Towards the Efficient and Privacy-Aware Diagnosis of Bladder Inflammations

Rashmi Ashtagi
Dr. Vishwanath Karad MIT World Peace University, Kothrud, Pune, India
rashmiashtagi@gmail.com (corresponding author)

Ranjeet Bidwe
Symbiosis Institute of Technology, Pune Campus, Symbiosis International (Deemed University) (SIU), Lavale, Pune, India
ranjeetbidwe@hotmail.com (corresponding author)

Sangita Jaybhaye
Vishwakarma Institute of Technology, Pune, India
sangita.jaybhaye@vit.edu

Nilesh J. Uke
Indira College of Engineering and Management, Pune, India
nilesh.uke@gmail.com

Jinay Jain
Vishwakarma Institute of Technology, Pune, India
jain.jinay21@vit.edu

Uday Jaju
Vishwakarma Institute of Technology, Pune, India
uday.jaju21@vit.edu

Khushi Agarwal
Vishwakarma Institute of Technology, Pune, India
khushi.agarwal21@vit.edu

Ajinkya Kalamkar
Vishwakarma Institute of Technology, Pune, India
ajinkya.kalamkar21@vit.edu

Received: 2 June 2025 | Revised: 16 July 2025, 13 August 2025, 3 September 2025, and 6 September 2025 | Accepted: 9 September 2025

 DOI: https://doi.org/10.48084/etasr.12519

ABSTRACT

Conventional centralized machine learning solutions for medical diagnosis face significant limitations, such as data privacy concerns, communication overhead, and poor scalability across healthcare institutions, restricting real-time and secure diagnostic capabilities for diseases such as acute bladder inflammation. To address these challenges, this work proposes a highly accurate and privacy-preserving model that ensures scalability and efficiency while reducing training time. A Federated Logistic Regression (FLR) approach uses data from multiple clinical institutions to collaboratively train a global model without exchanging raw patient data. The FedAvg algorithm is employed to aggregate locally trained models into a centralized global model, evaluated on a dataset of urinary tract inflammations. Experimental findings reveal that the federated architecture achieves diagnostic accuracy comparable to centralized frameworks while requiring significantly fewer training iterations, resulting in a ~20× improvement in convergence efficiency without

compromising patient privacy. The FLR model demonstrates a secure, realistic, and computationally feasible solution for the diagnosis of acute bladder inflammation, underscoring the transformative potential of federated learning to advance privacy-preserving medical diagnostic systems within distributed healthcare settings.

Keywords-nephritis; Federated Learning (FL); logistic regression; distributed systems; bladder inflammation

I. INTRODUCTION

Acute bladder inflammation, or bladder catarrh, is a common and disruptive condition characterized by symptoms such as urgency, frequency, dysuria, hematuria, and suprapubic discomfort. Although often self-limiting, severe or recurrent cases require further investigation to identify the underlying factors. Inflammation may extend to the renal pelvis and surrounding structures (nephritis), usually caused by bacterial infection in the lower urinary tract, potentially leading to serious complications such as sepsis or kidney damage if not treated.

Traditional centralized diagnostic approaches face limitations, including privacy concerns, data security risks, computational bottlenecks, and poor scalability, especially in resource-constrained settings. Federated Learning (FL) offers a decentralized alternative that allows efficient and privacy-preserving computation by leveraging distributed data sources. Previous studies have shown that FL can achieve better accuracy and efficiency compared to traditional models [3].

Centralized classical machine learning models used in medical diagnosis suffer significantly from data privacy, communication overhead, and scalability problems, especially when working with sensitive patient health data across institutions. Despite good accuracy, central models relying on aggregating data at one location exacerbate the threat of privacy breaches as well as hamper collective diagnostics in real-time. Existing studies have not effectively demonstrated a practical and privacy-preserving solution that offers comparable performance. This study aimed to fulfill the following objectives:

- Develop a highly accurate model while strictly adhering to privacy regulations.
- Propose a means of assuring scalability so that it can work in any healthcare setting to boost productivity and reduce training time.

This study presents a Federated Logistic Regression (FLR) framework using FedAvg to diagnose acute bladder inflammation across distributed data sources, addressing the privacy, efficiency, and scalability issues in centralized systems. The framework ensures data confidentiality, achieves high accuracy, and converges much faster than previous methods, demonstrating the practical potential of federated learning for secure and scalable medical diagnosis. The key contributions of this study are as follows:

- Proposes an FLR model for the diagnosis of acute bladder inflammation, enabling collaborative learning without compromising patient data privacy.
- Implements and evaluates the model using the FedAvg algorithm on multiple distributed nodes, demonstrating enhanced efficiency and scalability.
- Experimental results show that the proposed model achieves the same diagnostic accuracy as centralized models while significantly reducing the number of training iterations.

Machine learning models leveraging urine, blood, and demographic features offer a faster and more accurate diagnosis than conventional methods such as urine culture [4]. Fine-tuned large language models have shown a strong potential for diagnosing acute bladder inflammation and nephritis from symptom-based datasets, outperforming traditional neural networks, even with limited data [5]. Classical models, such as SVM, KNN, and ensembles, achieve high diagnostic accuracy, with techniques such as SVM-SMOTE and SHAP enhancing balance and interpretability [6]. Combining patient-reported outcomes with urinary biomarkers further improves the diagnostic accuracy of interstitial cystitis [7]. Comparative studies have shown that SVM outperforms random forests for the detection of bladder inflammation [8]. FL has been explored in various contexts, such as an IoT-enabled wheelchair system for secure and adaptive assistance to disabled pilgrims, limited by simulations and context-specific focus [9], a multimodal facial expression recognition method using GCNs with strong performance but lacking efficiency and occlusion robustness analysis [10], and AFRLS [11], which integrates federated and reinforcement learning in fog-cloud kidney image processing to reduce delays, though without real-world validation.

