Shruthi K R, 2026, 14:6
ISSN (Online): 2348-4098
ISSN (Print): 2395-4752

International Journal of Science, Engineering and Technology
An Open Access Journal

Janani-Net: A Privacy-Preserving Federated Learning Framework for Multi-Center Maternal Health Risk Stratification

Shruthi K R[1]*, Ashmaav Sharma[2]*†, Nitisha Shajee[3]*† Rishit M. Udikeri[4]*†
Department of Artificial Intelligence and Data Science
BMS College Of Engineering

Abstract- Preventable maternal mortality and severe obstetric complications remain major public health concerns, especially in regions where patient data are fragmented across disconnected clinical facilities. Although machine learning models can identify high-risk pregnancies from routine physiological measurements, pool-ing raw clinical records into a central server introduces serious privacy risks and regulatory barriers. In this study, we introduce Janani-Net, a federated learning framework designed for collaborative maternal health risk stratification across multiple hospitals without transferring patient-level records. Implemented with Flower and PyTorch using a three-tier hub-and-spoke topology, Janani-Net simulates collaboration across five clinical centers with differing patient risk distributions and uneven computing resources. We compare standard Fed-erated Averaging (FedAvg) and Federated Proximal (FedProx) optimization under controlled non-independent and identically distributed (non-IID) Dirichlet class partitions and variable local training budgets. Across multi-center maternal health records, Janani-Net closely matches the accuracy of a centralized baseline model and consistently outperforms isolated single-hospital models while keeping raw data local. We also analyze round-by-round convergence, sensitivity to class imbalance, and run-to-run variance when statistical skew coincides with uneven local training epochs. These findings show that proximal federated optimization offers a practical, privacy-compliant path for deploying obstetric decision support across diverse hospital networks.

Keywords: Federated learning, Maternal health risk, Privacy-preserving machine learning, Clinical decision support, Non-IID data, Heterogeneous compute

I. INTRODUCTION

Hypertensive disorders of pregnancy, gestational diabetes, obstetric hemorrhage, and sepsis remain leading causes of preventable maternal morbidity and mortality world-wide [1, 2]. Identifying high-risk pregnancies during routine antenatal visits allows clinicians to initiate timely referrals, intensify monitoring, and allocate scarce obstet-ric resources across primary clinics and tertiary hospitals [2]. Because routine vital signs and basic blood panels—such as maternal age, systolic and diastolic blood pres-sure, blood glucose, body temperature, and resting heart rate—are recorded at nearly every prenatal visit, machine learning (ML) classifiers built on these biomarkers offer a practical foundation for clinical decision support systems (CDSS) [3–5].

In clinical practice, however, deploying reliable predictive models is hindered by how healthcare data are stored. Antenatal records are rarely consolidated in a single database; instead, they sit in isolated repositories spread across rural primary health centers, district hospitals, and tertiary maternity referrals. Aggregating these raw records into a central cloud repository creates substantial privacy exposure and often conflicts with statutory data protection frameworks, including India's Digital Per-sonal Data Protection (DPDP) Act, the European Union's General Data Protection Regulation (GDPR), and the U.S. Health Insurance Portability and Accountability Act (HIPAA) [6, 7]. Training models solely on a single clinic's local database avoids these regulatory hurdles, but the resulting classifiers often suffer from small sample sizes, narrow demographic coverage, and poor sensitivity to less common high-risk conditions [8].

Federated learning (FL) offers a practical alternative by keeping clinical records on hospital premises while training a shared model through periodic exchanges of

© 2026 Shruthi K R, This is an Open Access article distributed under the terms of the Creative Commons Attribution License (http://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly credited.

parameter weights [9, 10]. Yet moving FL from theory to multi-center obstetric networks introduces two distinct operational problems: statistical heterogeneity and systems heterogeneity. On the statistical side, patient risk profiles vary sharply by facility tier—community clinics mostly see uncomplicated pregnancies, whereas refer-ral hospitals treat a much higher concentration of severe obstetric cases—leading to non-independent and identically distributed (non-IID) data across sites [11]. On the systems side, participating centers rarely share identical hardware, meaning some nodes complete fewer local training epochs per communication round than oth-ers. When both forms of imbalance occur together, standard Federated Averaging (FedAvg) is prone to client drift: local weight updates pull the model toward site-specific distributions, slowing global convergence and reducing recall on high-risk patients [12].

