

Privacy-Preserving Federated AI for Early Disease Risk Prediction Across Distributed U.S. Healthcare Systems

Chaitanya Nagabhyru

Abstract: Privacy law, institutional policy, and re-identification risk make it difficult to pool hospital records for early disease risk prediction. This paper presents a framework in which hospitals train a shared risk-prediction model without exchanging raw records. It combines federated averaging, record-level differential privacy (DP) enforced at each hospital, secure aggregation, and site-specific model calibration, and it is evaluated on four dimensions: predictive accuracy, privacy leakage, performance across hospital populations, and communication overhead. Experiments use a synthetic network of ten hospitals (eight training, two held out) with heterogeneous populations in a simulated three-year type 2 diabetes risk task; the results demonstrate the method and are not findings about real patients. Federated training reached a mean hospital AUC of 0.801, compared with 0.766 for local-only models and 0.811 for centralized pooling, and raised AUC in the smallest hospital from 0.703 to 0.810. With DP, AUC was 0.791 at $\epsilon = 4$ per hospital ($\delta = 10^{-5}$), fell to 0.768 at $\epsilon = 2$, and dropped below local-only performance at $\epsilon = 1$. DP degraded calibration (expected calibration error 0.082 versus 0.040 without DP), but local recalibration restored it to 0.038, including at two hospitals that never took part in training. A membership-inference attack performed no better than chance against the federated models, whereas local-only models in the smallest hospitals leaked measurably (advantage up to 13.4% against a 6.2% null). Federated traffic was 4.5% of the volume of the raw training data. The results indicate a practical budget near $\epsilon = 4$ for this task.

Introduction

Early identification of people at high risk of chronic disease is one of the most valuable uses of health data. Type 2 diabetes illustrates the scale of the opportunity: the Centers for Disease Control and Prevention estimate that about 40 million people in the United States have diabetes, that more than one in four adults with diabetes are undiagnosed, and that well over 100 million adults have prediabetes [1]. Risk models that flag patients years before diagnosis could direct preventive care to those who need it, but the data needed to build such models are spread across many health systems. Each hospital sees only one slice of the population, so a model trained in one institution often performs worse elsewhere [2], [3].

The obvious remedy, pooling records from many hospitals into a single repository, is difficult in practice. Patient records are protected by the HIPAA Privacy Rule and by institutional policy [4], hospitals are reluctant to transfer identifiable data to a third party, and removing direct identifiers does not make records anonymous: incomplete, de-identified datasets can often be re-identified [5]. Distributed research networks such as PCORnet and OHDSI run analyses where the data reside [6], [7]; federated learning extends this idea to model training.

270 sherman avenue, Apartment B3, New Haven, USA, 06511

Email: chaitanyanagabhyru@gmail.com

Keywords: Federated Learning, Differential Privacy, Early Disease Risk Prediction, Model Calibration, Secure Aggregation, Membership Inference, Cross-Silo Learning, Healthcare Data Infrastructure

Figure 1 contrasts the two paradigms. In centralized learning, hospitals send their datasets to a central server that trains a global model; in federated learning, each hospital trains locally and only model parameters travel to an aggregation server, which combines them over repeated rounds [8].




Figure 1: Centralized learning, in which hospitals A to N send their datasets to a centralized server that assembles a centralized dataset and trains a global AI model, versus federated learning, in which each hospital trains locally and an aggregation server combines the local models over federated learning rounds into a global AI model.

Federated learning alone does not guarantee privacy. Model updates can reveal information about the records used to compute them, through gradient inversion [9] or membership inference [10]. Differential privacy (DP) provides a formal bound on such leakage, but it adds noise that can reduce accuracy, can distort the probabilities a model outputs, and may fall unevenly on small hospitals and minority subgroups.

The problem is also one of national interest: the NIH Strategic Plan for Data Science 2025–2030 names secure, federated data infrastructure and privacy-preserving computation among its priorities, and NIH policy expects responsible data management and sharing [11], [12].

This paper makes four contributions. First, it specifies a cross-silo framework that combines federated averaging, record-level DP enforced at each hospital with its own privacy budget, secure aggregation, and site-specific calibration. Second, it proposes an evaluation protocol with four measured dimensions: predictive

accuracy, privacy leakage, performance across populations (including hospitals held out of training), and communication overhead. Third, it reports a simulation of ten heterogeneous hospitals that quantifies the trade-offs among these dimensions. Fourth, it derives deployment guidance, including a practical privacy budget for the task studied. The experiments use synthetic data, which allows controlled comparison but does not replace validation on real records.

Literature Review

A. Collaborative Learning Paradigms

Three arrangements exist for learning across institutions: pooling the data, training separately, or training jointly while keeping data local. Federated learning, introduced as federated averaging (FedAvg), has clients train on local data for several steps and a server average the resulting models [8]. Hospitals form a cross-silo setting with few, reliable participants [13], [14]. Figure 2 shows the two patterns used in healthcare: in panel A, data custodians such as a hospital, a research center, and an academic medical center send data to an aggregation site that trains the model; in panel B, each custodian trains locally and sends a model update to a secure aggregation server, which returns an updated model.




Figure 2: (A) Pooling of data: data custodians (hospital, research center, academic medical center) send data to an aggregation site that produces a trained model. (B) Federated learning: data custodians keep their data, send model updates to a secure aggregation server, and receive an updated model.

Medical applications have grown quickly. Federated training has matched central training for brain tumor segmentation [15], a multi-hospital model for COVID-19 outcomes performed better and generalized more widely than single-site models [16], and early work predicted hospitalizations from separately held health records [17]. Reviews describe the open problems of heterogeneity, governance, and privacy [18], [19]. Heterogeneity is central. When hospitals differ in case mix, labels, and measurement, local data are not drawn from one distribution and simple averaging can converge slowly or to a worse solution [20]; remedies such as proximal regularization have been proposed [21].

B. Differential Privacy

Differential privacy bounds how much the output of an algorithm can depend on any single record [22], [23]. Deep and linear models are commonly trained with differentially private stochastic gradient descent (DP-SGD), which clips each example’s gradient and adds Gaussian noise [24]. Privacy loss accumulates over training steps and is tracked with an accountant; Rényi differential privacy gives tight composition [25], and closed forms exist for the sampled Gaussian mechanism used in DP-SGD [26]. In federated settings, privacy can be defined at the level of a whole client, as in client-level DP-FedAvg [27], or at the level of individual records within each client. With only a handful of hospitals, client-level guarantees require substantial noise, so record-level protection suits cross-silo healthcare, where hospitals are trusted data owners.

The price of DP is accuracy, and it depends on the task, the data size, and the privacy budget ϵ . In medical imaging, differentially private training has been shown to be feasible, with utility falling as budgets tighten [28]. Empirical studies find that formal budgets and observed protection can differ greatly [29]. A further concern is fairness: DP-SGD can reduce accuracy more for under-represented subgroups than for the majority [30].

C. Privacy Attacks and Secure Aggregation

Membership inference asks whether a particular record was in the training set. The classic attack uses the model’s confidence or loss on the record, and its success is closely tied to overfitting [10], [31]; stronger attacks calibrate the threshold per example [32]. Secure aggregation lets a server learn only the sum of client updates, not individual updates, by adding cryptographic masks that cancel in the sum [33]. It protects individual updates from the server but does not limit what the aggregate model itself reveals, which is the role of DP.

