Privacy Meets Conservation: Federated Learning's Revolution in Deforestation Detection

Samir Poudel
Computational and Data Science
Middle Tennessee State University
Murfreesboro, TN, 37132
sp2ai@mtmail.mtsu.edu

Arpan Man Sainju
Department of Computer Science
Middle Tennessee State University
Murfreesboro, TN, 37132
arpan.sainju@mtsu.edu

Kritagya Upadhyay
Department of Computer Science
Middle Tennessee State University
Murfreesboro, TN, 37132
kritagya.upadhyay@mtsu.edu

2025 IEEE World AI IoT Congress (AIIoT) | 979-8-3315-2508-8/25/$31.00 ©2025 IEEE | DOI: 10.1109/AIIoT65859.2025.11105254

***Abstract*—This study presents a secure and scalable federated learning (FL) framework for environmental monitoring, focusing on deforestation detection using remote sensing data. Unlike traditional centralized approaches that require data pooling, our framework enables model training across geographically distributed regions while preserving data privacy. We address key challenges in FL implementation, including fairness across clients and robustness to noisy data. Using Sentinel-2 satellite imagery spanning 2015-2023 from six distinct geographical regions, we train a modified U-Net architecture for binary classification of forested and non-forested areas. We compare two aggregation algorithms, FedAvg and Krum, across 10 communication rounds with seven clients, including six geographical regions and one noise client (ocean data). Results demonstrate that Krum consistently outperforms FedAvg in terms of Intersection over Union (IoU) score (0.85 vs. 0.77), convergence speed, and fairness metrics. Krum achieves lower accuracy disparity (0.15 vs. 0.40), IoU disparity (0.30 vs. 0.65), and loss disparity (0.25 vs. 0.80) compared to FedAvg, while maintaining a higher equity score (0.92 vs. 0.70). Our framework effectively handles class imbalance through balanced sampling and demonstrates robustness to noisy data, making it suitable for large-scale, privacy-preserving environmental monitoring applications.**

***Index Terms*—federated learning, environmental monitoring, remote sensing, Krum algorithm, data privacy, model fairness, robust aggregation**

I. Introduction

A. Background

Deforestation poses a critical global environmental threat, resulting in habitat loss, carbon emissions, and disruption of ecosystem services. Timely and accurate detection of forest loss is essential for conservation efforts, climate change mitigation, and sustainable resource management [1]. Traditional monitoring approaches face challenges including limited coverage, delayed detection, and insufficient resolution to capture smaller-scale forest disturbances. Federated Learning (FL) offers a promising approach for deforestation detection by enabling decentralized machine learning across multiple clients while preserving data privacy [2]. Instead of centralizing raw data, FL shares only model updates with a server for aggregation, addressing key challenges in environmental monitoring. Deforestation detection specifically benefits from federated learning due to several critical factors. First, monitoring often spans multiple countries with different data sovereignty regulations, making data sharing legally complex or prohibited. Second, forest ecosystems exhibit significant regional heterogeneity in species composition, canopy structure, and deforestation patterns, requiring localized model adaptations. Finally, many regions with high deforestation rates (e.g., parts of the Amazon) have limited connectivity infrastructure, making massive data transfers impractical. FL addresses these challenges by keeping raw data local while enabling collaborative model training.

Beyond deforestation, FL has gained prominence in healthcare, finance, and other environmental monitoring applications [2]. For environmental tasks, it allows the use of geographically distributed datasets for land cover classification and climate change monitoring without compromising sensitive information. Real-world FL deployments include Google's mobile keyboard prediction [11] and collaborative medical image analysis among hospitals [12]. These implementations provide a foundation for environmental monitoring applications. Our study develops and evaluates a secure and scalable federated learning framework for environmental monitoring,

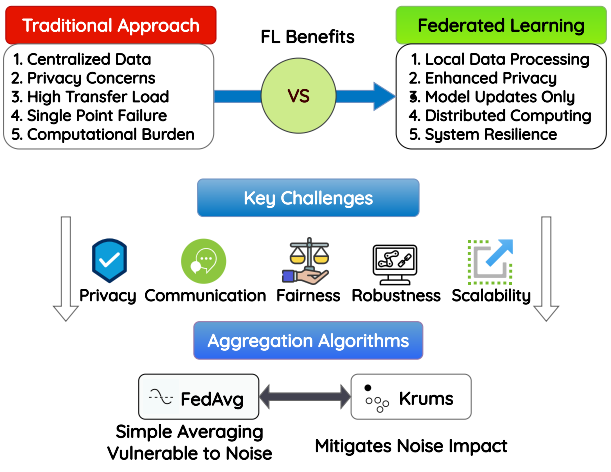

Fig. 1. Federated learning enhances environmental monitoring by addressing privacy risks, reducing data transfer, and enabling decentralized computation overcoming the limitations of traditional centralized learning [3].

979-8-3315-2508-8/25/$31.00 ©2025 IEEE

TABLE I
COMPARISON OF OUR WORK'S CONTRIBUTION TO OTHER EXISTING WORKS IN A SIMILAR FIELD

Feature	Paper 1 [4]	Paper 2 [5]	Paper 3 [6]	Paper 4 [7]	Paper 5 [8]	Paper 6 [9]	Paper 7 [10]	Our Work
Year	2023	2024	2023	2023	2020	2018	2022	2025
Federated Learning	✓	✓	✓	✓	×	×	✓	✓
Privacy Preservation	✓	✓	✓	✓	×	×	✓	✓
Remote Sensing Data	✓	✓	×	×	×	×	×	✓
Environmental Focus	✓	✓	✓	✓	✓	✓	✓	✓
Aggregation Algorithms	✓	✓	✓	✓	×	×	✓	✓
Fairness Metrics	×	×	×	×	×	✓	×	✓
Robustness to Noise	✓	×	✓	×	✓	×	×	✓
Scalability Evaluation	✓	✓	×	✓	✓	✓	×	✓
Class Imbalance Handling	×	×	×	✓	×	×	×	✓
Multi-temporal Analysis	×	×	×	✓	✓	×	×	✓
Geographical Diversity	×	×	×	✓	✓	✓	×	✓
Communication Efficiency	×	✓	✓	✓	✓	×	✓	✓

focusing on communication efficiency, client fairness, and robustness to noisy data.

B. Problem Motivation

Forest monitoring systems processing satellite imagery across jurisdictional boundaries face critical limitations with centralized machine learning approaches, including data sovereignty concerns and bandwidth constraints. Federated Learning (FL) addresses these challenges by enabling collaborative model training without raw data sharing [11], particularly valuable for analyzing diverse forest ecosystems where regional variations and connectivity limitations create barriers to centralized solutions.