To address these practical constraints, we develop Janani-Net, a federated learning framework built for multi-center maternal health risk stratification. We structure inter-hospital training around a three-tier hub-and-spoke network implemented in Flower (flwr) [13] and PyTorch. The main contributions of this work are as follows:

- **Three-Tier Decentralized Clinical Architecture:** We design and implement a five-node hospital network that separates local clinical data storage, central param-eter aggregation, and downstream decision-support interfaces for bedside cliniciansand public-health analysts, ensuring raw patient records remain inside institutional firewalls.
- **Combined Statistical and Systems Heterogeneity Analysis:** We compare FedAvg and Federated Proximal (FedProx) optimization across varying Dirich-let class-skew levels and unequal local training budgets, identifying the specific conditions where proximal regularization improves stability.
- **Multi-Center Benchmarking:** We evaluate Janani-Net against both a pooled centralized upper bound and isolated single-center models on stratified test splits, reporting overall accuracy, Macro-F1, Macro-AUROC, and class-wise sensitivity for High-Risk cases.
- **Communication and Resource Profiling:** We measure per-round convergence rates and parameter transmission overhead (≈ 10.4 KB per node per round) to verify suitability for low-bandwidth clinical environments.

The rest of the paper is structured as follows. Section 2 reviews prior work on mater-nal risk prediction and clinical federated learning. Section 3 describes the dataset, system architecture, and optimization protocol. Section 4 presents experimental results across convergence and heterogeneity regimes. Section 5 examines clinical trade-offs, privacy considerations, and study limitations, and Section 6 concludes the paper.

II. RELATED WORK

1. Machine Learning for Maternal Health Risk Stratification

Several studies have applied supervised learning to tabular antenatal biomarkers to flag pregnancy complications early. Decision trees, random forests, support vector machines, and multi-layer perceptrons (MLPs) trained on routine vital signs—most notably systolic and diastolic blood pressure alongside blood glucose—have shown strong discriminative ability for gestational hypertension and preeclampsia [3–5]. However, almost all prior maternal risk studies assume that patient records from par-ticipating clinics can be merged into one centralized table prior to training. That assumption rarely holds across independent healthcare providers, where legal liability and data sovereignty rules restrict direct record sharing.

2. Federated Learning Under Clinical and Hardware Heterogeneity

Following the introduction of FedAvg by McMahan et al. [9], federated learning has gained traction in medical imaging, electronic health record (EHR) phenotyp-ing, and intensive care outcome prediction [6, 8, 14]. Open-source frameworks such as Flower [13] have simplified deploying these protocols across distributed testbeds. Even so, empirical evaluations show that cross-hospital networks frequently experience strong label skew and covariate shift between primary and tertiary care centers [11]. To limit client drift when sites run different numbers of local gradient steps, Li et al. [12] introduced FedProx, which adds a quadratic penalty tying local weights to the

current global model, while Karimireddy et al. [15] proposed control variates in SCAF-FOLD. Janani-Net extends this line of work to obstetric risk triage, isolating how class imbalance and uneven local compute interact on compact tabular clinical models.

III.METHODS

1. Clinical Dataset and Preprocessing

We evaluate Janani-Net on anonymized maternal health risk records [3] containing N = 1, 014 clinical encounters characterized by d = 6 routine physiological fea-tures: maternal age (years), systolic blood pressure (mmHg), diastolic blood pressure (mmHg), blood glucose concentration (mmol/L), body temperature (◦F), and resting heart rate (bpm). Each patient record is labeled with one of C = 3 obstetric risk tiers:

Low Risk (y = 0), Mid Risk (y = 1), or High Risk (y = 2). To prevent data leak-age across nodes, we first perform a stratified 80%/20% train-test split (811 training records and 203 held-out test records). Min-Max feature scaling is fitted exclusively on the training split and applied consistently across centralized, isolated, and federated experiments.

