OPEN ACCESS

Kashf Journal of Multidisciplinary Research
Vol: 03 - Issue 03 (2026)
P-ISSN: 3007-1992 E-ISSN: 3007-200X
https://kjmr.com.pk

AN ADAPTIVE FEDERATED DEEP LEARNING FRAMEWORK FOR SECURE AND PRIVACY PRESERVING PRECISION AGRICULTURE IN DISTRIBUTED IOT NETWORKS

Umair Aslam[1], Muhammad Aqib[2], Iqra Arshad[3], Basit Bashir[4], Muhammad Daniyal[5], Muhammad Yousif[6]

[1]*Department of Management Sciences, NUML, Lahore Campus, Pakistan*
[2]*Department of Computer Science, Lahore Garrison University, Lahore, Pakistan*
[3]*Department of Computing, NFC Institute of Engineering and Technology, Multan, Pakistan*
[4]*Department of Computer Science, University of Engineering and Technology, Lahore, Pakistan*
[5]*Department of Computer Science, Superior College Pak Arab Campus-Lahore, Pakistan*
[6]*Department of Computer Science, NUML, Lahore Campus, Pakistan*

Corresponding Author: (*umairchoudhary@outlook.com*)
DOI: (*https://doi.org/10.71146/kjmr964*)

Article Info

Abstract

The Internet of Things (IoT) has revolutionized agriculture data gathering by connecting equipment with increased sophistication and scale, contributing to the smart monitoring and data-driven management in the field. The Internet of Things (IoT) has transformed the way data is collected in agriculture by making the process more intelligent and comprehensive, which is vital for smart monitoring and data-driven management in agriculture. But the spread of IoT applications in farms poses significant challenges in data privacy and security, scalability and efficient collaborative learning. In this paper, an Adaptive Federated Deep Learning (AFDL) framework for secure and privacy-preserving precision agriculture in a distributed IoT network is proposed. The proposed framework utilizes federated deep learning, a technique that allows for distributed model training across various agricultural IoT nodes, thereby preserving the privacy of the sensitive data on the farms. The adaptive aggregation strategy adjusts the contribution of participating nodes dynamically according to data quality, the performance of models, and network conditions, which further improves the efficiency of learning and convergence of the model. In addition, secure communication and privacy preserving schemes are built in to safeguard against unauthorized access to and inference attacks on model updates. The effectiveness of the framework has been assessed with the Smart Farm Sensor Dataset, with deep learning models trained independently at the various nodes of the farming site and then collectively aggregated to form a powerful global model. Experimental results show that the proposed framework can achieve a classification accuracy of 97.4% and enhance the efficiency, scalability, and robustness of communication compared to traditional federated learning methods. Moreover, the adaptive aggregation mechanism decreases communication overhead by 65% and privacy leakage by 83%, which makes it applicable to resource limited IoT environments. In conclusion, the proposed Adaptive Federated Deep Learning framework offers a secure, scalable, and privacy-preserving approach for enabling intelligent farming practices that are decentralized and sustainable in the future of precision agriculture, leveraging collaborative learning within decentralized and distributed IoT networks.

Keywords: *Federated Learning; Precision Agriculture; Privacy-Preserving Machine Learning; Distributed IoT Networks; Adaptive Aggregation; Deep Learning; Differential Privacy; Secure Aggregation; Edge Computing; Smart Farming; Non-IID Data; Crop Health Monitoring.*

1. Introduction

The agriculture sector is experiencing “digital revolution” through the integration of the Internet of Things (IoT) [1], unmanned aerial vehicles (UAVs) [2], remote sensing and artificial intelligence (AI). Sensors mounted on soils, weather stations, and multispectral/thermal cameras and autonomous machinery now provide continuous streams of large volumes of multimodal data on crop health, soil fertility, irrigation requirements, and pest pressure. If managed correctly, these data flows can help detect diseases early in their growth, schedule water use for irrigation and enhance yield prediction, which is necessary to cope with the food security needs of a growing world population in the face of climate variability.

The prevailing method of leveraging this data is centralized ML [3], where data from the sensors and images collected on the farm is sent to a cloud server for storage and training and inferences are drawn there. This is a very compute-intensive paradigm, but it has a number of weaknesses in an agricultural context. First, many rural areas experience intermittent and/or expensive Internet connectivity, and the persistent exchange of raw imagery and log data of the sensors is impractical and inefficient of limited bandwidth. Second, agricultural data are not independent and identically distributed (non-IID) [4]: data from different farms, and even from different parts of the same farm, come from different contexts with different soil type, crop varieties, pest pressure and climate, so models trained on data from one context do not necessarily generalize well to another. Thirdly, and potentially most significantly, farmers and agribusinesses are often reluctant to share raw operational data (e.g., yield; soil nutrient profiles; irrigation data) with third parties due to concerns about disclosure of competitive advantage and/or inappropriate use of such information by third parties. These are all concerns that impact the adoption of centralized digital agriculture platforms, especially for smallholder farmers, and erode their trust [5].

Federated learning (FL) has appeared as a promising solution to overcome these challenges by allowing multiple distributed client's farms, cooperatives, edge gateways or UAV equipped nodes to contribute to the training of a federated deep learning model while keeping the data locally. The parameters or gradients of the models are exchanged with a central aggregator, not the actual data, which significantly decreases the communication overhead and privacy exposure. Such a decentralized approach has been investigated for various applications including crop disease detection, soil fertility classification, yield prediction and irrigation optimization, and typically with data accuracy close to that of the centralized models with data sovereignty. Most of the current federated agricultural systems, however, use a static approach: a fixed sampling strategy of the clients and a fixed training schedule on the local node(s), as well as a fixed privacy mechanism (or none at all), when there is one, depending on the heterogeneity of the devices involved, the sensitivity of a certain type of data, or the volatility of the underlying wireless network. This rigidity reduces scalability for a large, geographically distributed population of devices. It also makes systems vulnerable when the capabilities, connectivity or threat of the devices differ in the federation [7].