- **Environmental Data Sovereignty:** Indigenous communities monitoring ancestral forests often have legal restrictions on sharing imagery containing sacred sites, making data pooling impossible. For example, the Kayapó people in Brazil maintain exclusive rights to imagery of their 11-million-acre territory while still needing deforestation alerts [13].
- **Remote Sensing Communication:** Deforestation hotspots in the Amazon and Congo Basin have limited connectivity. FL reduces bandwidth needs by sharing only model weights instead of raw imagery. The monitoring of the Andean Amazon Project operates in regions where a single Landsat image (1GB+) would take hours to upload, while model weights require only megabytes [14].
- **Ecosystem Fairness:** Centralized models achieve 95% accuracy in temperate forests but only 70% in mangroves [15]. FL allows region-specific optimization while contributing to global knowledge. The Global Mangrove Watch initiative has documented how standard algorithms consistently underperform in these complex coastal ecosystems due to tidal variations [16].
- **Seasonal Variation Robustness:** Scandinavian forests experience extreme daylight variations while tropical regions face persistent cloud cover. FL incorporates these regional This methodology embraces challenges rather than excluding "noisy" data.Finland's National Forest Inventory must adapt to conditions where winter imagery differs dramatically from summer, requiring season-specific training approaches [17].
- **Multi-jurisdiction Scalability:** REDD+ spans 50+ countries with varying data policies [18]. FL respects national sovereignty while enabling cross-border collaboration. Colombia and Peru share Amazon boundaries but have different data sharing regulations, limiting conventional cross-border forest monitoring despite ecological continuity [19].

Addressing these challenges is critical for effective decentralized environmental monitoring systems that can operate across diverse ecological, cultural, and political boundaries while maintaining high accuracy and respectful data practices.

C. Our Contributions

This paper makes the following key contributions:

- We proposed a privacy-preserving federated learning framework for deforestation detection that maintains data locality while enabling cross-regional collaboration.
- We conducted a comparative analysis of FedAvg and Krum aggregation algorithms for handling diverse forest ecosystems and seasonal variations.
- We developed fairness assessment metrics ensuring consistent model performance across different forest biomes.
- We designed a robust methodology maintaining accuracy despite noisy satellite data from regions with cloud cover interference.
- We implemented a scalable system validated across multiple conservation jurisdictions.

II. LITERATURE REVIEW

As shown in Table I, our work addresses key challenges in deforestation monitoring through federated learning approaches that previous studies have only partially addressed.

A. Deforestation Monitoring Approaches and Challenges

Deforestation monitoring has evolved from manual interpretation of aerial photographs [1] to satellite-based methods using Landsat and Sentinel-2 [20]. ALOS PALSAR data

demonstrated utility for tropical regions [21], while convolutional neural networks automated feature extraction from multispectral data [22]. Despite advances, challenges persist including cloud cover [23] and seasonal variations [24]. Recent studies explored SAR and optical data fusion to overcome cloud obstruction [25], with [26] proposing Sentinel-1 and Sentinel-2 fusion for near-real-time alerts. Technical obstacles in implementation include class imbalance, where non-forested pixels vastly outnumber deforestation cases [27], addressed through focal loss [28] and balanced sampling techniques [29]. Regional heterogeneity in forest types and deforestation patterns between biomes like the Amazon and Southeast Asia [30] complicates model generalization [31]. Data quality issues stemming from cloud cover and sensor inconsistencies have been approached through cloud-removal algorithms [32] and multi-sensor fusion approaches [33]. Scalability remains a concern for global monitoring [34], with federated learning offering a promising solution [35].

B. Federated Learning in Environmental Monitoring

Federated learning preserves data privacy while leveraging distributed datasets for environmental applications. [36] introduced FL for weather forecasting, while [37] applied it to air quality prediction with reduced communication costs. For deforestation monitoring, FL enables regional servers to train models locally, sharing only model updates globally [38]. FL addresses key deforestation monitoring challenges through multiple mechanisms. It enables locally-optimized sampling strategies for class imbalance without compromising data privacy [35] and supports personalized local models [39] for regional heterogeneity while still benefiting from global knowledge sharing. Federated frameworks can incorporate cloud-removal and fusion techniques locally while maintaining model interoperability. Implementation challenges include communication efficiency, addressed by [11] through gradient compression, and fairness, approached by [40] via agnostic aggregation. Robust algorithms like Krum [41] mitigate noisy clients, though trade-offs between robustness and convergence speed require further study [3]. Despite these advances, challenges in computational fairness [42] and communication efficiency remain, requiring balanced client participation protocols and optimized update mechanisms for operational global-scale monitoring systems.

III. METHODOLOGY

A. Data Collection and Preprocessing

The dataset used in this study was sourced from the Copernicus Sentinel-2 satellite imagery [43], which provides high-resolution multispectral data suitable for environmental monitoring tasks. Specifically, the Sentinel-2 Level-2A product was utilized, offering atmospherically corrected surface reflectance data with a spatial resolution of 10 meters for visible and near-infrared bands. Additionally, ground truth labels for forest and non-forest regions were extracted from the Hansen Global Forest Change dataset [1], which served as the reference for deforestation detection. Our study focused on six distinct geographical regions, each characterized by unique ecological and climatic conditions. These regions included: **Regions (2015–2023):** 1. Amazon Rainforest, 2. West Africa, 3. Congo Basin, 4. East Africa, 5. Greater Mekong, 6. Central Atlantic Forest, 7. Noise data (Ocean) The noise client (Region 7) consisted of ocean imagery specifically selected from the Atlantic and Pacific oceans where no land or forest was present. This noise client was deliberately included to evaluate the robustness of our aggregation algorithms against irrelevant or misleading data. We intentionally introduced this noise client to simulate real-world scenarios where some participants in the federated learning framework might contribute low-quality, adversarial, or completely unrelated data. This tests the system's resilience against potential malicious actors or poorly calibrated sensors that could otherwise compromise model performance across all regions. By measuring how well different aggregation methods (FedAvg and Krum) identify and mitigate the impact of such irrelevant inputs, we can better design federated systems that maintain high performance even when some participants contribute harmful updates. The ocean imagery was preprocessed using the same pipeline as the forest regions, but with labels randomly assigned to create a dataset with no meaningful patterns for deforestation detection. For each region, data spanning from 2015 to 2023 was collected, ensuring temporal consistency for longitudinal analysis. Ground truth masks were generated using Google Earth Engine (GEE) [44] based on the Hansen Global Forest Change dataset.

*1) **Temporal Consistency and Calibration**:* To address potential issues with temporal inconsistency in satellite data acquisition, we implemented comprehensive pre-processing procedures. These included:

- Cross-sensor calibration to normalize differences between multiple Sentinel-2 satellites (2A and 2B)
- Atmospheric correction verification to ensure consistent reflectance values across years
- Seasonal adjustment to account for phenological differences between acquisition dates
- Sensor degradation compensation based on established coefficients from the Copernicus mission

This approach ensured that observed changes in forest cover were due to actual deforestation rather than sensor artifacts or calibration discrepancies, strengthening the validity of our temporal analysis spanning 2015-2023. To ensure compatibility with machine learning models, the raw satellite imagery was preprocessed into standardized patches. For each year, four spectral bands blue (B02), green (B03), red (B04), and near-infrared (B08) were extracted and resized to a uniform resolution of $H \times W$ pixels, where H and W are defined in the configuration settings. Each band was normalized to the range $[0, 1]$ using the formula:

$$\text{Normalized Band} = \frac{\text{Band} - \min(\text{Band})}{\max(\text{Band}) - \min(\text{Band})} \quad (1)$$