II. PROPOSED MODEL

Traditional Logistic Regression (LR) requires frequent retraining on centralized datasets, causing delays and high processing costs. A federated approach, distributing data across multiple clients, overcomes this limitation by enabling faster and more efficient detection of nephritis and bladder inflammation. The proposed model employs four worker nodes, and each participating client device hosts a local LR model trained on its respective dataset. These client models autonomously update their parameters based on local data, capturing intricate features within their data distributions. Periodically, the client models communicate their updates to a central server, where federated averaging is employed to aggregate the model parameters across all clients. The global model on the server is then updated using the aggregated parameters, ensuring that it represents the combined knowledge of all participating clients. The global model can continuously improve its performance while adhering to the data privacy restrictions that come with FL environments because of this iterative training, aggregation, and updating procedure.

Figure 2 demonstrates the proposed design. Submodels A and B symbolize the client nodes that train the local models, which are then transmitted to the server node that uses FedAvg.

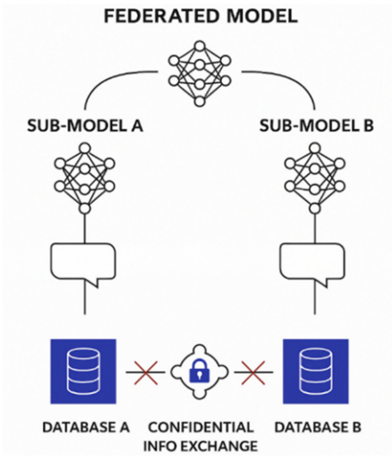

Fig. 1. Proposed architecture.

Although many FL healthcare frameworks—especially for image-based diagnostics, electronic health records, and multi-modal data—use complex architectures, such as CNNs or variational autoencoders, the proposed FLR model is lightweight, interpretable, and optimized for structured tabular clinical data to diagnose acute bladder inflammation and nephritis. Using LR preserves clinical transparency, allowing symptom-level interpretation. The model converges faster than centralized methods, ensuring high computational efficiency, full data privacy, and simple implementation. This makes it practical for resource-limited clinical settings and, to our knowledge, is the first FLR approach tailored to this diagnostic context.

A. Dataset

The experiments were conducted using the publicly available Acute Inflammations Dataset from the UCI Machine Learning Repository [12]. The dataset contains 120 patient records, each described by six independent variables representing clinically relevant attributes:

- Temperature of patient (continuous, in °C)
- Occurrence of nausea (binary: yes/no)
- Lumbar pain (binary: yes/no)
- Urine pushing (binary: yes/no)
- Micturition pains (binary: yes/no)
- Burning of the urethra, itch, swelling of the urethral outlet (binary: yes/no)

Two dependent binary target variables are provided:

- Diagnosis of acute inflammation of the urinary bladder
- Diagnosis of nephritis of renal pelvis origin.

Thus, the dataset supports two separate binary classification tasks.

B. Logistic Regression (LR)

LR is a statistical method for modeling the relationship between one or more independent variables and a binary outcome using the logistic, or sigmoid, function to convert a linear combination of predictors into a probability score between 0 and 1, unlike classic LR, which predicts continuous outcomes. The logistic function is defined as follows:

$$logit(p) = \ln\left(\frac{p}{1-p}\right) = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \cdots + \beta_n x_n$$

where p denotes the probability of the binary result (success or failure), $x_1, x_2, \ldots, x_{1n}$ are the independent variables, and $\beta_0, \beta_1, \ldots, \beta_n$ are the coefficients of independent factors contributing toward the log-odds of an outcome.

The linear combination of independent measures is transformed into log-odds with the logistic function, also called the logit [13]. The logit is then mapped back onto the probability scale with the logistic function:

$$p = \frac{1}{1+e^{-(\beta_0+\beta_1 x_1+\beta_2 x_2+\cdots+\beta_n x_n)}}$$

To maximize the possibilities of viewing the provided binary outputs with the occurrence of the predictor variables, LR applies maximum likelihood estimation or other optimization procedures to estimate the coefficients $\beta_0, \beta_1, \ldots, \beta_n$. As it applies the coefficient, it can yield the probability of binary outputs with the independent variable values.

C. Federated Learning (FL)

This approach trains models across decentralized devices or servers holding local data, sharing only model updates such as weights or gradients [14]. This approach enables on-device training, reduces data transmission, and addresses privacy concerns of centralized storage. Mathematically, FL can be formulated as follows:

Let $D_1, D_2, \ldots, D_n$ represent the datasets held by N devices or servers. Each dataset D_i contains samples (x_{i1}, y_{i1}), (x_{i2}, y_{i2}), ... , (x_{iM}, y_{iM}), where x_{ij} is the input data and y_{ij} is the corresponding label. The goal is to learn a global model θ that minimizes a global loss function $\zeta(\theta)$, defined as the average loss over all local datasets [15]:

$$min_\theta \frac{1}{N}\sum_{i=1}^{N} E_{(x,y)}[l(\theta; x; y)]$$

where $l(\theta; x; y)$ is the loss function for a single sample, and E denotes the expectation over the local data D_i.

