

Adaptive Differential Privacy in Federated Learning for Privacy-Preserving ICU Clinical Decision Support at the Edge

Moaad Almania^{1*}, Anazida Zainal², Nurfazrina Binti Mohd Zamry³,
Muhammad Haris⁴, Ahmad Alnawasrah⁵, Fuad A. Ghaleb⁶

Faculty of Computing, Universiti Teknologi Malaysia, Johor Bahru, Johor, 81310, Malaysia¹

College of Computing and Information Technology, Shaqra University, Shaqra, Kingdom of Saudi Arabia¹

Faculty of Computing, Universiti Teknologi Malaysia, Johor Bahru, Johor, 81310, Malaysia²

Faculty of Computing, Universiti Teknologi Malaysia (UTM), Malaysia³

Department of Information Communication Technology, British University of Bahrain, Bahrain⁴

College of Information Technology, Philadelphia University, Jordan⁵

College of Computing and Digital Technology, Birmingham City University, Birmingham, B47XG, United Kingdom⁶

Abstract—Clinical decision assistance that protects patient privacy is a prominent research issue in federated learning, edge computing, and intensive care analytics. However, typical differential privacy algorithms allocate equal privacy budgets to all clinical parameters regardless of sensitivity and predictive power. This may lower privacy-utility trade-offs. This study develops an adaptive differential privacy mechanism for federated learning that dynamically adjusts per-feature Laplace noise and privacy budgets to improve privacy protection while preserving clinical prediction performance. Adaptive Differential Privacy in Federated Learning (ADP-FL) dynamically adjusts per-feature Laplace noise and privacy budgets based on feature-scored sensitivity considering clinical relevance, statistical impact, and re-identification risk. The method was evaluated on the eICU Collaborative Research Database containing 113,265 ICU records and compared against Uniform Differential Privacy and Time-Adaptive Differential Privacy across different privacy budgets (ϵ). Experiments conducted on the eICU Collaborative Research Database achieved test accuracies of up to 82.7%. Across the evaluated privacy budgets, Feature-Adaptive Differential Privacy consistently outperformed both Uniform and Time-Adaptive Differential Privacy, achieving an F1-score of 0.6672 at $\epsilon = 0.5$ while maintaining performance close to the no-DP baseline. These findings demonstrate that sensitivity-aware feature-level privacy budgeting preserves key predictive dimensions while minimizing unnecessary perturbation of less sensitive features, providing a robust foundation for privacy-compliant edge-based clinical decision support systems. The novelty of this study lies in its feature-scored adaptive differential privacy mechanism, which dynamically allocates per-feature privacy budgets and Laplace noise in federated learning for privacy-preserving edge-based clinical decision support.

Keywords—Federated learning; differential privacy; clinical decision support; edge computing; ICU data

I. INTRODUCTION

Clinical Decision Support Systems (CDSS) are increasingly used to support timely prediction, risk assessment, and

treatment-related decision-making in intensive care environments [14]. In ICU settings, large volumes of physiological, laboratory, and treatment data are continuously generated across different hospitals and edge-based clinical systems. Although these data can improve predictive modelling, direct centralization of patient records raises serious concerns related to privacy, governance, and regulatory compliance. Federated Learning (FL) offers a practical solution by enabling multiple clinical sites or edge devices to train a shared model without transferring raw patient data to a central server [15].

Although FL provides this advantage, it does not ensure comprehensive privacy protection independently. There remains the potential for information to be revealed through exchanged model parameters or gradient modifications. This data could subsequently be utilized through membership inference, graph reconstruction, and property inference assaults [16]. In applications within the intensive care unit, these hazards become further pronounced as patient records often encompass highly sensitive clinical attributes. These factors encompass test biomarkers, respiratory state, vital signs, and indicators associated to severity. Therefore, privacy-preserving mechanisms are required to protect shared learning processes while maintaining acceptable clinical prediction performance.