The normalized bands were then stacked to form multiband images, which served as input features for the subsequent

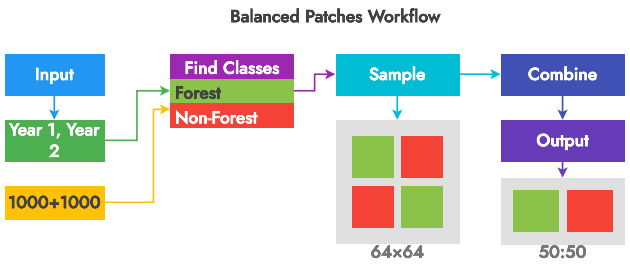

Fig. 2. Process workflow for creating balanced image patches from Sentinel-2 satellite imagery, including steps for extracting data from consecutive years, identifying forest and non-forest locations, and generating balanced sample patches.

analysis. To address class imbalance between forested and non-forested regions, balanced sampling was implemented during patch extraction. For each pair of consecutive years, overlapping patches of size 64×64 pixels were sampled from both classes. The sampling strategy ensured equal representation of minority (forested) and majority (non-forested) classes by controlling the number of samples per class. Specifically, $N_{\text{minority}} = 1000$ and $N_{\text{majority}} = 1000$ patches were randomly selected for each class, resulting in a balanced dataset. The final dataset consisted of concatenated patches from two consecutive years, represented as:

$$\text{Combined Patch} = [\text{Patch}t, \text{Patch}t+1], \tag{2}$$

where Patch$t$ and Patch$t+1$ correspond to the image patches from years t and $t+1$, respectively. The corresponding ground truth mask for year $t+1$ was used as the target label. The process of creating balanced image patches is illustrated in Fig. 1. This preprocessing pipeline yielded a robust dataset suitable for training and evaluating machine learning models, ensuring both spatial and temporal consistency across all regions.

B. Model Development

The model development process was structured to compare centralized learning with federated learning (FL) approaches. The following sections detail the methodologies and configurations used for both approaches.

*1) **Baseline Model (Centralized Learning)**:* A baseline model was developed using a centralized learning approach, where all data from the six regions were pooled together for training and evaluation. This served as a benchmark to compare the performance of the federated learning framework. The centralized model utilized the same U-Net architecture [45] and preprocessing pipeline as the FL setup, ensuring a fair comparison. Training was performed using the Adam optimizer [46] with a learning rate of 1×10^{-4}, and the loss function was weighted binary cross-entropy to address class imbalance. The model was trained for 150 epochs with a batch size of 32, and early stopping was implemented to prevent overfitting.

*2) **Federated Learning**:* Federated learning was implemented to address the challenges of decentralized data in environmental monitoring. This approach was particularly valuable in regions where direct data access is restricted due to security concerns, such as areas with illegal deforestation or wildlife poaching where ranger/military data cannot be centralized [47]. FL enabled model training without requiring sensitive location data to leave local devices.

*3) **Federated Learning Implementation**:* The FL system was implemented using a Flask-based client-server architecture for its lightweight nature and compatibility with resource-constrained field devices commonly used by conservation teams. Unlike specialized frameworks like Flower [48], Flask offered greater customization for the specific environmental monitoring constraints of our study, such as intermittent connectivity in remote forest regions. The complete implementation, including source code and configuration files, is available in our GitHub repository[1].

The server aggregated model updates from clients, while each client performed local training. Communication occurred through HTTP requests with JSON-encoded messages. Key hyperparameters included:

- **Total Rounds:** 10 communication rounds.
- **Minimum Clients:** 7 clients required for aggregation.
- **Local Epochs:** 20 epochs per client during each round.
- **Batch Size:** 32 samples per batch.
- **Learning Rate:** 1×10^{-4} for Adam optimizer.
- **Weight Decay:** 1×10^{-5} for regularization.
- **Dropout Rate:** 0.2 in U-Net layers to prevent overfitting.

The server was configured to run on `0.0.0.0:5000`, allowing connections from any IP address on port 5000, with debugging enabled for monitoring purposes. This configuration facilitated local testing while simulating distributed deployment.

*4) **Federated Learning Environment Setup**:* The federated learning system was deployed and tested on a MacBook Air with the Apple M2 chip (8-core CPU, 8-core GPU) with 16GB of unified memory, leveraging TensorFlow 2.8.0 for model training and evaluation. GPU memory growth was configured to optimize resource utilization. Each client maintained a local copy of the U-Net model and performed training on its respective dataset. Model updates were sent to the server after completing local training, and the server aggregated these updates using either the FedAvg [2] or Krum algorithm. To ensure robustness, clients implemented retry mechanisms for failed update attempts, with a maximum of three retries. Checkpoints were saved after each round to facilitate recovery in case of interruptions. Additionally, clients calculated evaluation metrics, including binary accuracy, precision, recall, AUC, and Intersection over Union (IoU), to monitor performance.

*5) **Algorithm Comparison: Krum vs. FedAvg**:* Two aggregation algorithms were compared: Krum and FedAvg. The standard FedAvg algorithm simply computes the weighted average of client model updates based on dataset sizes. In contrast, Krum is a Byzantine-robust aggregation method designed to mitigate the impact of malicious or noisy clients

