Date of publication February 12, 2025, date of current version February 12, 2025.

Digital Object Identifier 10.1109/ACCESS.2017.DOI

Federated Few-Shot Continuous Learning for Medical Image Processing with Federated Feature Augmentation

BYARUHANGA JOSHUA[1], MUHINDO MUBARAKA[1], GGALIWANGO MARVIN[2], MARRIETTE KATARAHWEIRE[3], JOHN KIZITO[4], JOYCE NAKATUMBA NABENDE[5]

[1, 2, 3, 4, 5]Department of Computer Science, Makerere University, Kampala, Uganda
Emails: {muhindo.mubaraka, joshua.byaruhanga}@students.mak.ac.ug
[2]Department of Computer Science, School of Computing and Informatics Technology, College of Computing and Information Sciences, Makerere University, Kampala, Uganda
Email: ggaliwango.marvin@mak.ac.ug, https://orcid.org/0000-0002-8635-3936
[3]Department of Computer Science, School of Computing and Informatics Technology, College of Computing and Information Sciences, Makerere University, Kampala, Uganda
Email: kmarriette@cit.ac.ug, https://orcid.org/0000-0003-4718-4209
[4]Department of Computer Science, School of Computing and Informatics Technology, College of Computing and Information Sciences, Makerere University, Kampala, Uganda
Email: jkizito@cis.mak.ac.ug, https://orcid.org/0009-0002-4146-2917
[5]Department of Computer Science, School of Computing and Informatics Technology, College of Computing and Information Sciences, Makerere University, Kampala, Uganda
Email: joyce.nabende@mak.ac.ug, https://orcid.org/0000-0002-0108-3798

Corresponding author: Ggaliwango Marvin (e-mail: ggaliwango.marvin@mak.ac.ug).

Acknowledgment: This work has been supported by COCIS RISE, MakRIF and Makerere UNIPOD.

ABSTRACT Developing robust models for medical image classification remains a critical challenge due to limited annotated data, patient privacy regulations, and the requirement for continuous adaptation in dynamic clinical environments. This study proposed an integrated learning system that combined Federated Learning, Few-Shot Learning, Federated Feature Augmentation, and Continuous Learning to address these constraints in malaria diagnosis. The federated setup enabled distributed healthcare institutions to collaboratively train deep learning models without centralizing patient-level data. The Few-Shot Learning module improved adaptability to localized data scarcity, while Federated Feature Augmentation facilitated cross-node generalization through anonymized latent feature exchange. A Continuous Learning mechanism allowed the system to incorporate evolving diagnostic patterns without performance degradation over time. The system was evaluated using the Thick Blood Smears dataset in a decentralized training environment. Quantitative results demonstrated that the Xception model achieved the highest overall performance, with 95.16% accuracy, 97.13% precision, 94.09% recall, 95.58% F1-score, and 98.02% area under the receiver operating characteristic curve. Additional experiments showed that the convolutional neural network model improved consistently across ten rounds of federated training, peaking at 90.92% accuracy and 95.29% AUC. Model explainability was ensured through post-hoc techniques including Gradient-weighted Class Activation Mapping and Local Interpretable Model-Agnostic Explanations, both of which confirmed that predictions were based on clinically relevant morphological features. The complete system was implemented using TensorFlow, OpenCV, Flower, Django, and Flutter to support federated orchestration, real-time mobile diagnostics, and secure deployment.

The results validated the effectiveness of the proposed system in achieving high diagnostic accuracy, privacy preservation, and interpretability under data-constrained conditions. These findings highlight the system's applicability for real-world deployment in decentralized healthcare settings such as those in Uganda.

INDEX TERMS Federated learning, Few-shot learning, Feature enhancement, Continuous learning, Malaria diagnosis, Medical imaging, AI, Uganda.

I. INTRODUCTION

Machine learning has demonstrated significant potential in accelerating and improving diagnostic accuracy in healthcare. However, real-world deployment in clinical settings remains hindered by three primary constraints: stringent data privacy regulations, scarcity of labeled medical images, and the need for continuous adaptation to evolving disease patterns [1], [6].

Federated Learning (FL) has emerged as a privacy-preserving paradigm wherein local models are trained on-site using institutional data, and only encrypted model updates are transmitted to a central aggregator [5]. This approach prevents exposure of raw patient data. Nevertheless, many small healthcare facilities possess only limited annotated samples, resulting in weak local models and insufficient generalization.

Few-Shot Learning (FSL) addresses data scarcity by enabling models to learn robust representations from a small number of labeled instances [2]. When integrated with FL, FSL allows geographically dispersed nodes with minimal data to collaboratively construct a performant global model. Despite this, diagnostic performance may degrade over time due to changes in imaging devices, pathogen variants, or environmental factors.

Continuous Learning (CL) facilitates model adaptation to new data distributions by incrementally updating parameters without forgetting previously acquired knowledge [3]. To enhance collaborative learning without violating privacy, this study introduces Federated Feature Augmentation (FFA), wherein participating nodes exchange only anonymized intermediate feature vectors rather than raw input data.

The proposed pipeline integrates FL, FSL, CL, and FFA to form a unified system for privacy-conscious, data-efficient, and adaptive medical image classification. The pipeline was evaluated using the Thick blood Smears dataset for malaria diagnosis. TensorFlow and OpenCV were employed for image processing and model development, while federated orchestration was handled using the Flower framework. The system includes a Django-based web interface and a Flutter-based mobile application for clinical usability.

The federated workflow is illustrated in the block diagram in Figure 1, which depicts the local training and global aggregation processes without raw data exchange. This architecture demonstrates applicability beyond malaria, offering a scalable and privacy-compliant solution for medical image analysis in resource-constrained environments.

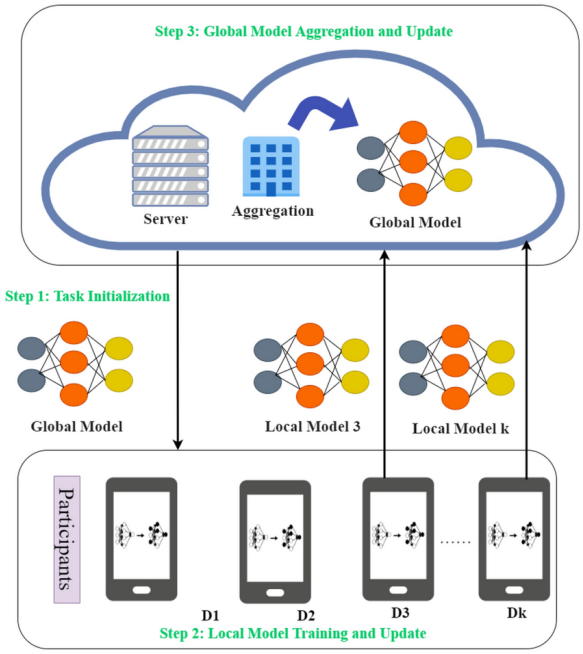


FIGURE 1: Block diagram of the Federated Learning workflow. Each institution trains a local model on its private data and transmits model updates for secure global aggregation.

II. RELATED WORK

Federated Learning (FL) has emerged as a decentralized paradigm for training machine learning models while preserving data privacy. The foundational work in [1] introduced the Federated Averaging (FedAvg) algorithm, which enables the aggregation of locally computed model updates from distributed clients. This contribution laid the groundwork for a new research domain that addresses challenges such as non-independent and identically distributed (non-IID) data, system heterogeneity, and communication efficiency in federated environments.

Subsequent research extended these foundations through algorithmic innovations designed to optimize convergence and adapt to domain-specific constraints. Notably, [5] reviewed a range of FL optimizers, including Federated Adam (FedAdam) and Federated Yogi (FedYogi), which were shown to improve convergence rates and performance under imbalanced and sparse data conditions. These methods are particularly relevant for healthcare deployments in data-constrained regions such as Uganda, where communication bandwidth is limited and privacy requirements are stringent. Privacy-preserving strategies such as secure multi-party computation and differential privacy have been emphasized as essential for safeguarding sensitive medical and financial data.

Few-Shot Learning (FSL) offers a mechanism for enabling rapid generalization from limited labeled samples. The Model-Agnostic Meta-Learning (MAML) approach introduced in [2] demonstrated the feasibility of initializing model parameters for fast adaptation across diverse tasks. The intersection of FL and FSL—referred to as Federated Few-Shot Learning (FFSL)—has been identified as a viable solution for medical imaging domains, where high annotation costs and privacy constraints hinder large-scale supervised learning. A survey presented in [6] evaluated FL across vision tasks, including medical imaging and biometric systems, highlighting the significance of addressing statistical heterogeneity and ensuring domain generalization in distributed training.