During training, each device computes a local update $\Delta\theta_i$ by minimizing its local loss function:

$$\Delta\theta_i = \text{argmin}\Delta\theta \sum_{j=1}^{M_i} l(\theta + \Delta\theta; x_{ij}, y_{ij})$$

The local updates $\Delta\theta_i$ are then aggregated at a central server using methods such as federated averaging (FedAvg) [17]. The server updates the global model θ by averaging the local updates:

$$\theta \leftarrow \theta - \frac{1}{N}\sum_{i=1}^{N} \Delta\theta_i$$

This process repeats until convergence, resulting in a global model that is trained on collective knowledge from all participating units, while data privacy is preserved. Stochastic Gradient Descent (SGD) is an efficient optimization algorithm that iteratively updates model parameters to minimize the loss, making it well-suited for large-scale datasets. For a given training example (x_i, y_i), the update for SGD is given by:

$$\theta_{(t+1)} = \theta_{(t)}\eta\nabla_{\theta}L(x_i, y_i, \theta_{(t)})$$

Here, parameter update is controlled by the learning rate η, and $L(x_i, y_i, \theta_{(t)})$ is the loss function that determines the deviation between the model conditions and the correct label.

In SGD, the gradient is computed using only a single training example at each iteration. This makes SGD highly efficient and scalable for large datasets, as it avoids computing gradients for the entire dataset in each iteration. However, due to the randomness introduced by the selection of individual training examples, SGD may exhibit noisy updates and slower convergence compared to batch gradient descent methods.

D. Pseudocode

Algorithm 1 outlines the training of a basic LR model. It begins by initializing the model parameters, defining the loss function, and setting up the optimizer. In each epoch, the model iterates over minibatches of training data, making predictions, computing loss, and updating parameters. Training and validation metrics are saved for examination. The model is evaluated using different test data.

```
Algorithm 1: LR Model
Initialize the model parameters (weights
and biases).
Define the loss function using cross-
entropy loss (e.g.,
torch.nn.CrossEntropyLoss()).
Define the optimizer (e.g., SGD or Adam)
with the chosen learning rate.
For each epoch, repeat the following
steps:
  Iterate over mini-batches of the
  training data:
  a. Perform the forward pass to get model
     predictions.
  b. Compute the loss by comparing
     predictions with true labels.
  c. Perform the backward pass to
     calculate gradients.
  d. Update the model parameters using the
     optimizer's step() function.
  Evaluate the model's performance using
  validation data.
  Plot training and validation accuracy
  and loss curves across epochs.
Evaluate the trained model using the test
dataset.
Report final test accuracy and any other
relevant evaluation metrics.
```

Algorithm 2 describes the FL approach, where the training process involves distributing the global model to participating units, training local models using unit-specific data, aggregating local model parameters to update the global model, and repeating the process over multiple federated rounds. It includes steps for distributing, training, aggregating, evaluating, plotting, testing, reporting, and discussing the FL experiment.

```
Algorithm 2: Federated LR Model
Initialize global model parameters and
necessary data structures.
Define the architecture of the global LR
model.
Define the loss function (e.g., cross-
entropy) and optimizer for local training.
For each federated round, repeat the
following:
  Distribute the current global model
  to all participating hospitals.
  At each hospital, perform the
  following steps:
    a. Train the local model using that
       hospital's private dataset.
    b. Compute gradients and update the
       local model parameters.
    c. Send the updated local model
       parameters back to the central
       server.
  Aggregate all local model parameters
  (e.g., using weighted averaging) to
  update the global model.
  Evaluate the global model's
  performance on a central validation
  dataset.
  Store federated training metrics such as
  loss, accuracy, or communication cost.
  Plot the stored federated training
  metrics across rounds.
  Test the final global model using held-
  out test data from a separate source.
  Report the final test accuracy and any
  additional observations or performance
  metrics.
```

Tables I and II show the parameters of the LR and FL algorithms.

TABLE I. PARAMETERS OF LOGISTIC REGRESSION

	Hyperparameter	Value
1	Learning rate	0.01
2	Number of epochs	20
3	Batch size	32
4	Optimizer	Stochastic Gradient Descent
5	Loss function	CrossEntropyLoss for multiclass scenarios and BCELoss for binary classification
6	Regularization	L2 penalty with $\lambda = 0.001$ to avoid overfitting

TABLE II. PARAMETERS OF FEDERATED LEARNING

	Hyperparameter	Value
1	Number of clients	4
2	Local epochs per round	5
3	Global communication rounds	20
4	Learning rate	0.01
5	Batch size	32
6	Optimizer	Stochastic Gradient Descent
7	Loss Function	CrossEntropyLoss for multiclass scenarios and BCELoss for binary classification
8	Regularization	L2 penalty with λ = 0.001 to avoid overfitting

III. PERFORMANCE EVALUATION

A. *Number of Iterations*

The number of iterations required to reduce the loss of a model is a key performance metric in machine learning, reflecting training efficiency and convergence speed [16]. Fewer iterations indicate faster convergence, reduced computational time, and optimized resource use. Iteration counts also provide insight into model complexity, as more complex models with larger parameters often require more iterations to learn data patterns. Convergence stability can be assessed by the consistency of required iterations, and significant variation may indicate instability or sensitivity to initialization. Determining the optimal iteration count supports early stopping, prevents overfitting, and conserves computational resources.