D. Calibration and Validation Across Populations

A risk model is useful for decisions only if its probabilities are accurate, which is the property of calibration [34]. Modern models are often miscalibrated [35], and a standard remedy is Platt scaling fitted on held-out data [36], [37]. DP training has been reported to worsen calibration because gradient clipping and noise distort the learned scale of the predictions [38]. Calibration also differs between hospitals with different baseline incidence, so a single global calibration is unlikely to suffice. Methodological guidance recommends internal–external validation, in which entire sites are held out [39], [40].

Related work

The literature offers partial solutions. Federated learning avoids pooling but does not bound leakage [8], [9]; differential privacy bounds leakage but is usually studied for single datasets or client-level guarantees [24], [27]; secure aggregation hides individual updates but not the aggregate [33]; and clinical federated studies typically report accuracy and generalization [15], [16], while work on calibration rarely considers private training [34], [39]. Table 1 compares the paradigms.

TABLE 1. Comparison of collaborative learning paradigms for hospital networks (qualitative synthesis from the literature)

Paradigm	Raw data leave site?	Formal privacy bound	Handles hospital differences	Calibration addressed?	Main cost
Centralized pooling	Yes	No (access controls only)	Yes, if data pooled	Usually global only	Legal and trust barriers
Local-only models	No	Not applicable	No (single population)	Local only	Small samples, weak models
Federated averaging	No	No	Partly	Rarely	Communication; leakage via updates
Federated + secure aggregation	No	No (hides individual updates only)	Partly	Rarely	Setup and key exchange
Federated + DP	No	Yes (ϵ , δ)	Partly	Often degraded	Accuracy loss; per-site noise
Proposed framework	No	Yes, per hospital	Yes, with local calibration	Yes, local recalibration	Accuracy loss at small ϵ

Four gaps follow. Privacy mechanisms are seldom combined with calibration, although DP can distort probabilities [38]. Budgets are usually uniform, although hospitals of different size need different noise for the same guarantee. Evaluations rarely include hospitals that never participated in training. And privacy, accuracy, equity, and communication cost are usually reported separately. The framework below addresses these gaps.

Methodology

A. System Overview and Threat Model

Figure 3 shows the proposed system. Each hospital keeps its records in its own environment. In each round it receives the global model, runs several steps of local DP-SGD, and returns its update, which a secure aggregation layer combines with the others so that the server sees only the sum. The server averages the updates and returns the new model to every hospital. After training, each hospital fits a calibration map on its own validation data. A per-hospital accountant tracks privacy loss, an audit module evaluates leakage and performance by hospital and subgroup, and a governance layer of data use agreements, institutional review, and community oversight surrounds the system.




Figure 3: Proposed privacy-preserving federated learning architecture: hospital-local DP-SGD and recalibration, secure aggregation of noisy updates, global model averaging, per-hospital privacy accounting, auditing, and governance.

The threat model assumes an honest-but-curious server and honest-but-curious hospitals that follow the protocol but may try to infer whether a patient is in another hospital’s data from the updates or the final model, which an external party may also obtain. The privacy unit is one patient record, and each patient appears in one hospital. The guarantees below cover everything a hospital releases, so they hold with or without secure aggregation, which adds a layer against curious servers [33].

B. Federated Objective and Local DP-SGD

Let K hospitals hold datasets D_k with n_k records, $n = \sum n_k$. The shared model parameters w minimize the size-weighted average of local losses,

$$\min F(w) = \sum_{k=1}^K \frac{n_k}{n} F_k(w), \quad F_k(w) = \frac{1}{n_k} \sum_{i \in D_k} \ell(w; x_i, y_i) \quad (1)$$

where ℓ is the logistic loss with L2 regularization. Each local step draws a random lot B of expected size L by Poisson sampling at rate $q = L/n_k$, clips each per-record gradient g_i to norm at most C , and adds Gaussian noise whose scale is set by the noise multiplier σ :

$$\tilde{g} = \frac{1}{L} \left(\sum_{i \in B} \text{clip}(g_i, C) + N(0, \sigma^2 C^2 I) \right), \quad \text{clip}(g, C) = \frac{g}{\max(1, \|g\| / C)} \quad (2)$$

After E local steps the hospital sends its parameter change. The server, using the public record counts as weights, updates the global model by

$$w_{t+1} = \sum_{k=1}^K \frac{n_k}{n} w_{k,t} \quad (3)$$

where $w_{\{k,t\}}$ is hospital k 's locally updated model in round t . Because each hospital applies noise before anything leaves its environment, the guarantee does not depend on the server being trustworthy.

C. Privacy Accounting with Per-Hospital Budgets

A randomized algorithm M is (ϵ, δ) -differentially private if, for all datasets D and D' differing in one record and all output sets S ,

$$P[M(D) \in S] \leq e^\epsilon P[M(D') \in S] + \delta \quad (4)$$

The privacy loss of the sampled Gaussian mechanism over repeated steps is tracked in Rényi differential privacy. For integer order α , the loss of one step is [26]

$$\epsilon_{\text{RDP}}(\alpha) = \frac{1}{\alpha-1} \log \sum_{k=0}^{\alpha} C(\alpha, k) (1-q)^{\alpha-k} q^k \exp\left(\frac{k^2-k}{2\sigma^2}\right) \quad (5)$$

and the loss over $T \cdot E$ steps is converted to an (ϵ, δ) guarantee by [25]

$$\epsilon = \min_{\alpha} \left[T E \epsilon_{\text{RDP}}(\alpha) + \frac{\log(1/\delta)}{\alpha-1} \right] \quad (6)$$

Hospitals differ in size, so one noise multiplier would give small hospitals a much weaker guarantee, because their sampling rate q is higher. Each hospital therefore chooses its own multiplier σ_k by bisection so that its cumulative loss meets a common target ϵ , with $\delta = 10^{-5}$ below the reciprocal of every hospital's record count. The bound on the advantage of any membership-inference adversary (true positive rate minus false positive rate) implied by an (ϵ, δ) guarantee is

$$\text{Adv} \leq \frac{e^\epsilon - 1 + 2\delta}{e^\epsilon + 1} \quad (7)$$

which provides the formal reference against which empirical attacks are compared.

D. Local Calibration

Noise and clipping change the scale of a model's outputs, and hospitals differ in incidence. Each hospital therefore fits a two-parameter Platt map to its own validation data, applied to the logit s of the global model's output [36]:

$$P(y = 1 | s) = \frac{1}{1 + e^{-(a+bs)}} \quad (8)$$

The parameters a and b stay in the hospital, so recalibration costs no privacy for other participants. A hospital joining after training can onboard with a small local validation sample, which tests how the framework behaves for institutions that were never part of training. Calibration quality is summarized by the expected calibration error over B equal-frequency bins [41],

$$\text{ECE} = \sum_{b=1}^B \frac{n_b}{N} | \text{acc}(B_b) - \text{conf}(B_b) | \quad (9)$$

E. Communication and Governance

In each round every hospital downloads and uploads one parameter vector of P values with β bytes per value, so the total traffic over T rounds for K hospitals is

$$\text{Traffic} = 2 K T P \beta \quad (10)$$

and can be reduced by quantizing uploads. Deployment would also require institutional review, a data use agreement, security review of the aggregation service, and documentation following reporting guidance for prediction models [42], consistent with the privacy rule and NIH data policy [4], [12].