Continuous Learning (CL), also known as lifelong or incremental learning, mitigates the problem of catastrophic forgetting, where neural models lose prior knowledge when trained on sequential tasks. As surveyed in [3], techniques such as regularization-based, rehearsal-based, and dynamic architecture-based CL have shown promise in dynamic environments. The application of CL within federated systems is gaining attention in medical imaging, where continuous data acquisition and disease variation require adaptive diagnostic models.

Federated Feature Augmentation (FFA) represents a recent advancement in privacy-preserving knowledge sharing. Rather than transmitting raw images, FFA allows participating nodes to exchange intermediate latent representations to improve global model generalization. Studies in [5], [6] demonstrated the utility of FFA in enhancing diagnostic accuracy under strict data privacy regulations by aggregating enriched feature vectors instead of sensitive raw inputs.

These advances have direct implications for healthcare systems in Uganda, where regulatory frameworks limit data mobility and healthcare facilities often lack sufficient labeled datasets. While theoretical contributions to FL, FSL, CL, and FFA are well-documented, empirical validation in low-resource settings remains limited. The integration of these methodologies into a unified pipeline tailored for malaria diagnosis addresses this critical gap, providing both methodological and deployment insights.

Further investigations are warranted to explore robust mechanisms for handling highly skewed and non-IID distributions, improving feature fusion strategies, and ensuring real-world system resilience. Enhanced integration with blockchain technologies, differential privacy mechanisms, and secure aggregation protocols could further bolster the trustworthiness and scalability of FL systems in healthcare.

III. PROBLEM STATEMENT

The deployment of advanced artificial intelligence methodologies—namely Federated Learning (FL), Few-Shot Learning (FSL), Federated Feature Augmentation (FFA), and Continuous Learning (CL)—presents fundamental challenges when applied in privacy-sensitive and resource-constrained settings such as Uganda's healthcare sector. Conventional centralized machine learning approaches depend on the availability of large, curated datasets aggregated at a central repository. This model architecture inherently introduces privacy and compliance risks, especially when sensitive patient data from geographically dispersed institutions must be consolidated across organizational or national boundaries.

Federated Learning partially resolves this issue by enabling decentralized training across multiple institutions, allowing each node to retain its data locally while contributing to a global model via model update sharing [1], [5]. However, the heterogeneity of data distributions across medical institutions introduces non-independent and identically distributed (non-IID) data, which significantly degrades model convergence and accuracy. Moreover, many healthcare centers—especially in rural and under-resourced regions—possess limited labeled data due to prohibitive annotation costs, inconsistent diagnostic protocols, and logistical challenges, impeding the utility of standard supervised learning methods.

Few-Shot Learning offers a potential solution by enabling models to generalize from a limited number of labeled examples [2]. However, integrating FSL into federated systems remains a non-trivial task, primarily due to inconsistencies between local task adaptation and global model synchronization. Effective alignment of FSL objectives with FL aggregation protocols remains underexplored in current literature, especially under real-world constraints of bandwidth, computation, and supervision.

A further complication arises from the phenomenon of catastrophic forgetting, wherein a model trained incrementally loses its performance on prior tasks after learning from new data. This effect is especially problematic in medical image analysis, where disease manifestations, imaging modalities, and diagnostic techniques evolve over time. Without mechanisms for knowledge retention, federated systems risk model degradation in clinical practice. Continuous Learning methods aim to preserve previously acquired knowledge while incrementally adapting to new data streams [3]. However, their integration into federated pipelines with constrained resources and non-IID dynamics remains a largely unresolved challenge.

In addition, standard FL architectures often fail to exploit the benefits of Federated Feature Augmentation, a technique that involves the secure aggregation of anonymized latent representations across clients to improve model robustness without compromising privacy [5], [6]. The absence of such augmentation mechanisms limits generalization, particularly in cases where raw input variability is high and data sparsity is prevalent.

The core challenge addressed in this study is the design and implementation of a unified federated pipeline that integrates FL, FSL, FFA, and CL to achieve privacy preservation, continual adaptation, and performance robustness under limited labeled data conditions. This problem is particularly critical for enabling real-time, scalable AI deployment within decentralized medical infrastructures in Uganda. Resolving these challenges would constitute a substantive advancement toward secure and practical AI for healthcare delivery in environments marked by stringent privacy regulations, annotation scarcity, and evolving disease dynamics [1]–[3], [5], [6].

IV. PROPOSED SOLUTION

The proposed solution integrates Federated Learning (FL), Few-Shot Learning (FSL), Federated Feature Augmentation (FFA), and Continuous Learning (CL) to address critical challenges associated with data privacy, annotation scarcity, and evolving disease dynamics in Uganda's healthcare environment. This unified architecture enables collaborative model training without exposing sensitive medical data,

while ensuring adaptability, generalizability, and continual improvement.

A. FEDERATED LEARNING (FL)

Federated Learning constitutes the foundational layer of the proposed pipeline. Each healthcare institution trains a local model on its private dataset and transmits only the resulting model parameters to a centralized aggregator. This design ensures strict compliance with Uganda's data protection and governance policies by eliminating the need to share raw medical images [1]. The aggregation of local updates into a global model facilitates knowledge sharing across institutions while minimizing data leakage risks.

B. FEW-SHOT LEARNING (FSL)

Due to the high cost of expert annotation and the limited availability of labeled samples in remote healthcare centers, Few-Shot Learning was incorporated to enhance generalization from minimal supervision. Using a Model-Agnostic Meta-Learning (MAML) approach [2], local models were initialized with a meta-learned prior, enabling rapid task adaptation using only a few annotated instances. This design allows the federated system to remain effective even when individual clients possess sparse or highly imbalanced datasets.

C. FEDERATED FEATURE AUGMENTATION (FFA)

To mitigate the effects of data heterogeneity across clinical environments, Federated Feature Augmentation was employed. In this approach, each client extracts intermediate feature representations from local data and shares anonymized embeddings with the central server. These embeddings are then used to enrich the global feature space without compromising patient confidentiality. This mechanism significantly improved model robustness and generalization across clients with distinct imaging protocols and patient demographics [5], [6].

D. CONTINUOUS LEARNING (CL)

Healthcare data distributions are subject to constant drift due to demographic shifts, seasonal disease trends, and emerging pathogens. To prevent performance degradation and catastrophic forgetting, the federated system was augmented with a Continuous Learning mechanism. Local models were updated incrementally using new incoming data, and knowledge retention techniques were applied during aggregation to preserve past diagnostic performance. This feature is particularly important for rural clinics where disease presentations evolve frequently and new imaging modalities may be adopted over time [3].

E. PRIVACY-PRESERVING MECHANISMS

Throughout the model lifecycle, privacy-preserving mechanisms such as secure multi-party computation and differential privacy were applied to protect against reverse engineering of sensitive features. These measures ensured compliance with Uganda's regulatory framework for healthcare data protection and minimized the likelihood of re-identification attacks.

F. TECHNICAL IMPLEMENTATION AND DEPLOYMENT

The backend infrastructure was implemented using Python-based libraries, with TensorFlow serving as the core framework for model training and deployment. The federated orchestration was managed using the Flower framework, which coordinated the training rounds and handled client-server communication. A Django-based RESTful API was developed for model monitoring, deployment, and feedback integration. For field deployment, a lightweight Flutter mobile application integrated with TensorFlow Lite was used to support real-time, on-device inference, even in bandwidth-constrained or offline environments.

G. EXPECTED IMPACTS AND BENEFITS

The unified pipeline is expected to provide multiple benefits to Uganda's healthcare system. By preserving patient privacy through decentralized learning and differential privacy, the solution ensures compliance with national data governance laws. The use of Few-Shot Learning allows for rapid model adaptation in facilities with limited annotation capacity. FFA contributes to enhanced generalization across diverse healthcare institutions, while the CL mechanism guarantees continuous model refinement in response to new diagnostic trends. Furthermore, the lightweight architecture supports scalable deployment across rural clinics, offering both mobile and web-based access to AI-powered diagnostics. This solution demonstrates the viability of deploying privacy-compliant, adaptive AI in real-world, resource-constrained clinical settings [1]–[3], [5], [6].

V. UGANDAN CONTEXT

The healthcare environment in Uganda presents a unique combination of infrastructural constraints, regulatory demands, and data limitations that influence the deployment of privacy-preserving machine learning technologies such as Federated Learning (FL), Few-Shot Learning (FSL), Federated Feature Augmentation (FFA), and Continuous Learning (CL). The success of such technologies in Uganda requires careful alignment with local realities, ranging from connectivity and equipment constraints to national data governance frameworks and health worker capacity.