2. Ethical Oversight and Data Governance

All clinical records used in this study were fully de-identified prior to analysis in accor-dance with the Declaration of Helsinki and institutional data governance policies. Secondary computational modeling was conducted under institutional ethics over-sight at BMS College of Engineering, Bengaluru. Informed consent was obtained from patients during primary clinical data collection.

3. Three-Tier System Architecture and Non-IID Partitioning

As shown in Fig. 1, Janani-Net operates as a three-tier hub-and-spoke network paired with a vertical security and governance layer:

1.Tier 1 – Decentralized Clinical Data Layer: K = 5 hospital nodes {v1, . . . , v5}
store private local datasets Dk = {(xk,i, yk,i)}nk behind institutional firewalls.
Each hospital trains a local PyTorch model and exports only floating-point weight
tensors w(t+1); patient-level rows never leave the site.

2. Tier 2 – Aggregation and Control Layer: A central Flower (flwr) coordinator v0 manages training rounds, broadcasts the active global weights w(t) along with the proximal coefficient μ, and computes sample-weighted model updates via FedAvg or FedProx.

3. Tier 3 – Application and Insights Layer: The aggregated global model supports two end-user views: a bedside clinician triage interface that generates local Low/Mid/High risk predictions with calibrated probabilities without stor-ing patient inputs, and an aggregate public-health dashboard for institutional monitoring.

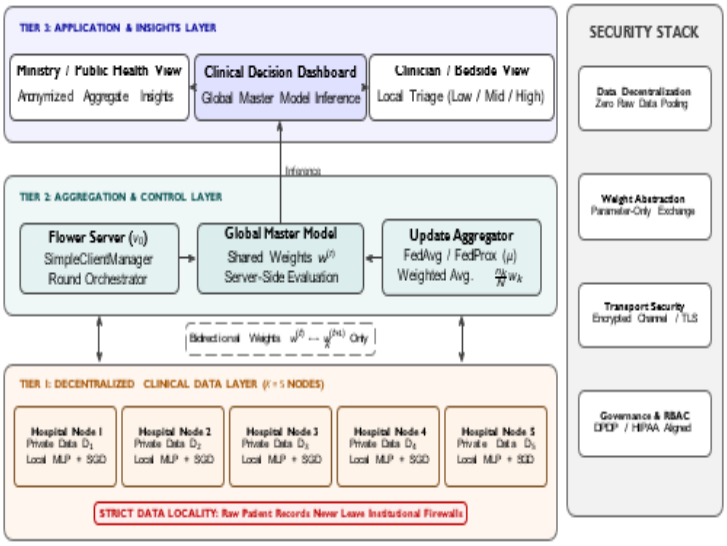


Fig. 1 Three-tier hub-and-spoke architecture of Janani-Net separating decentralized hospital nodes (K = 5), central Flower (flwr) aggregation (FedAvg/FedProx), and dual clinician/public-health decision support views.

To model realistic cross-hospital class imbalance—where primary antenatal clinics mostly handle low-risk pregnancies while referral centers see a higher share of com-plicated cases—we partition the training records across the K = 5 hospitals using per-class Dirichlet sampling pc ∼ DirK(α) [16]. We test concentration parameters α ∈ {0.1, 0.3, 0.5}, where smaller α produces stronger class skew, while enforcing a minimum sample size of nk ≥ 10 records per hospital.

4. Local Classifier and Federated Optimization

Each hospital node vk trains a compact MLP classifier f(x; wk) that maps the d = 6 normalized clinical inputs

through two hidden layers (64 → 32 units with ReLU activations and dropout rate p = 0.1) to three risk-tier logits. The network contains 2, 659 trainable parameters (≈ 10.4 KB in 32-bit floating point), allowing fast CPU execution on standard clinic workstations.