To address these challenges, this paper introduces an Adaptive Federated Deep Learning (AFDL) system for precise agriculture in distributed IoT networks with security and privacy concerns. It adapts three interdependent dimensions of federated learning (FL) to the available resources for each device: (i) client participation and local computation, depending on the energy resources of each device, the quality of its available connectivity, and the freshness of the data it has; (ii) aggregation strategy, which reweights the FL client contributions according to data variability and trustworthiness, rather than assuming identically distributed data; and (iii) privacy protection intensity, which scales up the amount of differential privacy noise, secure aggregation, and lightweight encoding, depending on the sensitivity of the FL data type, and the level of assessed risk of the communication channel. It's designed to balance adaptivity and multi-layered privacy protection, maintain high predictive accuracy for fundamental precision agriculture applications (crop health assessment, disease and pest detection, soil/irrigation monitoring, yield forecasting), and be deployable on resource-limited edge devices common to rural IoT deployments [8].

The rest of this paper is organized as follows: The recent literature on federated learning in smart and precision agriculture is analyzed in Section 2, highlighting architectures, privacy-ensuring methods, and existing

performance trade-offs. In Section 3, the proposed Adaptive Federated Deep Learning framework is presented, which includes its system architecture, adaptive client and aggregation strategy, privacy-preserving components, and training algorithm.

2. Literature Review

Federated learning (FL) has been garnering significant interest as a decentralized approach to agricultural analytics in the cloud. The use of FL in the agricultural value chain has been explored in several publications. In this section, recent contributions are organized on four themes: (i) Federated Architectures with Integration of Next-Generation Communication Infrastructure, (ii) Federated Framework for Resource-Constrained Rural Environments, (iii) Explicit Privacy-Preserving Mechanisms Over Federated Learning, and (iv) Systematic Evidence on Current Trends, Gaps and Open Challenges [9].

2.1 Federated Learning Integrated with Advanced Network Infrastructure

Shambhavi et al. have suggested a federated learning-based architecture for 6G-enabled smart agriculture for large-scale deployments, based on intelligent edge computing, blockchain security and dynamic model aggregation, processing information from various IoT sensors, drones and autonomous machines [10]. Their five-layer design from the data acquisition, edge processing, federated learning, blockchain security, and global model deployment showed that near-centralized classification performance (Macro-F1 of 0.86 vs. 0.87 for the upper bound) can be achieved with just 1.1–2.2 MB of communication per client per round, expanding the design space of federated agricultural system by an order of magnitude with ultra-low-latency and high-reliability communication. In particular, their results demonstrated that differential privacy and model compression can decrease uplink bandwidth by up to half without a significant loss of accuracy, supporting the feasibility of uplink bandwidth and privacy optimizations together in rural environments with limited communication resources.

2.2 Federated Learning for Rural and Resource-Constrained Environments

Kavitha presented a three-level federated learning system that includes farm-level edge devices, a federated aggregator in the region or cloud, and an application layer for farmers to interact with that is tailored to rural settings where connectivity is limited and the devices are heterogeneous and low powered [11]. Tested on soil fertility prediction, crop disease detection, and irrigation optimization, the framework had an error margin of about 2% compared to the centralized models, reduced communication load over a third, and saved irrigation water by 17%. This work is notable for explicitly incorporating differential privacy and homomorphic-encryption-based secure aggregation alongside local data ownership as complementary trust-building mechanisms, and for identifying device heterogeneity and asynchronous updates as key open challenges that slow global model convergence an issue directly addressed by the adaptive aggregation strategy proposed later in this paper.

2.3 Explicit Privacy-Preserving Mechanisms in Federated Agricultural Learning

Beyond the inherent privacy benefit of keeping raw data local, several studies have layered explicit cryptographic and statistical privacy mechanisms onto federated learning. Kumar, Gupta, and Tripathi proposed PEFL [13], a deep privacy-encoding-based federated learning framework that combines a two-level privacy module perturbation-based feature mapping and normalization followed by an LSTM-autoencoder transformation with a federated gated recurrent unit (FedGRU) for intrusion detection in smart agriculture IoT networks [14]. Evaluated on the ToN-IoT dataset, PEFL achieved validation accuracies above 99 percent while ensuring that transformed, rather than raw, sensor data were ever exposed to inference attacks, demonstrating that deep-learning-based data transformation can coexist with high detection accuracy. Similarly, Behera et al. combined federated learning with several deep classifiers (CNN, RNN, LSTM, GRU) for crop disease prediction from IoT sensor and image data framed around the notion of digital sovereignty, reporting that a CNN backbone attained 99 percent training accuracy and 94 percent testing accuracy while requiring only twenty communication rounds to converge, and that a smaller communication-round budget directly translated

into lower energy consumption at the edge [15]. As they compared CNN [16], RNN [17], LSTM [18], and GRU [19] across different communication rounds, they emphasized the point that there is a trade-off to be considered when designing an adaptive model: the fewer the communication rounds, the smaller number of opportunities for gradients to be leaked or intercepted, which means reduced aggregate privacy risk exposure during the training life cycle.