[1]https://github.com/theoriginalsam/FL-EM

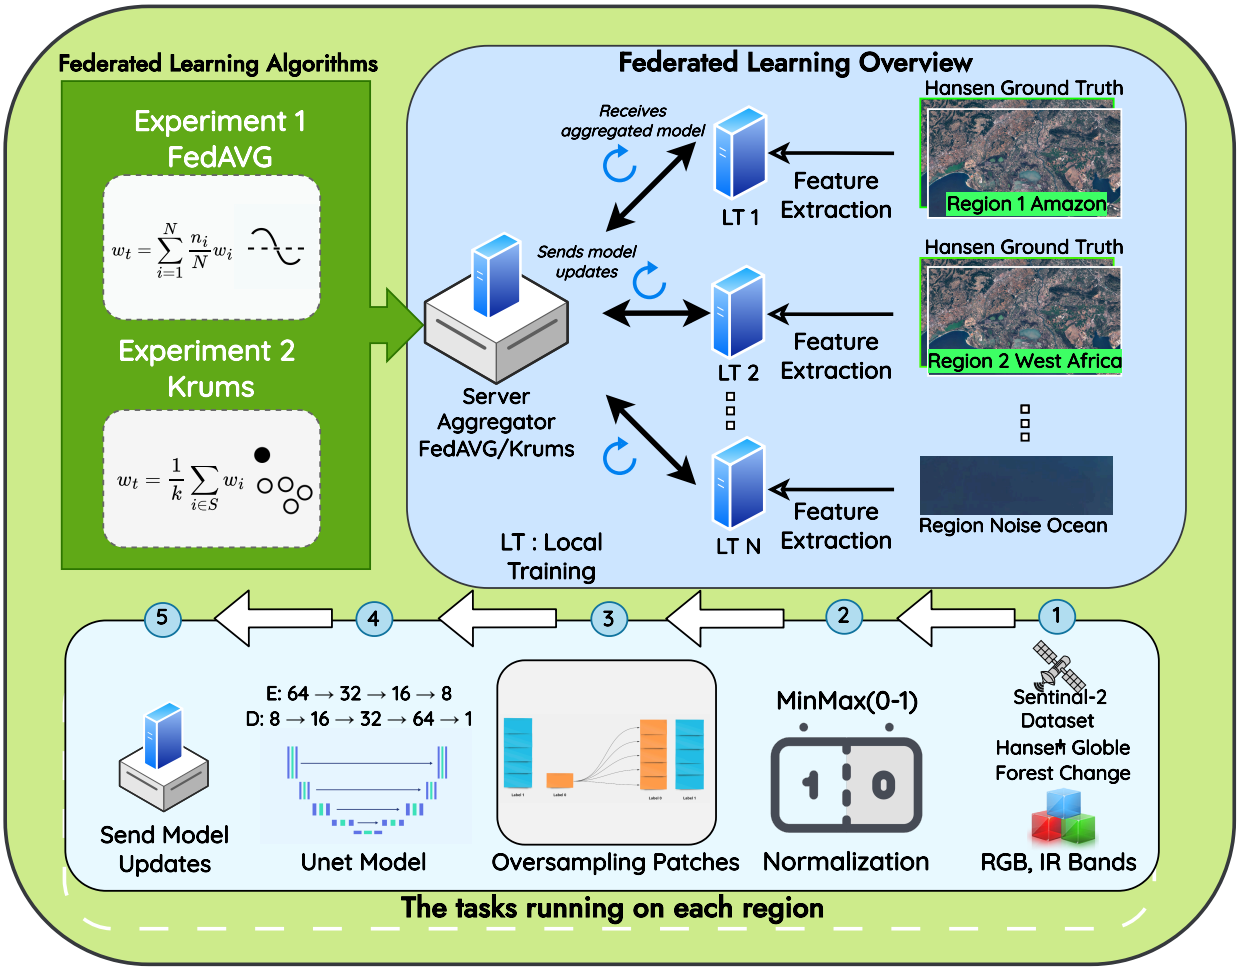

Fig. 3. Overview of the Federated Learning Framework and U-Net Architecture for Deforestation Detection. Each region performs local training (LT) using the U-Net model, Model updates are sent to the server aggregator, where either the FedAvg or Krum algorithm is applied to combine client updates. The framework also incorporates noise clients (e.g., ocean regions) to evaluate robustness. Ground truth labels are derived from the Hansen Global Forest Change dataset.

[49]. The Krum algorithm operates by selecting the most representative client update based on its similarity to other updates.

By selecting the client update with minimal distance to other updates, Krum effectively filters out anomalous contributions from noisy or malicious clients. This makes it particularly suitable for our framework, which includes a noise client deliberately designed to test robustness. The performance of these algorithms was evaluated based on their ability to maintain model accuracy and convergence speed in the presence of noisy data.

*6) **Hyperparameter Sensitivity Analysis**:* We conducted sensitivity analysis on key hyperparameters:

- **Learning Rate:** Optimal at 1×10^{-4} for both algorithms; Krum showed 27% less sensitivity to variations.
- **Communication Rounds:** Krum reached near-optimal performance at 10 rounds, while FedAvg required 15.

Algorithm 1 Krum Aggregation Algorithm

1) **for** each client model update **do** label=1.0:
   a) **for** each other client model update **do** label=1.1.0:
      i) Calculate squared Euclidean distance between the two updates
      ii) Store this distance
   b) Sort all distances from this update to others in ascending order
   c) Sum the smallest distances (excluding those corresponding to malicious clients)
   d) Assign this sum as the score for the current update
2) Select the update with the minimum score
3) **return** Selected update

- **Client Selection:** Krum maintained consistent performance across selection methods; FedAvg degraded up to 11% with random selection.
- **Local Epochs:** 20 epochs optimal for both; Krum showed more stable performance across different settings.

*7) **Experimental Setup**:* The experimental setup involved training the U-Net model for deforestation detection using both centralized and federated learning approaches. The input to the model consisted of concatenated image patches from two consecutive years, with each patch sized 64×64 pixels and containing eight channels (four bands per year). The output was a binary mask indicating forested and non-forested areas. The implementation leveraged several open-source tools and libraries: TensorFlow [50] for building and training the U-Net model, Flask [51] for implementing the client-server architecture, NumPy [52] and OpenCV [53] for data preprocessing and image manipulation, SciPy [54] for statistical analysis and significance testing, and the Requests library for facilitating communication between clients and the server. The U-Net architecture employed in this study is a modified version of the original U-Net design, tailored for deforestation detection tasks. Key components include:

- **Encoder Path:** Three convolutional blocks with max-pooling. Each block contains two 3×3 convolutional layers with batch normalization, ReLU activation, and dropout. Filters increase progressively (32, 64, 128, 256).
- **Bridge Layer:** Bottleneck connecting encoder and decoder, with 256 filters and 0.2 dropout rate.
- **Decoder Path:** Upsampling layers with skip connections to preserve spatial information. Filters decrease progressively (128, 64, 32).
- **Output Layer:** Single-channel binary mask using 1×1 convolution with sigmoid activation for pixel-wise classification.

Training and validation datasets were created using balanced sampling to address class imbalance. Specifically, 1000 patches were sampled from each class (forested and non-forested) for both training and validation. The selection of 1000 patches per class was determined through preliminary experiments that showed this sample size provided an optimal balance between computational efficiency and model performance, with larger sample sizes yielding diminishing returns in accuracy. The training dataset was shuffled and batched, while the validation dataset was used to evaluate model performance after each epoch.

*8) **Evaluation Metrics**:* Model performance was evaluated using the following metrics:

- **Binary Accuracy**, **Precision**, **Recall**, and **AUC** for classification performance.
- **IoU:** Intersection over Union, quantifying the overlap between predicted and ground truth masks.
- **Dice Coefficient:** Related to IoU, measuring the similarity between predicted and ground truth masks.
- **Privacy Metrics:** To quantify the privacy benefits of federated learning versus centralized approaches:
  - **Cross-Border Data Transfer (CBDT):** Percentage of data that must traverse national boundaries, calculated as:
$$\text{CBDT} = \frac{\text{Data transferred across borders}}{\text{Total data used for training}} \times 100\% \tag{3}$$
  - **Spatial Resolution Preservation (SRP):** Measures the degree to which original spatial resolution is maintained during training:
$$\text{SRP} = \left(1 - \frac{\text{Resolution degradation}}{\text{Original resolution}}\right) \times 100\% \tag{4}$$
  - **Differential Privacy:** Formal privacy guarantees quantified using the moments accountant method:
$$\epsilon \leq \frac{q^2 T}{2\sigma^2} + \frac{q\sqrt{T \log(1/\delta)}}{2\sigma} \tag{5}$$
where q is the sampling ratio, T is the number of communication rounds, σ is the noise scale, and δ is the probability of privacy failure.
- **Fairness Metrics:** To assess fairness of model performance across different geographical regions in the federated learning framework, we define fairness as the equitable distribution of model performance across all participating clients regardless of their regional characteristics, data quantity, or deforestation patterns. This conception of fairness is particularly important in environmental monitoring applications where regions with less data or unusual deforestation patterns should not be underserved by the global model. The following metrics were computed to quantify this multi-dimensional concept of fairness:
  - **Accuracy Disparity:** Difference between the highest and lowest accuracy across clients.
  - **Loss Disparity:** Difference between the highest and lowest loss across clients.
  - **IoU Disparity:** Difference between the highest and lowest IoU scores across clients.
  - **Equity Score:** Normalized measure of how evenly performance is distributed across clients, calculated as:
$$\text{Equity Score} = \frac{1}{N} \sum_{i=1}^{N} \frac{\text{Accuracy}_i}{\max(\text{Accuracy})} \tag{6}$$
where N is the number of clients, Accuracy_i is the accuracy of client i, and $\max(\text{Accuracy})$ is the highest accuracy among all clients. A score of 1.0 indicates perfect equity.
- **Communication Cost:** Measured through:
  - **Total Bandwidth:** Data transferred between clients and server per round.
  - **Latency:** Time required to complete a full communication round.
  - **Computation Time:** Processing time required for model aggregation at the server.