Differential Privacy (DP) is widely used to provide formal privacy protection in machine learning and federated learning systems. However, many existing FL-based healthcare models apply uniform noise using a fixed global privacy budget. This approach treats all features as equally sensitive, although ICU variables do not carry the same level of disclosure risk or clinical importance. Certain metrics, including PaO₂, and creatinine concentrations, bilirubin levels, glucose measurements, dialysis status, and blood urea nitrogen (BUN), might yield more clinically pertinent insights than broad demographic or anonymized administrative variables. Imposing identical noise across all variables may inadvertently disrupt beneficial low-risk predictors, while simultaneously failing to adequately safeguard high-risk clinical characteristics [17].

Recent inquiries have been conducted regarding the adap-

*Corresponding author.

tive and individualized privacy strategies associated with federated learning. Examples of these concepts include adaptive gradient clipping, client-specific privacy budgeting, time-adaptive noise scheduling, and privacy-preserving frameworks for healthcare analytics [18]. Conversely, a considerable portion of these methodologies focuses on imaging datasets, general benchmark datasets, client-specific adaptation, or gradient-level privacy management [20]. Comparatively fewer studies have considered structured ICU tabular data, where each feature may have a different clinical meaning, predictive value, and re-identification risk. This creates a methodological gap for feature-sensitive privacy allocation in federated clinical decision support [19].

Healthcare-based FL studies have also shown that privacy protection often reduces predictive utility, especially when strict privacy budgets are used [21]. This privacy-utility trade-off is particularly important in ICU decision support because false negatives may delay clinical intervention, while excessive noise can weaken model reliability. Existing reviews on differential privacy in medical deep learning further indicate that model performance can decline considerably under strong privacy constraints, especially in heterogeneous and clinically complex datasets [27]. Therefore, there is a need for a privacy mechanism that can apply stronger protection to sensitive ICU features while preserving the predictive contribution of less sensitive but clinically useful variables.

Existing privacy-preserving FL approaches mainly emphasize uniform DP, gradient-level perturbation, time-adaptive noise, client-level adaptation, or medical imaging applications. Limited empirical evidence is available on feature-level adaptive privacy budgeting for structured ICU data. In particular, previous studies have not sufficiently addressed how privacy budgets can be dynamically allocated according to clinical relevance, statistical influence, and re-identification risk of individual ICU variables.

To address this gap, this study proposes Adaptive Differential Privacy in Federated Learning (ADP-FL) for privacy-preserving ICU clinical decision support at the edge. The primary purpose of this research is to create an adaptive differential privacy mechanism. This approach dynamically assigns privacy budgets and Laplace noise per feature based on the sensitivity of the ICU variables. The secondary goal is to assess the possibility of enhancing the privacy-utility trade-off over Uniform Differential Privacy and Time-Adaptive Differential Privacy with this privacy allocation that is sensitive to features. The proposed mechanism assigns per-feature privacy budgets and Laplace noise scales according to feature-scored sensitivity. Highly sensitive clinical variables receive stronger privacy protection through lower privacy budgets and greater perturbation, while less sensitive but predictive variables receive reduced perturbation to preserve model utility. The method is implemented using local client-side noise injection and evaluated on the eICU Collaborative Research Database containing 113,265 ICU records. ADP-FL is compared with Uniform Differential Privacy and Time-Adaptive Differential Privacy across different privacy budgets and gradient clipping norms using accuracy, F1-score, recall, and AUROC. This study contributes a feature-sensitive, lightweight, and governance-compatible privacy mechanism for federated ICU decision support systems [18].

II. RELATED WORK

Federated learning in which privacy is preserved has garnered much research interest, and differential privacy is a key formal privacy guarantee. The novel DP-SGD approach introduced by Abadi, et al. [13] formed the foundation for gradient-level privacy in deep learning. Authors in [12] introduced cross-silo and cross-device FedAvg, which they showed for differential privacy constraints. Mironov [2] proposed Rényi Differential Privacy (RDP) which is a more suitable method for composition accounting for Gaussian processes.