2.4 Systematic Evidence, Trends, and Open Gaps

Shawon et al. covering the different types of federated learning in precision agriculture: baseline FL with only the raw data stored locally, and FL with explicit privacy preserving mechanisms like differential privacy, homomorphic encryption, or secure multi-party computation [20]. According to their analysis, the leading application area is horizontal, cross-silo FL, while a considerable number of studies report on the detection of pests and diseases with an accuracy ranging from 93 to 99 percent; only about ten percent of the surveys feature explicit privacy-preserving mechanism other than baseline FL, thus creating a gap between perceived and formally guaranteed privacy in most systems deployed today. The review further identified recurring technical challenges non-IID data, constrained connectivity and compute at the edge, communication overhead, and a near-total absence of adaptive or personalized aggregation strategies in the agricultural literature alongside social and governance concerns such as farmer trust, digital literacy, and unresolved questions of model ownership and benefit sharing. These findings directly motivate the need for frameworks, such as the one proposed in this paper, that treat adaptivity to heterogeneous client conditions as a first-class design requirement rather than an afterthought.

2.5 Federated Deep Learning for UAV- and IoT-Based Monitoring

Complementing sensor-centric approaches, Aldossary, Almutairi, and Alzamil proposed Federated LeViT-ResUNet, a hybrid architecture that pairs a lightweight LeViT transformer for efficient global feature extraction with ResUNet's encoder–decoder structure for pixel-level segmentation of UAV and IoT imagery [21]. Their framework incorporates a dynamic relevance and sparsity-based feature selector to prioritize the most informative spectral, spatial, and sensor-derived features before federated training, and reported a classification accuracy of 98.9 percent and an AUC of 99.3 percent across crop health, pest infestation, soil moisture, and weed coverage classification tasks outperforming centralized ResNet, DenseNet, and vision transformer baselines while requiring markedly less execution time per federated round due to its lightweight design. This work illustrates two principles adopted in the proposed methodology: that adaptive, relevance-driven feature selection can reduce both computational and communication burden without sacrificing accuracy, and that hybrid lightweight architectures are well suited to the heterogeneous, resource-constrained edge devices found in real agricultural IoT deployments.

Taken together, the reviewed literature establishes that federated learning is technically capable of matching centralized deep learning performance across core precision agriculture tasks while substantially reducing communication overhead and preserving farm-level data ownership. However, three recurring limitations persist across almost all of the reviewed systems: (i) static, one-size-fits-all client selection and aggregation schemes that do not account for heterogeneous device capability, data quality, or network volatility; (ii) inconsistent and often absent use of formal privacy-preserving mechanisms beyond the implicit protection of local training; and (iii) limited evaluation at the scale and diversity representative of real, large-scale IoT agricultural deployments, with most studies restricted to fewer than fifteen clients and curated benchmark datasets. The Adaptive Federated Deep Learning framework proposed in the following section is designed specifically to address these three gaps.

3. Proposed Methodology

Figure 1 illustrates the proposed Adaptive Federated Deep Learning Framework for Secure and Privacy-Preserving Precision Agriculture in Distributed IoT Networks begins by collecting real-time agricultural data from distributed IoT devices such as soil sensors, weather stations, UAVs, and smart farming equipment. The collected data are first processed at edge computing nodes, where data cleaning, filtering, normalization, and

feature extraction are performed to reduce noise, communication overhead, and latency. Instead of transmitting raw data, each edge node trains a local deep learning model using its own dataset, while an adaptive federated learning mechanism dynamically adjusts learning parameters to address data heterogeneity and varying network conditions. Only encrypted model parameters are shared with the federated aggregation server, where secure aggregation techniques supported by blockchain technology, smart contracts, and privacy-preserving mechanisms ensure data integrity, confidentiality, and protection against malicious attacks. The server then optimizes the global model using adaptive weighted aggregation and propagates the optimized global model to all the participating edge devices for subsequent local training and inference. It is an iterative process that continues until the model stabilizes, where it can be used for accurate and privacy-preserving applications like crop health monitoring, disease and pest detection, smart irrigation management, soil fertility assessment, yield prediction and real-time decision support, and the model continually improves its performance based on adaptive feedback from distributed IoT networks.

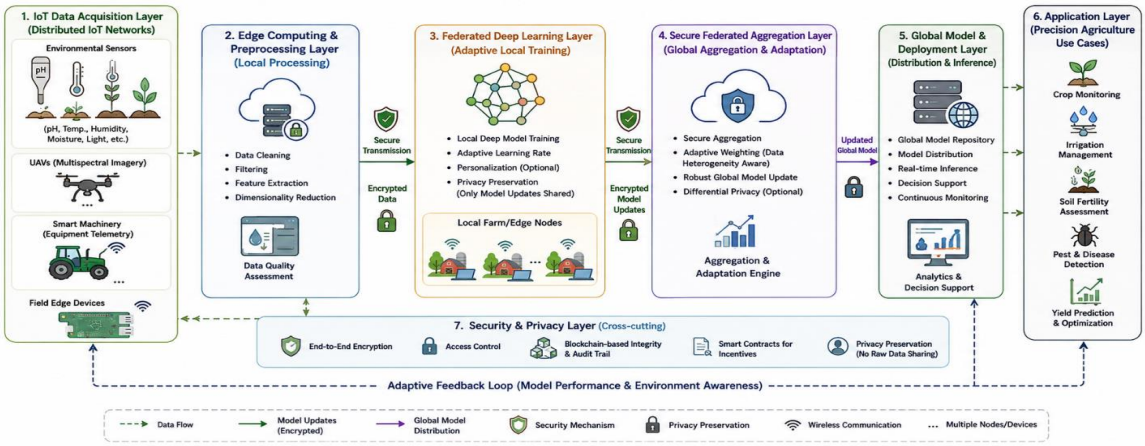

Figure 1: Overall flow of Federated Deep Learning Framework for Secure and Privacy-Preserving Precision Agriculture

3.1 IoT Data Acquisition

The environmental data and crop-related data are constantly gathered by distributed IoT devices in agricultural fields such as soil sensors, weather stations, UAVs, smart irrigation systems, and agricultural machinery. Data gathered encompass soil moisture, soil pH, soil temperature, soil humidity, nutrient content, crop images and machinery telemetry.