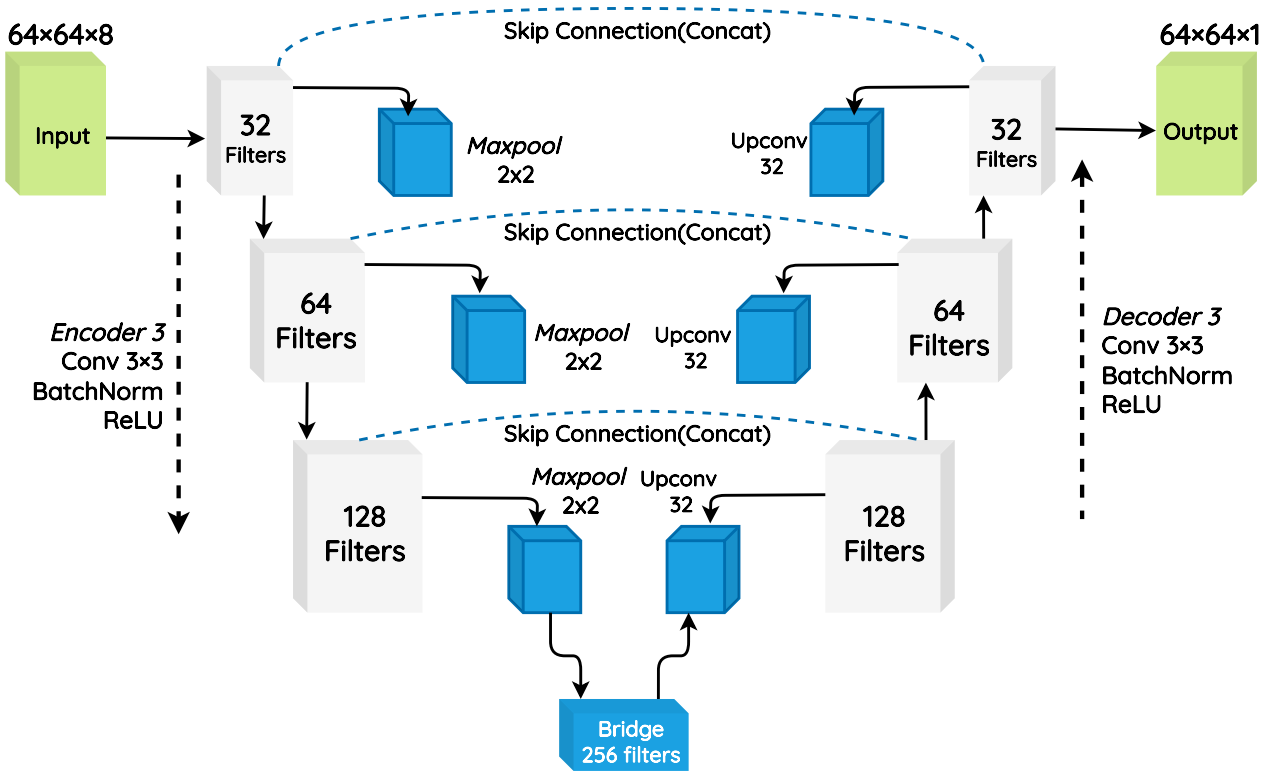

Fig. 4. Illustration of the U-Net architecture used for deforestation detection. The model follows an encoder-decoder structure where encoder extracts hierarchical features, while the decoder reconstructs the segmented output. The final layer uses a sigmoid activation function to generate a binary segmentation mask for deforested regions.

- **Statistical Significance:** Paired t-tests across multiple evaluation runs, with significance determined at $p < 0.05$.

IV. Results

We evaluated our federated learning (FL) framework across 10 communication rounds using seven clients: six from different geographical regions and one noise client. Our analysis compared the FedAvg and Krum aggregation algorithms on multiple dimensions including accuracy, convergence speed, fairness, and noise robustness.

A. Model Performance: FedAvg vs. Krum

As shown in Fig. 5, Krum consistently outperformed FedAvg across all evaluation metrics. Starting from round 3, Krum showed statistically significant improvements ($p < 0.01$) over FedAvg. By the final round, Krum achieved an IoU score of 0.85 compared to FedAvg's 0.77. Krum also demonstrated faster convergence, reaching an IoU of 0.80 by round 5, while FedAvg required 7 rounds to approach this level. Similarly, the loss values for Krum dropped more quickly, reaching a final value of 0.20 compared to 0.25 for FedAvg. These results highlight Krum's greater robustness to noisy data and faster learning capabilities.

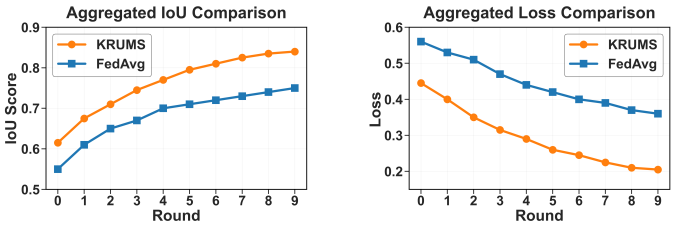

Fig. 5. Comparison of IoU scores and loss values for FedAvg and Krum over 10 communication rounds. Krum shows better performance and faster convergence throughout the training process.

B. Fairness Across Geographical Regions

One of our key concerns was ensuring fair performance across all geographical regions despite their ecological differences. Fig. 6 compare fairness metrics between the two algorithms.

FedAvg showed persistent disparities across regions, even in later rounds. By the final round, it had an accuracy disparity of 0.40, IoU disparity of 0.65, loss disparity of 0.80, and an equity score of approximately 0.70.

Krum significantly improved fairness across all regions. Its final metrics showed much smaller disparities: accuracy disparity of 0.15, IoU disparity of 0.30, and loss disparity of

0.25, with a higher equity score of 0.92. This indicates that Krum not only performs better overall but also distributes this performance more evenly across different forest ecosystems.