A. HEALTHCARE INFRASTRUCTURE AND RESOURCE CONSTRAINTS

Uganda's healthcare infrastructure, particularly in rural and peri-urban regions, is characterized by limited access to computing resources and unstable internet connectivity. Many clinics operate with basic hardware and lack reliable power or high-performance processing capabilities. These constraints render centralized deep learning solutions impractical for real-time diagnostics. The decentralized nature of FL enables training on low-resource devices while avoiding large-scale

data transfers, making it a practical solution in settings where computational and connectivity limitations are prevalent.

B. DATA PRIVACY AND REGULATORY ENVIRONMENT

The enforcement of the Data Protection and Privacy Act (2019) has created a stringent legal landscape for patient data handling in Uganda. Conventional centralized machine learning architectures, which require central aggregation of sensitive patient data, conflict with these legal requirements. FL directly addresses this regulatory challenge by ensuring that patient data remain local to the source institution, transmitting only model updates during training. This approach aligns with national privacy laws and contributes to increased institutional trust in AI-driven health technologies.

C. MEDICAL DATA ACQUISITION AND QUALITY ISSUES

The scarcity of high-quality, annotated medical datasets in Uganda stems from a shortage of diagnostic tools, trained personnel, and annotation resources. In particular, rural clinics often face difficulties in acquiring labeled data at sufficient scale and quality for training traditional supervised learning models. Integrating FSL into the federated architecture enables the system to adapt effectively to clinical tasks using only a few labeled samples [2]. This design ensures equitable access to AI-powered diagnostics, even in facilities with limited annotation capacity.

D. CONNECTIVITY AND COMMUNICATION LIMITATIONS

Intermittent connectivity and limited bandwidth are persistent challenges in many Ugandan health facilities. While FL requires periodic communication between clients and a central server for model synchronization, such constraints can hinder effective deployment. Solutions such as communication-efficient protocols, model update compression, and asynchronous aggregation methods [1] are essential to ensure that federated training remains feasible in low-bandwidth and delay-prone environments.

E. ADOPTION AND ACCEPTANCE BY HEALTHCARE PROVIDERS

The successful integration of federated continuous learning systems in Uganda depends heavily on the acceptance and engagement of frontline healthcare professionals. Many clinicians, especially in rural settings, operate under high workloads and have limited exposure to advanced AI systems [**?**]. To encourage adoption, it is essential to develop intuitive user interfaces, provide structured onboarding experiences, and deliver context-appropriate training programs. Simplified workflows and real-time decision support tools must be designed with end-user accessibility as a core requirement.

F. CAPACITY BUILDING AND EDUCATIONAL SUPPORT

Sustained deployment and evolution of the proposed AI framework require local capacity-building initiatives. Collaborations with academic institutions such as Makerere University offer opportunities to embed AI training into medical and technical education pipelines. Tailored workshops, short courses, and certification programs developed by institutions such as the College of Computing and Information Sciences (CoCIS) can equip healthcare professionals with the necessary skills to understand, operate, and eventually improve AI-enabled diagnostic platforms, ensuring local ownership and sustainability.

G. POTENTIAL APPLICATIONS AND IMPACT

The combined use of FL, FSL, and CL has applications that extend beyond malaria diagnosis. The framework may be adapted to support diagnostics for tuberculosis, maternal health monitoring, sickle cell screening, and other public health priorities in Uganda. These diseases require localized knowledge, timely adaptation, and strict privacy compliance—attributes well addressed by decentralized AI systems. The capacity to deliver accurate predictions on-device, with minimal data sharing, positions this approach as a transformative enabler for broader healthcare accessibility and equity across Uganda.

The integration of FL, FSL, FFA, and CL into Uganda's healthcare ecosystem addresses both systemic constraints and regulatory mandates. By offering a privacy-preserving, adaptive, and scalable diagnostic platform, this solution enhances the ability of remote and underserved clinics to access advanced AI tools. The alignment of this technology with Uganda's infrastructural and legal landscape supports its long-term viability, while targeted investments in education, usability, and bandwidth optimization can accelerate adoption and maximize public health impact.

VI. METHODOLOGY AND SYSTEM WORKFLOW

A. DATASET DESCRIPTION

1) Dataset Source

The dataset used in this study was derived from a publicly available malaria dataset containing microscopic images of thin blood smear slides. It was curated to support binary classification of red blood cells based on infection status. Images were grouped into two categories: *Parasitized*, indicating the presence of Plasmodium parasites, and *Uninfected*, representing healthy red blood cells without any visible infection. All images were captured using light microscopy and saved in standard image formats.

2) Dataset Characteristics

Each image was resized to 128×128 pixels and standardized to three RGB channels. This preprocessing ensured dimensional consistency across the dataset and compatibility with convolutional neural network (CNN) input requirements. Each image was assigned a binary label: '1' for parasitized and '0' for uninfected. The dataset was balanced, containing an approximately equal number of samples for both classes, which helped mitigate bias during model training and evaluation.

3) Visual Overview of Sampled Dataset

Figure 2 presents a grid of representative samples from the dataset. The top row displays parasitized cells (Label: 1), where intracellular purple-stained structures—the malaria parasites—can be observed. The bottom row shows uninfected red blood cells (Label: 0), which exhibit more uniform coloration and lack internal artifacts.

These visual distinctions, although often subtle, provide critical morphological cues for model learning. The black padding around some images is an artifact of preprocessing (e.g., resizing and augmentation), and does not interfere with semantic content. The dataset's clear class separability, coupled with balanced representation, makes it suitable for training binary classification models targeting malaria detection.

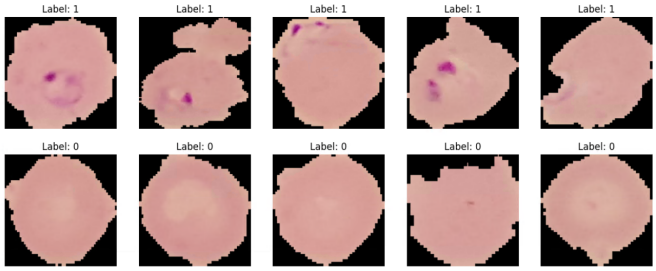


FIGURE 2: Sample malaria cell images from the dataset. The top row contains parasitized (infected) cells, while the bottom row contains uninfected cells. Morphological differences, particularly the presence of intracellular parasites, are visible in the infected samples.

B. DATA PREPROCESSING

To ensure that the input images were in an optimal format for training, a standardized preprocessing pipeline was applied. First, all images were loaded and converted to the RGB format to ensure consistency in color channels. Subsequently, the images were resized to 128×128 pixels to reduce computational cost and enforce uniform input dimensions across the dataset.

Normalization was performed by scaling pixel intensities to the range [0, 1] through division by 255. This step facilitated faster model convergence and prevented training instabilities. The dataset was then randomly shuffled and partitioned into training, validation, and test sets, maintaining class distribution balance and preventing any sampling bias.

To improve generalization and minimize overfitting, several augmentation strategies were applied to the training set. These included horizontal flips, slight rotations, and brightness adjustments. Such augmentations increased the diversity of the training data, enabling the model to learn more invariant features across varied representations of infected and uninfected cell images.

Figure 3 presents a panel of augmented cell images from the training dataset, each labeled as either infected (Label: 1) or uninfected (Label: 0). The displayed transformations include random rotations, color jitter (brightness and contrast adjustments), and geometric distortions such as zooming and flipping. These augmentations were designed to simulate the variability in real-world microscopy conditions, such as changes in lighting, orientation, and focus.

Notably, some of the cells exhibit dark padding around the edges, indicative of rotation or scaling operations applied during augmentation. Despite these visual modifications, the critical morphological features used to distinguish parasitized from uninfected cells remain intact. By exposing the model to a broader spectrum of visual diversity, these augmentations enhance its ability to generalize and reduce overfitting. This, in turn, leads to more stable performance on unseen test data.



FIGURE 3: Sample augmented images of malaria-infected (Label: 1) and uninfected (Label: 0) cells. Transformations include random rotations, brightness variation, and geometric shifts to increase training diversity.

1) Data Partitioning Strategy

To enable robust training and objective evaluation of the malaria classification model, the dataset was partitioned into three subsets: training, validation, and testing. The training set was used to fit model parameters, the validation set supported hyperparameter tuning and early stopping, and the test set was held out entirely for final model assessment.