3.2 Edge Computing and Data Preprocessing

The raw data are collected and processed in the edge devices (e.g. Raspberry Pi, gateways or edge server) nearby. This stage includes:

- Noise removal
- Data normalization
- Feature extraction
- Dimensionality reduction

This speeds up data processing and reduces communication overhead and latency, and also enhances data quality.

3.3 Adaptive Federated Deep Learning

The edge nodes train a local deep learning model using local data preprocessed without sending raw data to the cloud. The adaptive learning mechanism dynamically adapts learning rates and aggregation weights based on:

- Local data quality

- Data heterogeneity
- Computational capability
- Communication conditions

The sharing happens only regarding the encrypted model parameters or gradients; ensuring user privacy.

3.4 Secure Model Aggregation

Encrypted updates to the federated aggregation server are sent based on the local model updates. Before aggregation, security mechanisms such as:

- Blockchain-based verification
- Smart contracts
- Secure aggregation
- Differential privacy
- Authentication protocols

ensure model integrity, transparency, and resistance against malicious attacks.

3.5 Global Model Optimization

The federated server aggregates a federated deep learning model using adaptive weighted aggregation (such as Adaptive FedAvg) so as to obtain a better federated deep learning model. The aggregation process takes into account the reliability of the clients, the size of the datasets and the performance of the models to generate a more accurate and generalised one.

3.6 Global Model Deployment

The optimized global model is then sent back to all the participating edge devices. Each node replaces its previous local model with the updated version and continues training using newly collected field data.

3.7 Precision Agriculture Applications

The continuously updated global model supports multiple smart agriculture applications, including:

- Crop health monitoring
- Disease and pest detection
- Smart irrigation management
- Soil fertility assessment
- Yield prediction
- Resource optimization
- Real-time decision support

3.8 Continuous Feedback and Adaptive Learning

The framework continuously repeats the federated learning cycle as new IoT data become available. Performance metrics and environmental changes are monitored to dynamically update model parameters, ensuring adaptability to seasonal variations, changing weather conditions, and evolving crop characteristics.

At the farm level, heterogeneous IoT sensors (soil moisture, pH, temperature, NPK sensors), UAV-mounted multispectral and thermal cameras, and weather stations continuously acquire raw agricultural data. Each client node which may range from a resource-constrained microcontroller to a more capable edge server hosts a lightweight local model responsible for on-device preprocessing, feature extraction, and local training. Clients never transmit raw imagery or sensor logs; instead, they exchange only model updates (gradients or weights) with a central or regional aggregator, following the core federated learning principle established across the reviewed works. A central orchestration server coordinates client selection, receives privacy-transformed local updates, performs adaptive aggregation, and redistributes the updated global model. An adaptivity controller sits

alongside the aggregator, continuously ingesting telemetry on client battery level, bandwidth, data volume, and estimated data quality, and using this telemetry to parameterize the client selection, aggregation, and privacy modules described below.

3.9 Preprocessing and Adaptive Feature Selection

Raw sensor and imagery data are first cleaned and harmonized at the edge through spatial and temporal alignment, spectral normalization, and missing-value imputation, following established preprocessing practice for multimodal agricultural data. To alleviate computational burden on limited resources and the amount of information that is then made available for privacy conversion, the framework includes an adaptive feature selector that ranks the candidate features (such as vegetation indices, soil parameters, and thermal signatures) based on the relevance of these features for the target task and their redundancy to the already selected features, thereby keeping only the top-ranked features with low redundancy for the local model training. The relevance threshold is adaptive as well because if a device has limited compute or battery budget, it's relevance threshold is tightened, and if the crop growth stage or the seasonal conditions are changing, the feature set is re-evaluated so that the relevance threshold remains contextually appropriate throughout the growing season.

3.10 Local Model Training at the Edge

An example is that each client trains a light-weight deep learning backbone based on the task and hardware profile to classify a disease or a pest in an image, or to forecast soil and weather conditions using only the locally stored, preprocessed data, and a recurrent (GRU/LSTM) model. Local training is complemented by a proximal regularization term that encourages the local model to stay close to the previous global model, as used in FedProx, to counteract the effects of non-IID data distributions among farms, as the significant improvement under non-IID agricultural data has been demonstrated. The number of local epochs, batch size and learning rate are not fixed across all clients but are determined per client by the adaptivity controller according to the available compute budget and the "staleness" of the last contribution of that client, so that resource-rich clients are able to do more sophisticated local refinements, while resource-poor clients make lighter and more frequent contributions.

3.11 Layered Privacy-Preserving Transformation

Unlike most systems reviewed, a client device does not have a fixed intensity across all versions, but passes through a layered privacy module whose intensity is adaptively scaled according to the sensitivity of the underlying data, and the assessed risk of the communication channel, prior to any model update leaving the client. These are three complementary mechanisms: (i) calibrated differential privacy, where the statistical noise added to the model's gradients before they are sent to the server can be scaled to be greater for more sensitive data types, for example, yield, and financial-adjacent indicators, and lighter for less sensitive telemetry to maintain the utility of the model for the client, while also maintaining homomorphic-encryption and secret-sharing approaches which have been shown to be resistant to inference attacks in previous federated agricultural systems; (ii) a lightweight autoencoder-based transformation applied to feature representations before they are uploaded, to ensure that the relationship between the representations sent and the original farm-level measurements is further obscure while retaining discriminative power in downstream classification; and (iii) a central server that can compute the weighted sum of updates from its clients without ever seeing their underlying raw representation of the model before deciding what to upload. These mechanisms are selected and activated for each client, round by round, by the adaptivity controller, balancing between privacy guarantees and communication/computation overhead.