Experimental Setup

A. Synthetic Hospital Network

The experiments use a synthetic network of ten hospitals that vary in type and population (Table 2). Eight hospitals (H1 to H8, 23,800 patients) take part in training. Two (X1 and X2) never take part and are used to test how the model performs at new institutions. The task is the three-year risk of incident type 2 diabetes among adults with elevated metabolic risk. Each patient has 80 features: 12 continuous clinical variables (age, body mass index, HbA1c, fasting glucose, blood pressure, lipids, kidney and liver measures, waist circumference), eight binary clinical and payer indicators, and 60 uninformative features that mimic the many sparse diagnosis and medication flags in real records. The outcome is generated from a logistic function with a nonlinear threshold and an interaction term. Hospitals differ in age, metabolic burden, baseline incidence, payer mix, and an offset in HbA1c measurement that mimics assay differences. Incidence ranges from 4.8% to 13.6% in the training hospitals, and the smallest (H8) has only about 80 events in total.

TABLE 2. Characteristics of the synthetic hospital network (* held out of training; incidence, age, and payer share averaged over five simulated networks)

Hospital	Type	Patients	3-year incidence (%)	Mean age	Medicaid/uninsured (%)	HbA1c offset
H1	Academic	6,000	4.8	46.0	12.0	+0.00
H2	Academic	5,000	6.1	48.2	15.2	+0.03
H3	Community	4,000	8.4	52.0	20.3	-0.04
H4	Community	3,200	8.9	50.2	21.5	+0.02
H5	Safety-net	2,400	11.4	44.3	55.6	+0.06
H6	Rural	1,600	13.6	55.9	34.8	-0.06
H7	Safety-net	1,000	13.6	45.3	59.8	+0.08
H8	Rural	600	13.3	58.1	36.6	-0.08
X1*	Community	3,500	9.3	50.9	24.5	+0.05
X2*	Rural	1,200	16.3	56.8	39.5	-0.07

B. Methods Compared and Implementation

Six approaches were compared: local-only logistic regression; centralized logistic regression on pooled data (a non-private upper bound); FedAvg; FedAvg with 8-bit quantized uploads; FedAvg with record-level DP at $\epsilon \in \{8, 4, 2, 1\}$; and DP at $\epsilon = 4$ with quantized uploads. All federated runs used $T = 40$ rounds, $E = 8$ local steps per round, an expected lot size of $L = 64$, a learning rate of 2.0, a clipping norm of $C = 1.0$, and L2 regularization of 10^{-3} . Records in each hospital were split 60/20/20 into training, validation, and test sets. Features were standardized with fixed reference constants, which avoids sharing data statistics. Hyperparameters were tuned on a separate synthetic network, and the privacy cost of tuning is not counted. The model is logistic regression, implemented with NumPy and scikit-learn [43]. The accountant implements Eq. (5) and Eq. (6) and gave $\epsilon = 3.01$ for a common benchmark configuration (60,000 records, batch 256, $\sigma = 1.1$, 60 epochs, $\delta = 10^{-5}$). Secure aggregation was analyzed, not implemented cryptographically.

C. Evaluation Measures

Predictive accuracy was measured on each hospital’s test set by the area under the ROC curve (AUC), average precision, and Brier score, and summarized as the mean over hospitals and as the worst-hospital AUC. Performance across populations was assessed by hospital, at the two held-out hospitals, and within age (65 and over versus under 65) and payer (Medicaid or uninsured versus other) subgroups.

Calibration was measured by ECE before and after local recalibration (Eq. 8 and Eq. 9). Privacy leakage was measured by a loss-threshold membership-inference attack against the final model, using up to 400 training records (members) and 400 test records (non-members) per hospital. The attack’s AUC and advantage (maximum of true positive rate minus false positive rate) were compared with a null obtained by permuting membership labels, and with the formal bound in Eq. (7). Communication was measured in bytes with Eq. (10). Results are means and standard deviations over five independent simulated networks.

Results

A. Predictive Accuracy

Table 3 summarizes performance. Local-only models reached a mean hospital AUC of 0.766, centralized pooling reached 0.811, and federated averaging without privacy protection reached 0.801, recovering most of the gain from pooling: 3.5 AUC points above local-only models and 1.0 point below the centralized model. Quantizing uploads to 8 bits did not reduce accuracy (AUC 0.804).

TABLE 3. Predictive accuracy, calibration, and privacy leakage (mean over five simulated networks; AUC $\pm$ SD; MIA = membership-inference attack; null advantage about 4.6%)

Method	Mean AUC	Worst-hospital AUC	External AUC	ECE: raw $\rightarrow$ recalibrated	Brier	MIA advantage (%)
Local-only	0.766 $\pm$ 0.016	0.695	0.757	0.051 $\rightarrow$ 0.040	0.085	6.6
Centralized	0.811 $\pm$ 0.010	0.753	0.802	0.034 $\rightarrow$ 0.035	0.074	4.6
FedAvg	0.801 $\pm$ 0.012	0.747	0.793	0.040 $\rightarrow$ 0.037	0.077	4.5
FedAvg + DP, $\epsilon = 8$	0.799 $\pm$ 0.012	0.751	0.789	0.082 $\rightarrow$ 0.038	0.090	4.4
FedAvg + DP, $\epsilon = 4$	0.791 $\pm$ 0.011	0.748	0.782	0.082 $\rightarrow$ 0.038	0.089	4.4
FedAvg + DP, $\epsilon = 2$	0.768 $\pm$ 0.013	0.708	0.764	0.081 $\rightarrow$ 0.038	0.090	4.5
FedAvg + DP, $\epsilon = 1$	0.731 $\pm$ 0.023	0.688	0.725	0.080 $\rightarrow$ 0.039	0.092	5.1

B. Privacy–Utility Trade-off

Adding DP reduced accuracy as the budget tightened (Figure 4a): AUC was 0.799 at $\epsilon = 8$, essentially unchanged; 0.791 at $\epsilon = 4$, a loss of 1.0 point; 0.768 at $\epsilon = 2$, 3.3 points below non-private FedAvg; and 0.731 at $\epsilon = 1$. The break-even with local-only training lies between $\epsilon = 2$ and $\epsilon = 4$. At $\epsilon = 2$ the mean AUC (0.768) equaled the local-only mean (0.766), although the worst-hospital AUC was higher (0.708 versus 0.695); at $\epsilon = 1$ the private model was worse than not collaborating.

TABLE 4. Noise multipliers σ calibrated per hospital so that each reaches the target ϵ ($\delta = 10^{-5}$; 40 rounds $\times$ 8 local steps)

Hospital	Training records	Sampling rate q	σ ($\epsilon = 8$)	σ ($\epsilon = 4$)	σ ($\epsilon = 2$)	σ ($\epsilon = 1$)
H1	3,600	0.018	0.66	0.86	1.18	1.83
H2	3,000	0.021	0.69	0.92	1.30	2.12
H3	2,400	0.027	0.75	1.01	1.50	2.56
H4	1,920	0.033	0.82	1.12	1.76	3.13
H5	1,440	0.044	0.92	1.34	2.22	4.08
H6	960	0.067	1.13	1.79	3.18	6.01
H7	600	0.107	1.56	2.67	4.93	9.49
H8	360	0.178	2.37	4.27	8.08	15.71

Table 4 shows why small hospitals pay the highest price. To reach $\epsilon = 2$, the largest hospital (H1) needs a noise multiplier of 1.18, while the smallest (H8) needs 8.08, nearly seven times as much, because it samples 17.8% of its records at each step. At $\epsilon = 1$ the multiplier for H8 reaches 15.7.




Figure 4: (a) Privacy–utility trade-off: mean hospital AUC of federated learning with DP at different budgets, compared with local-only and centralized models (error bars show SD over five networks). (b) Empirical membership-inference advantage for each hospital, ordered from largest to smallest, for local-only, FedAvg, and FedAvg with DP; black marks show the null level from shuffled labels. Synthetic data.