B. *Binary Cross-Entropy Loss (BCELoss)*

BCELoss is among the most widely utilized loss functions in binary classification issues, particularly in neural network-based models. It measures the loss between the expected probability and the true binary labels of the target variable.

$$BCELoss(p, y) = -\frac{1}{N}\sum_{i=1}^{N}[y_i.\log(p_i) + (1 - y_i).\log(1p_i)]$$

where p is the model's anticipated probability output, y is the target variable's actual binary label (0 or 1), and N represents the dataset's sample count. For the i-th sample, y_i and p_i represent the real label and expected probability, respectively.

For binary classification, p is typically the output of a sigmoid activation function, such that output probabilities are between 0 and 1. BCELoss is a suitable loss function for training binary classification models since it captures agreement between true labels and output probabilities, guiding the optimization toward accurate classification performance [16].

IV. RESULTS AND DISCUSSION

This study implemented a classic LR and an FLR model to compare their efficiency in terms of iterations and training loss. Table III presents the resulting values recorded at intervals of 1000 iterations for the training model for inflammation of the urinary bladder, and Table IV shows the same for nephritis of the renal pelvis for a classic LR model.

TABLE III. LR MODEL FOR INFLAMMATION OF THE URINARY BLADDER

Number of Iterations	Training loss	Training accuracy (%)
0	3.0699	47.92
500	1.506	47.92
1500	0.2249	92.71
2500	0.1627	100
3500	0.137	100
4500	0.1185	100
5500	0.1046	100
6500	0.0936	100
7500	0.0848	100
8500	0.0775	100
9500	0.0714	100
10500	0.0662	100
11500	0.0617	100
12500	0.0578	100
13500	0.0544	100
14500	0.0513	100
15500	0.0486	100
16500	0.0462	100
17500	0.044	100
18500	0.042	100
19500	0.0401	100

TABLE IV. LR MODEL FOR NEPHRITIS OF THE RENAL PELVIS

Number of iterations	Training loss	Training accuracy (%)
0	2.6804	69.79
500	0.877	60.42
1500	0.5576	60.42
2500	0.2121	60.42
3500	0.1793	86.46
4500	0.1562	100
5500	0.1389	100
6500	0.1253	100
7500	0.1142	100
8500	0.105	100
9500	0.0973	100
10500	0.0906	100
11500	0.0848	100
12500	0.0798	100
13500	0.0753	100
14500	0.0713	100
15500	0.0677	100
16500	0.0644	100
17500	0.0615	100
18500	0.0588	100
19500	0.0563	100

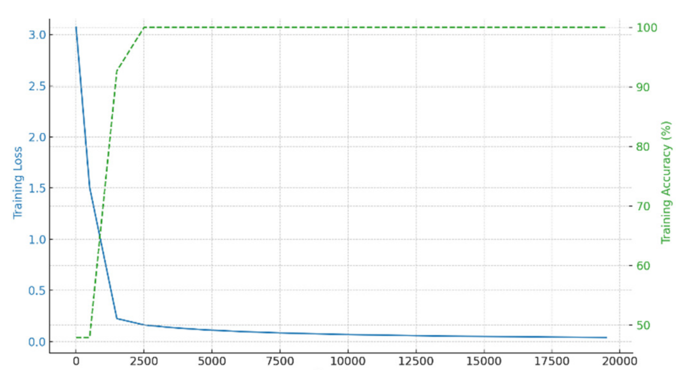

Fig. 2. Training loss and accuracy vs iterations for the LR on urinary bladder inflammation.

The bladder inflammation LR model achieved a BCELoss of 0.0401, and the nephritis LR model a BCELoss of 0.0563 at 19500 iterations. Both tasks showed 100% training and testing accuracy. However, there is an important distinction in iteration counts: Tables III and IV show the iteration at which training accuracy first reached 100% (bladder: 2,500 iterations; nephritis: 4,500 iterations). For clarity, training continued until 19,500 iterations to observe the loss trajectory and ensure stability of the learned parameters. When comparing centralized and federated approaches, the number of iterations to first reach 100% training accuracy and the number of iterations required to reach the final recorded BCELoss value are reported. The federated model reached comparable BCELoss values (≤0.06) in approximately 900 iterations, while the centralized models needed 19,500 iterations to reach similar BCELoss values. The proposed federated model was trained and tested for both diseases, with its accuracy and loss for bladder inflammation shown in Tables V and VI, and for nephritis in Tables VII and VIII.

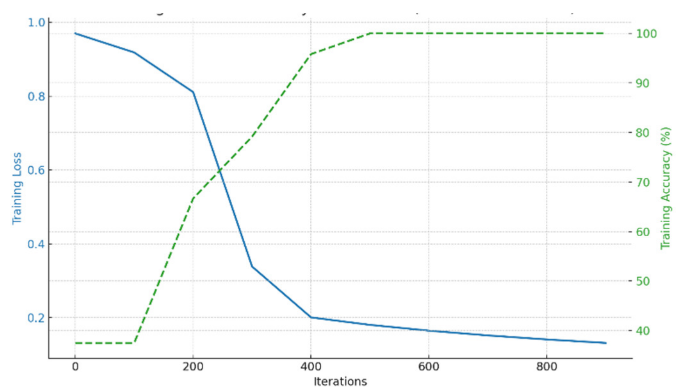

Fig. 3. Combined convergence analysis of the FLR model.

