Governance-Aware Privacy-Preserving AI Infrastructure for Ghana's Maternal Health Ecosystem

Abstract

The Maternal health data in Ghana is sensitive and difficult to share because of privacy and governance concerns. This study proposes a governance-aware privacy-preserving AI infrastructure using federated learning and differential privacy for maternal health risk prediction. A neural network model was trained using centralized learning, federated learning, and federated learning with Gaussian noise. The centralized model achieved 68.97% accuracy. The federated model improved to 70.94% after 20 rounds. When differential privacy was added, the accuracy was 68.47%, showing only a small performance drop. High-risk cases were still well detected under privacy protection. These findings are consistent with existing healthcare federated learning studies. This research connects data governance, privacy, AI, and healthcare data infrastructure, and shows how Ghana can enable collaborative health analytics without sharing raw patient data.

keywords: Federated Learning, Differential Privacy, Data Governance, Maternal Health, AI Infrastructure

Introduction

Maternal health data is sensitive and difficult to share across hospitals because of privacy and governance concerns. In decentralized healthcare systems, data are stored in separate institutions, which limits centralized machine learning. A systematic review by Alam et al. (2025) explains that data silos and regulatory restrictions are major barriers to collaborative healthcare AI. Federated learning addresses this problem by allowing hospitals to train a shared model without transferring raw patient data (Rieke et al., 2022). However, federated learning alone may still expose sensitive information through model updates.

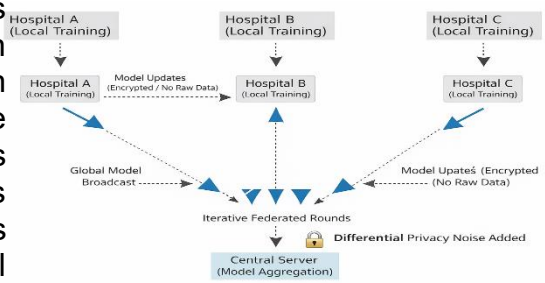

Figure 1: Privacy-Preserving Federated Learning Workflow

To improve privacy protection, differential privacy adds controlled noise during training. Studies using clinical datasets show that federated learning combined with differential privacy can maintain useful performance while reducing privacy risks (Nguyen et al., 2023). In this study, we design a governance-aware privacy-preserving AI infrastructure for maternal health risk prediction in Ghana. We compare centralized learning, federated learning, and federated learning with differential privacy to evaluate the performance–privacy trade-off in a distributed healthcare setting.

Methods

Experimental Setup

This study implemented a neural network model for maternal health risk classification using three training approaches: centralized learning, federated learning, and federated learning with differential privacy. The dataset was divided into training and testing sets using standard preprocessing steps in Python and PyTorch. Features were normalized before training.

```
print("Training shape:", X_train.shape)
print("Test shape:", X_test.shape)
```

```
Training shape: (811, 6)
Test shape: (203, 6)
```

[6] ✓ 7s

```
import torch
import torch.nn as nn
```

The neural network consisted of an input layer, hidden layers with ReLU activation function, and an output layer for multi-class classification. The model was trained using cross-entropy loss and the Adam optimizer.For centralized learning, all training data were combined and used to train one global model.

For federated learning, the training data were split into three simulated hospital clients. Each client trained the model locally, and the central server aggregated model weights using federated averaging (FedAvg), which is a standard federated learning approach (Rieke et al., 2022). The training was performed for 20 communication rounds.

```
    results[r] = acc
    print(f"Rounds: {r}, Accuracy: {acc:.4f}")

Rounds: 1, Accuracy: 0.5616
Rounds: 5, Accuracy: 0.6798
Rounds: 10, Accuracy: 0.6798
Rounds: 20, Accuracy: 0.7094
```

To enhance privacy, differential privacy was introduced by adding Gaussian noise to model updates during federated training. Prior healthcare studies show that adding noise to gradients can reduce privacy risks while maintaining acceptable model performance (Nguyen et al., 2023)

```
import matplotlib.pyplot as plt

def add_gaussian_noise(model, sigma=0.01):
    noisy_model = MaternalNN()
    noisy_model.load_state_dict(model.state_dict())

    for param in noisy_model.parameters():
        noise = torch.normal(0, sigma, size=param.size())
        param.data += noise

    return noisy_model
```

```
private_model, private_accuracy = run_federated_private(rounds=20, sigma=0.01)
print(f"Federated + Differential Privacy Accuracy: {private_accuracy:.4f}")

Federated + Differential Privacy Accuracy: 0.6897
```

Model performance was evaluated using accuracy, precision, recall, F1-score, and confusion matrices. These metrics were used to compare centralized, federated, and privacy-preserving federated models and to analyze the privacy–utility trade-off.

Results

The performance of the three models was evaluated using accuracy and class-level metrics. The centralized model achieved 68.97% accuracy. The federated model improved performance to 70.94% after 20 communication rounds. When differential privacy was introduced using Gaussian noise, the model achieved 68.47% accuracy. This shows that federated learning improved performance compared to centralized learning, while differential privacy caused only a small reduction in accuracy.

Parameter	Centralized	Federated (20 Rounds)	Federated + DP
Accuracy	68.97%	70.94%	68.47%
Macro F1-Score	0.69	0.68	0.70
High-Risk Recall	0.89	0.85	0.81

Table 1: Performance Comparison of Centralized, Federated, and Privacy-Preserving Federated Models

Class-level evaluation showed that high-risk cases maintained strong recall even after adding privacy noise. The federated model also showed stable convergence across communication rounds. These results demonstrate that privacy-preserving federated learning can maintain useful performance while improving data governance and privacy protection.