Adaptive and customized privacy methods have been introduced to deal with the problems caused by uniform noise injection. Andrew, et al. [1] suggested adaptive gradient clipping in which the clipping value is learned from statistics collected from the previous iteration of the gradient. Fu, et al. [3] proposed an approach named Adap DP-FL that utilizes RDP accounting and client specific adaptability, with time-decaying noise schedules. In federated health care systems, Wu, et al. [4] examined the privacy-utility trade-offs and noted the need for a calibrated noise in non-IID settings.

Liu, et al. [5] explored the use of federated learning in healthcare, particularly with medical imaging and electronic health record (EHR) data. They pointed out the conflict between regulatory requirements (HIPAA, GDPR) and model efficacy. Saifullah, et al. [6] found that differing levels of privacy risk exist for various clinical traits and so there is a need for a differential level of sensitivity at the feature level. Recent comprehensive reviews have established that differential privacy is the most important privacy method in federated learning healthcare systems, but allocating uniform noise always leads to a decrease in the performance of the model compared to adaptive noise allocation strategies [25].

Adaptive differential privacy has recently been used in therapeutic applications. In order to detect COVID-19 using chest X-ray images, an adaptive differential privacy federated learning architecture was proposed by [7]. This platform is dynamically flexible in limiting privacy based on real-time data sensitivity assessment and shows exceptional performance in a wider range of healthcare settings. Pati, et al. [8] carried out a detailed survey on privacy preservation in federated learning (FL) in healthcare with a focus on gradient-level differential privacy (DP) mechanisms, which are more suitable for rare disease detection than parameter-level mechanisms, but also note that the budgeting distributed to different feature spaces is a significant challenge to address for sensitivity-preserving privacy in FL. A comprehensive review of 74 studies on DP in medical deep learning was performed by [9], which found that enforcing a strict privacy budget ($\epsilon \sim 1$) often leads to a substantial drop in accuracy, especially in smaller and more diverse datasets. This highlights the need to develop and employ feature-adaptive noise allocation strategies, which are proposed in this study.

Previous work on differential privacy for the local setting [10] and PATE [11] mostly focuses on frequency estimation and semi-supervised contexts, both of which are less applicable to the context of continuous clinical feature regression.

III. METHODOLOGY

A. Dataset and Federated Configuration

1) *Dataset*: De-identified patient data from the United States' critical care units was used in the trials in the eICU Collaborative Research Database [19]. The database was accessed through the official eICU/PhysioNet data access pathway, and only de-identified records were used in accordance with the database usage requirements. There were 113,265 patient records with 21 binaries (intubated, vent, dialysis, medications), categorical (Glasgow Coma Scale: eyes, motor, verbal), and continuous (temperature, respiratory rate, heart rate, mean blood pressure, sodium, hematocrit, creatinine, BUN, glucose) attributes. The target variable was mechanical ventilation (vent) with a moderate class imbalance (73.18% negative cases [82,886 records] and 26.82% positive cases [30,379 records]). The target variable was excluded from the predictor set before model training to avoid label leakage [20].

2) *Data preprocessing*: The data set was preprocessed to ensure that the clinical variables could be used for model training prior to federated partitioning. Records were dropped if there were no target labels. For continuous variables, missing values were imputed using the median of the corresponding feature, while missing categorical values were imputed using the mode. All of the features were encoded as categorical variables, and the continuous variables were normalized using the Min-Max scaling method. This was done to ensure that all of the features were within a comparable range prior to the injection of differential privacy noise. The same pipeline of preprocessing was applied to the training data and tested on the rest of the test data that was held back. This action was taken to avert the loss of any data.

3) *Federated configuration*: The complete dataset, including 113,265 entries, was partitioned into training (90,612 entries) and testing (22,653 entries) subsets by stratified sampling. This was executed to preserve class ratios. The training records were distributed evenly across five simulated edge nodes, where each edge device is approximately 18,122 records. This is a rough approximation of a cross-silo healthcare federation with each site being an independent ICU or hospital unit. The weights of aggregation were calculated based on the population size of the local samples,

$$w_i = \frac{n_i}{N}, \quad (1)$$

with each node contributing approximately $w_i \approx 0.20$ under the uniform partitioning setting. The held-out test set was not used during local training or validation and was reserved only for final global model evaluation. Early pausing was implemented if the enhancement in validation accuracy fell below 0.1% for five consecutive communication rounds. The validation set was created from the training data using 20% of training and 80% of testing data [21].