TABLE V. TRAINING ACCURACY FOR THE FEDERATED INFLAMMATION MODEL

Iterations	W1	W2	W3	W4
0	37.5	33.33	50	54.17
100	37.5	91.67	50	54.17
200	66.67	66.67	50	54.17
300	79.17	70.83	83.33	75
400	95.83	100	100	100
500	100	100	100	100
600	100	100	100	100
700	100	100	100	100
800	100	100	100	100
900	100	100	100	100

TABLE VI. TRAINING LOSS PER WORKER

Iterations	W1	W2	W3	W4
0	0.9696	0.8328	1.9309	1.3678
100	0.9181	0.4144	1.5114	1.1034
200	0.8108	0.4497	1.0155	0.8722
300	0.3383	0.3264	0.3469	0.4698
400	0.2006	0.1812	0.1624	0.2024
500	0.1804	0.1632	0.1447	0.1817
600	0.1647	0.1488	0.1309	0.1655
700	0.1518	0.1369	0.1197	0.1523
800	0.1409	0.127	0.1103	0.1411
900	0.1315	0.1184	0.1024	0.1315

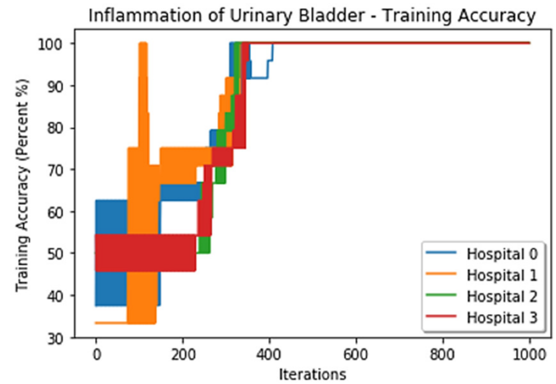

Fig. 4. Training accuracy of the Federated model on urinary bladder.

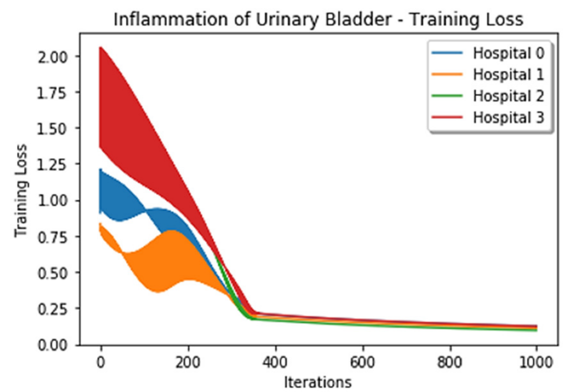

Fig. 5. Loss for the federated inflammation model on urinary bladder.

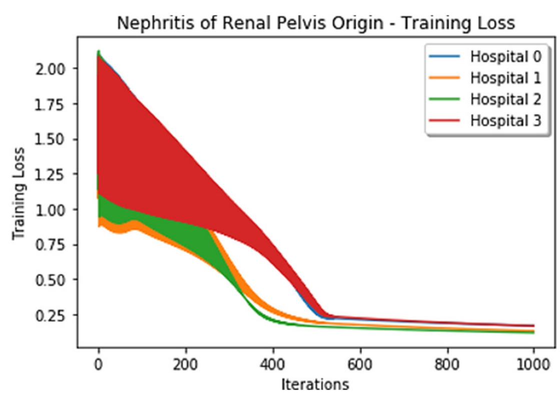

Fig. 6. Training loss of the federated model on nephritis.

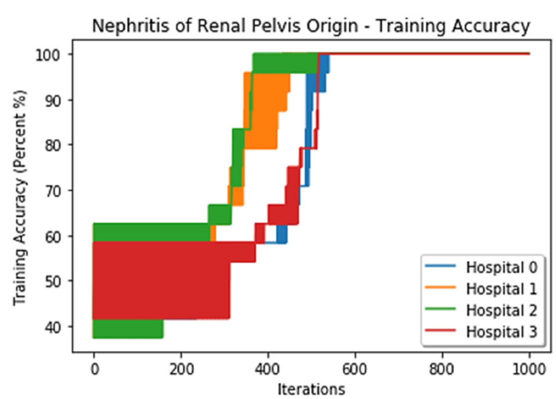

Fig. 7. Accuracy of the federated model on nephritis.