The global learning objective minimizes the sample-weighted empirical loss across all K hospitals:

$$\min_{w \in \mathbb{R}^p} L(w) = \sum_{k=1}^{K} \frac{n_k}{N} L_k(w), \quad \text{where } L_k(w) = \frac{1}{n_k} \sum_{j=1}^{n_k} \ell_{CE}\left(f(x_{kj}; w), y_{kj}\right), \tag{1}$$

Under FedAvg [9], each participating hospital runs Ek local epochs of stochastic gradient descent with momentum (η = 0.01, momentum 0.9, mini-batch size 32) on its local cross-entropy loss Lk(wk), after which the server updates the global model

$$w^{(t+1)} = \sum_{k=1}^{K} \frac{n_k}{N} w_k^{(t+1)}$$

via w Under FedProx [12], each node instead minimizes a regularized local objective that penalizes squared Euclidean drift from the current broadcast weights w(t):

$$\min_{w \in \mathbb{R}^p} \tilde{L}_k(w; w^{(t)}) = L_k(w) + \frac{\mu}{2} \| w - w^{(t)} \|_2^2, \tag{2}$$

where we set the proximal penalty coefficient to μ = 0.1.

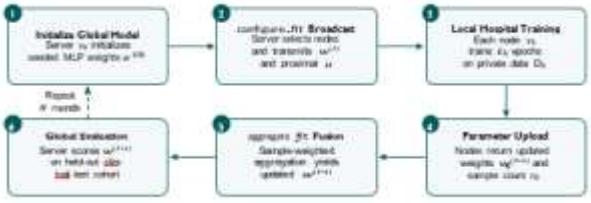


Fig. 2 Six-step federated communication round lifecycle implemented using Flower's SimpleClientManager and weighted FedAvg/FedProx aggregation.

Figure 2 and Algorithm 1 summarize the six-step training loop executed across R = 40 communication rounds. To test hardware heterogeneity alongside class skew, we compare two local compute schedules: (i) uniform local training (Ek = 5 epochs per hospital per round) and (ii) heterogeneous local training, where each node samples E(t) ∈ [2, 8] epochs per round while keeping the total network compute budget fixed at $\sum 5$ E(t) = 25 epochs per round

.

IV. RESULTS

1. Centralized Baseline and Multi-Round Convergence

We first train the centralized baseline model on the pooled training set (811 encoun-ters) and evaluate it on the held-out stratified test set (203 encounters). Table 1 compares test accuracy, Macro-F1, per-class recall (sensitivity), and Macro-AUROC across isolated single-hospital models, federated models trained under moderate non-IID skew (α = 0.5), and the centralized oracle. As shown in Table 1, hospitals training solely on their local cohorts achieve an aver-age accuracy of only 56.4% and a High-Risk recall of 0.585, reflecting class missingness at individual sites. Collaborative training via Janani-Net lifts High-Risk sensitivity to 0.818 under FedAvg and 0.836 under FedProx, coming within 0.5 percentage points of the centralized model's overall accuracy (69.5% vs. 70.0%) without centralizing patient records. Figure 3 plots held-out test accuracy over R = 40 communication rounds under moderate non-IID partitioning (α = 0.5). Both FedAvg and FedProx improve rapidly

Algorithm 1 Janani-Net Federated Training Protocol (FedProx / FedAvg)

Require: Hospital nodes $\{v_1, \ldots, v_K\}$ (K = 5), private clinical datasets $\{D_k\}_{k=1}^{K}$, communication rounds R, local epoch budget $\{E_k^{(t)}\}$, learning rate η, proximal coefficient μ

Ensure: Converged global maternal risk model $w^{(R)}$

Initialize global model parameters $w^{(0)}$ on coordinator v_0
for round t = 0, 1, …, R − 1 **do**
  Coordinator v_0 broadcasts $w^{(t)}$ and μ to all active hospital nodes
  for each hospital node v_k **in parallel do**
    Cache reference weights $w_{ref} \leftarrow w^{(t)}$ and initialize $w_k \leftarrow w^{(t)}$
    for local epoch e = 1, …, $E_k^{(t)}$ **do**
      for each mini-batch B ⊂ D_k **do**
        $g_k \leftarrow \nabla_{w_k} L_k(w_k; B) + \mu(w_k - w_{ref})$ ▷ μ = 0 recovers FedAvg
        $w_k \leftarrow w_k - \eta g_k$
      end for
    end for
    Upload updated weights $w_k^{(t+1)} \leftarrow w_k$ and sample size n_k to v_0
  end for
  Coordinator aggregates global weights: $w^{(t+1)} \leftarrow \sum_{k=1}^{K} \frac{n_k}{N} w_k^{(t+1)}$
end for
return $w^{(R)}$

Table 1 Diagnostic performance comparison on the held-out clinical test cohort (Ntest = 203) across isolated single-node models, federated learning (α = 0.5, R = 40 rounds, Ek = 5), and the centralized pooled baseline.