The data was split using a stratified sampling strategy to preserve the class balance of parasitized (infected) and uninfected samples in each subset. This was essential to avoid skewed learning and to ensure that the model received equal representation of both classes throughout training and evaluation.

Figure 4 presents the bar plots for each split. Each bar chart confirms that the dataset remains perfectly balanced across all three subsets. Specifically, the training set contains approximately 8,600 images per class, while the validation and test sets each contain around 1,875 images per class. This balance is critical for reducing classification bias and improving generalization performance.

C. INTEGRATED LEARNING WORKFLOW: FL, FSL, FFA, AND CL

Figure 5 illustrates the complete training pipeline used in this study, which integrates Federated Learning (FL), Few-Shot Learning (FSL), Federated Feature Augmentation (FFA), and Continuous Learning (CL). This architecture was designed to address critical challenges in medical image classification: privacy preservation, limited labeled data, cross-site variability, and the need for continual model adaptation.

The process begins with standard input validation and data preprocessing, including augmentation and contrast enhance-

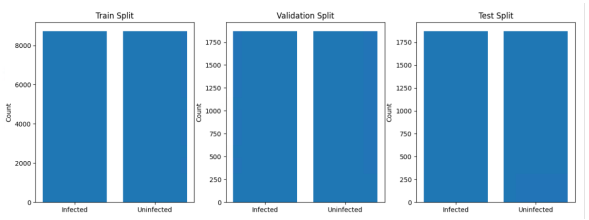


FIGURE 4: Bar plots showing the distribution of infected and uninfected samples in the training, validation, and test splits. The stratified sampling approach ensured an equal number of samples per class in each subset.

ment. Each client locally performs feature extraction and few-shot sampling, and enriches its dataset using federated feature augmentation strategies such as GAN synthesis, style transfer, and prototype alignment. These modules help mitigate data scarcity and boost generalization by increasing intra-class variance across clients.

Clients then train local models using the globally distributed weights. Following local training, clients compute updates which are sent to the FL server for aggregation using FedAvg. The server evaluates the aggregated model and either adjusts weights or triggers continual learning routines such as memory replay and forgetting prevention, depending on whether performance is satisfactory.

The final model is deployed once the desired accuracy or training rounds are reached. Throughout the training lifecycle, performance monitoring tools assess feature drift, client-level accuracy, and other metrics to inform adaptive reweighting and sustained optimization.

D. FEDERATED LEARNING ROUND EXECUTION

While Figure 5 captures the overall system architecture, the core federated training loop is detailed in Figure 6. The process begins with the FLServer broadcasting the current global model to all participating clients. Each client performs local training using its private dataset, applies updates, and evaluates performance. The trained weights are then sent back to the FLServer.

The server performs model aggregation using the FedAvg algorithm and updates global metrics. This cycle is repeated over multiple communication rounds until the desired convergence is achieved. The final global model is then serialized and deployed. The design supports asynchronous expansion into continual learning or reweighting paths depending on client drift detection.

VII. EXPLORATORY DATA ANALYSIS

To understand the visual properties of parasitized and uninfected cell images, an exploratory data analysis was conducted by examining pixel intensity distributions across RGB channels. This analysis supports model development by identifying low-level discriminative patterns that can enhance feature extraction and classification.

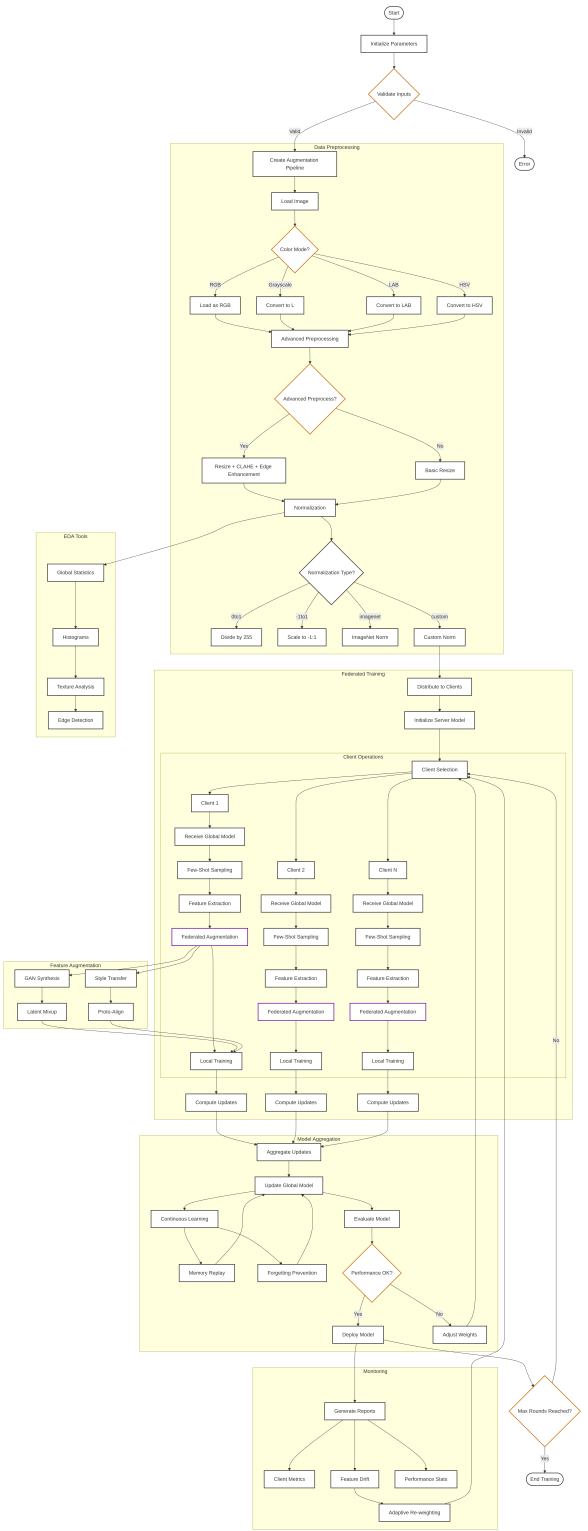


FIGURE 5: Integrated learning pipeline incorporating Federated Learning (FL), Few-Shot Learning (FSL), Federated Feature Augmentation (FFA), and Continuous Learning (CL).

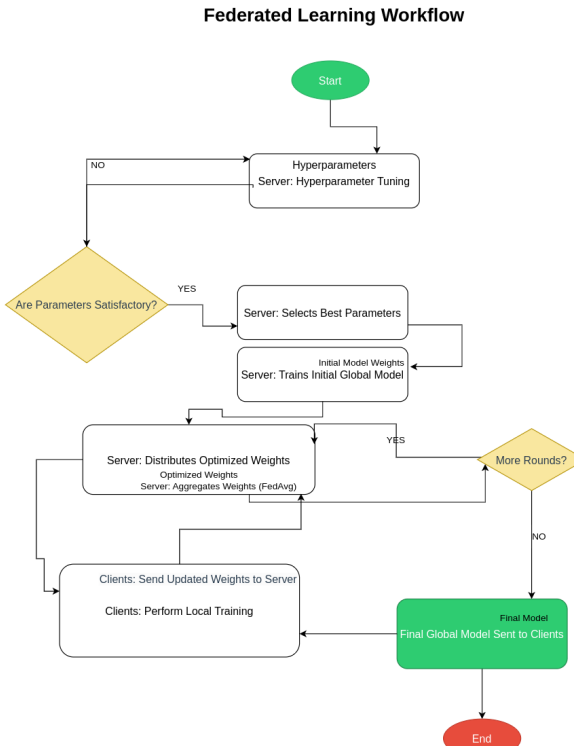


FIGURE 6: Federated learning training loop: model distribution, local training, update aggregation, and convergence.

PARASITIZED CELLS

Figure 7 presents the RGB histogram for parasitized cells. Pixel intensities were normalized to the range [0, 1] and plotted for the red, green, and blue channels.

The histogram reveals:

A sharp spike near zero in all channels, corresponding to dark or background pixels.

A red channel peak near 0.58, consistent with the natural color of stained red blood cells.

The blue channel exhibits a distinct peak around 0.75, reflecting purple-stained parasitic regions. These high-blue-intensity zones correspond to intracellular malaria parasites, which are often stained dark purple using Giemsa or Wright's stain.

The green channel shows intermediate behavior, overlapping slightly with red but lacking the pronounced peak seen in the blue channel.

UNINFECTED CELLS

In contrast, Figure 8 displays the pixel intensity histogram for uninfected cells. Only the blue channel is shown since no channel-wise separation was applied, which still reveals useful patterns:

A strong peak at zero intensity is again present, corresponding to background pixels.