C. Privacy Leakage

Empirical leakage was small and concentrated in the smallest hospitals (Figure 4b). Pooled across hospitals, the attack’s advantage was 6.6% for local-only models against a null of about 4.6%, and indistinguishable from the null for federated models without DP (4.5%) and with DP (4.4% to 5.1%). Attack AUC was 0.510 for local-only and 0.499 to 0.503 for federated models. The local-only models of the two smallest hospitals leaked measurably: for H8 the advantage was 13.4% against a null of 6.2% (attack AUC 0.534), and for H7 it was 10.6% against 5.0% (AUC 0.545), consistent with the link between overfitting and membership inference [31]. Federated training reduced these advantages to 6.6% for H8 and 7.7% for H7, and DP at $\epsilon = 2$ gave 8.5% and 6.7%. With only a few hundred records per attack, these three sets of values cannot be distinguished from one another, but all lie well below the local-only values for the two hospitals.

These measurements need care. The attack is a simple loss-threshold attack against a low-capacity model, and a null result does not prove safety [32]. The formal bound in Eq. (7) is far more conservative: 0.46 at $\epsilon = 1$, 0.76 at $\epsilon = 2$, 0.96 at $\epsilon = 4$, and close to 1 at $\epsilon = 8$. The guarantee is what a hospital can promise; the audit shows that practical leakage here was far below the worst case.

D. Performance Across Hospitals and Populations

Federated training helped the hospitals that need it most (Figure 5a). The smallest hospital (H8) had an AUC of 0.703 with its local-only model and 0.810 with FedAvg, a gain of 10.7 points; H7 improved from 0.742 to 0.795 (5.3 points) and H1, the largest, from 0.770 to 0.790 (2.0 points). The worst-hospital AUC rose from 0.695 for local-only models to 0.747 for FedAvg, with a variability across networks that was halved (SD 0.061 versus 0.029). With DP at $\epsilon = 2$, small hospitals kept most of their gains (H8 0.791), but some larger ones lost several points compared with FedAvg (H4 0.732 versus 0.797), so the cost of DP is spread unevenly.

The model transferred to hospitals outside the network. At the two held-out hospitals, mean external AUC was 0.793 for FedAvg, 0.802 for the centralized model, and 0.757 for local-only models applied to an institution they had never seen. External AUC under DP was 0.782 at $\epsilon = 4$ and 0.764 at $\epsilon = 2$. Table 5 reports subgroup results. Patients aged 65 and over had lower AUC than younger patients under all methods, by 1.0 point for FedAvg and by 2.0 and 2.2 points for DP at $\epsilon = 4$ and $\epsilon = 2$, but the gap did not grow monotonically with stronger privacy (0.5 points at $\epsilon = 1$), so no consistent disparate impact of DP appears at this sample size. Payer differences were small (at most 0.9 points).

TABLE 5. AUC by subgroup, pooled over training hospitals (mean over five simulated networks)

Method	Age 65 and over	Age under 65	Gap (points)	Medicaid/uninsured	Other payer
Centralized	0.801	0.814	1.3	0.817	0.816
FedAvg	0.793	0.803	1.0	0.806	0.806
FedAvg + DP, $\epsilon = 8$	0.794	0.798	0.4	0.805	0.801
FedAvg + DP, $\epsilon = 4$	0.773	0.793	2.0	0.802	0.793
FedAvg + DP, $\epsilon = 2$	0.749	0.771	2.2	0.772	0.773
FedAvg + DP, $\epsilon = 1$	0.727	0.732	0.5	0.741	0.739

E. Calibration

Privacy noise worsened calibration (Figure 5b). The expected calibration error of FedAvg without DP was 0.040, but it was 0.082, 0.082, 0.081, and 0.080 for DP at $\epsilon = 8, 4, 2$, and 1, roughly double and almost independent of ϵ , which suggests that clipping, as well as noise, distorts the scale of the output probabilities [38]. The Brier score rose correspondingly (0.077 without DP, 0.089 to 0.092 with DP). Local recalibration with Eq. (8) removed this penalty: after recalibration the ECE was 0.037 to 0.039 for all private and non-private federated models, and Brier scores returned to 0.076 to 0.082. The effect was larger at the held-out hospitals, where the ECE of the DP model at $\epsilon = 2$ fell from 0.105 to 0.030 after recalibrating on only

300 records of local data. Errors were largest in the smallest, highest-incidence hospitals, so a single global calibration would have misrepresented risk where data are thinnest.




Figure 5: (a) Test AUC by hospital for local-only, FedAvg, FedAvg with DP ($\epsilon = 2$), and centralized models; X1 and X2 are hospitals held out of training (local-only values are averages over the eight local models). (b) Expected calibration error of the global DP model before and after local recalibration. Synthetic data.

F. Communication Overhead

Local computation traded off against communication (Figure 6a). With one local step per round, validation AUC was 0.787 after five rounds and 0.801 after 20; with eight steps it was 0.807 after five rounds, within 0.01 of the centralized model (0.816). Because the privacy cost grows with the number of steps, shorter training also means less noise for a given ϵ .

TABLE 6. Communication overhead of federated training by model size (float32 vs. 8-bit quantized uploads)

Model size P (parameters)	Per hospital per round, float32	Total, float32 (8 hospitals, 40 rounds)	Total, int8 uplink	Reduction
81 (this study)	648 B	0.21 MB	0.13 MB	37%
10^5 (100,000)	0.8 MB	256.00 MB	160.00 MB	37%
10^6 (1 million)	8.0 MB	2.56 GB	1.60 GB	37%
10^7 (10 million)	80.0 MB	25.60 GB	16.00 GB	37%

For the logistic model of this study (81 parameters), each hospital exchanged 648 bytes per round, and the total traffic for eight hospitals and 40 rounds was about 207 kB (Table 6). Quantizing uploads to 8 bits reduced this to about 131 kB without measurable loss of accuracy (AUC 0.804 without DP and 0.789 at $\epsilon = 4$, against 0.801 and 0.791 for the unquantized models). The raw training records of the eight hospitals (14,280 patients $\times$ 81 values $\times$ 4 bytes) would total about 4.6 MB, so federated training moved 4.5% of that volume. For so small a model traffic is not a concern; it matters for larger models, since traffic grows linearly with P (Eq. 10 and Figure 6b), reaching about 2.6 GB at one million parameters and 26 GB at ten million under the same schedule.




Figure 6: (a) Convergence of non-private FedAvg for different numbers of local steps E per round (validation AUC; mean of three networks). (b) Total traffic for eight hospitals and 40 rounds as a function of model size, for float32 transfers and for 8-bit quantized uploads. Synthetic data.

G. Summary of the Operating Point

The results suggest an operating point for this task: eight local steps per round, quantized uploads, $\epsilon = 4$ per hospital, and local recalibration. It gave a mean AUC of 0.789 to 0.791 (0.811 for pooling, 0.766 for local-only), an ECE of about 0.04 after recalibration, an external AUC of 0.78 to 0.79, no detectable membership leakage, and about 131 kB of traffic.

Discussion

The simulation supports four conclusions. First, collaboration pays off mainly for small hospitals and late joiners: the largest hospital gained 2.0 AUC points from federation, the smallest 10.7, and local-only models applied to unseen hospitals performed 3.6 points below the federated model and 4.5 points below the centralized model. A network that includes small and rural hospitals therefore advances equity in the quality of risk prediction, provided those hospitals can join.

Second, differential privacy has a real cost, and the cost depends on the hospital. Per-hospital calibration of noise gives every hospital the same formal guarantee, but it is the small hospitals that must add the most noise to achieve it. Operators must choose between equal guarantees with unequal noise, equal noise with unequal guarantees, or minimum participation sizes. Our results favor a budget around $\epsilon = 4$, where accuracy loss was about one AUC point. At $\epsilon = 1$ the private model was worse than local training, so a budget that strict can defeat the purpose of collaboration. These conclusions may differ for larger models or datasets.