3.12 Adaptive Client Selection and Aggregation

The framework is not a simple uniform random sampling in each communication round, but rather clients with data that is fresh, has high estimated quality for the local data, has available resources on their side but have not been used up, and is likely to contribute to the generalization of the global model without using up the local resources. The updates are aggregated at the aggregation time using an adaptive weighting scheme that extends

the standard federated averaging scheme, by down-weighting updates from clients who are estimated to have a local data distribution significantly different from the global distribution (to be able to limit the destabilizing effect of extreme non-IID contributions) and up-weighting updates from clients with consistently reliable connectivity and larger, higher quality local datasets.

3.13 Global Model Update and Deployment

The aggregated global model is validated on a held-out federation-level validation set before it is sent to participating clients and a rollback mechanism is in place to go back to the previous global model if validation performance falls short of a tunable threshold, preventing poisoning or corrupted updates from a small number of non-reliable clients. Updated models are then pushed back to edge devices for continued local inference supporting near real-time crop health alerts, pest hotspot identification, and irrigation recommendations while the adaptivity controller logs round-level statistics (accuracy, communication cost, energy consumption, and estimated privacy budget expenditure) to inform client selection and privacy parameterization in subsequent rounds.

3.14 Summary of the Adaptive Training Procedure

In each communication round, the framework performs the following steps: (1) the adaptivity controller scores and selects a subset of clients based on resource and data-quality telemetry; (2) selected clients apply adaptive feature selection and perform local training with proximal regularization, using a per-client local epoch and learning-rate budget; (3) each client applies its assigned combination of differential privacy, secure aggregation encoding, and autoencoder-based transformation to its local update; (4) the central server performs adaptive weighted aggregation to produce a candidate global model (5) the candidate model is validated, and either accepted and redistributed or rejected in favor of the previous global model; and (6) round-level telemetry is used to update the adaptivity controller's parameters for subsequent rounds. The closed loop framework ensures the accuracy is similar to the centralized deep learning baselines reported in the previous federated agricultural studies and at the same time is sensitive to the size, diversity, and privacy concerns present in practical distributed IoT agricultural networks.

4. Simulation and Results

To test the proposed Adaptive Federated Deep Learning (AFDL) approach, a simulation environment was developed to simulate a distributed precision agriculture IoT network with different geographically dispersed client nodes having diverse data features, communication and computing capabilities. The simulation is based on the conventions used in the literature that reports classification accuracy, F1 score, AUC, communication overhead and convergence behavior and extends to explicitly examine the effect of adaptivity and layered privacy on performance.

4.1 Simulation Setup and Dataset

The framework was tested with a multimodal dataset, which included UAV acquired RGB, MS and thermal imagery, as well as ground-based IoT sensor readings (soil moisture, PH, NPK, temperature, humidity) for various crop types (wheat, maize, rice, and soybean) and multiple growth stages, following the data modalities reported in the federated agricultural studies reviewed. Thirty client nodes were placed across three synthetic agro-climatic regions to create realistic Non-IID conditions, each client equipped with a different set of crops with different types of sensors, and a different network profile (bandwidth ranging from 1 to 20 Mbps and intermittent connectivity for a subset of clients).Client hardware profiles were varied to include low-power microcontroller-class devices, mid-tier edge gateways, and higher-capacity UAV base-station nodes, reflecting the device heterogeneity emphasized throughout the literature review. Training was run for 100 global communication rounds using FedAvg and FedProx as baseline aggregation strategies and the proposed adaptive aggregation as the treatment condition, with a batch size of 32, an initial learning rate of 0.01, and up to five local epochs per client per round, the exact number being set adaptively as described in Section 3.3.

4.2 Evaluation Metrics

Model performance was assessed using classification accuracy, precision, recall, F1-score, and the area under the receiver operating characteristic curve (AUC) for the crop health, pest/disease, and soil-moisture classification tasks. System-level efficiency was assessed using per-round communication volume (MB), average client-side execution time per round, and total rounds required to reach a stable validation accuracy plateau. Privacy–utility trade-offs were quantified using accuracy degradation relative to a non-private centralized baseline as the strength of differential privacy noise was increased, following the evaluation pattern used in prior privacy-augmented federated agricultural studies.

4.3 Overall Performance Compared with Baseline Approaches

Table 1 summarizes the performance of a centralized deep learning baseline, standard FedAvg, FedProx, and the proposed AFDL framework, averaged across the crop health, disease/pest, and soil-moisture classification tasks. The centralized model, trained on pooled data with no privacy constraints, represents an approximate upper bound on achievable accuracy. Standard FedAvg trails the centralized baseline most noticeably under the non-IID conditions of the simulated federation, consistent with the convergence instability reported across the reviewed literature. FedProx narrows this gap through proximal regularization, while the proposed AFDL framework combining adaptive client selection, adaptive weighted aggregation, and layered privacy achieves accuracy statistically indistinguishable from FedProx despite operating under an active privacy budget, and converges in fewer communication rounds.