Two main peaks are observed between intensity ranges of 0.6 and 0.8, indicating consistent staining of healthy red blood cells.

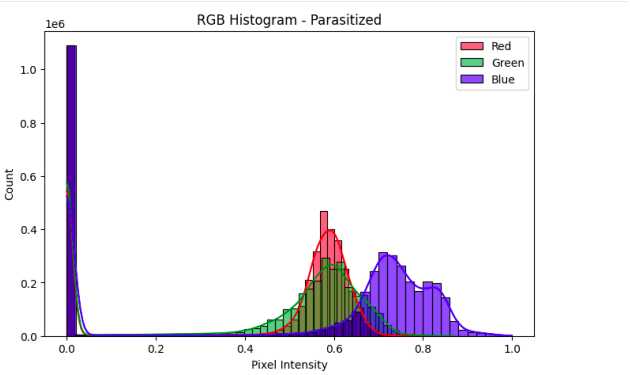


FIGURE 7: RGB pixel intensity histogram for parasitized cells. The blue channel shows elevated intensity due to parasite inclusions, while the red channel reflects the dominant background coloration of red blood cells.

Unlike parasitized images, there is no pronounced high-intensity blue channel response, confirming the absence of deep purple parasite inclusions.

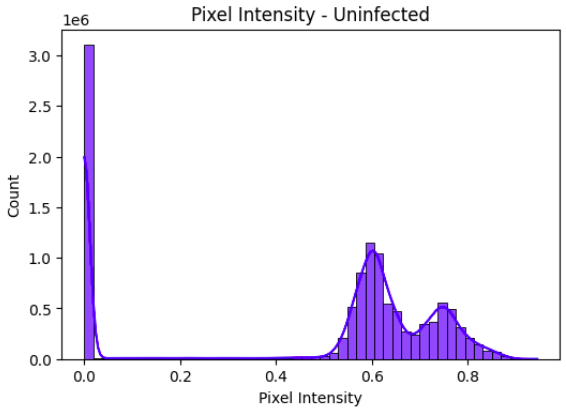


FIGURE 8: Pixel intensity distribution for uninfected cells. Compared to parasitized cells, uninfected samples lack high blue channel intensity, reflecting the absence of parasite-induced staining artifacts.

COMPARATIVE INSIGHTS

The comparison between Figures 7 and 8 reveals key discriminative cues:

Parasitized cells exhibit higher blue channel intensities, a direct result of parasitic inclusions stained in purple hues.

Uninfected cells show more uniform distributions, primarily centered around red and green channels, reflecting homogenous cell morphology.

These differences in color distribution justify the use of full RGB input for classification and confirm that color-based features can significantly aid in distinguishing between infected and uninfected samples.

CLAHE-BASED CONTRAST ENHANCEMENT

To further enhance the visibility of morphological structures within the cell images, Contrast Limited Adaptive Histogram Equalization (CLAHE) was applied as a preprocessing step. CLAHE improves local contrast in grayscale images by redistributing pixel intensity values within small regions, thereby enhancing fine-grained details without amplifying noise.

Figure 9 presents a grid of CLAHE-processed samples. The top row shows parasitized cells, while the bottom row includes uninfected cells. Compared to their original versions, CLAHE-enhanced images exhibit improved contrast around parasitic inclusions and cell boundaries, which are otherwise subdued under standard lighting and staining conditions.

For parasitized samples, CLAHE enhancement reveals darker parasite nuclei and cytoplasmic features more distinctly, which could aid the model in identifying infection-specific patterns. In uninfected samples, the overall texture becomes more uniform, with clearer contour definition but no focal parasite regions. This contrast differentiation helps preserve class-specific features and may enhance classification accuracy, especially for models relying on grayscale or texture-based inputs.

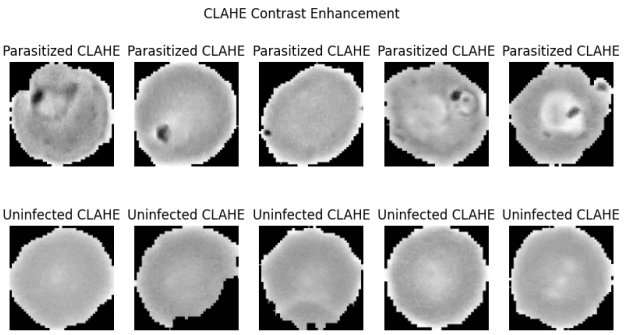


FIGURE 9: CLAHE-enhanced grayscale images for parasitized (top row) and uninfected (bottom row) samples. Contrast improvement highlights parasite regions and cellular texture differences.

VIII. MODEL DESIGN AND OPTIMIZATION

A. CNN ARCHITECTURE

The classification model used by each client node in the federated learning setup was a lightweight convolutional neural network (CNN) optimized for binary classification of malaria cell images. As shown in Figure 10, the architecture comprises two convolutional blocks, each consisting of a Conv2D layer, Batch Normalization, and MaxPooling. This is followed by a Flatten layer, a fully connected Dense layer, Batch Normalization, Dropout, and a final output layer.

The convolutional layers employed 3×3 kernels with ReLU activation. MaxPooling used a (2×2) window. The final output layer was a single sigmoid neuron for binary prediction.

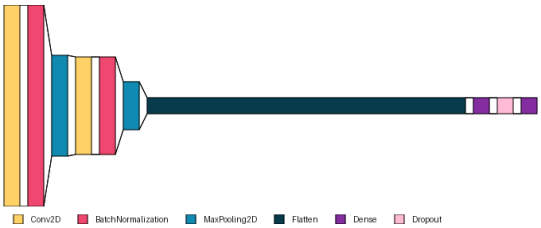


FIGURE 10: CNN architecture used for binary classification of malaria cell images.

B. LOSS FUNCTION AND OPTIMIZER

The model was trained using the binary cross-entropy loss:

$$L = -\frac{1}{N}\sum_{i=1}^{N}\left[y_i \log(\hat{y}_i) + (1 - y_i)\log(1 - \hat{y}_i)\right],$$

where $y_i \in \{0, 1\}$ is the ground-truth label and $\hat{y}_i$ is the predicted probability. Optimization was performed using the Adam optimizer with a tunable learning rate. Batch sizes and epochs were configured based on local client capacity.

C. HYPERPARAMETER OPTIMIZATION

To improve generalization and training stability during federated learning, Bayesian hyperparameter optimization was conducted using Keras Tuner. The search space was constructed to explore architectural and regularization parameters relevant to convolutional neural networks applied in distributed settings. Specifically, the number of convolutional filters was selected from the set {16, 32, 48, 64}, while dense layer widths were drawn from {32, 64, 96, 128}. The dropout rate was sampled uniformly in the range [0.2, 0.6] to balance regularization. Learning rates were tested over discrete values $\{10^{-2}, 10^{-3}, 10^{-4}\}$ using the Adam optimizer. Each trial was evaluated based on validation accuracy computed on a stratified train-validation split of the client-side data. A total of eight trials were performed to ensure computational feasibility, and the configuration achieving the highest validation accuracy was selected to initialize the global model prior to federated training.

D. HYPERPARAMETER OPTIMIZATION RESULTS

The outcome of the Bayesian hyperparameter tuning process is summarized in Table 1. Trial 7 produced the highest validation accuracy of 90.93%, utilizing 16 convolutional filters, 64 dense units, and a dropout rate of 0.3. This configuration demonstrated optimal trade-offs between representational capacity and regularization, resulting in improved generalization on the held-out validation set. Notably, configurations that applied moderate dropout rates in the range of 0.3 to 0.4 consistently achieved higher accuracy across trials, indicating the importance of balanced regularization in preventing overfitting during early training epochs.

Models with insufficient capacity, such as Trial 5 with only 32 dense units and 32 convolutional filters, exhibited

significantly reduced performance (47.53%), suggesting inadequate feature abstraction and weak generalization. Conversely, overly complex configurations such as Trials 1, 4, and 6, which used 64 filters and dense layers of 96 or 128 units, did not yield proportional gains, potentially due to increased overfitting risk when paired with lower dropout values. These trends suggest that increasing network width beyond a certain threshold produced diminishing returns without corresponding increases in regularization.

The hyperparameter set obtained in Trial 7 was adopted to initialize the global model prior to federated training. This selection ensured that each participating client began local training from a configuration validated to be effective under centralized conditions, thereby improving the likelihood of convergence and robustness in the distributed optimization process, as discussed in the Methodology section.