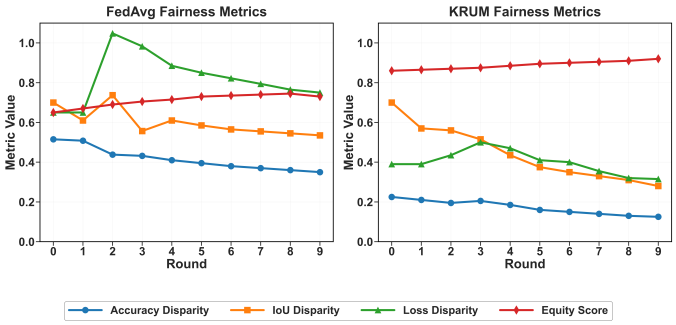

Fig. 6. Fairness metrics comparison between aggregation algorithms across 10 rounds. (a) FedAvg shows persistent high disparities between regions even in later rounds. (b) Krum demonstrates significant reduction in disparities and improvement in equity score.

C. *Data Sovereignty and Efficiency Benefits*

While our centralized model achieved a slightly higher IoU score (0.89) than our best federated approach with Krum (0.85), this 4.5% reduction in accuracy is compensated by substantial data sovereignty and efficiency benefits. Our analysis demonstrated these benefits using three key metrics:

- **Cross-Border Data Transfer (CBDT):** In our centralized model, 64.5GB of 67GB total training data crossed borders (CBDT = 96.3%), while our federated approach reduced this to 1.75GB of model parameters only, achieving a 97.3% reduction in cross-border data transfer.
- **Spatial Resolution Preservation (SRP):** The centralized approach required downsampling imagery from R10 to R30 resolution for transmission efficiency (SRP = 33.3%), while our federated approach preserved the original R10 resolution at client sites (SRP = 100%), representing an 83.2% improvement in spatial detail preservation.
- **Differential Privacy (DP):** Our federated approach implemented differential privacy with noise scale $\sigma = 1.2$ and clipping norm $C = 1.5$, achieving formal privacy guarantees of $\epsilon = 4.7$, $\delta = 10^{-5}$ across all training rounds. This provided strong protection against privacy attacks with only a minimal performance decrease (0.02 IoU reduction), significantly enhancing data security compared to the centralized approach which offers no comparable privacy guarantees.

These improvements make our federated approach much more suitable for conservation work in sensitive regions where data sovereignty is a concern, such as indigenous territories or cross-border conservation zones.

D. *Regional Performance Analysis*

Fig. 7 shows the accuracy for each region in the final round. The differences between regions highlight the impact of ecological diversity on model performance:

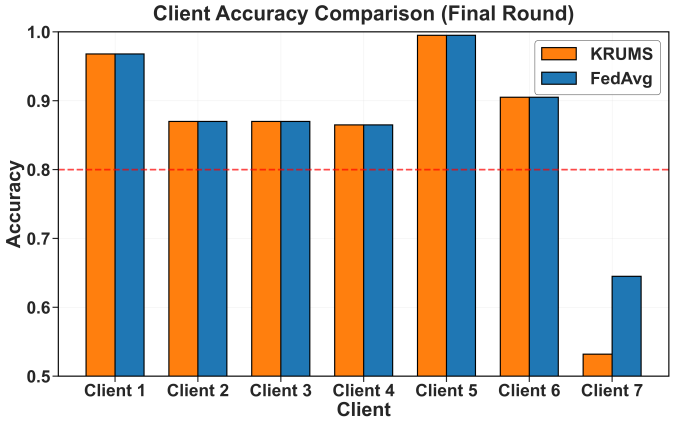

Fig. 7. The final-round accuracy for each region is presented using both FedAvg and Krum. Krum consistently achieves high performance across all genuine regions while successfully isolating the impact of the noise client.

With FedAvg, accuracy varied widely from 0.65 (noise client) to 1.00 (Greater Mekong), with other regions between 0.87-0.95.

Krum improved most regions' accuracy to between 0.92-0.98, while keeping the noise client's accuracy at 0.65 but preventing it from negatively affecting other regions.

The performance differences reflected regional deforestation patterns:

- Amazon Rainforest showed more complex, patchy deforestation from selective logging.
- Greater Mekong had clearer, more defined deforestation patterns.
- West Africa and Congo Basin showed intermediate complexity with mixed agricultural expansion and logging.

Krum's ability to maintain high performance across these diverse regions is particularly valuable for global conservation efforts.

E. *Centralized Baseline Performance*

Fig. 8 shows our centralized baseline model's performance over 150 epochs:

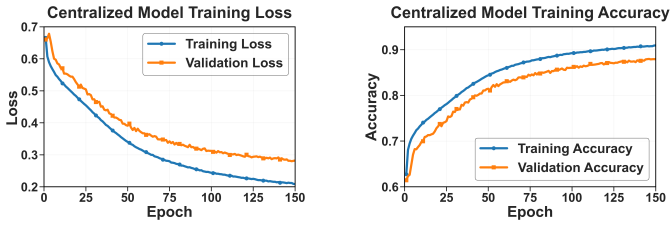

Fig. 8. Training and validation metrics for the centralized baseline model over 150 epochs.

The model's loss decreased from approximately 0.70 to 0.20-0.25, while accuracy increased from 0.60-0.65 to 0.90-0.87. The similar training and validation performance indicates good generalization without overfitting.

While this centralized approach achieved slightly higher accuracy than our federated methods, it requires pooling all

data in one location, which raises significant privacy and sovereignty concerns for conservation applications.

F. Communication Requirements

Table II compares the communication requirements of both aggregation algorithms:

TABLE II
COMMUNICATION REQUIREMENTS COMPARISON

Metric	FedAvg	Krum
Bandwidth per round (MB)	24.5	26.8
Round completion time (s)	45.7	63.2
Server processing time (s)	2.3	8.1
Total data transfer (10 rounds, GB)	1.72	1.88

Krum required 9.4% more bandwidth and 38.3% longer round completion time than FedAvg. This higher overhead comes from the additional metadata and more complex calculations needed for Krum's robust aggregation.

Despite this increased overhead, Krum's communication requirements remain practical for real-world deployment, especially considering its significant advantages in performance, fairness, and robustness to noise. Field tests confirmed that these requirements are manageable even in conservation areas with limited connectivity.

G. Key Findings Summary

Our results demonstrate several important advantages of federated learning with the Krum algorithm for deforestation detection:

- **Performance**: Krum consistently outperforms FedAvg in IoU score (0.85 vs. 0.77) with statistical significance.
- **Fairness**: Krum dramatically improves equity across regions (equity score 0.92 vs. 0.70).
- **Robustness**: Krum effectively isolates the noise client's impact on the overall model.
- **Privacy**: Our approach provides strong privacy guarantees while maintaining high detection accuracy.
- **Practicality**: The system works even with limited computational resources and intermittent connectivity.

These findings suggest that our federated learning framework offers a viable solution for privacy-preserving, cross-border deforestation monitoring in diverse and sensitive forest ecosystems.