TABLE VII. TRAINING ACCURACY OVER ITERATIONS FOR EACH WORKER ON NEPHRITIS

Iterations	W1	W2	W3	W4
0	41.67	37.5	37.5	41.67
100	41.67	37.5	37.5	41.67
200	41.67	45.83	54.17	41.67
300	54.17	66.67	66.67	41.67
400	58.33	95.83	100	62.5
500	100	100	100	79.1
600	100	100	100	100
700	100	100	100	100
800	100	100	100	100
900	100	100	100	100

TABLE VIII. TRAINING LOSS FOR FEDERATED NEPHRITIS MODEL

Iterations	W1	W2	W3	W4
0	1.9586	1.0755	1.1401	1.9368
100	1.7387	1.5665	1.5511	1.7416
200	1.3884	1.1495	1.1017	1.4136
300	1.0380	0.6462	0.5630	1.0771
400	0.6920	0.2910	0.2020	0.7449
500	0.2535	0.1947	0.1616	0.3370
600	0.2098	0.1723	0.1495	0.2196
700	0.1957	0.1590	0.1395	0.2021
800	0.1841	0.1476	0.1309	0.1877
900	0.1741	0.1378	0.1234	0.1754

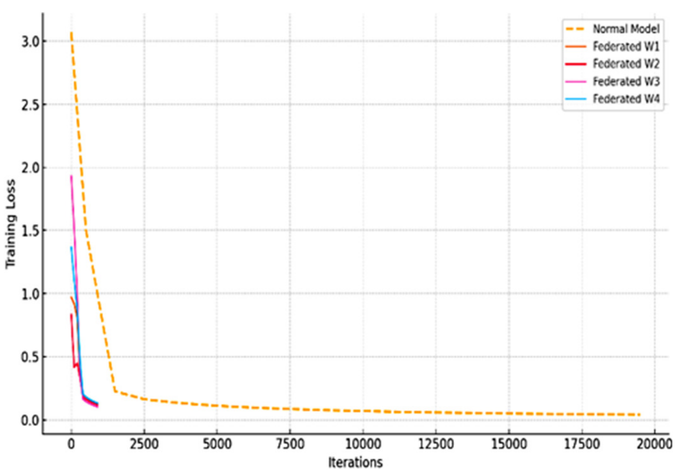

Fig. 8. Training loss vs iterations graph comparing the RL and FLR (all workers) models for bladder inflammation.

For the diagnosis of nephritis, the centralized model required approximately 19,500 iterations to reach a loss of 0.0563, while the federated model achieved comparable results in less than 1,000 iterations, confirming its ability to preserve accuracy while greatly accelerating training. For bladder inflammation, all four federated model workers reached and maintained 100% training accuracy by the 500th iteration, demonstrating stability and consistency despite varying initial performance levels. For nephritis detection, all federated model workers reached 100% accuracy by 900 iterations, even those starting as low as 37.5%, due to effective local training and FedAvg. Overall, the results show that federated learning matches centralized model accuracy while greatly reducing training time, offering a privacy-preserving, scalable, and efficient solution for distributed clinical settings.

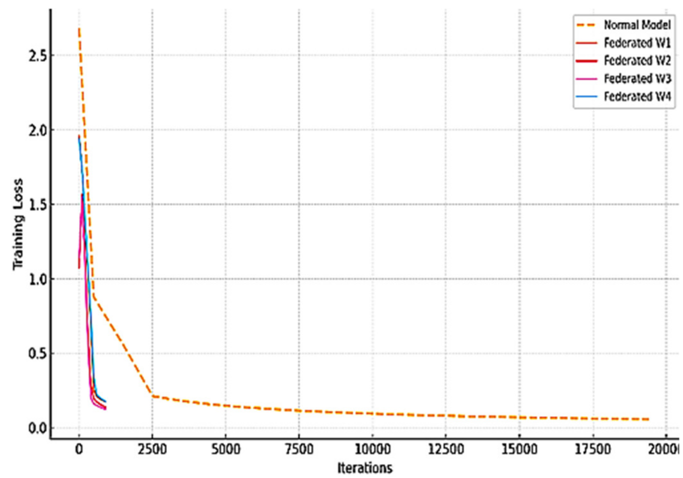

Fig. 9. Training loss vs. iterations for the Normal and FLR models on nephritis diagnosis.

A. *Analysis of Results*

Training loss and accuracy over iterations were plotted to compare the centralized and FLR models. For bladder inflammation detection, the centralized model reached a loss of 0.0401 after ~19,500 iterations, whereas the federated model achieved a similar loss in just 900 iterations across all worker nodes, demonstrating far greater computational efficiency without loss of performance. For the diagnosis of nephritis, the centralized model required approximately 19,500 iterations to reach a loss of 0.0563, while the federated model achieved comparable results in less than 1,000 iterations, confirming its ability to preserve accuracy while greatly accelerating training. For bladder inflammation, all four federated model workers reached and maintained 100% training accuracy by the 500th iteration, demonstrating stability and consistency despite varying initial performance levels. For nephritis detection, all federated model workers reached 100% accuracy by or before 900 iterations, even those starting as low as 37.5%, due to effective local training and FedAvg. Overall, the results show that federated learning matches centralized model accuracy while greatly reducing training time, offering a privacy-preserving, scalable, and efficient solution for distributed clinical settings.

The FLR model not only achieves high accuracy but also improves convergence speed and training efficiency. For both bladder inflammation and nephritis, the FLR model reached 100% accuracy in under 900 iterations, offering a great speed-up. Convergence across nodes further demonstrates the stability of the federated setup. Table IX summarizes this comparison, contrasting the centralized and federated approaches in terms of final accuracy, training iterations, and communication rounds. While the centralized model performs training on a monolithic dataset, the federated model leverages local training followed by periodic aggregation using FedAvg, with approximately 20 communication rounds required in total. This significantly reduces the communication overhead, as only model parameters are exchanged, preserving data privacy and reducing bandwidth usage. For example, sharing a few hundred model weights per round imposes far less communication cost than transmitting the full dataset.