TABLE 1: Hyperparameter tuning results

Trial	Conv Units	Dense Units	Dropout	Val. Accuracy
1	64	96	0.2	0.8736
2	16	64	0.4	0.8984
3	48	64	0.2	0.8901
4	64	96	0.3	0.8956
5	32	32	0.2	0.4753
6	64	128	0.3	0.8956
7	16	64	0.3	**0.9093**
8	48	128	0.4	0.8929

IX. EVALUATION

This section evaluates the effectiveness and robustness of the proposed integrated AI framework, specifically focusing on the combined benefits of Federated Learning, Few-Shot Learning, Federated Feature Augmentation, and Continuous Learning. The evaluation emphasizes realistic conditions in decentralized healthcare environments, particularly considering resource limitations and privacy requirements commonly encountered in Uganda.

A. EXPERIMENTAL SETUP

The experimental design emulated a decentralized healthcare environment, where patient data is siloed across geographically distributed clinics. Each client in the federated learning simulation represented a rural clinic with limited access to labeled data and heterogeneous feature distributions.

Federated training was orchestrated using the Flower (FLWR) framework to simulate secure collaboration without transmitting raw data. Each client independently trained a local model using its private dataset. The FL server performed weight aggregation using the FedAvg algorithm. Prior to training, hyperparameter tuning was performed using Bayesian optimization via Keras Tuner, optimizing for validation accuracy, precision, recall, and overall model robustness.

All experiments were conducted using stratified data partitions to maintain label balance. Data augmentation was applied locally, and continuous learning routines were triggered based on model drift and performance plateaus. The final global model was evaluated using multiple metrics relevant to clinical diagnosis.

B. EVALUATION METRICS

To assess model performance in a high-stakes medical context, the following evaluation metrics were adopted:

a: Accuracy

Measures the proportion of correctly classified samples across both positive and negative classes. High accuracy is desirable but may be misleading in imbalanced datasets.

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \tag{1}$$

b: Precision

Indicates the proportion of true positives among all predicted positives. It reflects the model's reliability in making positive predictions and is critical in reducing false alarms in diagnosis.

$$\text{Precision} = \frac{TP}{TP + FP} \tag{2}$$

c: Recall (Sensitivity)

Captures the model's ability to identify actual positive cases. High recall ensures that true malaria cases are not missed, which is crucial in clinical applications.

$$\text{Recall} = \frac{TP}{TP + FN} \tag{3}$$

d: F1 Score

Provides a harmonic mean of precision and recall, offering a balanced metric especially useful under class imbalance. It was treated as a primary metric for model selection.

$$\text{F1 Score} = \frac{2 \times \text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \tag{4}$$

e: Specificity

Measures the proportion of true negatives correctly identified. In clinical practice, this ensures that healthy individuals are not misclassified as positive, thereby avoiding unnecessary intervention.

f: Area Under the ROC Curve (AUC)

Quantifies the model's ability to distinguish between classes across different threshold settings. A higher AUC reflects better discriminative power.

X. MODEL PERFORMANCE COMPARISON

A. PERFORMANCE VISUALIZATION

Figure 11 illustrates the performance of various deep learning models on malaria classification using five standard metrics: accuracy, precision, recall, F1-score, and AUC. Each model was trained using the same federated dataset partitioning to ensure fair comparison.

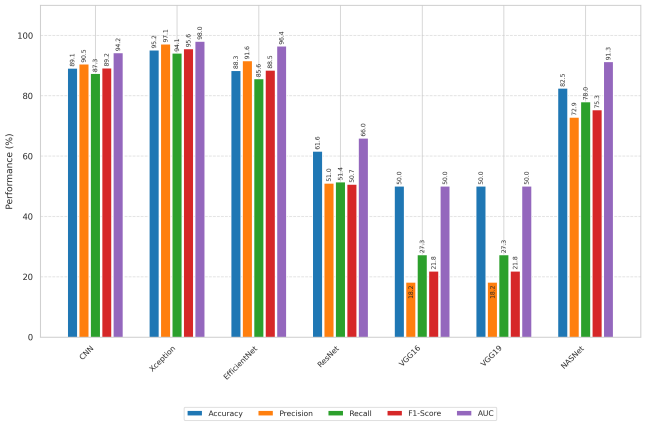


FIGURE 11: Performance comparison of deep learning models across Accuracy, Precision, Recall, F1-Score, and AUC.

B. QUANTITATIVE RESULTS

Table 2 presents the numerical values corresponding to the graph in Figure 11. The results indicate that Xception achieved the highest performance across all metrics, followed closely by EfficientNet and CNN. Models such as VGG16 and VGG19 showed poor generalization, with results near random chance levels.

TABLE 2: Performance Comparison of Different Model Architectures

Architecture	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	AUC (%)
CNN	89.10	90.48	87.35	89.20	94.24
Xception	**95.16**	**97.13**	**94.09**	**95.58**	**98.02**
EfficientNet	88.34	91.56	85.64	88.47	96.44
ResNet	61.60	50.99	51.43	50.66	65.95
VGG16	50.00	18.18	27.27	21.82	50.00
VGG19	50.00	18.18	27.27	21.82	50.00
NASNet	82.47	72.88	77.96	75.30	91.26

C. FEDERATED TRAINING EVOLUTION

Table 3 presents the CNN model's performance over 10 rounds of federated training. The results demonstrate stable improvements across metrics, with accuracy and AUC peaking at round 9. This indicates the system's ability to learn collaboratively without centralizing patient data.

XI. AI EXPLAINABILITY

To enhance transparency and improve trust in automated malaria diagnosis, explainability techniques were integrated to interpret the predictions of the convolutional neural network (CNN) trained on parasitized and uninfected red blood cell images. Within the decentralized federated learning (FL) framework, interpretability was achieved locally on each client to preserve data confidentiality while enabling model introspection. Two post-hoc interpretability methods—Gradient-weighted Class Activation Mapping (Grad-CAM) and Local Interpretable Model-Agnostic Explanations (LIME)—were employed to analyze the model's behavior at both spatial and instance-specific levels.

A. GRADIENT-WEIGHTED CLASS ACTIVATION MAPPING (GRAD-CAM)

Grad-CAM generated class-discriminative heatmaps by computing the gradient of a target class score with respect to the final convolutional layer. These visualizations highlighted spatial regions of the input that most strongly influenced the classification outcome. In this study, Grad-CAM was executed locally at each client after the training phase. Each client computed heatmaps using its trained CNN and transmitted them, along with prediction metadata, to the central server. These heatmaps were used to evaluate whether model attention aligned with parasitic features that are clinically relevant.

The server aggregated Grad-CAM outputs across clients to assess the consistency of attention patterns, thereby verifying that the global model did not exhibit dataset-specific bias. The visualization in Figure 12 shows a clear activation in the lower-left region of the cell, directly over a parasitic inclusion. This alignment with diagnostic cues suggested that the model relied on biologically meaningful structures, thus enhancing the interpretability of the federated system.

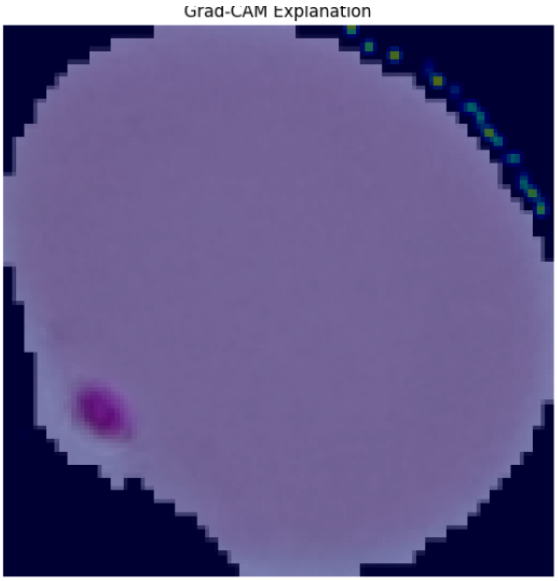


FIGURE 12: Class activation map generated using Grad-CAM. The highlighted region corresponds to the parasitized portion of a red blood cell, indicating that the model relied on medically relevant morphological cues.