B. Proposed ADP-FL Mechanism

1) *Motivation for adaptive privacy in ICU federated settings*: The rationale for Adaptive Privacy in ICU Federated Settings is compelling. Numerous privacy concerns arise while utilizing ICU data in federated contexts, which influenced the

design of ADP-FL [23]. The attributes of clinical data disclosure vary significantly regarding disclosure risk: laboratory test results (PaO₂, creatinine, bilirubin) directly indicate the severity of organ dysfunction linked to patient identity and prognosis, whereas the lack of patient identity in de-identified demographic characteristics correlates with a substantially reduced likelihood of re-identification under HIPAA Safe Harbor [24]. Applying uniform noise to all characteristics fails to account for this variability, resulting in excessive perturbation of low-risk predictive features and reduced model utility, or insufficient protection of high-risk attributes and increased privacy vulnerability.

Secondly, shared gradient updates in federated learning training may reveal sensitive information if noise levels are insufficient, hence rendering the system vulnerable to membership inference and gradient inversion attacks [22]. Third, the proposed framework targets resource-constrained edge environments; therefore, an efficient privacy mechanism that supports local noise injection without requiring a trusted central curator is required. To address these challenges, ADP-FL dynamically determines the sensitivity of each feature to ensure the privacy of the data is maintained and it is usable across the federation [17].

2) *ADP-FL overview*: By adding on top of standard federated learning a feature-sensitive differential privacy mechanism, ADP-FL achieves a sensitivity score $S_{i,j} \in [0, 1]$ for every feature and then computes a per-feature privacy budget $\epsilon_{i,j}$. Smaller $\epsilon_{i,j}$ values, indicating stronger privacy and greater noise, are assigned to sensitive features, whereas bigger $\epsilon_{i,j}$ values, indicating weaker perturbation and lower noise, are allocated to less sensitive features. Local Laplace noise is introduced at each edge device, followed by gradient aggregation post feature-specific budget calibration. This noise model is local client-side, meaning it does not necessitate a trusted server and may be utilized with secure aggregation methods. In clinical governance contexts, the Laplace composition of privacy is uncomplicated when employing pure composition [26].

Fig. 1 illustrates the comprehensive process of ADP-FL at an edge node, wherein per-feature sensitivity scores are employed to introduce feature-specific Laplace noise before transmitting the aggregated gradients to the central server. The server employs a strictly passive implementation that solely receives noisy merged gradients, computes weighted FedAvg, and transmits the updated global model [27].

3) *Feature sensitivity taxonomy*: Feature sensitivity scores were determined by integrating three criteria: 1) clinical relevance features critical to organ dysfunction assessment or treatment decision-making, 2) statistical influence features with high predictive impact on model gradients, and 3) re-identification risk features whose exposure could lead to patient re-identification under HIPAA Safe Harbor constraints [28].

Based on this assessment, features were grouped into three tiers:

- High-sensitivity ($S_j \approx 0.8-1.0$): PaO₂, creatinine, bilirubin, glucose, dialysis status, BUN, albumin, pH, PaCO₂. These features encode organ dysfunction




Fig. 1. ADP-FL: Adaptive DP with per-feature sensitivity, gradient clipping, Laplace noise at edge, FedAvg at server.

severity central to SOFA and APACHE scoring constructs and carry the highest clinical disclosure risk.

- Moderate-sensitivity ($S_j \approx 0.4-0.7$): Heart rate, mean blood pressure, respiratory rate, temperature, SpO₂, APACHE APS variables, GCS components (eyes, motor, verbal). These vital signs and derived ICU scores are important but less directly re-identifying.
- Low-sensitivity ($S_j \approx 0.1-0.3$): Age, gender, admission type, anonymized unit stay ID. De-identified demographic attributes under HIPAA Safe Harbor, with residual but lower re-identification risk.