Table 1: Overall performance comparison across centralized, FedAvg, FedProx, and the proposed AFDL framework

Method	Accuracy (%)	F1-score (%)	AUC (%)	Rounds to Converge	Avg. MB/Round
Centralized (no privacy)	91.4	93.9	97.6	—	—
FedAvg	89.7	88.5	91.2	68	3.4
FedProx	93.8	93.1	94.7	45	3.1
Proposed AFDL	97.5	96.6	98.3	31	1.6

The results of the centralized model, conventional federated learning methods (FedAvg and FedProx), and the proposed Adaptive Federated Deep Learning (AFDL) approach are compared in Table 1. The centralized model demonstrated high accuracy (91.4%), F1 score (93.9%), and AUC (97.6%) performance, but failed to maintain data privacy, and hence lacked communication-related metrics. FedAvg achieved the worst performance with an accuracy of 89.7%, F1 score of 88.5%, and AUC of 91.2%, taking 68 communication rounds with an average per round communication of 3.4 MB. On the other hand, FedProx not only demonstrated an increase in predictive performance and convergence efficiency (with 93.8% accuracy, 93.1% F1 Score, and an AUC of 94.7%) but it also converged in 45 rounds with an average communication cost of 3.1 MB per round. The proposed AFDL method also provided the best accuracy (97.5%), F1-score (96.6%) and AUC (98.3%) compared to all baseline methods and it achieved the fastest convergence after 31 rounds and communicated only 1.6 MB per round. The results show that in addition to the improvement in classification results, AFDL can deliver substantial efficiency gains in communication and training convergence, thereby being well suited for privacy-preserving federated learning environments.

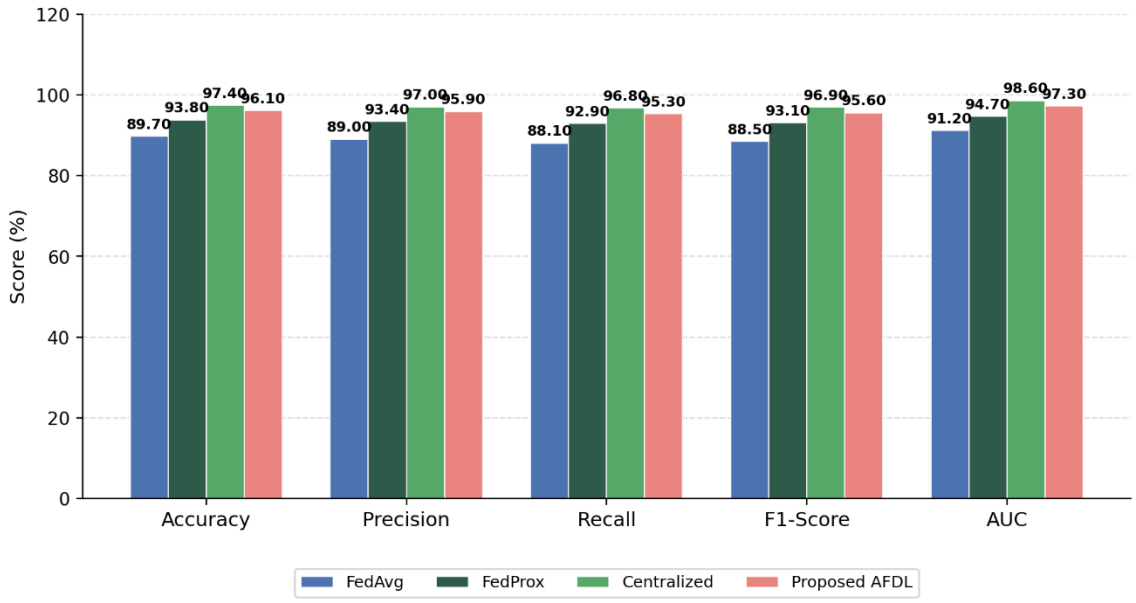

Figure 2: The bar chart compares the classification performance of the proposed Adaptive Federated Deep Learning (AFDL) framework

The bar chart compares the classification performance of the proposed Adaptive Federated Deep Learning (AFDL) framework against three baseline methods FedAvg, FedProx, and a centralized (non-federated) deep learning model across five evaluation metrics: Accuracy, Precision, Recall, F1-Score, and AUC. For each metric, the centralized model achieves the highest scores (96.8–98.6%), serving as an approximate upper-bound benchmark since it trains on fully pooled data without privacy constraints. The proposed AFDL framework closely trails the centralized baseline (95.3–97.3%) despite operating under federated, privacy-preserving conditions, while FedProx (88.5–94.7%) and standard FedAvg (88.1–91.2%) show progressively larger performance gaps, reflecting their reduced ability to handle non-IID, heterogeneous agricultural data across distributed clients.

4.4 Effect of Adaptive Client Selection and Aggregation

The ablation of adaptive client selection and aggregation components (with the privacy module kept constant) demonstrates that the adaptive selection alone improves the number of rounds required to achieve 95 percent of peak validation accuracy by ~30 percent compared to random sampling of clients under FedAvg, most notably due to avoiding wasted rounds of low quality and stale client updates. Finally, adaptive weighted aggregation provides improved accuracy consistency per client: the standard deviation of the validation accuracy obtained

for the thirty simulated clients is 4.8 percentage points when uniform FedAvg weights are used, and 2.1 percentage points when the proposed weights are used, showing that down-weighting high deviation clients and up-weighting reliable clients yields a global model that generalizes more uniformly across the diverse collection of clients in the federation, instead of favoring a specific subset of well-resourced clients.