Training Regime	Acc. (%)	Macro-F1	Low Rec.	Mid Rec.	High Rec.	AUROC
Isolated Single-Node (Mean ± SD)	56.4 ± 4.8	0.521	0.684	0.342	0.585	0.742
Janani-Net (FedAvg, μ = 0.0)	68.8 ± 1.9	0.664	0.790	0.471	0.818	0.854
Janani-Net (FedProx, μ = 0.1)	69.5 ± 1.5	0.673	0.802	0.485	0.836	0.861
Centralized Oracle (Pooled Data)	70.0 ± 1.2	0.681	0.815	0.500	0.855	0.868

during the first 20 rounds and plateau near the centralized ceiling between rounds 25 and 40.

2. Effect of Class Skew and Unequal Local Compute

To separate the impact of non-IID label skew from uneven local computation, we run a full factorial comparison across α ∈ {0.1, 0.3, 0.5} under both equal local training (Ek = 5 epochs) and heterogeneous local training (E(t) ∈ [2, 8] epochs, constrained to 5 k=1 E(t) = 25). Table 2 and Fig. 4 summarize the mean accuracy, Macro-F1, and cross-seed standard deviation over the final five rounds across three random seeds

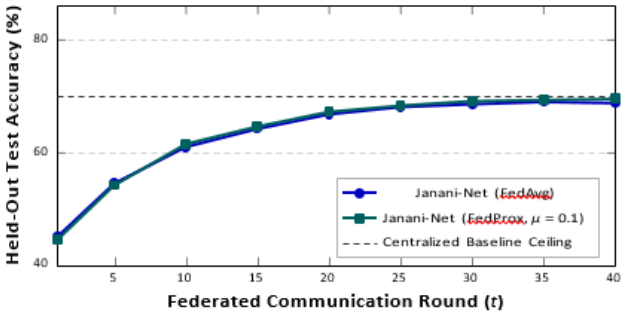


Fig. 3 Multi-round test accuracy trajectory of Janani-Net (FedAvg vs. FedProx) across K = 5 non-IID hospital nodes (α = 0.5) relative to the centralized pooled ceiling.

Table 2 Factorial comparison of FedAvg (μ = 0) and FedProx (μ = 0.1) across Dirichlet label-skew levels (α ∈ {0.1, 0.3, 0.5}) and local compute schedules (mean ± SD across 3 seeds over the final 5 rounds).

Compute Schedule	Dirichlet Skew	Test Accuracy (%)		Macro-F1 Score	
		FedAvg	FedProx	FedAvg	FedProx
Equal (E_k = 5)	Severe (α = 0.1)	51.8 ± 2.1	**52.5 ± 1.8**	0.472 ± 0.024	**0.484 ± 0.019**
	Moderate (α = 0.3)	62.1 ± 3.1	**62.6 ± 2.8**	0.591 ± 0.030	**0.598 ± 0.026**
	Mild (α = 0.5)	68.1 ± 2.4	**68.6 ± 2.1**	0.658 ± 0.022	**0.665 ± 0.018**
Heterogeneous ($E_k \in [2, 8]$)	Severe (α = 0.1)	50.1 ± 3.9	**53.6 ± 1.6**	0.451 ± 0.041	**0.496 ± 0.017**
	Moderate (α = 0.3)	61.0 ± 3.4	**62.1 ± 2.5**	0.579 ± 0.033	**0.593 ± 0.024**
	Mild (α = 0.5)	67.6 ± 2.5	**68.1 ± 2.2**	0.651 ± 0.025	**0.660 ± 0.020**