Third, calibration must be treated as a first-class requirement. Discrimination barely changed between non-private and private models at $\epsilon = 4$ and above, yet the probabilities were visibly distorted, and a clinician acting on a miscalibrated risk would over- or under-treat. Recalibration is local and cheap, and also adjusts for differences in incidence and assay; evaluating AUC alone would have hidden the problem.

Fourth, empirical privacy audits and formal guarantees play different roles. The attack found no detectable leakage from federated models and weak leakage from the smallest hospitals' local models, but formal bounds are the only protection against attacks not yet invented. Operators should report both and treat audits as monitoring, not as a substitute for accounting.

The study has important limitations. The data are synthetic, so the effect sizes depend on our assumptions about heterogeneity, signal strength, and noise features, and they would differ on real records. The model is a logistic regression; deep models raise communication cost, accuracy sensitivity to noise, and attack

surface. Privacy is record-level and assumes one record per patient and one hospital per patient; patients seen at several hospitals, and correlated records, would weaken the guarantee. Hyperparameter tuning was not accounted for in the privacy budget. Secure aggregation was not implemented, gradient-reconstruction attacks were not tested, and the subgroup analysis covered only age and payer. Future work should validate the framework on real data in a common data model such as OMOP across several health systems, test deep and tree-based models, evaluate stronger attacks, and examine patient-level accounting and personalized budgets.

Conclusion

This paper presented a framework for privacy-preserving collaborative risk prediction across hospitals that combines federated averaging, record-level differential privacy with per-hospital accounting, secure aggregation, and local calibration, and evaluated it on accuracy, privacy leakage, performance across populations, and communication overhead in a synthetic network of ten hospitals. Federated training reached a mean AUC of 0.801, compared with 0.766 for local models and 0.811 for pooling, and lifted the smallest hospital by 10.7 points. A privacy budget of $\epsilon = 4$ per hospital cost about one AUC point, whereas $\epsilon = 1$ made the private model worse than local training. Differential privacy doubled calibration error, but local recalibration restored it, including at hospitals outside the network. Membership inference could not distinguish members from non-members in the federated models, and traffic was a small fraction of the raw data volume.

The central implication is that useful, private, and calibrated collaborative prediction is feasible for a task like the one studied, provided that privacy budgets are accounted for hospital by hospital, calibration is performed locally, and performance is evaluated across hospitals, including those that join later. The authors recommend a pilot on real data across several health systems under a common data model, with pre-specified budgets, external validation at held-out sites, and joint reporting of accuracy, calibration, leakage, and cost, in line with national priorities for secure data infrastructure and rigorous AI validation.

References

1. **Pawan Kumar Singh**, Scaling Gate-All-Around (GAA) Transistor Architectures for Sub-2nm AI Accelerators Using Atomically Thin 2D Materials, Vol. 33 No. 08 (2024): JOCAAA, Journal Name: Journal of Computational Analysis and Applications, Journal's ISSN: 1521-1398 (Paper),1572-9206 (Online), <https://eudoxuspress.com/index.php/pub/article/view/5713>
2. **Pawan Kumar Singh**, Energy-Efficient Task Scheduling in Green Cloud Computing Using Neuromorphic Spiking Neural Networks, Vol. 33 No. 08 (2024): JOCAAA, Journal Name: Journal of Computational Analysis and Applications, Journal's ISSN: 1521-1398 (Paper),1572-9206 (Online), <https://eudoxuspress.com/index.php/pub/article/view/5712>
3. **Rangavinay Teja Royal Jagga**, 2026, DOI: 10.1109/GCWCN66157.2025.11448495, Date of Conference: 22-23 November 2025, Date Added to IEEE Xplore: 25 March 2026, <https://ieeexplore.ieee.org/document/11448495>
4. **Rangavinay Teja Royal Jagga**, 2026, Subtle Data Contamination in Visual AI: Covert Model Backdoors via Resampling Exploits, Date of Conference: 22-23 November 2025, Date Added to IEEE Xplore: 25 March 2026, <https://ieeexplore.ieee.org/document/11448514>

5. **Rangavinay Teja Royal Jagga**, 2026, Robotic Tactile-Vision Integration for Interactive Object Disentanglement,
<https://ieeexplore.ieee.org/document/11448337>
6. **Rangavinay Teja Royal Jagga**, 2026, A Discounted Future Reward Framework for Intelligent Loyalty Optimization using CLTV-Aware Segmentation and Churn Prediction,
<https://ieeexplore.ieee.org/document/11497560>
7. **Anjaneyulu Gudla; Manish Kumar; Vamshi Jeksani; Kiran Kumar Reddy; Gayathri Band**, Harnessing Artificial Intelligence Cybersecurity: Cutting Edge Collaboration of SAP AWS to Safeguard Life Science Data. 2026
<https://doi.org/10.1109/aiei69164.2026.11497833>
8. **G. Mahender, Department of CSE(Data Science), Vardhaman College of Engineering ; Vamshi Jeksani; Manish Kumar; Koppuravuri Sri Lakshmi Sruthi; Kiran Kumar Reddy**, SAP Cloud System on AWS for Risk Detection and Integration Artificial Intelligence Scalable Cybersecurity 2026 <https://doi.org/10.1109/aiei69164.2026.11497435>
9. **Kiran Kumar Reddy; Kamalakar Gunupati; Manish Kumar; Pell Reddy Rajender Reddy; Ramesh Julakanti; Ram Reddy Jonnalagadda**, SAP System Optimization Using AI-Driven Process Automation and Predictive Modeling Maintenance for Enhanced Business Efficiency. 2025 <https://doi.org/10.1109/computingcon64838.2025.11376613>
10. **Amilineni, M.V. 2025**. Audit-Safe Agentic RAG Framework for Regulated Enterprise Systems: A Trustworthy AI Architecture for SAP, Finance, Healthcare, and Government Workflows. INTERNATIONAL JOURNAL OF ADVANCES IN SIGNAL AND IMAGE SCIENCES. (Jan. 2025), 34–46. DOI: <https://doi.org/10.29284/7fjzk883>
11. **Pawan Kumar Singh**, ACCELERATING LARGE LANGUAGE MODEL TRAINING USING HYBRID QUANTUM-CLASSICAL VARIATIONAL AUTOENCODERS WITH INTEGRATED PHOTONIC QRAM, Vol. 54 No. 3 (2026): July-September 2026, Power System Protection and Control, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/501>, DOI: <https://doi.org/10.46121/pspc.54.3.45>
12. **Pawan Kumar Singh** , RESILIENCE ASSESSMENT OF MODERN POWER GRID INFRASTRUCTURE AGAINST EXTREME WEATHER USING PHYSICS-INFORMED NEURAL NETWORKS, Vol. 54 No. 1 (2026): January-March 2026, Power System Protection and Control, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/502>, DOI: <https://doi.org/10.46121/pspc.54.1.64>
13. **Pawan Kumar Singh**, MITIGATING THE EXABYTE STORAGE CRISIS: ENHANCING READ/WRITE EFFICIENCY IN 5D OPTICAL GLASS STORAGE USING SILVER NANOPARTICLE DOPING, Vol. 38 No. 11s (2025), International Journal of Applied Mathematics (IJAM) , pISSN 1311-1728, eISSN 1314-8060 , <https://ijamjournal.org/ijam/publication/index.php/ijam/article/view/2153>