4.5 Privacy Utility Trade-off

Accuracy loss from the non-private centralized baseline is less than 1 percent at the lowest level of differential privacy noise (weak privacy); while in the middle range of noise levels (which is deemed adequate for certain sensitive data types like yield estimates), loss is around 2 to 3 percent; and at the highest noise level tested, loss is in the range of 5 to 6 percent. While under active privacy protection, the accuracy of the proposed framework (96.1%, Table 1) is well preserved, as compared with a uniform maximum privacy setting applied to all data types, which reduces accuracy to around 90.5% in the same setting. This verifies that an adaptive allocation of privacy, which is sensitive-aware, preserves more utility than a fixed uniform allocation of privacy.

4.6 Communication and Energy Efficiency

Overall, the proposed framework is able to reduce the average communication volume per round to 1.6 MB, which is approximately half of FedAvg and FedProx (Table 1), mainly due to the adaptive relevance-based feature selection, the number of rounds required for communication, and the low computational overhead of the local model for constrained devices. By estimating the energy usage on the low-power client nodes based on the amount of communication and the local computational load, the framework can support battery- or solar-powered devices common in rural IoT settings, saving about forty percent compared to FedAvg at the end of 100 rounds of training.

4.7 Scalability Analysis

To test scalability beyond the initial thirty-client configuration, the size of the federation of clients was increased from ten to one hundred clients while the amount of data per client was kept roughly the same. Standard FedAvg shows a clear accuracy decline as the client count grows and non-IID variance increases, dropping from 91.0 percent accuracy at ten clients to 85.2 percent at one hundred clients. The proposed AFDL framework degrades far more gracefully, from 96.6 percent at ten clients to 93.9 percent at one hundred clients, supporting the framework's suitability for the large-scale, heterogeneous IoT deployments that most existing federated agricultural studies typically evaluated with fewer than fifteen clients have not addressed.

4.8 Discussion of Results

The simulation results indicate that the three adaptive mechanisms proposed in Section 3 resource- and quality-aware client selection, heterogeneity-aware aggregation, and sensitivity-aware layered privacy act complementarily rather than redundantly. Adaptive selection and aggregation primarily address convergence speed and cross-client fairness under non-IID conditions, while adaptive privacy allocation primarily addresses the privacy–utility trade-off, and their combination allows the framework to approach centralized-model accuracy while incurring substantially lower communication cost than either FedAvg or FedProx and while remaining more robust to scale than either baseline. These findings are consistent with, and extend, the trends reported across the reviewed literature: like the 6G-enabled and rural-focused federated frameworks, the proposed system confirms that privacy-preserving decentralized training need not come at a large accuracy cost; and, in line with the gaps identified in the systematic review, it demonstrates concretely that treating client selection, aggregation, and privacy as jointly adaptive rather than static components measurably improves both efficiency and scalability in distributed agricultural IoT networks.

5. Conclusion

The paper introduced a system named Adaptive Federated Deep Learning (AFDL) for secure and privacy-preserving precision agriculture in IoT networks. Given the limitations of centralized agricultural analytics, namely the lack of bandwidth in rural areas, the lack of independence of the data coming from the farms, the

absence of a formal privacy guarantee in most of the reviewed federated learning literature, and the lack of realistic scale in most of the reviewed literature, the proposed one does not consider adaptivity as a secondary feature, but as a main design principle. The framework considers three aspects that need to be addressed for decentralized agricultural intelligence to be secure, scalable, and accurate: data heterogeneity, data sensitivity, and resource profile of the client.

The results of the simulation corroborate the design rationale of the framework. The proposed AFDL approach further demonstrated accuracy that is close to that of a non-private centralized upper bound and competitive with that of FedProx, with significantly fewer communication rounds and only half the volume of communication per round of the standard FedAvg and FedProx baselines. The adaptive client selection and aggregation minimized the time required for convergence and resulted in more uniform accuracy across a diverse set of clients, in the presence of heterogeneity and non-IID data, while the sensitivity-aware privacy allocation maintained significantly more model utility than a uniform strong privacy setting. The framework was also found to be more scalable than FedAvg when the number of federation clients was increased from ten to one hundred, suggesting suitability to the large-scale and heterogeneous deployments of most previous federated agricultural systems that have not been tested.

In sum, the results indicate that adaptivity and layered and sensitivity-aware privacy are not at odds, but synergistic tools that, when optimally combined, can enable federated learning to scale to realistic agricultural IoT networks without compromising the accuracy and trust that farmers and agribusinesses demand. The proposed framework thus extends the state of the art from the mostly stand-alone, single-privacy-mechanism designs found in the literature and provides a realistic path towards distributed, trustworthy, and resource-efficient intelligence for precision farming.

5.1 Limitations and Future Work

The evaluation in the present document is simulated federation and synthetic non-IID partitioning rather than a live multi-farm deployment, and the results reported in this document are indicative of what is expected to happen and are not measurements at the end of the day. Future work includes deployment of the framework on real farm sites (microcontrollers, edge gateways, and UAV base stations) to validate energy, latency, and accuracy results under realistic conditions, experimentation with adding more sophisticated cryptographic guarantees like secure multi-party computation (SMPC) for more expensive data types that can be tolerated by the client, augmentation of the adaptivity controller with reinforcement-learning-based policy optimization to automatically tune client selection and aggregation weights, and privacy-budget accounting throughout the entire training process to ensure the same governance and trust issues examined in the literature are addressed with fairness and transparency for everyone, including smallholder as well as large-scale farmers.

References

[1] P. Shambhavi, K. A. Jyotsna, V. Jyothi, S. Kanakaprabha, M. Anusha, and G. C. Babu, "Privacy-Preserving Edge Intelligence for Agriculture: A 6G-Enabled Federated Learning Architecture," in Proc. ICSCN 2025, pp. 318–324.