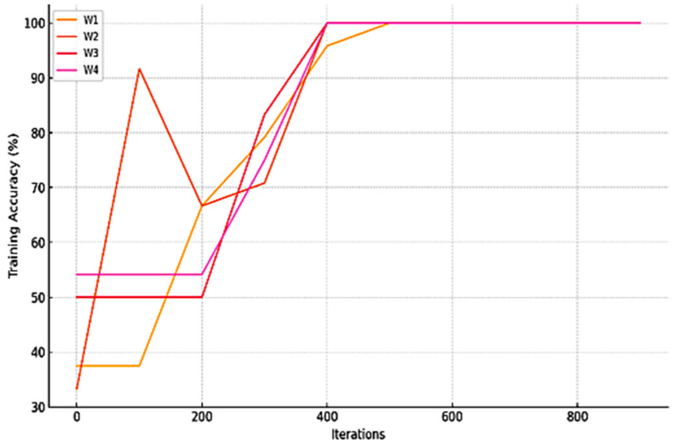

Fig. 10. Federated training accuracy vs. iterations for bladder inflammation.

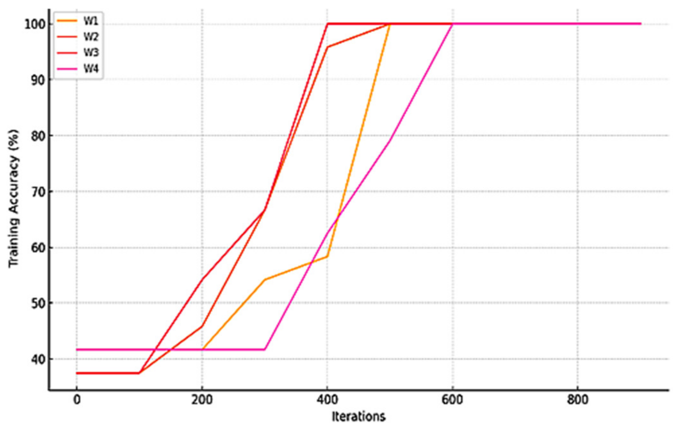

Fig. 11. Federated training accuracy vs. iterations for Nephritis diagnosis.

TABLE IX. TRAINING ACCURACY OVER ITERATIONS FOR EACH WORKER

Model type	Final accuracy	Iterations to converge	Communication rounds	Observations
Centralized	100%	~19,500	N/A	Slow convergence
Federated (W1–W4 avg)	100%	~900	~20	Fast, consistent learning

Furthermore, due to its reliance on local computation of small data partitions, such a federated model is particularly simple to deploy in low-resource environments, e.g., small hospitals or rural clinics, with limited computing capacity. Although this experiment used four worker nodes, this architecture of underlying aggregation is scalable up to larger federated networks with just minimal modifications. The federated method is accurate and privacy-preserving, and is computationally and communicationally efficient as well, providing a practicable substitute for central learning in decentralized health settings.

The proposed FLR framework is systematically scalable and can be ideally used in large-scale real-world healthcare environments with a large number of institutions. As the number of participating nodes increases, the communication and computational load per node remains manageable since each institution processes only its local data and transmits lightweight model updates (e.g., weights or gradients) to the central aggregator. This makes the system highly adaptable to settings where data is distributed between hospitals, clinics, or diagnostic laboratories. Furthermore, the use of the FedAvg aggregation strategy ensures that the global model benefits from the collective knowledge of all nodes, even in cases of data heterogeneity or varying sample sizes. From a practical point of view, the FLR model can be deployed with minimal computational requirements on edge devices or hospital servers, as it uses a lightweight and interpretable model. By keeping sensitive patient data within institutional boundaries and only exchanging model parameters, the framework aligns well with regulatory requirements (such as HIPAA or GDPR) and data governance policies commonly enforced in healthcare systems. This positions the FLR model as a viable, privacy-preserving, and scalable solution for collaborative learning in real-world medical environments.

V. CONCLUSION

This study developed a novel FLR model for the diagnosis of acute bladder inflammation and nephritis, focusing on maintaining data privacy, computational efficiency, and scalability in decentralized healthcare settings. By integrating LR into a federated framework and applying it to structured clinical data, this research contributes to both the theoretical understanding of FL in low-complexity models and the practical deployment in resource-constrained clinical settings. The key contribution of this study is the development of a privacy-preserving FLR model to diagnose acute bladder inflammation. Implementation using the FedAvg algorithm ensures scalability and efficiency across distributed nodes, while maintaining diagnostic accuracy comparable to centralized models with fewer training iterations.

The proposed FLR model has several advantages. Due to its low computational requirements, the model is appropriate for processing-limited hospitals. Due to the dependency on interpretable LR, the model can be easily validated and understood by physicians. Additionally, due to the centralization of patient data, which is only distributed for model updates, the model can ensure complete data confidentiality and compliance with healthcare regulations.

A limitation of this study is that it used a relatively small, highly structured dataset, which does not necessarily capture the complete variability of clinical cases from the real world. In addition, the simulation was performed under a controlled environment where the hardware and software settings were homogeneous, which might not best represent the diversity of real-world hospital systems. In addition, it is assumed that the data distribution remains stable at the client level, which might not be the case for federated settings that include diverse demographics or clinical practices.

