintaining lightweight but sufficient security. Specifically, we explore how energy-aware client participation combined with low-overhead cryptographic protection can achieve a balance between efficiency, accuracy, and trustworthiness in real-world scenarios such as smart irrigation. Our main contributions are:

- A threat-aware lightweight security framework for FL across the IoT–edge continuum.
- Integration of AES-128 + SHA-256 to protect against tampering and unauthorized modification with $< 1\%$ overhead.
- Energy–security trade-off evaluation on a real agricultural IoT dataset with heterogeneous clients (case study: smart irrigation).
- Docker-based testbed replicating adversarially vulnerable communication between clients and server.

The rest of this paper is organized as follows: Section 2 reviews related work, Section 3 introduces the dataset, Section 4 presents the framework, Section 5 reports implementation and results, and Section 6 concludes the paper.

2 Related Work

FL enables collaborative training in IoT and edge environments without centralizing sensitive data, but deployments face two key challenges: limited resources and security. Prior works have mainly optimized efficiency. Nishio and Yonetani [7] proposed client selection for heterogeneous devices, Wang et al. [8] studied adaptive scheduling, and Shi et al. [9] developed communication-efficient updates. These improve scalability but do not explicitly address security. On the other hand, security-focused methods such as secure aggregation [10] and differential privacy [11] provide strong guarantees but incur high computational and communication costs, making them less practical for constrained IoT–edge systems. Recent works have continued exploring these trade-offs. Zhang et al. [12] introduced a lightweight hierarchical FL approach for edge–cloud coordination with reduced encryption cost. Li et al. [13] proposed an efficient secure aggregation protocol tailored for IoT devices, achieving improved energy–accuracy balance. Moreover, Rahman et al. [14] presented an energy-aware FL framework for agricultural IoT, emphasizing sustainability and device longevity. Our work fills this gap by introducing a lightweight cryptographic framework (AES-128 with SHA-256) that secures model updates with under 1% overhead, combined with energy-aware client participation to achieve a balance of efficiency, accuracy, and trustworthiness in IoT–edge continuum scenarios.

3 Dataset Description

The dataset was collected from six agricultural plots in the Region of Murcia, Spain, where citrus and peach crops are cultivated under semi-arid conditions (Figure 1). Since 2020, soil probes and meteorological stations have continuously monitored underground and atmospheric variables. Soil probes (Sentek Drill & Drop) measure moisture, temperature, and salinity at

multiple depths, while meteorological stations record air temperature, humidity, solar radiation, and precipitation. Data are sampled every 10–15 minutes and aggregated monthly, with each month corresponding to one FL round. The dataset includes:

- Inputs: water meter readings, air humidity, rainfall, and air temperature.
- Outputs: soil moisture at depths of 10, 20, 30, 40, and 50 cm.

Despite challenges such as sensor heterogeneity and missing values, the dataset provides a realistic basis for evaluating energy-aware FL in smart irrigation. Table 1 presents an excerpt of the

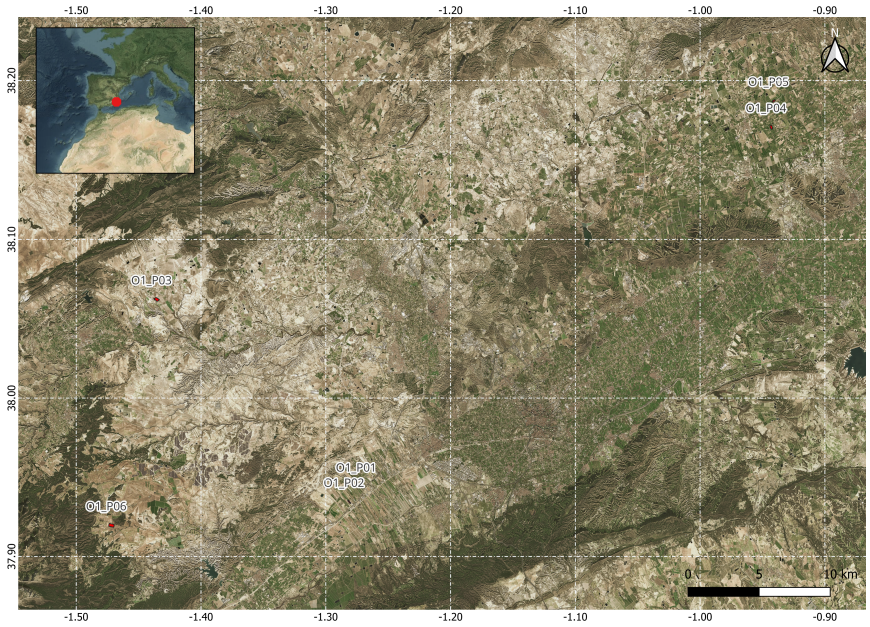

Figure 1: Locations of the six experimental plots in southeastern Spain.