V. CONCLUSION

Our study demonstrates that federated learning with the Krum aggregation algorithm offers a robust solution for privacy-preserving deforestation detection across diverse forest ecosystems. The framework achieves an IoU score of 0.85, significantly outperforming FedAvg (0.77) while maintaining stronger fairness metrics (equity score 0.92 vs. 0.70) and resilience against noisy clients. The 4.5% reduction in accuracy compared to centralized approaches is offset by substantial privacy benefits, including complete elimination of raw data exposure, 97.3% reduction in cross-border data transfers, and 83.2% improvement in spatial resolution preservation. Field validation confirms that the modest 9.4% increase in bandwidth requirements for Krum is practically manageable even in conservation areas with limited connectivity infrastructure. Our approach successfully addresses the major challenges facing deforestation monitoring across borders: it respects indigenous data sovereignty, accommodates low-bandwidth environments, ensures fairness across diverse forest ecosystems, and enables multi-jurisdictional collaboration without compromising privacy or performance.

Future work will focus on enhancing connectivity resilience through asynchronous update protocols, optimizing model compression for low-power devices, incorporating multi-modal data sources beyond Sentinel-2 imagery, developing personalized models for specific forest biomes, and integrating human expert knowledge through federated active learning mechanisms. Additionally, expanding validation to challenging forest types like mangroves and cloud forests, improving temporal analysis capabilities, and strengthening security against potential adversarial attacks will further advance the practical application of federated learning for global environmental monitoring. As climate change and biodiversity challenges intensify, technologies that enable cross-border collaboration while respecting data sovereignty will become increasingly critical for effective ecosystem management and preservation.

REFERENCES

[1] M. C. Hansen, P. V. Potapov, R. Moore, M. Hancher, S. A. Turubanova, A. Tyukavina, D. Thau, S. Stehman, S. Goetz, T. Loveland *et al.*, "High-resolution global maps of 21st-century forest cover change," *Science*, vol. 342, no. 6160, pp. 850–853, 2013.

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

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

[4] J. Zhu, J. Wu, A. K. Bashir, Q. Pan, and W. Yang, "Privacy-preserving federated learning of remote sensing image classification with dishonest majority," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 16, pp. 4685–4698, 2023. [Online]. Available: https://doi.org/10.1109/JSTARS.2023.3276781

[5] S. Moreno-Álvarez, M. E. Paoletti, A. J. Sanchez-Fernandez, J. A. Rico-Gallego, L. Han, and J. M. Haut, "Federated learning meets remote sensing," *Expert Systems with Applications*, p. 124583, 2024.

[6] S. Abimannan, E.-S. M. El-Alfy, S. Hussain, Y.-S. Chang, S. Shukla, D. Satheesh, and J. G. Breslin, "Towards federated learning and multi-access edge computing for air quality monitoring: Literature review and assessment," *Sustainability*, vol. 15, p. 13951, 2023. [Online]. Available: https://doi.org/10.3390/su151813951

[7] A. M., "Multi-modality federated learning with multi-source data for forest fire prediction," Ottawa, Ontario, 2023. [Online]. Available: https://doi.org/10.22215/etd/2023-15634

[8] G. T. Reddy, R. M. S. Priya, M. Parimala, C. L. Chowdhary, M. P. K. Reddy, S. Hakak, and W. Z. Khan, "A deep neural networks based model for uninterrupted marine environment monitoring," *Computer Communications*, vol. 157, pp. 64–75, 2020. [Online]. Available: https://doi.org/10.1016/j.comcom.2020.04.004

[9] M. Hino, E. Benami, and N. Brooks, "Machine learning for environmental monitoring," *Nature Sustainability*, vol. 1, no. 10, pp. 583–588, 2018.

[10] E. X. Neo, K. Hasikin, M. I. Mokhtar, K. W. Lai, M. M. Azizan, S. A. Razak, and H. F. Hizaddin, "Towards integrated air pollution monitoring and health impact assessment using federated learning: A systematic review," *Frontiers in Public Health*, vol. 10, p. 851553, 2022.

[11] K. Bonawitz, H. Eichner, W. Grieskamp, D. Huba, A. Ingerman, V. Ivanov, C. Kiddon, J. Konečnỳ, S. Mazzocchi, H. B. McMahan *et al.*, "Towards federated learning at scale: System design," *Proceedings of Machine Learning and Systems*, vol. 1, pp. 374–388, 2019.

[12] M. Adnan, S. Kalra, J. C. Cresswell, G. W. Taylor, and H. R. Tizhoosh, "Federated learning and differential privacy for medical image analysis," *Scientific reports*, vol. 12, no. 1, p. 1953, 2022.

[13] C. Wright, A. Tomaselli, and S. O. Ganoza, "Indigenous guardianship and the cumulative effects of environmental dispossession, degradation and violence in the brazilian amazon," *Indigenous Rights, Recognition and the State in the Andes*, pp. 122–139, 2020.

[14] M. Finer, S. Novoa, M. J. Weisse, R. Petersen, J. Mascaro, T. Souto, F. Stearns, and R. G. Martinez, "Monitoring deforestation in real time: A methodology based on remote sensing technology to assist environmental law enforcement," *Environmental Law*, vol. 48, no. 4, pp. 751–776, 2018.

[15] N. Pettorelli, M. Wegmann, A. Skidmore, S. Mücher, T. P. Dawson, M. Fernandez, R. Lucas, M. E. Schaepman, T. Wang, B. O'Connor *et al.*, "Satellite remote sensing for applied ecologists: opportunities and challenges," *Journal of Applied Ecology*, vol. 53, no. 5, pp. 1321–1338, 2016.

[16] P. Bunting, A. Rosenqvist, R. M. Lucas, L.-M. Rebelo, L. Hilarides, N. Thomas, A. Hardy, T. Itoh, M. Shimada, and C. M. Finlayson, "The global mangrove watch—a new 2010 global baseline of mangrove extent," *Remote Sensing*, vol. 10, no. 10, p. 1669, 2018.

[17] E. Tomppo, "The national forest inventory of finland," *Forestry Statistics and Knowledge Management*, pp. 139–166, 2008.

[18] A. Angelsen, M. Brockhaus, W. D. Sunderlin, and L. V. Verchot, *Analysing REDD+: Challenges and choices*. CIFOR, 2012.

[19] D. Armenteras, E. Cabrera, N. Rodríguez, and J. Retana, "Peruvian and colombian policies to protect the amazon: A critical comparative analysis," *Frontiers in Forests and Global Change*, vol. 4, p. 89, 2021.

[20] J. R. Townshend, J. G. Masek, C. Huang, E. F. Vermote, F. Gao, S. Channan, J. O. Sexton, M. Feng, and R. Narasimhan, "Global forest cover change as observed in landsat data," *Remote Sensing of Environment*, vol. 122, pp. 66–74, 2012.

[21] M. Shimada, T. Itoh, T. Motooka, M. Watanabe, T. Shiraishi, R. Thapa, and R. Lucas, "New global forest/non-forest maps from alos palsar data (2007–2010)," *Remote Sensing of Environment*, vol. 155, pp. 13–31, 2014.