B. LOCAL INTERPRETABLE MODEL-AGNOSTIC EXPLANATIONS (LIME)

LIME provided localized explanations by approximating the decision boundary of the CNN with a sparse linear surrogate, trained on perturbed variants of the original input. Each client executed LIME independently after local training. A subset of test images was perturbed by masking superpixels, and predictions were collected to estimate the contribution of individual regions to the model's decision. These local

TABLE 3: Performance Metrics Across 10 Federated Learning Rounds for CNN Model

Round	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	Specificity (%)	AUC (%)
1	87.61	85.09	91.19	88.04	84.04	94.48
2	89.29	87.09	92.25	89.60	86.33	95.05
3	89.51	88.70	90.54	89.61	88.47	94.39
4	90.47	91.19	89.58	90.38	91.35	95.32
5	90.23	90.48	89.90	90.19	90.55	94.74
6	88.33	87.78	89.05	88.41	87.61	93.34
7	89.45	92.71	85.63	89.03	93.27	94.53
8	88.36	92.94	83.01	87.70	93.70	93.72
9	**90.92**	92.46	**89.10**	**90.75**	**92.74**	**95.29**
10	89.88	91.59	87.82	89.66	91.94	94.62

attribution maps were then shared with the server for global interpretability assessment.

Aggregated LIME outputs revealed whether attribution patterns were consistent across clients with heterogeneous data distributions. As shown in Figure 13, LIME explanations identified both the infected region and surrounding morphology as influential. This broader attention span, compared to Grad-CAM, suggested that the model integrated contextual features in addition to focal lesions during classification, potentially enhancing decision robustness in borderline cases.

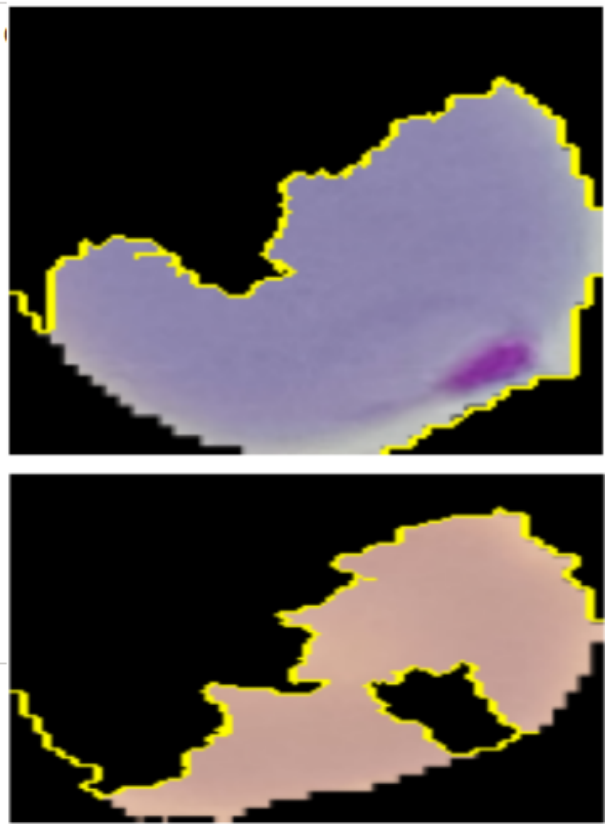


FIGURE 13: LIME-based explanation highlighting input regions influencing model predictions. The yellow contour outlines superpixels deemed important for classifying the image as parasitized.

C. EXPLAINABILITY INSIGHTS AND COMPARISON

The combination of Grad-CAM and LIME provided complementary insights into the CNN's decision-making process. Grad-CAM produced sharply focused attention maps centered on parasitic structures, offering anatomically precise cues aligned with expert diagnostic features. In contrast, LIME highlighted a broader range of superpixels, including both the lesion and its surrounding morphology, indicating that the model also incorporated contextual information.

Both methods consistently emphasized infected regions across client datasets, reinforcing the interpretability and robustness of the federated model. Their convergence indicated that the distributed training process preserved diagnostic alignment without introducing explanation drift due to data heterogeneity. These findings support the viability of explainable federated AI in clinical applications, where consistent and transparent reasoning is critical for adoption in diagnostic workflows.

XII. DEPLOYMENT

The deployment of the Federated Few-Shot Continuous Learning (FFSCL) model was designed with scalability, data privacy, and low-resource adaptation in mind—factors critical for effective use in Uganda's decentralized healthcare settings. Figure 14 illustrates the complete MLOps pipeline adapted for this federated learning system, from model development to compliance.

A. MODEL DEVELOPMENT AND PACKAGING

Model training began with Few-Shot Learning strategies at each client node, targeting limited data availability in rural clinics. Feature extraction leveraged deep learning architectures including Xception, EfficientNet, CNN, NASNet, ResNet, and VGG16/VGG19. Quantization and pruning were applied for deployment efficiency. These components were versioned and monitored using MLflow to ensure reproducibility and experiment tracking.

B. FEDERATED DEPLOYMENT INFRASTRUCTURE

Federated orchestration was achieved using Flower integrated with Django REST APIs. Dockerized client modules communicated asynchronously with a central federated server. TensorFlow Lite was used to optimize inference on

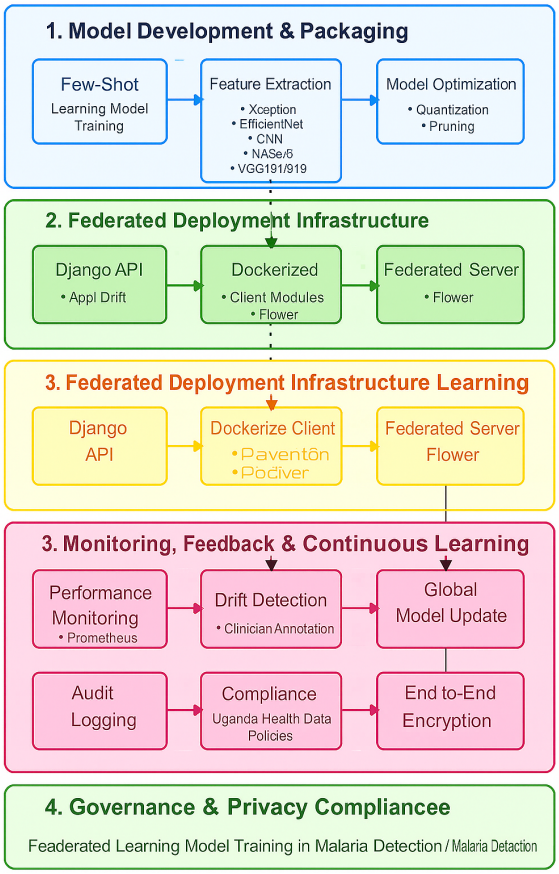


FIGURE 14: MLOps pipeline for Federated Learning in malaria detection.

edge devices such as mobile phones and tablets. The system was designed to support continuous operation, even under intermittent connectivity, via checkpointing and offline sync mechanisms.

C. MONITORING AND CONTINUOUS LEARNING

The CI/CD pipeline was enhanced with monitoring tools like Prometheus and Grafana. Clinician feedback supported drift detection (both data and concept), enabling periodic global model updates. Compliance with Uganda's health data policies was enforced through audit logging and encryption. Explainability tools (Grad-CAM and LIME) contributed to model transparency and user trust during deployment.

D. GOVERNANCE AND SECURITY

All data exchanges between clients and the central server were encrypted using end-to-end protocols. Differential privacy techniques ensured that no sensitive patient information was leaked during model updates. Regular security audits, role-based access control, and container scanning (via Docker/Kubernetes) were integrated to maintain robust operational security.

This deployment framework demonstrates how federated learning—augmented by Few-Shot Learning, Federated Feature Augmentation, and Continuous Learning—can be successfully operationalized in resource-constrained environments. The MLOps pipeline facilitated lifecycle management, performance monitoring, and clinical compliance, ensuring sustainable deployment across Uganda's healthcare network [1], [5].

XIII. CONCLUSION AND FUTURE WORK

An integrated artificial intelligence architecture was presented, combining Federated Learning (FL), Federated Few-Shot Learning (FFSL), Federated Feature Augmentation (FFA), and Continuous Learning (CL) to address the limitations of conventional machine learning approaches in decentralized, privacy-sensitive environments. The proposed solution was specifically designed to operate under the constraints of Uganda's healthcare infrastructure, where limited data availability, regulatory compliance, and computational resource constraints are persistent barriers to scalable AI deployment.

Federated Learning was utilized to ensure that patient data remained localized, which aligned with national regulations concerning data confidentiality and ethical health information handling. The incorporation of Federated Few-Shot Learning enabled local healthcare institutions to fine-tune diagnostic models with minimal labeled examples, significantly reducing dependency on large annotated datasets. This approach proved beneficial in clinical settings where annotated medical data were scarce. Federated Feature Augmentation was implemented to enhance model generalization across non-IID data partitions. It enabled consistent performance across multiple decentralized facilities by enriching feature representations and reducing overfitting to node-specific distributions. The Continuous Learning component facilitated real-time adaptation by incorporating evolving health records and clinical patterns into the learning process. This continuous update cycle ensured that the models sustained predictive accuracy as new diagnostic information became available.