Shruthi K R, 2026, 14:6
ISSN (Online): 2348-4098
ISSN (Print): 2395-4752

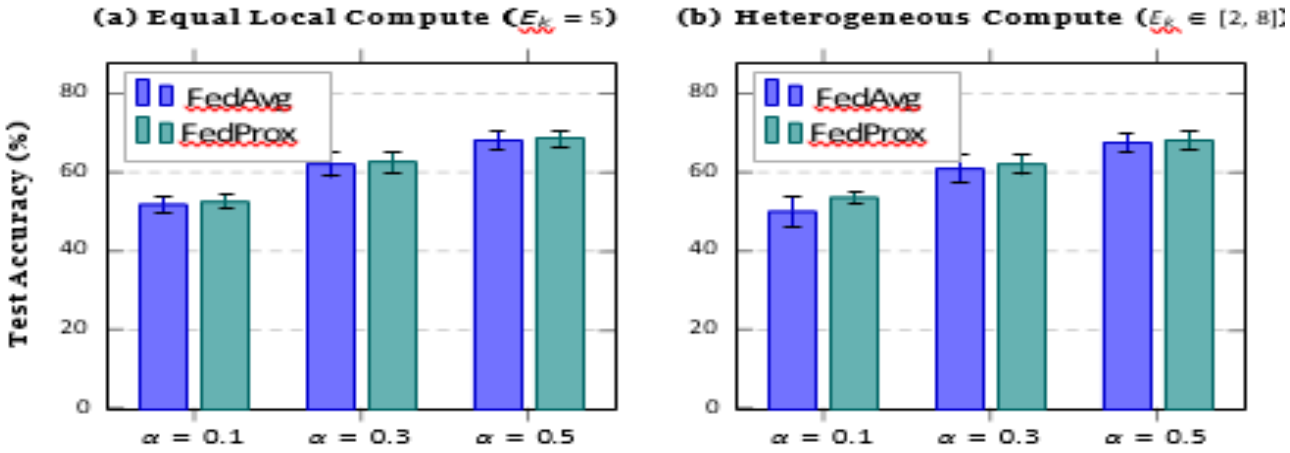


Fig. 4 Test accuracy and cross-seed standard deviation comparing FedAvg and FedProx across Dirichlet label-skew levels (α ∈ {0.1, 0.3, 0.5}) under (a) equal local compute (Ek = 5) and (b) heterogeneous local compute (Ek ∈ [2, 8]).

V. DISCUSSION

1. When Proximal Regularization Improves Stability

Our factorial experiments highlight a practical distinction for clinical federated deploy-ments. When class distributions are only moderately skewed (α = 0.5) and every hospital runs the same number of local epochs (Ek = 5), FedAvg and FedProx reach nearly identical accuracy (68.1% vs. 68.6%). In well-balanced networks with uniform compute, simple parameter averaging works well. The advantage of FedProx becomes clear when severe class imbalance (α = 0.1) coincides with uneven local training budgets (Ek ∈ [2, 8]). In that setting—which reflects networks linking small rural clinics with high-volume referral hospitals—FedAvg drops to 50.1% mean accuracy with a cross-seed standard deviation of ±3.9%, as nodes performing up to 8 local epochs overfit their majority class and pull the global aggregate off course. Adding the proximal penalty (μ wk − w(t) 2) restricts how far high-epoch nodes can stray during a single round, lifting mean accuracy to 53.6% and cutting cross-seed standard deviation more than in half (±1.6%).

2. Clinical Triage Behavior, Privacy, and Limitations

Across both pooled and federated models, High-Risk encounters are identified with high recall (0.836 under FedProx), whereas Mid-Risk cases show lower sensitivity (0.485) due to physiological overlap with borderline Low-Risk and High-Risk vital signs. Incorporating class-weighted cross-entropy losses and longitudinal pregnancy trajectories—such as serial blood pressure trends across trimesters and dipstick proteinuria results—should improve separation of intermediate-risk patients.