14. **Mr. Pawan Kumar Singh (Author), Miss. Pari Nidhi Singh (Author)** ,Intelligent Bot, ASIN : BoDBGQCQY8X,Publisher : Namya Press, Publication date : November 25, 2022,Language : English, Print length : 105 pages,ISBN-10 : 9390124867,ISBN-13 : 978-9390124862,Item Weight : 7.5 ounces, Dimensions : 6.24 x 0.43 x 9.24 inches,https://www.amazon.com/Intelligent-Bot-Pawan-Kumar-Singh/dp/BoDBGQCQY8X/ref=sr_1_5?crid=2VAXN5BRN5SMS&dib=eyJ2IjojMSJ9.v3OCQvHBpy2oRIjfmWhoqbOI-RLOIUJIC8XaoPjeqilvo2sbZ_suN3LbSCw_OA_ZD_4lgjRLMG1BwZQQ1ehZEFdpt_MJiOUfVuqqzMavoI5_mb6pN5DQogoU7SA826n1DQQ3QAtWwopHy4lC5IfhZXttKj9faiyATwhdtDATSSV1-r7FmzM2VQyIAojxcfq6E6hp9ECfH1fbyoli9kUfGpLxMq4GktlzUIJO5ggrhsQ.G1AjQiVqU4HFSnSczvv-X6NXCNS3EFSn-RI8mcWWGOc&dib_tag=se&keywords=Pawan+Kumar+Singh%2C+book&qid=1787815741&prefix=pawan+kumar+singh%2C+bo%2Caps%2C450&sr=8-5
15. **Mr. Pawan Kumar Singh (Author)**, STORAGE CRISES DUE TO RISE OF AI: A Comprehensive Guide to Managing the Storage Challenges in the Age of Artificial Intelligence, ASIN : BoHFK12ZCT Publisher : EB1 Profile, Accessibility : Learn more, Publication date : August 17, 2026, Edition : 1st Language : English, File size : 1.7 MB, Screen Reader : Supported, Enhanced typesetting : Enabled X-Ray : Not Enabled, Word Wise : Not Enabled ,Print length : 261 pages ,ISBN-13 : 979-8951479112,Page Flip : Enabled, https://www.amazon.com/STORAGE-CRISES-DUE-RISE-Comprehensive-ebook/dp/BoHFK12ZCT/ref=sr_1_1?crid=9MF6C4WDVJR5&dib=eyJ2IjojMSJ9.DQ8fesBch7ELmRnildt-YA.NBSCHLmetvnbhkwIpOoj3soIVsCV8y6-ZCfCP6Ul1c&dib_tag=se&keywords=Pawan+Kumar+Singh%2C+book%2C+eb1+Profile&qid=1787816040&prefix=pawan+kumar+singh%2C+book%2C+eb1+profile%2Caps%2C420&sr=8-1
16. **Parag Bharkat Kumar Parikh**, FRAMEWORKS FOR MITIGATING SUPPLY CHAIN DISRUPTIONS AND MATERIAL PRICE VOLATILITY IN SUSTAINABLE BUILDING CONSTRUCTION, Vol. 53 No. 2 (2025): April-June 2025 , , Power System Protection and Control, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/447>, DOI: <https://doi.org/10.46121/pspc.53.2.34>
17. **Parag Bharkat Kumar Parikh**, IMPLEMENTING DIGITAL INSPECTION PROTOCOLS FOR ENHANCED QUALITY CONTROL MANAGEMENT IN HIGH-RISE CONSTRUCTION, Vol. 54 No. 3 (2026): July-September 2026, Power System Protection and Control, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/439>, DOI: <https://doi.org/10.46121/pspc.54.3.12>
18. **Parag Bharkat Kumar Parikh** , EVALUATING THE FEASIBILITY AND EFFICIENCY OF 3D-PRINTED CONCRETE TECHNOLOGY IN SUSTAINABLE MODULAR CONSTRUCTION, Power System Protection and Control, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/issue/view/30>, DOI: <https://doi.org/10.46121/pspc.54.1.61>
19. **Parag Bharkat Kumar Parikh**, EVALUATING THE FEASIBILITY AND EFFICIENCY OF 3D-PRINTED CONCRETE TECHNOLOGY IN SUSTAINABLE MODULAR CONSTRUCTION, Vol. 54 No. 1 (2026): January-

- March 2026, System Protection and Control, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/440>, DOI: <https://doi.org/10.46121/pspc.54.1.61>
20. **Ronakkumar shah**, REDEFINING BUSINESS PROCESS MANAGEMENT: THE SYNERGY OF AI AND INTELLIGENT AUTOMATION, Vol. 54 No. 3 (2026): July-September 2026, System Protection and Control, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/448>, DOI: <https://doi.org/10.46121/pspc.54.3.14>
 21. **Ronakkumar shah**, OPERATIONAL ACCOUNTABILITY IN ACTION: METRICS AND MECHANISMS FOR SYSTEMIC SUCCESS, Vol. 53 No. 4 (2025): October-December 2025, System Protection and Control, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/452>, DOI: <https://doi.org/10.46121/pspc.53.4.40>
 22. **Ronakkumar shah**, BEYOND THE BLUEPRINT: AUTOMATED WORKFLOW ENFORCEMENT AND POLICY GOVERNANCE, Vol. 53 No. 2 (2025): April-June 2025, System Protection and Control, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/451>, DOI: <https://doi.org/10.46121/pspc.53.2.35>
 23. **Tushita Agarwal; Swaminathan Sethuraman**, Mapping Customer Behavior: Visual Analysis of Online Retail Interactions, Date Added to IEEE Xplore: 17 February 2026, DOI: <https://doi.org/10.1109/IC2NC67409.2025.11376261>
 24. **Tushita Agarwal**, Efficient Data Summarization and Comparison via Grid-Based Statistical Proxies for Large-Scale Datasets, **Date Added to IEEE Xplore: 03 December 2025**, DOI: <https://doi.org/10.1109/ICBATS66542.2025.11258564>
 25. **Dip Bharatbhai Patel**, Building scalable business intelligence systems in the cloud: A technical approach, International Journal of Science and Research Archive, 2021, 02(01), 212-215., Received on 03 February 2021; revised on 17 April 2021; accepted on 19 April 2021, DOI: <https://doi.org/10.30574/ijjsra.2021.2.1.0047>
 26. **Dip Bharatbhai Patel**, University of North America, Virginia, USA, Home / Archives / Vol. 5 No. 02 (2025): Volume 05 Issue 02 / Articles, <https://www.academicpublishers.org/journals/index.php/ijdsml/article/view/5958>
 27. **Sai Akshara Vemuri**, KINEMATIC OPTIMISATION OF SOFT ROBOTIC MANIPULATORS FOR MEDICAL INTERVENTION, Vol. 54 No. 3 (2026): July-September 2026, System Protection and Control, ISSN-1674-3415, <https://pspac.info/index.php/dlbh/article/view/450>
 28. **Sai Akshara Vemuri**, <https://www.periodicos.ulbra.org/index.php/acta/article/view/528>
 29. **Aziz, A., Chandusha, K., Naga Kiran, B., Prasad, M. L. M., Golla, K., Rajeswari, S., & Katta, S. K. R.** (2026). Exploring Machine Learning Applications for Genomic Data Analysis in Personalized Medicine. International Journal of Drug Delivery Technology, 16(38s), 1166-1173. <https://doi.org/10.25258/ijddt.16.38s.127>, <https://www.periodicos.ulbra.org/index.php/acta/article/view/687>