dataset collected from one agricultural patch. Each record corresponds to a 15–30 minute interval and includes irrigation (water meter readings), meteorological, and soil parameters. The table captures transitions between dry, irrigated, and rainy conditions, illustrating the natural variability of the environmental data. Similar datasets were gathered from six distinct agricultural plots in Murcia (Spain), each with different soil and crop characteristics, and collectively used to train and evaluate the proposed federated learning framework.

Table 1: Example data excerpt from one agricultural patch showing time progression across dry, irrigated, and rainy conditions.

Timestamp	Water (L)	Hum. (%)	Temp (°C)	Rain (mm)	SM10 (%)	SM20 (%)	SM30 (%)	SM40 (%)	SM50 (%)
2022-11-09 09:00	0.0	41.0	19.3	0.0	34.85	37.92	38.40	38.10	36.70
2022-11-09 09:30	0.0	42.6	19.8	0.0	34.90	37.95	38.42	38.11	36.72
2022-11-09 10:00	10.8	45.3	20.5	0.0	35.40	38.20	38.60	38.20	36.80
2022-11-09 17:15	12.4	49.9	21.3	0.2	35.10	38.05	38.60	38.20	36.85
2022-11-10 11:00	0.0	50.8	21.4	0.0	36.21	39.28	39.72	39.01	36.96
2022-11-10 12:00	8.7	51.4	22.8	0.0	36.10	39.05	39.64	39.01	36.92
2022-11-11 00:00	0.0	91.3	13.6	1.4	35.39	38.06	39.17	38.98	36.96

4 System Model and Proposed Framework

This section presents the overall system model and the proposed energy-aware FL framework for smart irrigation. The framework integrates five main components: clients, server, communication, dataset, and learning model. In addition, we define a threat model that captures potential adversarial actions against the communication and update process, and explain how lightweight cryptographic mechanisms are integrated to mitigate these risks.

4.1 System Overview

As shown in Figure 2, the proposed FL setup reflects the IoT–edge continuum. Multiple heterogeneous IoT clients collaborate with a containerized **edge server** that coordinates training. Clients represent constrained devices in agricultural fields, while the server runs as a Docker container at the edge, reducing latency and preserving data locality.

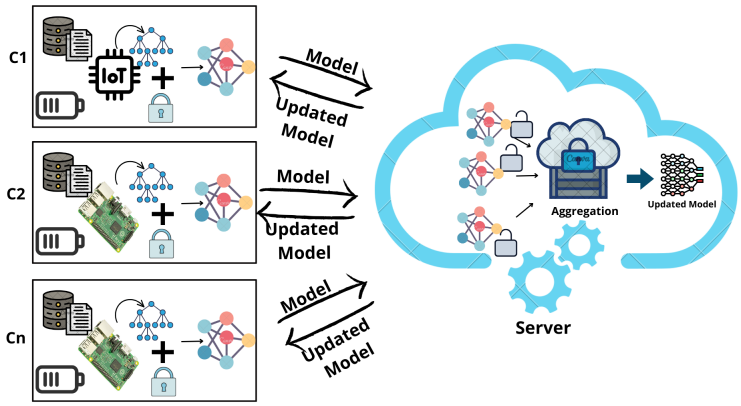

Figure 2: Energy-aware FL setup in the IoT–edge continuum with a Docker-based edge server and encrypted client updates.

The setup includes five elements:

Clients: Docker containers simulate heterogeneous devices with varying energy and computation levels. Each trains locally and contributes according to its energy budget.

Edge Server: Coordinates FL by initializing the global model, selecting clients, aggregating updates, and broadcasting results. Running at the edge ensures lower latency and privacy.

Communication: Training proceeds in rounds. The edge server sends the model, clients return AES-128 encrypted updates with SHA-256 hashes, and the server verifies, decrypts, and aggregates them. This protects against tampering and eavesdropping with negligible overhead.

Dataset: Collected from six agricultural patches in Murcia (since 2020), partitioned by month to align with FL rounds and seasonal variations.

Learning Model: A lightweight neural network with two hidden layers predicts soil moisture from environmental and irrigation features, balancing accuracy and efficiency for constrained devices.