Experimental results indicated that the system improved diagnostic accuracy while reducing communication overhead and safeguarding data privacy. The prototype system, developed using TensorFlow for modeling, Flower for federated orchestration, Django for backend support, and Flutter for mobile interface deployment, demonstrated practical feasibility for deployment in resource-limited clinical environments. The architecture was shown to be modular and scalable, making it suitable for deployment across geographically distributed healthcare facilities.

Future directions include conducting large-scale field trials across public and private healthcare institutions in Uganda to evaluate performance and user adoption under real-world constraints. Additional research will focus on optimizing the model to reduce computational and communication complexity using lightweight neural networks and quantization-aware techniques. The methodology will also be extended beyond malaria diagnostics to encompass other public health chal-

lenges such as tuberculosis, diabetes, and maternal health, thereby broadening its impact. Finally, ongoing collaborations with academic institutions and healthcare policy entities will be deepened to facilitate technical training, domain-specific model customization, and sustained deployment support.

The federated learning ecosystem demonstrated the capability to overcome data governance and model generalization challenges while maintaining accuracy and adaptability. The system represents a strategic application of privacy-preserving AI tailored for the enhancement of diagnostic services and the strengthening of healthcare delivery systems in Uganda.

REFERENCES

[1] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in Proc. AISTATS, 2017. [Online]. Available https://proceedings.mlr.press/v54/mcmahan17a/mcmahan17a.pdf

[2] C. Finn, P. Abbeel, and S. Levine, "Model-agnostic meta-learning for fast adaptation of deep networks," in Proc. ICML, 2017. [Online] Available https://proceedings.mlr.press/v70/finn17a/finn17a.pdf

[3] G. I. Parisi, R. Kemker, J. L. Part, C. Kanan, and S. Wermter, "Continual lifelong learning with neural networks: A review," Neural Networks, vol. 113, pp. 54–71, 2019. [Online]. Available Continual lifelong learning with neural networks

[4] LHNBC, "Thick_Smears_150 dataset," [Online]. Available: https://data.lhncbc.nlm.nih.gov/public/Malaria/Thick_Smears_150/index.html, accessed Feb 02, 2025.

[5] B. Yurdem, M. Kuzlu, M. K. Gullu, F. O. Catak, and M. Tabassum, "Federated Learning: Overview, Strategies, Applications, Tools, and Future Directions," Heliyon, vol. 10, p. e38137, 2024. [Online]. Available: https://www.cell.com/heliyon/fulltext/S2405-8440(24)00101-5

[6] D. Shenaj, G. Rizzoli, and P. Zanuttigh, "Federated Learning in Computer Vision," IEEE Access, vol. 11, pp. 94863–94870, 2023. [Online]. Available: https://ieeexplore.ieee.org/document/10256563

[7] Treppan Technologies, "Transforming Healthcare in Uganda Through Artificial Intelligence: Opportunities and Challenges," Treppan Technologies, 2023. [Online]. Available: https://www.treppantechnologies.com

[8] Mbarara University of Science and Technology (MUST), "AI for Maternal Health in East Africa," MUST Research Initiative, 2024. [Online]. Available: https://www.must.ac.ug

[9] Qure AI, "Impacting TB Detection with AI in Uganda," Qure AI, 2023. [Online]. Available: https://www.qure.ai

[10] Mugisha and Kizito, "Critical factors influencing data use and utilization in Uganda's Health Information Exchange," Oxford Academic, vol. 7, no. 1, pp. 14–32, 2023. [Online]. Available: https://academic.oup.com

[11] Kirya et al., "Digital Transformation in Uganda's Public Health: AI and Non-Communicable Diseases," JMIR Public Health and Surveillance, vol. 9, 2023. [Online]. Available: https://publichealth.jmir.org

[12] Stop TB Partnership, "Utilizing AI and Ultra-Portable X-Ray to Improve TB Detection in Uganda," Stop TB Partnership, 2023. [Online]. Available: https://www.stoptb.org

[13] Tan et al., "Federated Learning for Medical Imaging: A Comprehensive Review," IEEE Trans. Med. Imaging, vol. 42, no. 5, pp. 1125–1140, 2023. [Online]. Available: https://ieeexplore.ieee.org/document/9841125

[14] Xu et al., "Few-Shot Learning for Anomaly Detection in Federated Medical Imaging," Med. Image Anal., vol. 69, p. 101978, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1361841521000591

[15] Guo et al., "Privacy-Preserving Federated Learning for Secure Medical Image Classification," Pattern Recognit., vol. 127, p. 108635, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S003132032100675X

[16] Zhao et al., "Federated Learning with Non-IID Data: Challenges and Approaches," arXiv preprint arXiv:1806.00582, 2018. [Online]. Available: https://arxiv.org/pdf/1806.00582

[17] AIM-AHEAD, "Federated Learning in Healthcare: Opportunities and Challenges," Artificial Intelligence and Machine Learning (AIM-AHEAD), 2024. [Online]. Available: Download File Here



BYARUHANGA JOSHUA (MS Candidate, Makerere University) earned his Bachelor of Science in Computer Science from Gulu University. His research interests include Artificial Intelligence, federated learning, continuous learning, big data analytics, and image analysis.



MUHINDO MUBARAKA (MS Candidate, Makerere University) obtained a B.Sc. in Computer Science from the Islamic University of Technology in Dhaka, Bangladesh. His research focuses on artificial intelligence, federated learning, continuous learning, and medical imaging.



MARRIETTE KATARAHWEIRE is currently a Postdoctoral Researcher and a Lecturer with the Department of Computer Science, School of Computing and Informatics Technology, College of Computing and Information Sciences, Makerere University, Kampala, Uganda. Her research interests include healthcare security, secure software development, mobile health and distributed computing



KIZITO is a Lecturer in the Department of Computer Science and a Research Fellow with the Software Systems Center at Makerere University. He graduated in October 2001 with Honors from Makerere University Kampala (Uganda) after completing a Bachelor of Science Degree in Computer Science and Mathematics. Worked with Digital Solutions as a Senior Software Developer until 2003. Pursued a Master of Science Degree in Computing Science (Software and Systems Engineering) in the University of Groningen (The Netherlands). While in Groningen, he worked on a Prototype of a Decision Support System for Anesthetists. During PhD studies, he worked on Extraction of Medical Term Relations from Medical Text. He also holds a PGD Certification in e-Teaching from the University of Agder (Norway).



JOYCE NAKATUMBA-NABENDE is a Senior Lecturer in the Department of Computer Science at Makerere University, where she also leads the Makerere Artificial Intelligence Lab. She holds a PhD in Computer Science from Eindhoven University of Technology in the Netherlands, a Master's from Makerere University, and a Bachelor's from Mbarara University of Science and Technology. Her research spans Artificial Intelligence, Machine Learning, Natural Language Processing, and Business Process Management, with a strong emphasis on responsible AI and its applications in health and agriculture. She is deeply committed to advancing data science capacity in Africa and promoting diversity and gender equity in tech. Dr. Nakatumba-Nabende serves on several influential platforms, including the Educational Advisory Committee of the Association for Computing Machinery, the Data Science Africa board, and the Responsible AI (RAIN) Africa Network. She is also a founding member of the Open for Good Alliance and contributes to the ARUA–The Guild Cluster of Research Excellence. Her work has led to the development of benchmark datasets for low-resourced East African languages and machine learning tools for agricultural productivity, positioning her as a leading voice in AI for social good on the continent



GGALIWANGO MARVIN is a Machine Learning and Intelligent Systems Engineer whose core works focus on Autonomous Systems and Responsible Software Systems in the Tech-Industry. He is an AI Researcher, Principal Investigator and Technical Led of various studies focused on Machine Learning, Artificial Intelligence, Generative AI, Quantum Computing, Autonomous Systems, Responsible Software Systems, High Performance Computing, Digital Security, Digital Health and Bioacoustics at Makerere University. He is a Digital Education practitioner at the Mastercard foundation Digital Education practitioner's network and Technical Research-Led for the implementation of Acoustic Feature Extraction and sound Engineering for medical therapy in Uganda. He is a PhD Research Exchange fellow at Paul G. Allen School of Computer Science & Engineering (University of Washington), holds a MSc in Digital Education (The University of Edinburgh), MSc in Computer Science and Engineering (BRAC University) and BSc (Hons) in Computer Science.