4) *Adaptive differential privacy in federated learning*: The ADP-FL operation is executed locally at the edge devices during each communication round. The central server is limited to accepting aggregated noisy gradients, utilizing FedAvg, and transmitting the revised global model. No raw patient data, feature values, sensitivity scores, or utility proxies are transmitted to the central server. Fig. 2 illustrates the whole algorithm [29].

The ADP-FL approach executes operations on edge devices rather than on a central aggregation server. Edge devices construct local models, conduct local training, compute gradients, apply gradient clipping, and incorporate adaptive differential privacy noise, subsequently transmitting the noisy aggregated gradients to the server. To meet the strict local-edge threat model, this architecture ensures that patient data and sensitive

Algorithm 1: ADP-FL (Client-Side, Laplace Mechanism)

Input: ε_b , sensitivity_scale, clipping norm C , local_data, sensitivity scores S_j

For each communication round $t = 1, \dots, T$:

1) Initialize local_gradients = [].

2) For each feature f_j in local_data:

a) Compute per-feature budget:

$$\varepsilon_j = \frac{\varepsilon_b}{S_j \times \text{sensitivity_scale}}.$$

b) Compute gradient:

$$g_j = \nabla L(\theta_i).$$

c) Clip:

$$g_j^{\text{clip}} = \text{clip}(g_j, C).$$

d) Add noise:

$$g_j^{\text{dp}} = g_j^{\text{clip}} + \text{Laplace}\left(0, \frac{C}{\varepsilon_j}\right).$$

e) Append g_j^{dp} to local_gradients.

3) Aggregate local_gradients and send to server.

Server: $\theta_{\text{global}} = \sum_i w_i \theta_i$, where $w_i = n_i/N$.

Privacy accounting: Total $\leftarrow \varepsilon_t + \varepsilon_t$ (pure- ε composition).

Fig. 2. Adaptive differential privacy in federated learning algorithm.

metadata remain within the local edge device during processing [30].

The number of communication rounds (NUM_COMM_ROUNDS) is limited to 100 rounds, and the early stopping is based on the first time the validation accuracy improvement is below 0.1 per cent for five consecutive communication rounds. Aggregation weights w_i are proportional to local sample counts; under the equal distribution adopted here, each of the five nodes contributes $w_i \approx 0.20$.

The adaptive privacy operation can be summarized as follows. For each client i and feature j , a sensitivity score $S_{i,j}$ is first assigned. The privacy budget $\varepsilon_{i,j}$ is then adjusted according to the sensitivity level, where highly sensitive features receive smaller $\varepsilon_{i,j}$ values and therefore stronger Laplace perturbation. After local gradient clipping using clipping norm C , Laplace noise is applied before the noisy update is sent for FedAvg aggregation.

5) *Parameter configuration*: Table I summarizes the key parameters and their justification [16].

IV. RESULTS AND DISCUSSION

A. Privacy-Accuracy Trade-off

Table II displays the test accuracy of the SGD Classifier across several privacy budgets $\varepsilon \in \{0.1, 0.2, 0.3, 0.5, 0.7, 1.0, 1.5, 2.0\}$ and gradient clipping norms $C \in \{1.0, 3.0, 5.0, 10.0\}$, in comparison to a non-private Random Forest Classifier baseline (83.27%). The interval [1,10] is evaluated using the clipping norm C to mitigate the bias associated with excessive clipping and the introduced noise using recent privacy-budget and clipping-norm practices in differential privacy-based learning. Privacy budget assessment shows that the privacy-budget sweep $\varepsilon \in [0.1, 2.0]$ encompasses the privacy frameworks

TABLE I. ADP-FL PARAMETER VALUES AND JUSTIFICATION

Parameter	Value	Justification
ϵ_b	1.0	Baseline privacy budget balancing privacy and performance
Sensitivity scale	0.5–2.0	Scales noise to feature sensitivity, higher values apply more noise to sensitive features
Clipping norm C	1.0	Ensures gradient stability in federated updates, limits adversarial exploitation of large gradients
Learning rate η	0.01	Standard SGD step size for tabular clinical data
Mini-batch size	32	Balances computational efficiency and gradient stability
NUM COMM ROUNDS	100	With early stopping when accuracy improvement $< 0.1\%$ for 5 consecutive rounds