4.2 Proposed Framework

The proposed framework implements energy-aware and secure FL in the IoT–edge continuum, allowing heterogeneous clients to contribute effectively without compromising overall training.

- **Privacy and Security in the Continuum:** FL ensures that raw data remains local to IoT devices, while updates transmitted to the edge server are encrypted with AES-128 and protected by SHA-256 hashes. This lightweight mechanism safeguards against tampering and eavesdropping during communication across potentially insecure networks.
- **Energy-Aware Client Participation:** Clients with higher energy budgets participate more frequently and perform longer local training, while low-energy clients reduce their workload or skip rounds.
- **Dynamic Scheduling at the Edge:** The edge server adaptively selects clients based on their reported energy, ensuring continuous training across the continuum while preventing weaker devices from being exhausted.

4.3 Threat Model and Assumptions

We consider a FL system over heterogeneous IoT devices in smart irrigation, where communication between clients and the edge server is exposed to several threats. Although privacy risks are relatively moderate in such scenarios, certain farms may belong to competing agricultural stakeholders whose local data (e.g., irrigation schedules or soil conditions) remain sensitive. Therefore, our lightweight cryptographic layer ensures both data integrity and confidentiality, even under semi-collaborative or multi-owner deployments. The main risks are: (i) tampering with updates during transmission, (ii) model poisoning by compromised clients, and (iii) eavesdropping on exchanged data. To counter these, clients encrypt updates with *AES-128* (CBC) and attach a *SHA-256* hash, which the server verifies before aggregation. This ensures confidentiality and integrity with negligible computational cost. AES-128 and SHA-256 were selected because they provide a strong balance between security and computational efficiency. AES-128 is hardware-accelerated on most embedded processors, ensuring minimal delay, while SHA-256 remains the de facto standard for data integrity verification in lightweight IoT communications

5 Implementation and Evaluation

This section outlines the implementation of the proposed framework in `Docker` and its evaluation over **30 communication rounds**, corresponding to five years of monthly data.

5.1 Implementation Setup

The framework was implemented in Python using Docker containers. Six client containers simulated heterogeneous agricultural patches, while an edge server container coordinated training. The dataset (2020–2025) was partitioned by patch, with one month per round. For security, each client applied AES-128 encryption with a SHA-256 hash, which the server verified before aggregation. This added less than 1% overhead in total computation and energy cost per round, measured as the difference in CPU time and power consumption with and without encryption.

5.2 Simulated Device Specifications

Each Docker container was configured to emulate a different IoT device profile, reflecting the heterogeneity of hardware typically deployed in smart irrigation systems. The assigned CPU and memory resources replicate microcontrollers (e.g., Arm Cortex-M7) and single-board computers (e.g., Raspberry Pi 3B). This heterogeneity spans CPU capacity, memory, and energy availability, representing the diversity of IoT nodes commonly found in agricultural deployments. Clients 1–4 represent low-power MCUs, while clients 5–6 emulate higher-capacity edge devices. Such heterogeneity directly affects participation frequency and local computation effort. Table 2 summarizes the specification of each simulated client and the central server. In this setup, each Docker container acted as a client device contributing to the FL process with its assigned computational and energy constraints.

Table 2: Specifications of simulated client and server devices in the Docker testbed.

Container	**Simulated device**	**Assigned memory**	**Assigned CPU**
server	Central aggregation server	7 GB	4 vCPU
client01	Arm Cortex-M7 MCU (512 KB Flash, SDRAM)	900 MB	0.25 vCPU
client02	Arm Cortex-M7 MCU (512 KB Flash, SDRAM)	900 MB	0.25 vCPU
client03	Arm Cortex-M7 MCU (1 MB Flash, SDRAM + TFT)	900 MB	0.5 vCPU
client04	Arm Cortex-M7 MCU (1 MB Flash, SDRAM + TFT)	900 MB	0.5 vCPU
client05	Raspberry Pi 3B (1 GB RAM, Quad-core ARM Cortex-A53 1.2 GHz)	1 GB	1 vCPU
client06	Raspberry Pi 3B (1 GB RAM, Quad-core ARM Cortex-A53 1.2 GHz)	1 GB	1 vCPU

5.3 Energy-Aware FL

Each client was initialized with an **energy budget**. At round t, energy was updated as a function of local computation and communication cost. Energy profiles were selected based on the specifications of representative IoT devices deployed in agricultural monitoring ranging from low-power MCUs (0.6 MJ initial energy) to mid-range edge nodes (1.2 MJ).

$$E_i^t = E_i^{t-1} - \Delta E_i \quad \text{if client } i \text{ participates}, \qquad E_i^t = E_i^{t-1} \quad \text{otherwise.}$$

The term ΔE_i represents the per-round energy cost associated with client i, reflecting its hardware characteristics and computational capacity. Clients with larger budgets and higher ΔE_i values sustained more frequent participation, whereas those with smaller budgets and lower ΔE_i values contributed less often in order to preserve resources.