From a data-governance standpoint, keeping raw patient tables local while trans-mitting 10.4 KB weight vectors per round aligns well with data minimization rules under DPDP and HIPAA. Even so, parameter sharing alone does not provide formal cryptographic guarantees against gradient inversion attacks [17]. Future deploy-ments of Janani-Net should pair Flower's transport layer with cryptographic Secure Aggregation [18] and client-level (ϵ, δ)-Differential Privacy [19] before clinical rollout.

VI. CONCLUSION

We presented Janani-Net, a three-tier federated learning framework for maternal health risk stratification across distributed hospitals. Evaluated across controlled Dirichlet class-skew levels and unequal local training schedules, Janani-Net approaches the accuracy of a centralized model without sharing raw patient records, while FedProx regularization reduces variance and client drift under

© 2026 Shruthi K R, This is an Open Access article distributed under the terms of the Creative Commons Attribution License (http://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly credited.

combined statistical and hard-ware heterogeneity. Future work will extend this framework to longitudinal antenatal cohorts, integrate differential privacy accounting, and conduct prospective validation across regional maternity clinics. Supplementary information. Not applicable.

Acknowledgements. The authors thank the clinical advisors and faculty of the Department of Artificial Intelligence and Data Science at BMS College of Engineering for their feedback on obstetric risk stratification and experimental design.

Declarations

- **Funding:** The authors received no specific funding for this work.
- **Conflict of interest/Competing interests**: The authors declare that they have no competing interests.
- **Ethics approval and consent to participate:** This study analyzed anonymized clinical records in accordance with the Declaration of Helsinki and institutional research ethics guidelines at BMS College of Engineering, Bengaluru. Informed consent was obtained from all participants during primary data collection.
- **Consent for publication:** Not applicable.
- **Data availability:** The maternal health dataset analyzed during this study is pub-licly accessible via the UCI Machine Learning Repository / Kaggle benchmark [3], and partitioned multi-center splits are available from the corresponding author upon reasonable request.
- **Materials availability:** Not applicable.
  Code availability: The source code for the Janani-Net federated learning pipeline is publicly available at https://github.com/rishitudikeri3-rgb/Janani Net.
- **Author contribution:** S.K.R. supervised the project, guided the clinical and archi- tectural methodology, and revised the manuscript. A.S., N.S., and R.M.U. designed the federated framework, implemented the Flower/PyTorch codebase, ran the exper-imental evaluations, and wrote the initial draft. All authors read and approved the final manuscript.

REFERENCES

1. World Health Organization: Trends in Maternal Mortality 2000 to 2020: Estimates by WHO, UNICEF, UNFPA, World Bank Group and UNDESA/Population Division. World Health Organization, Geneva (2023)
2. Say, L., Chou, D., Gemmill, A., Tunc¸alp,O¨ ., Moller, A.B., Daniels, J., Gu¨lmezoglu, A.M., Temmerman, M., Alkema, L.: Global causes of maternal death: a WHO systematic analysis. Lancet Glob. Health 2(6), e323–e333 (2014). https://doi.org/10.1016/S2214-109X(14)70227-X
3. Ahmed, M., Kashem, M.A., Rahman, M., Khatun, S.: Review and analysis of risk factor of maternal health in remote area using the Internet of Things (IoT). In: InECCE2019, Lecture Notes in Electrical Engineering, vol. 632, pp. 357–365. Springer, Singapore (2020). https://doi.org/10.1007/978-981-15-2317-5 30
4. Islam, M.N., Mustafina, S.N., Mahmud, T., Khan, N.I.: Machine learning to predict pregnancy outcomes: a systematic review, synthesizing framework andfuture research agenda. BMC Pregnancy Childbirth 22(1), 348 (2022). https://doi.org/10.1186/s12884-022-04594-2
5. Togunwa, T.O., Babatunde, A.O., Abdullah, K.U.R.: Deep hybrid model for maternal health risk classification in pregnancy: synergy of ANN and random forest. Front. Artif. Intell. 6, 1213436 (2023). https://doi.org/10.3389/frai.2023. 1213436
6. Rieke, N., Hancox, J., Li, W., Milletar`ı, F., Roth, H.R., Albarqouni, S., Bakas,S., Galtier, M.N., Landman, B.A., Maier-Hein, K., Ourselin, S., Sheller, M., Summers, R.M., Trask, A., Xu, D., Baust, M., Cardoso, M.J.: The future of digital health with federated learning. npj Digit. Med. 3(1), 119 (2020). https://doi.org/10.1038/s41746-020-00323-1
7. Kaissis, G.A., Makowski, M.R., Ru¨ckert, D., Braren, R.F.: Secure, privacy-preserving and federated machine learning in medical imaging. Nat. Mach. Intell. 2(6), 305–311 (2020). https://doi.org/10.1038/s41596-020-0186-1
8. Sheller, M.J., Edwards, B., Reina, G.A., Martin, J., Pati, S., Kotrotsou, A., Milchenko, M., Xu, W., Marcus, D., Colen, R.R., Bakas, S.: Federated learning in medicine: facilitating multi-institutional