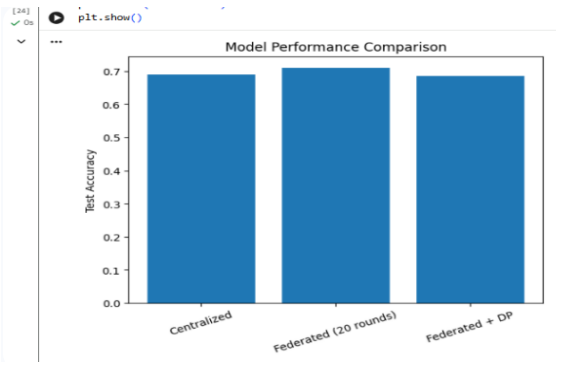

Figure 2: Model Accuracy Comparison Bar Chart.

This shows that federated learning achieved the highest accuracy, while privacy-preserving federated learning maintained comparable performance.

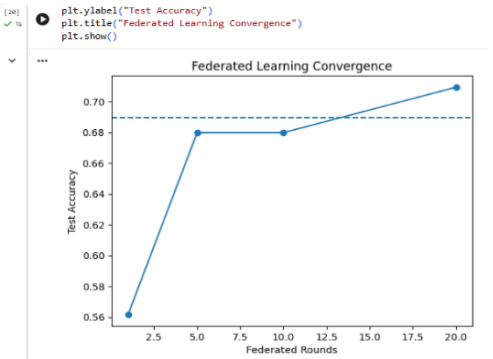

Figure 3: Federated Learning Convergence Curve (20 Rounds)

This shows that model performance improved gradually across communication rounds and stabilized near 70% accuracy.

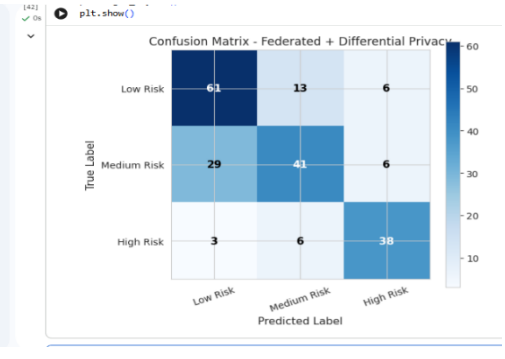

Figure 4: Confusion Matrix for Federated + Differential Privacy Model

This shows that high-risk cases were still correctly identified even after adding privacy noise.

These results show the performance–privacy trade-off in a distributed healthcare setting. When differential privacy was added, accuracy reduced slightly from 70.94% to 68.47%, showing that privacy protection comes with a small performance cost (i.e that small drop = cost of privacy protection.) which is about 2.47% from federated.

Discussion and Conclusion

The aim of this study was to design a governance-aware and privacy-preserving AI system for maternal health risk prediction in Ghana. We compared centralized learning, federated learning, and federated learning with differential privacy. The results show that federated learning improved the model accuracy from 68.97% to 70.94% after 20 communication rounds. When differential privacy was added, the accuracy reduced slightly to 68.47%. This confirms the performance–privacy trade-off. Stronger privacy protection slightly reduces model accuracy, but the drop was small. This means privacy can be improved without losing much predictive performance in Ghana's healthcare system.

From the confusion matrix, high-risk cases were still correctly detected even after adding Gaussian noise. This is very important in healthcare because high-risk patients must not be missed. The results are similar to previous studies which show that federated learning with differential privacy can protect patient data while still giving useful results (Nguyen et al., 2023; Rieke et al., 2022).

This study shows that Ghana can use federated AI systems to allow hospitals to work together without sharing raw patient data. This supports data governance, privacy, and responsible AI in healthcare.

However, this study has some limitations. The hospitals were simulated and not real hospitals. The dataset was not very large. Only one model type (neural network architecture) was tested. Also, different privacy levels were not tested in detail.

In future work, real hospital data in Ghana can be used. Different privacy settings can also be tested. More advanced models and secure aggregation methods can also be explored. This will help improve the AI system for national healthcare use.

References

[1] Jha, K. J., Ameta, G. K., Panchal, E., & Jani, K. A. Differential privacy-enhanced federated learning in medical data environments. International Journal of Environmental Sciences, 11(20), 2025.

[2] Sathish Kumar, K. A., Nelson, L., & Jibinsingh, B. R. Systematic review of privacy-preserving federated learning in decentralized healthcare systems. Franklin Open, 13, 100440, 2025.

[3] Fares, M. H., & Saad, A. M. S. E. Towards privacy-preserving medical imaging: Federated learning with differential privacy and secure aggregation using a modified ResNet architecture. arXiv preprint arXiv:2412.00687v1, 2024.

[4] Horvath, A. N., Berchier, M., Nooralahzadeh, F., Allam, A., & Krauthammer, M. Exploratory analysis of federated learning methods with differential privacy on MIMIC-III. arXiv preprint arXiv:2302.04208v1, 2023.

[5] Adnan, M., Karla, S., Cresswell, J. C., Taylor, G. W., & Tizhoosh, H. R. Federated learning and differential privacy for medical image analysis. Scientific Reports, 12, 1953, 2022.