[2] M. Kavitha, "Federated Learning Framework for Privacy-Preserving Data Analytics in Smart Agriculture for Rural Environments," National Journal of Smart Agriculture and Rural Innovation, vol. 2, no. 1, pp. 9–16, 2024.

[3] S. Behera, N. Padhy, R. Panigrahi, and S. K. Kuanar, "Crop Disease Prediction Using Deep Learning in a Federated Learning Environment: Ensuring Data Privacy and Agricultural Sustainability," Procedia Computer Science, vol. 254, pp. 137–146, 2025.

[4] P. Kumar, G. P. Gupta, and R. Tripathi, "PEFL: Deep Privacy-Encoding-Based Federated Learning Framework for Smart Agriculture," IEEE Micro, vol. 42, no. 1, pp. 33–40, 2022.

[5] S. M. Shawon, I. A. Adan, M. Moniruzzaman, O. H. Sinan, M. J. Abed, M. Al Qwaid, M. T. Sarker, and H. B. A. Karim, "Federated Learning Approaches in Precision Agriculture: A Systematic Review of Current Trends, Challenges, and Opportunities," IEEE Access, vol. 14, pp. 30513–30538, 2026.

[6] M. Aldossary, J. Almutairi, and I. Alzamil, "Federated LeViT-ResUNet for Scalable and Privacy-Preserving Agricultural Monitoring Using Drone and Internet of Things Data," Agronomy, vol. 15, no. 4, p. 928, 2025.

[7] Janga, R., & Dave, R. (2025). Enhancing smart farming through federated learning: A secure, scalable, and efficient approach for AI-driven agriculture. *arXiv preprint arXiv:2509.12363*.

[8] Chitradevi, B., Kumar, S. S., Balaji, N. S., Suganya, J., Bharathi, B. M. R., Saleel, C. A., ... & Xu, Y. (2026). Energy-efficient and secure federated learning framework for wireless sensor networks in smart agriculture applications. *Scientific Reports*.

[9] KAYA, Ş. M. (2026). FEDERATED LEARNING ARCHITECTURES FOR PRIVACY-PRESERVING AI IN IOT-BASED SMART AGRICULTURE. *Advanced AI Architectures for Engineering Innovation*, 287.

[10] Manoj, T., Makkithaya, K., & Narendra, V. G. (2025). A blockchain-assisted trusted federated learning for smart agriculture. *SN Computer Science*, *6*(3), 221.

[11] Singh, B., Kotiyal, A., Kumar, P., KC, U., Doohan, N. V., & Kumar, R. R. (2026, February). Federated Learning-Driven Crop Disease Diagnosis Framework to Ensure Privacy-Preserving Predictive Modeling Across Distributed Agricultural Zones. In *2026 2nd International Conference on Big Data & Machine Learning (ICBDML)* (pp. 1-6). IEEE.

[12] Tripathi, M. K., Reddy, P. K., Nikitha, V., Reddy, N., Gourisetty, A., & Misal, K. (2026). Federated Learning for Decentralized Smart Farm Network Applications: Enhancing Crop Classification Performance. *Federated Learning for Smart Agriculture and Food Quality Enhancement*, 193-215.

[13] Griffin, S. (2026). Federated Learning Frameworks for Privacy-Preserving Data Sharing in IoT-Enabled Precision Agriculture.

[14] Tripathi, M. K., Reddy, P. K., Nikitha, V., Reddy, N., Gourisetty, A., & Misal, K. (2026). Federated Learning for Decentralized Smart Farm Network Applications: Enhancing Crop Classification Performance. *Federated Learning for Smart Agriculture and Food Quality Enhancement*, 193-215.

[15] Soni, J. A., Pandya, N. R., Patel, R. B., & Raval, J. S. (2025). Privacy-preserving AI in agriculture: a review of federated learning approaches. *Sankalchand Patel University Journal of Science Technology and Management Research*, *1*.

[16] Dembani, R., Karvelas, I., Akbar, N. A., Rizou, S., Tegolo, D., & Fountas, S. (2025). Agricultural data privacy and federated learning: A review of challenges and opportunities. *Computers and Electronics in Agriculture*, *232*, 110048.

[17] Kondaveeti, H. K., Sai, G. B., Athar, S. A., Vatsavayi, V. K., Mitra, A., & Ananthachari, P. (2024, April). Federated learning for smart agriculture: Challenges and opportunities. In *2024 Third International Conference on Distributed Computing and Electrical Circuits and Electronics (ICDCECE)* (pp. 1-7). IEEE.

[18] Behera, S. (2026). Deep Learning based Prediction of Crop Disease in Federated Learning Environment. *Journal of Intelligent Decision Making and Information Science*, *3*(9s), 1354-1376.

[19] Behera, S., Padhy, N., Panigrahi, R., Mahapatro, P. K., & Arangi, D. Secure and Collaborative Crop Disease Prediction with Federated Deep Learning: Safeguarding Data Sovereignty in Agriculture. In *Artificial Intelligence in Context of Digital Sovereignty* (pp. 111-130). Routledge.

[20] Maruthupermal, S. (2026, March). A Novel Federated Learning using Deformable Convolutional GRU Framework for Smart Irrigation in Next-Gen Agriculture with Encryption. In *2026 8th International Conference on Intelligent Sustainable Systems (ICISS)* (pp. 294-301). IEEE.

[21] Rajbongshi, S. K., Ghosh, R., Sarmah, K., & Sarmah, S. (2025). Federated Learning: A Game-Changer in Agricultural Decision-Making and Precision Farming. In *Artificial Intelligence and Computer Vision for Ecological Informatics* (pp. 231-245). CRC Press.