5.4 Evaluation Metrics and Results

The framework was evaluated in terms of predictive accuracy (RMSE, R^2), client energy usage, and security overhead. Results cover client-side performance, global server performance, comparison with centralized training, and the impact of encryption and hashing.

1. Client-Side Results: Clients with larger initial energy budgets (1.2 MJ) participated in all 30 rounds, while lower-energy clients (0.6 MJ) joined fewer rounds. All clients achieved acceptable accuracy (Table 3).

Table 3: Client energy and performance (capped at 30 rounds).

Client	Initial Energy (MJ)	ΔE_i (per round)	Rounds	RMSE	R^2
1	1.2	0.020	30	0.072	0.91
2	0.9	0.017	27	0.081	0.89
3	1.2	0.019	30	0.072	0.91
4	0.6	0.013	23	0.086	0.87
5	0.9	0.0016	30	0.072	0.91
6	0.7	0.002	28	0.079	0.89

2. Server-Side Results: The global model converged by round 30 and stabilized thereafter (Table 4).

Table 4: Global model performance at the server.

Round	RMSE	R^2
10	0.128	0.76
20	0.087	0.87
30	0.072	0.91

We chose 30 training rounds because the global model reached convergence around round 25, and extending training beyond 30 rounds yielded negligible accuracy improvement (less than 0.3%). Conversely, with fewer than 20 rounds, the R^2 score dropped below 0.85, indicating underfitting. Therefore, 30 rounds provided a balanced trade-off between predictive accuracy and computational cost for the considered dataset and device configurations.

3. Comparison with Centralized Training: Centralized training achieved slightly better accuracy, but FL provided close performance while accounting for heterogeneous energy constraints (Table 5).

Table 5: Centralized vs. FL.

Method	RMSE	R^2	Energy-Aware
Centralized	0.068	0.93	No
Federated (ours)	0.072	0.91	Yes

4. **Impact of Security Mechanism:** We evaluated the overhead of encryption and hash-based verification during client-to-server communication. As shown in Table 6, the added cost was under 1% per round, with no effect on RMSE or R^2. This confirms that lightweight AES-128 and SHA-256 secure updates without compromising efficiency or accuracy.

Table 6: Impact of security mechanism on performance.

Setting	RMSE	R^2	Overhead
Without Security	0.072	0.91	–
With Security	0.072	0.91	$< 1\%$

5.5 Comparison with Alternative Security Mechanisms

Compared to heavier security mechanisms, our approach based on AES-128 and SHA-256 offers a practical balance: it preserves model accuracy, introduces negligible energy overhead, and still protects the integrity and confidentiality of updates. The following comparison Table 7 highlights why this lightweight design is better suited for constrained IoT–edge environments than more complex alternatives.

Table 7: Conceptual comparison of security mechanisms in FL.

Method	**Energy/Computation Cost**	**Impact on Accuracy**	**Privacy/Security Guarantees**
AES-128 + SHA-256 (ours)	Very low ($< 1\%$ overhead)	None (no accuracy loss)	Protects integrity and confidentiality of updates during transmission; suitable for constrained IoT devices.
Secure Aggregation [10]	Moderate to high (extra rounds of communication)	Minimal accuracy impact	Provides confidentiality of individual updates against server inspection; higher communication cost.
Differential Privacy [11]	Low to moderate (noise addition, extra computations)	Potential reduction in accuracy depending on privacy budget	Provides formal privacy guarantees against data reconstruction, but may degrade model performance.

The framework is designed to scale to larger IoT deployments. With 100 simulated nodes, communication time grows linearly, while encryption overhead remains constant (¡1%). Thus, the main limitation is bandwidth rather than computation.