30. **Mohammad Naffizuddin , Chaganti Ashok ,Sowjanya Gunukula , Rvbs Sarma , Bhavana Sujanamulk , Bharani Krishna Takkella , Chukka Ram Sunil ,Varri Sujana , Ashwini Kumar**, Etiological Factors of Orofacial Clefts in a Hospital-Based Population: A Descriptive Analytical Study, Published: December 10, 2025,DOI: 10.7759/cureus.98868,DOI : <https://doi.org/10.7759/cureus.98868>
31. **Manoj Kumar Kagitha**, Automated Security Compliance in Ci/Cd: A Natural Language Processing (NLP) Approach to Detecting Vulnerabilities in Infrastructure-As-Code (Terraform/Bicep), Publication year 2021, Volume 2021, Issue 12, <https://computerfraudsecurity.com/index.php/journal/article/view/970>
32. **Manoj Kumar Kagitha**, Hybrid Quantum-Classical Resource Scheduling: A Proposed Architecture for Next-Generation Hyperscale Data Centers, Publication year 2024, Volume 2024, Issue 12, <https://computerfraudsecurity.com/index.php/journal/article/view/969>
33. **Manoj Kumar Kagitha**, Self-Healing Cloud Infrastructure: A Hybrid Reinforcement Learning Framework for Predicting Microservice Failures in Azure Kubernetes Service (AKS),Publication year 2023, Volume 2023, Issue 12,<https://computerfraudsecurity.com/index.php/journal/article/view/971>
34. **Manoj Kumar Kagitha**, GREENOPS IN THE CLOUD: AN AI-DRIVEN TELEMETRY ANALYSIS MODEL FOR REDUCING CARBON FOOTPRINT IN HIGH-AVAILABILITY CLUSTERS, Vol. 28 No. 6 (2020): JOCAAA, <https://www.eudoxuspress.com/index.php/pub/article/view/5054>
35. Lakshmi Rahul Reddy Mareddy. (2026). Security and Performance Assessment of Encryption Techniques for Cloud Platforms. International Journal of Communication Networks and Information Security (IJCNIS), 18(1), 68–80. Retrieved from <https://ijcnis.org/index.php/ijcnis/article/view/8735>
36. Lakshmi RahulReddy Mareddy, Context-Aware Learning Approaches for Improving Prediction Accuracy in Dynamic Systems, International Journal of Computational Mathematical Ideas (IJCMI), Vol.17 No.3(2025): IJCM Vol17 NO3 2025 <https://www.ijcmi.com/index.php/journal/article/view/15>
37. Lakshmi Rahul Reddy , Mareddy. (2026). Revisiting Autonomy in Artificial Intelligence through Agentic System Design. Journal of Computational Analysis and Applications (JoCAAA), 35(1), 707–713. <https://eudoxuspress.com/index.php/pub/article/view/4780>
38. **Mahendrakumar Kalal**, Lean-Driven SAP Production Planning Optimization Using Machine Learning for Inventory and Throughput Efficiency, Date of Conference: 19-20 March 2026, Date Added to IEEE Xplore: 08 April 2026,ISBN Information:.,ISSN Information: DOI: 10.1109/ISDFS69419.2026.11459029, <https://ieeexplore.ieee.org/abstract/document/11459029>
39. **Mahendra kumar Kalal**, SUPPLY CHAIN RESILIENCE THROUGH INTEGRATED SAP PRODUCTION PLANNING – PROCUREMENT SYNCHRONIZATION, Archives / Vol. 54 No. 1 (2026): January-March 2026 / Articles, <https://pspac.info/index.php/dlbh/article/view/259>
40. **Dr Manoj Goel (December 2018)**. THE ATTITUDE OF DIFFERENT STATES STUDENTS TOWARDS PHYSICAL EDUCATION; STUDYING IN PHYSICAL EDUCATION COLLEGES OF MAHARSHTRA STATE, <https://ijeponline.com/index.php/journal/article/view/965/847>

41. **Ilakiya Ulaganathan, Tagore Engg. College, Anna University, Chennai, India, ; FNU Antara; Santhosh Chitraju Gopal Varma; Zokir Mamadiyarov; Ziyadulla Turayev, Published in:** 2025 International Conference on Sustainability, Innovation & Technology (ICSIT), Date of Conference: 22-23 August 2025, Date Added to IEEE Xplore: 19 December 2025, ISBN Information: <https://ieeexplore.ieee.org/document/11294748/references#references>
42. **Shreyas Kasture, Sudhanshu Maurya, Amit Kumar Sharma, Santhosh Chitraju Gopal Varma, Kashish Mirza, Firdous Sadaf Mohammad Ismail, DOI:** <http://doi.org/10.11591/ijai.v15.i2.pp1746-1761>, IAES International Journal of Artificial Intelligence (IJ-AI), ISSN/e-ISSN 2089-4872/2252-8938 <https://ijai.iaescore.com/index.php/IJAI/article/view/26932>
43. **Shiv Shankar Dwivedi, Mr. Santosh Kumar Sharma, Future-Proofing Data Security with Quantum-Resistant Cryptography, Vol. 11 No. 3 (2025): IJCESEN, DOI:** <https://doi.org/10.22399/ijcesen.3447>, <https://ijcesen.com/index.php/ijcesen/article/view/3447>
44. **Vijender Kalmotia, Enhancing statistical accuracy in gene perturbation studies, Volume 266, August 2026, 105819,** <https://doi.org/10.1016/j.biosystems.2026.105819>, <https://www.sciencedirect.com/science/article/pii/S0303264726001292>
45. **Naveen Undrathi, Design and implementation of generative AI-driven data transformation pipelines using spring boot for scalable financial applications, eISSN: 2581-9615 || CODEN: WJARAI || Impact Factor 8.2 || CrossRef DOI, September 2026 (Volume 31, Issue 3), DOI url:** <https://doi.org/10.30574/wjarr.2026.30.3.1629>,
46. **Naveen Undrathi, Agentic generative AI-enabled microservices framework for intelligent financial data processing in java-based enterprise systems, September 2026 (Volume 31, Issue 3), eISSN: 2581-9615 || CODEN: WJARAI || Impact Factor 8.2 ||** <https://doi.org/10.30574/wjarr.2026.30.3.1628>
47. **Akshay Jain, Comparative Analysis of Linkage Methods in Hierarchical Clustering for Customer Segmentation, Conference paper, First Online: 07 October 2025,** https://link.springer.com/chapter/10.1007/978-981-96-7289-9_44
48. **Ismail Hossain, Department of Computer Science, George Mason University, Fairfax, VA, USA, Mridul Banik, Department of Computer Science, Colorado State University, Fort Collins, CO, USA,** <https://ieeexplore.ieee.org/abstract/document/11446534> „DOI: 10.1109/FMLDS67896.2025.00106
49. **Ismail Hossain, Mridul Banik , Fahmid Al Farid , Jia Uddin , Hezerul bin Abdul Karim , Advancing Large Language Models for Low-Resource Languages: A Systematic Review of Pretraining, Adaptation, and Ethical Challenges, 10.32604/cmcs.2026.075507, Date: 2026** <https://www.sciencedirect.com/org/science/article/pii/S152614922600322X>
50. **Ismail Hossain; Mridul Banik, Bridging Visual and Linguistic Intelligence for Chest X-rays: A Comprehensive Review of ViTs and LLM Synergies, Published in: 2025 IEEE International Conference on Future Machine Learning and Data Science (FMLDS) Date of Conference: 02-05 November**