REFERENCES

[1] F. Del Ben *et al.*, "A fully interpretable machine learning model for increasing the effectiveness of urine screening," *American Journal of Clinical Pathology*, vol. 160, no. 6, pp. 620–632, Dec. 2023, https://doi.org/10.1093/ajcp/aqad099.

[2] R. Magherini, E. Mussi, Y. Volpe, R. Furferi, F. Buonamici, and M. Servi, "Machine Learning for Renal Pathologies: An Updated Survey," *Sensors*, vol. 22, no. 13, Jul. 2022, Art. no. 4989, https://doi.org/10.3390/s22134989.

[3] J. Czerniak and H. Zarzycki, "Application of rough sets in the presumptive diagnosis of urinary system diseases," in *Artificial Intelligence and Security in Computing Systems*, Boston, MA, USA, 2003, pp. 41–51, https://doi.org/10.1007/978-1-4419-9226-0_5.

[4] S. Farashi and H. E. Momtaz, "Prediction of urinary tract infection using machine learning methods: a study for finding the most-informative variables," *BMC Medical Informatics and Decision Making*, vol. 25, no. 1, Jan. 2025, Art. no. 13, https://doi.org/10.1186/s12911-024-02819-2.

[5] M. K. S. Ma'aitah, A. Helwan, and A. Radwan, "Urinary Bladder Acute Inflammations and Nephritis of the Renal Pelvis: Diagnosis Using Fine-Tuned Large Language Models," *Journal of Personalized Medicine*, vol. 15, no. 2, Jan. 2025, Art. no. 45, https://doi.org/10.3390/jpm15020045.

[6] S. Li, X. Li, Y. Wang, J. Shen, G. U. Nneji, and H. N. Monday, "Dependable AI Machine Learning Models for the Prediction of Urinary System Diseases," in *2024 17th International Conference on Advanced Computer Theory and Engineering (ICACTE)*, Hefei, China, Sep. 2024, pp. 366–370, https://doi.org/10.1109/ICACTE62428.2024.10871567.

[7] L. E. Lamb *et al.*, "Risk Classification for Interstitial Cystitis/Bladder Pain Syndrome Using Machine Learning Based Predictions," *Urology*, vol. 189, pp. 19–26, Jul. 2024, https://doi.org/10.1016/j.urology.2024.03.043.

[8] S. Baijwan and A. Dhyani, "Performance Analysis of an Inflammatory Process in the Bladder Cystitis Development using Machine Learning Approach," in *2023 IEEE International Conference on Integrated Circuits and Communication Systems (ICICACS)*, Raichur, India, Feb. 2023, pp. 01–06, https://doi.org/10.1109/ICICACS57338.2023.10099872.

[9] M. A. Mohammed, M. K. A. Ghani, A. Lakhan, B. AL-Attar, and W. Khaled, "Federated Learning-Driven IoT and Edge Cloud Networks for Smart Wheelchair Systems in Assistive Robotics," *Iraqi Journal for Computer Science and Mathematics*, vol. 6, no. 1, Mar. 2025, https://doi.org/10.52866/2788-7421.1241.

[10] A. Lakhan *et al.*, "FDCNN-AS: Federated deep convolutional neural network Alzheimer detection schemes for different age groups," *Information Sciences*, vol. 677, Aug. 2024, Art. no. 120833, https://doi.org/10.1016/j.ins.2024.120833.

[11] M. A. Mohammed *et al.*, "Federated-Reinforcement Learning-Assisted IoT Consumers System for Kidney Disease Images," *IEEE Transactions on Consumer Electronics*, vol. 70, no. 4, pp. 7163–7173, Nov. 2024, https://doi.org/10.1109/TCE.2024.3384455.

[12] J. Czerniak, "Acute Inflammations." UCI Machine Learning Repository, 2003, https://doi.org/10.24432/C5V59S.

[13] T. Haifley, "Linear logistic regression: an introduction," in *IEEE International Integrated Reliability Workshop Final Report, 2002.*, Lake Tahoe, CA, USA, 2002, pp. 184–187, https://doi.org/10.1109/IRWS.2002.1194264.

[14] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, "Federated Learning: Challenges, Methods, and Future Directions," *IEEE Signal Processing Magazine*, vol. 37, no. 3, pp. 50–60, May 2020, https://doi.org/10.1109/MSP.2020.2975749.

[15] L. Guo, M. Li, S. Xu, and F. Yang, "Application of Stochastic Gradient Descent Technique for Method of Moments," in *2020 IEEE International Conference on Computational Electromagnetics (ICCEM)*, Singapore, Aug. 2020, pp. 97–98, https://doi.org/10.1109/ICCEM47450.2020.9219400.

[16] S. Gunter and H. Bunke, "Optimizing the number of states, training iterations and Gaussians in an HMM-based handwritten word recognizer," in *Seventh International Conference on Document Analysis and Recognition, 2003, Proceedings.*, Edinburgh, UK, 2003, vol. 1, pp. 472–476, https://doi.org/10.1109/ICDAR.2003.1227710.

[17] U. Ruby, P. Theerthagiri, I. J. Jacob, and Y. Vamsidhar, "Binary cross entropy with deep learning technique for Image classification," *International Journal of Advanced Trends in Computer Science and Engineering*, vol. 9, no. 4, pp. 5393–5397, Aug. 2020, https://doi.org/10.30534/ijatcse/2020/175942020.