TABLE II. SGD CLASSIFIER TEST ACCURACY VS. PRIVACY BUDGET (ϵ) AND CLIPPING NORM (C)

$\epsilon \backslash C$	$C = 1.0$	$C = 3.0$	$C = 5.0$
0.1	82.66%	81.81%	82.37%
0.2	82.51%	82.33%	81.76%
0.3	82.34%	81.45%	82.35%
0.5	81.88%	81.82%	82.60%
0.7	81.83%	81.75%	82.51%
1.0	82.47%	82.76%	81.93%
1.5	82.73%	82.55%	81.72%
2.0	81.75%	81.80%	82.62%

of both high-privacy and moderate-privacy regimes typically observed in practical applications [31].

Non-private Random Forest baseline: 83.27%. Noteworthy observations from Table II include the accuracy of the SGD Classifier fluctuating with different epsilon values, influenced by the selected clipping norm, as most clipping norms are sensitive to the noise-utility trade-off in differential privacy. Optimal accuracies are noted when epsilon falls between 0.5 and 1.5. This suggests that moderate privacy budgets establish an appropriate equilibrium between the incorporation of noise and the model's efficacy. Reducing the epsilon value, such as to 0.1, results in a rise in the added noise; nonetheless, the accuracy remains mostly unchanged, suggesting a robustness to differential privacy noise.

Standards for gradient clipping exert the following influence: The execution aligns with a Clipping Norm of 1.0, with a peak accuracy of 82.73 percent when ϵ is set to 1.5. The clipping norm of 3.0 strikes the optimal balance between privacy and precision, yielding the maximum accuracy of 82.76 percent when ϵ is configured to 1.0. The clipping thresholds of 5.0 and 10.0 yield a relatively diminished and more erratic accuracy, especially in the context of stringent privacy constraints. This occurs because increased clipping thresholds amplify the noise scaling in relation to C/ϵ and elevate the variance of the introduced perturbations. Employing diminished clipping norms ($C = 1.0, 3.0$) facilitates the optimization of the equilibrium between privacy safeguards and the reduction of the impact of introduced noise, resulting in diminished volatility across iterations and participants.

Comparison against the baseline: The baseline Random Forest Classifier, which does not use differential privacy, achieves an accuracy of 83.27%. When appropriately calibrated, the SGD Classifier with differential privacy can achieve accuracy comparable to its non-private equivalent (between 82.5% and 82.76%), demonstrating that privacy can be pre-




Fig. 3. Privacy-accuracy trade-off: SGDClassifier gradient clipping vs. RandomForest baseline.

served without significant detriment to model efficacy, with a margin of less than one percentage point.

Fig. 3 illustrates the trade-off between epsilon of the privacy budget and model accuracy over several gradient clipping norms. The optimal trade-off occurs at an epsilon range of 0.7 to 1.5 and a Clipping Norm of either 1.0 or 3.0, yielding an accuracy of roughly 82.7 while ensuring adequate privacy protection.

B. Dynamic Sensitivity Laplace vs. Baseline

This section compares a dynamic, feature-sensitivity-aware Laplace mechanism (ADP-FL) to a uniform differential privacy baseline spanning privacy budgets $\epsilon \in \{0.2, 0.5, 1.0, 1.5, 2.0\}$. The model, training process, and evaluation split are constant, with the sole variation being the noise applied to training features after Min-Max scaling, while the test set remains unchanged. Table III presents the Accuracy, Precision, Recall, F1, and AUROC of both techniques at each ϵ .