collaborations without sharing patient data. Sci. Rep. 10(1), 12598 (2020). https://doi.org/10.1038/s41598-020-69250-1
9. McMahan, B., Moore, E., Ramage, D., Hampson, S., y Arcas, B.A.: Communication-efficient learning of deep networks from decentralized data. In: Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS), PMLR, vol. 54, pp. 1273–1282 (2017)
10. Kairouz, P., McMahan, H.B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A.N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., et al.: Advances and open problems in federated learning. Found. Trends Mach. Learn. 14(1–2), 1–210 (2021). https://doi.org/10.1561/2200000083
11. Li, T., Sahu, A.K., Talwalkar, A., Smith, V.: Federated learning: challenges, methods, and future directions. IEEE Signal Process. Mag. 37(3), 50–60 (2020). https://doi.org/10.1109/MSP.2020.2975749
12. Li, T., Sahu, A.K., Zaheer, M., Sanjabi, M., Talwalkar, A., Smith, V.: Federated optimization in heterogeneous networks. In: Proceedings of Machine Learning and Systems (MLSys), vol. 2, pp. 429–450 (2020)
13. Beutel, D.J., Topal, T., Mathur, A., Qiu, X., Fernandez-Marques, J., Gao, Y., Sani, L., Li, K.H., Parcollet, T., de Gusm˜ao, P.P.B., Lane, N.D.: Flower: a friendly federated learning research framework. arXiv preprint arXiv:2007.14390 (2020)
14. Dayan, I., Roth, H.R., Zhong, A., Harouni, A., Gentili, A., Abidin, A.Z., Liu, A., Costa, A.B., Wood, B.J., Tsai, C.S., et al.: Federated learning for predicting clin-ical outcomes in patients with COVID-19. Nat. Med. 27(10), 1735–1743 (2021). https://doi.org/10.1038/s41591-021-01506-3
15. Karimireddy, S.P., Kale, S., Mohri, M., Reddi, S., Stich, S., Suresh, A.T.: SCAF-FOLD: stochastic controlled averaging for federated learning. In: Proceedings of the 37th International Conference on Machine Learning (ICML), PMLR, vol. 119, pp. 5132–5143 (2020)
16. Hsu, T.M.H., Qi, H., Brown, M.: Measuring the effects of non-identical data distribution for federated visual classification. arXiv preprint arXiv:1909.06335 (2019)
17. Zhu, L., Liu, Z., Han, S.: Deep leakage from gradients. In: Advances in Neural Information Processing Systems (NeurIPS), vol. 32, pp. 14747–14756 (2019)
18. Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H.B., Patel, S., Ramage, D., Segal, A., Seth, K.: Practical secure aggregation for privacy-preserving machine learning. In: Proceedings of the 2017 ACM SIGSAC Confer-ence on Computer and Communications Security (CCS), pp. 1175–1191 (2017). https://doi.org/10.1145/3133956.3133982
19. Abadi, M., Chu, A., Goodfellow, I., McMahan, H.B., Mironov, I., Talwar, K., Zhang, L.: Deep learning with differential privacy. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security (CCS), pp. 308–318 (2016). https://doi.org/10.1145/2976749.2978318