6 Conclusion and Future Work

This work presented an energy-aware FL framework designed for heterogeneous IoT devices within the IoT–edge continuum, with smart irrigation used as a representative case study. The framework integrates a threat model that addresses tampering, poisoning, and eavesdropping attacks, and mitigates these risks through AES-128 encryption and SHA-256 integrity verification. Using Dockerized clients with heterogeneous energy budgets and a real dataset from six agricultural patches in Murcia, Spain, we simulated 30 training rounds under realistic conditions. The results demonstrate that energy-aware scheduling reduces consumption, sustains client participation, and achieves reliable convergence, while the lightweight cryptographic layer secures communications with less than 1% overhead. Overall, the framework offers a practical balance between predictive accuracy, energy efficiency, and security, making it suitable for constrained IoT–edge environments. Future work will extend the threat model, investigate advanced privacy-preserving techniques, and validate the framework through real-world deployment on energy-limited IoT devices.

Acknowledgments

This work has been funded by the project **OSIRIS (TSI-100921-2023-1)** under the *Cátedras ENIA*, promoted by the Ministry for Digital Transformation and Public Function and co-funded by the European Union – NextGenerationEU, through the PRTR. It was also supported by grant **RYC2023-043553-I**, funded by MICIU/AEI/10.13039/501100011033 and ESF+. This study was also carried out within the project **"Ecosistema TECH4YOU – Technologies for climate change adaptation and quality of life improvement"**, and received funding from the European Union NextGenerationEU (*National Recovery and Resilience Plan (PNRR) - M4C2 - Investment 1.5 - "Innovation Ecosystems" - D.D. 3277 of 30 December 2021*). This manuscript reflects only the authors' views and opinions; neither the European Union nor the European Commission can be considered responsible for them.

References

[1] Z. Dakhia, A. Lazzaro, M. R. Sebti, D. Iero, M. Russo, and M. Merenda. Preliminary analysis of federated learning in agrifood: From mnist to real-world iot applications in a heterogeneous environment. In *Proceedings of the 2025 10th International Conference on Smart and Sustainable Technologies (SpliTech)*, pages 1–7. IEEE, 2025.

[2] A. Ahmadi, M. Keshavarz, and F. Ejlali. Resilience to climate change in agricultural water-scarce areas: The major obstacles and adaptive strategies. *Water Resources Management*, 39(3):1195–1214, 2025.

[3] E. M. Campos, P. F. Saura, A. González-Vidal, J. L. Hernández-Ramos, J. B. Bernabe, G. Baldini, and A. Skarmeta. Evaluating federated learning for intrusion detection in internet of things: Review and challenges. *Computer Networks*, 203:108661, 2022.

[4] Z. Dakhia and M. Merenda. Client selection in federated learning on resource-constrained devices: A game theory approach. *Applied Sciences*, 15(13):7556, 2025.

[5] A. González-Vidal, J. Fernández-García, and A. F. Skarmeta. A combination of multi and univariate anomaly detection in urban irrigation systems. In *Proceedings of the 3rd International Conference on Embedded & Distributed Systems (EDiS)*, pages 31–36. IEEE, 2022.

[6] Z. Dakhia, M. Russo, and M. Merenda. Ai-enabled iot for food computing: Challenges, opportunities, and future directions. *Sensors*, 25(7):2147, 2025.

[7] T. Nishio and R. Yonetani. Client selection for federated learning with heterogeneous resources in mobile edge. In *IEEE International Conference on Communications (ICC)*, pages 1–7, 2019.

[8] H. Wang et al. Adaptive federated learning in resource-constrained edge computing systems. *IEEE Journal on Selected Areas in Communications*, 39(1):254–267, 2021.

[9] S. Shi et al. Communication-efficient federated learning with asynchronous model updates. In *IEEE International Conference on Distributed Computing Systems (ICDCS)*, pages 246–256, 2020.

[10] K. Bonawitz, V. Ivanov, B. Kreuter, A. Marcedone, H. B. McMahan, S. Patel, D. Ramage, A. Segal, and K. Seth. Practical secure aggregation for privacy-preserving machine learning. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 1175–1191. ACM, 2017.

[11] R. C. Geyer, T. Klein, and M. Nabi. Differentially private federated learning: A client level perspective. In *Proceedings of the 2017 NIPS Workshop on Privacy and Security in Machine Learning*, 2017.

[12] Y. Zhang, T. Nguyen, and K. Lee. Lightweight hierarchical federated learning for edge–cloud collaboration in iot networks. *IEEE Internet of Things Journal*, 12(4):5231–5243, 2025.

[13] H. Li, M. Chen, and J. Luo. Secure and efficient aggregation protocol for federated learning in resource-constrained iot. *IEEE Transactions on Network and Service Management*, 19(6):834–847, 2025.

[14] M. Rahman, A. Al-Hamadi, and L. Zhang. Energy-aware federated learning for sustainable smart agriculture. *Computers and Electronics in Agriculture*, 222:109–118, 2024.