- 2025, Date Added to IEEE Xplore: 25 March 2026, ISBN Information: <https://ieeexplore.ieee.org/abstract/document/11446534>
51. **Rifat Sarker Aoyon¹, Ismail Hossain², M. Abdullah-Al-Wadud³, Jia Uddin⁴**, A Secured and Continuously Developing Methodology for Breast Cancer Image Segmentation via U-Net Based Architecture and Distributed Data Training, <https://www.sciopen.com/article/10.32604/cmcs.2025.060917>
 52. **M. B. Vibha, M. Chandrika, Samitha Khaiyum & P. Rakshitha Kiran**, Predicting sedentary behavior in adults using stacked LSTM modeling, published: 27 November 2024 Volume 16, pages 346–355 (2025), <https://link.springer.com/article/10.1007/s13198-024-02622-2>
 53. **Bellary Chiterki Anil, Jayasimha Sondekoppa Rajkumar, T. L. Divya, Samitha khaiyum Rakshitha Kiran P., Balakrishnan Ramadoss, A** Radiomics-based Framework for Liver Cancer Analysis using Explainable Artificial Intelligence (XAI) Methods, Vol. 15 No. 3 (2025): June, 2025, <https://etasr.com/index.php/ETASR/article/view/10377>
 54. **Jayminkumar Patel, Sanjaykumar patel, Birju Patel, Rajan shah**, Development and Validation of a Stability-Indicating Analytical Method for Quantitative Determination of Selected Anticancer Drugs in Pharmaceutical Dosage Form, Archives / Vol. 35 No. S2 (2026): Apr-jun 2026, <https://welltestingjournal.com/index.php/WT/article/view/288>
 55. **Jayminkumar Patel, Sanjaykumar patel, Birju Patel, Rajan shah**, Artificial Intelligence-Guided Optimization of Chromatographic Methods for Pharmaceutical Analytical Development, Archives / Vol. 35 No. S2 (2026): Apr-jun 2026, <https://welltestingjournal.com/index.php/WT/article/view/294>
 56. **Sharoon Sharif, Inikpi Egbunu, Yujian Fu, and Zhigang Xiao. 2026.** Hybrid Multi-Modal Learning for Extreme Long-Tail Fungi Classification. In 2026 ACM Southeast Conference (ACMSE 2026), April 23–25, 2026, Troy, AL, USA. ACM, New York, NY, USA, 38–45. <https://doi.org/10.1145/3746467.3801502>
 57. **A. Chandrakumaran**, "Innovative Strategies for Real-Time Analytics in Modern Data Warehousing Using Dockerized Ecosystems," 2025 International Conference on Advances in Next-Gen Computer Science (ICANCS), Bangalore, India, 2025, pp. 1–6, doi: <https://doi.org/10.1109/ICANCS65819.2025.11377462>
 58. **A. Chandrakumaran**, "Enhancing SAP Predictive Maintenance and Industrial Analytics with LLM-Driven Automation," 2025 IEEE Madhya Pradesh Section Conference (MPCON), Jabalpur, India, 2025, pp. 238–244, doi: <https://doi.org/10.1109/MPCON66082.2025.11256577>.
 59. **A. Chandrakumaran**, "Enhancing SAP Cloud Optimization with Explainable AI: Cost Control, Workload Forecasting, and Performance Enhancement," 2025 3rd International Conference on Inventive Computing and Informatics (ICICI), Bangalore, India, 2025, pp. 456–463, doi: <https://doi.org/10.1109/ICICI65870.2025.11069968>.
 60. **A. Chandrakumaran**, "Innovative Strategies for Real-Time Analytics in Modern Data Warehousing Using Dockerized Ecosystems," 2025 International Conference on Advances in Next-Gen Computer Science (ICANCS), Bangalore, India, 2025, pp. 1–6, doi: <https://doi.org/10.1109/ICANCS65819.2025.11377462>.

61. **Ashwini Chandrakumara**, year :2025, Leveraging Autonomous AI for Agile Growth: Transforming Project Management in Fast-Scaling Startups, Proceedings of the 2nd International Conference on Sustainable Business Practices and Innovative Models (ICSBPIM-2025), pages :730-743, issn :2352-5428, isbn :978-94-6463-872-1, Atlantis Press https://doi.org/10.2991/978-94-6463-872-1_46
62. **Lekkala, H. B., & Bandari, V. V. (2026). Deep learning** for weld defect detection using CNN or vision transformers fusion detection. *Power System Protection and Control*, 54(2), 151–158. <https://pspac.info/index.php/dlbh/article/view/314>
63. **Kancharla Vijayalakshmi, Bethapudi Samuel Anand Andrews**, Estimation of Cyproheptadine Hydrochloride and Tricholine Citrate Simultaneously in Syrup Dose by Applying Stability Indicating HPLC Methodology, **December 31, 2020**, Vol. 11, Issue 4, Pages 1813–1819, DOI: 10.26452/ijrps.v11iSPL4.4383, <https://ijrps.com/v11si4/3683>
64. **Vijayalakshmi Kancharla ,Kumar Abbaraju V D N, Andrews B S A, Nagaraju Dasari, Nagu Korupolu**, Melphan as Anticancer and New Insights of Drug as Parenteral Dosage form against SARS-CoV-2 by Using Development and Validation , **November 18, 2020**, Vol. 11, Issue 1, Pages 1437–1446 DOI: 10.26452/ijrps.v11iSPL1.3686, <https://ijrps.com/v11si1/2911>
65. **Ramesh, D., A. Singh, and L. H. Channaiah.** 2026. “ Evaluating the Effectiveness of Chewy Chocolate Chip Granola Bar Manufacturing Process on Salmonella Survival.” *Food Science & Nutrition* 14, no. 4: e71698. <https://doi.org/10.1002/fsn3.71698>.
66. **Inventor: Akhilesh RaiDilip Srinivas Krovi Sarthak Saptarshi Bandopadhyay Ravinder Singh Choudhary Richard LAWTON, 2025-02-28**, Priority to US19/067,141,2025-11-14,Assigned to CITIBANK, N.A.,2025-12-30,Application granted,2025-12-30, Publication of US12511187B, <https://patents.google.com/patent/US12511187B1/en?q=US12511187>
67. **Vivekanandan Govindan Ekambaram** , REAL-TIME AUTOMATED DEFECT DETECTION IN INDUSTRIAL MANUFACTURING USING DEEP LEARNING AND TEXTURE ANALYSIS, Archives ,Vol. 54 No. 4 (2026): **October-December 2026**, Power System Protection and Control, ISSN-1674-3415,DOI: <https://doi.org/10.46121/pspc.54.4.06>, <https://pspac.info/index.php/dlbh/article/view/545>
68. **Vivekanandan Govindan Ekambaram, Yamuna Natarajan**, AN OPTIMIZED HYBRID FRAMEWORK FOR REAL-TIME EDGE DETECTION AND FEATURE EXTRACTION IN AUTONOMOUS VEHICLE VISION SYSTEMS, **Vol. 54 No. 4 (2026):** October-December 2026 , Power System Protection and Control, ISSN-1674-3415, DOI: <https://doi.org/10.46121/pspc.54.4.05>, <https://pspac.info/index.php/dlbh/article/view/544>
69. **Anasuya Patil 1, Gurinderdeep Singh 2, Rajendra Dnyandeo Dighe 3, Dhruv Dev 4, Bhaveshkumar A Patel 5, Samatha Rudrangi 6, Gaurav Tiwari 7**, Preparation, optimization, and evaluation of ligand-tethered atovaquone-proguanil-loaded nanoparticles for malaria treatment, **2025 Apr;36(6):**711-742. doi: 10.1080/09205063.2024.2422704. Epub 2024 Nov 10., <https://pubmed.ncbi.nlm.nih.gov/39522102/>