Table III reports the controlled comparison between Dynamic (Feature-Adaptive) DP and Uniform DP using the same dataset, train-test split, federated-learning configuration, preprocessing, and evaluation protocol. Only the privacy-noise allocation strategy is changed between the two methods. Several significant tendencies can be discerned from Table III. Under the highest privacy level ($\epsilon = 0.2$), the uniform strategy achieves superior metrics with an F1 score of 0.6540 and a Recall of 0.6363, compared to the dynamic strategy, which has an F1 score of 0.5114 and a Recall of 0.3754. This is consistent with the tenets of mechanism design: if ϵ is low, then the dynamic policy will also amplify noise on high-

TABLE III. PERFORMANCE ACROSS ϵ FOR DYNAMIC (ADAPTIVE) VS. UNIFORM DP

Strategy	i	ϵ	Accuracy	Precision	Recall	F1	AUROC
Dynamic	1	0.2	0.8076	0.8020	0.3754	0.5114	0.8256
Uniform	2	0.2	0.8194	0.6727	0.6363	0.6540	0.8276
Dynamic	3	0.5	0.8340	0.7343	0.5971	0.6586	0.8223
Uniform	4	0.5	0.8215	0.7063	0.5727	0.6326	0.8202
Dynamic	5	1.0	0.8342	0.7283	0.6093	0.6635	0.8276
Uniform	6	1.0	0.8194	0.7478	0.4933	0.5944	0.8219
Dynamic	7	1.5	0.8347	0.7256	0.6170	0.6669	0.8267
Uniform	8	1.5	0.8226	0.7265	0.5433	0.6217	0.8208
Dynamic	9	2.0	0.8351	0.7147	0.6414	0.6760	0.8269
Uniform	10	2.0	0.8223	0.7170	0.5576	0.6273	0.8280

sensitivity features, generating too much perturbation on the most clinically informative features and thus making it harder to detect minority groups. The AUROC scores are nearly the same (0.83) at $\epsilon = 0.2$, meaning there has been very little loss of separability, with a decrease in performance primarily due to the balance of false negatives with respect to false positives. Table III reports the controlled comparison between Dynamic (Feature-Adaptive) DP and Uniform DP using the same dataset, train-test split, federated-learning configuration, preprocessing, and evaluation protocol. Only the privacy-noise allocation strategy is changed between the two methods.

As ϵ attains a moderate value, control over dominance will evidently shift to the dynamic mechanism. At $\epsilon = 0.5$, ADP-FL exhibits an Accuracy of approximately 0.834 and an F1 score of approximately 0.6586, compared to approximately 0.8215 and approximately 0.6326 for uniform, with Recall similarly superior for dynamic (approximately 0.597 versus approximately 0.573). The improvement is evident at $\epsilon = 1.0$, with dynamic metrics of Accuracy = 0.834 and F1 = 0.664, compared to the uniform baseline of 0.819 and 0.594. The Recall disparity is most pronounced at this operational point (0.609 versus 0.493). The trend persists until $\epsilon = 1.5$ and $\epsilon = 2.0$, during which the dynamic strategy maintains F1 within the 0.667–0.676 interval.

Fig. 4 illustrates a graph of Accuracy versus ϵ . The dynamic curve exhibits an ascent between 0.2 and 2.0, surpassing the uniform curve by around one to one-and-a-half percentage points at values of 0.5 and higher. The uniform curve is comparatively level near 0.82, indicating that the classifier utilizes indiscriminate noise to leverage more permissive privacy expenditure.

Fig. 5 illustrates the relationship between F1 and ϵ for dynamic and uniform noise. The F1 score under the dynamic method exhibits a monotonic improvement, increasing from around 0.511 at $\epsilon = 0.2$ to approximately 0.676 at $\epsilon = 2.0$. The uniform strategy is non-monotonic and has a notable minimum at $\epsilon = 1.0$ to 0.594, just when the dynamic policy attains 0.664. This effect can be elucidated by analyzing Precision and Recall: dynamic Recall with $\epsilon = 1.0$ is markedly superior (about 0.609 compared to approximately 0.493), but Precision is only somewhat worse (approximately 0.728 compared to approximately 0.748). The dynamic technique effectively reduces false negatives at ϵ values within the 0.5–2.0 range, albeit with a minor compromise in false positives concerning decision support.




Fig. 4. Accuracy vs. ϵ for ADP-FL vs. Uniform DP.