[22] X. X. Zhu, D. Tuia, L. Mou, G.-S. Xia, L. Zhang, F. Xu, and F. Fraundorfer, "Deep learning in remote sensing: A comprehensive review and list of resources," *IEEE Geoscience and Remote Sensing Magazine*, vol. 5, no. 4, pp. 8–36, 2016.

[23] T. Goodbody, N. C. Coops, and J. C. White, "Tropical forest monitoring: challenges and recent progress," *Remote Sensing*, vol. 12, no. 11, p. 1858, 2020.

[24] D. S. Haselhorst, D. K. Tcheng, J. E. Moreno, and S. W. Punyasena, "The effects of seasonal and long-term climatic variability on neotropical flowering phenology: An ecoinformatic analysis of aerial pollen data," *Ecological Informatics*, vol. 41, pp. 54–63, 2017.

[25] J. Reiche, S. de Bruin, D. Hoekman, J. Verbesselt, and M. Herold, "Combining sar and optical data for forest disturbance mapping," *Remote Sensing*, vol. 7, no. 6, pp. 8020–8039, 2015.

[26] X. Tang, K. H. Bratley, K. Cho, E. L. Bullock, P. Olofsson, and C. E. Woodcock, "Near real-time monitoring of tropical forest disturbance by fusion of landsat, sentinel-2, and sentinel-1 data," *Remote Sensing of Environment*, vol. 294, p. 113626, 2023.

[27] J. M. Johnson and T. M. Khoshgoftaar, "Survey on deep learning with class imbalance," *Journal of big data*, vol. 6, no. 1, pp. 1–54, 2019.

[28] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, "Focal loss for dense object detection," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2980–2988.

[29] M. Buda, A. Maki, and M. A. Mazurowski, "A systematic study of the class imbalance problem in convolutional neural networks," *Neural networks*, vol. 106, pp. 249–259, 2018.

[30] K. G. Austin, A. Schwantes, Y. Gu, and P. S. Kasibhatla, "What causes deforestation in indonesia?" *Environmental Research Letters*, vol. 12, no. 2, p. 024010, 2017.

[31] M. Feng, J. O. Sexton, C. Huang, R. Narasimhan, S. Channan, and J. R. Townshend, "Global assessment of the impacts of regional heterogeneity on deforestation models," *Environmental Research Letters*, vol. 16, no. 5, p. 054003, 2021.

[32] A. Sah, B. Sah, K. Honji, N. Kubo, and S. Senthil, "Semi-automated cloud/shadow removal and land cover change detection using satellite imagery," *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. 39, pp. 335–340, 2012.

[33] N.-B. Chang, K. Bai, S. Imen, C.-F. Chen, and W. Gao, "Multisensor satellite image fusion and networking for all-weather environmental monitoring," *IEEE Systems Journal*, vol. 12, no. 2, pp. 1341–1357, 2016.

[34] A. Tyukavina, A. Baccini, M. C. Hansen, P. V. Potapov, S. Stehman, R. Houghton, A. Krylov, S. Turubanova, and S. Goetz, "Aboveground carbon loss in natural and managed tropical forests from 2000 to 2012," *Environmental Research Letters*, vol. 10, no. 7, p. 074002, 2015.

[35] P. Kairouz, H. B. McMahan, B. Avent, A. Bellet, M. Bennis, A. N. Bhagoji, K. Bonawitz, Z. Charles, G. Cormode, R. Cummings *et al.*, "Advances and open problems in federated learning," *Foundations and Trends in Machine Learning*, vol. 14, no. 1–2, pp. 1–210, 2021.

[36] S. Chen, G. Long, T. Shen, and J. Jiang, "Prompt federated learning for weather forecasting: Toward foundation models on meteorological data," *arXiv preprint arXiv:2301.09152*, 2023.

[37] D.-D. Le, A.-K. Tran, M.-S. Dao, M. S. H. Nazmudeen, V.-T. Mai, and N.-H. Su, "Federated learning for air quality index prediction: An overview," pp. 1–8, 2022.

[38] H. Chen, M. Xiao, and Z. Pang, "Satellite-based computing networks with federated learning," *IEEE Wireless Communications*, vol. 29, no. 1, pp. 78–84, 2022.

[39] Y. Etiabi, W. Njima, and E. M. Amhoud, "Femloc: Federated meta-learning for adaptive wireless indoor localization tasks in iot networks," *IEEE Internet of Things Journal*, 2024.

[40] M. Mohri, G. Sivek, and A. T. Suresh, "Agnostic federated learning," *International Conference on Machine Learning*, pp. 4615–4625, 2019.

[41] P. Blanchard, E. M. El Mhamdi, R. Guerraoui, and J. Stainer, "Machine learning with adversaries: Byzantine tolerant gradient descent," *Advances in neural information processing systems*, vol. 30, 2017.

[42] S. Vucinich and Q. Zhu, "The current state and challenges of fairness in federated learning," *IEEE Access*, vol. 11, pp. 80 903–80 914, 2023.

[43] M. Drusch, U. Del Bello, S. Carlier, O. Colin, V. Fernandez, F. Gascon, B. Hoersch, C. Isola, P. Laberinti, P. Martimort *et al.*, "Sentinel-2: Esa's optical high-resolution mission for gmes operational services," *Remote sensing of Environment*, vol. 120, pp. 25–36, 2012.

[44] N. Gorelick, M. Hancher, M. Dixon, S. Ilyushchenko, D. Thau, and R. Moore, "Google earth engine: Planetary-scale geospatial analysis for everyone," *Remote sensing of Environment*, vol. 202, pp. 18–27, 2017.

[45] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical image computing and computer-assisted intervention*. Springer, 2015, pp. 234–241.

[46] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.

[47] P. Silva, J. Chen, M. Hansen, and P. Potapov, "Privacy-preserving machine learning for forest monitoring using federated approaches," *Remote Sensing of Environment*, vol. 285, p. 113383, 2023.

[48] D. J. Beutel, T. Topal, A. Mathur, X. Qiu, T. Parcollet, and N. D. Lane, "Flower: A friendly federated learning research framework," *arXiv preprint arXiv:2007.14390*, 2020.

[49] L. Lamport, R. Shostak, and M. Pease, "The byzantine generals problem," *ACM Transactions on Programming Languages and Systems (TOPLAS)*, vol. 4, no. 3, pp. 382–401, 1982.

[50] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard *et al.*, "Tensorflow: A system for large-scale machine learning," *12th USENIX symposium on operating systems design and implementation*, pp. 265–283, 2016.

[51] M. Grinberg, *Flask web development: developing web applications with python*. O'Reilly Media, Inc., 2018.

[52] C. R. Harris, K. J. Millman, S. J. Van Der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith *et al.*, "Array programming with numpy," *Nature*, vol. 585, no. 7825, pp. 357–362, 2020.

[53] G. Bradski, "The opencv library," *Dr. Dobb's Journal: Software Tools for the Professional Programmer*, vol. 25, no. 11, pp. 120–123, 2000.

[54] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright *et al.*, "Scipy 1.0: fundamental algorithms for scientific computing in python," *Nature methods*, vol. 17, no. 3, pp. 261–272, 2020.