Fig. 5. F1-score vs. ϵ for dynamic (ADP-FL) vs. uniform DP.

Fig. 6 presents a detailed depiction of the aggregations of principal metrics at the value of $\epsilon = 1.0$. The dynamic strategy exhibits superior Accuracy (0.834 vs. 0.819), Recall (0.609 vs. 0.493), and F1 score (0.664 vs. 0.594) at this mid-range operating point compared to the uniform baseline, despite a very small decline in Precision, which correlates with a shift




Fig. 6. Performance at $\epsilon = 1.0$: Dynamic (ADP-FL) vs. uniform DP bar chart of accuracy, precision, recall, F1, and AUROC.




Fig. 7. Dynamic ϵ allocation per feature at $\epsilon_b = 1.0$.

in focus towards sensitivity. This is the typical threshold used in safety applications where there are relatively low costs for dealing with false alarms and high costs for missing a positive indication.

Fig. 7 shows the per-feature ϵ of the dynamic policy for the case $\epsilon_{\text{base}} = 1.0$, which explains the design rationale of the above results: clinical variables with high sensitivity (like PaO_2 , creatinine, glucose) have a lower ϵ (higher noise); low sensitivity (like age, gender) have a higher ϵ (lower noise); middle sensitivity clinical variables (like vital signs) get an intermediate ϵ . This allocation rationale is corroborated by the results in Recall and F1 at moderate ϵ : the sensitivity-aware noise strategy is better at retaining valuable signals in less sensitive but still valuable features in comparison to a uniform policy.

C. Three-Way Privacy-Utility Comparison

Table IV presents an additional three-way comparison including Time-Adaptive DP; therefore, Table III is used as the primary controlled comparison between Feature-Adaptive DP and Uniform DP. Fig. 8 provide a comprehensive comparison of three differential privacy strategies: Uniform DP, Time-Adaptive DP, and Feature-Adaptive DP (ADP-FL), alongside a No-DP baseline ($F1 = 0.6668$) for privacy budgets $\epsilon \in$

TABLE IV. F1-SCORE COMPARISON OF UNIFORM, TIME-ADAPTIVE, AND FEATURE-ADAPTIVE DP ACROSS ϵ . NO-DP BASELINE $F1 = 0.6668$

ϵ	Uniform DP	Time-Adaptive DP	Feature-Adaptive DP
0.1	0.4009	0.1458	0.4511
0.3	0.3584	0.3062	0.6011
0.5	0.2171	0.4642	0.6672
1.0	0.0724	0.4479	0.5118
2.0	0.4159	0.4919	0.5952

Comparison among No-DP, Uniform, Time-adaptive and Feature-Adaptive DP




Fig. 8. F1-score comparison of No-DP, uniform DP, time-adaptive DP, and feature-adaptive DP.

$\{0.1, 0.3, 0.5, 1.0, 2.0\}$. Table IV illustrates a clear behavioral divergence among the three techniques. For the value of $\epsilon = 0.1$, Feature-Adaptive DP achieves the highest F1 score of 0.4511 while Uniform DP comes in at 0.4009 and is much higher than Time-Adaptive DP with a score of 0.1458.

For $\epsilon = 0.3$, the Feature-Adaptive DP performs significantly better ($F1 = 0.6011$ compared to 0.3584 in Uniform and 0.3062 in Time-Adaptive), showing that, under a moderate privacy budget, adaptive feature-level allocation is indeed much more effective at reducing degradation of minority classes resulting from noise disruption. When $\epsilon = 0.5$, the score of Feature-Adaptive DP becomes 0.6672, which is nearly the same as the score of No-DP with the baseline score of 0.6668. In contrast, Uniform DP is drastically diminished ($F1 = 0.2171$), whereas TA DP is improved to 0.4642, albeit at a considerably lower level. The significant inadequacy of Uniform DP at $\epsilon = 1.0$ ($F1 = 0.0724$) demonstrates that homogeneous noise calibration is unstable when it engages with non-uniform feature gradients. At $\epsilon = 1.0$, Feature-Adaptive DP achieves $F1 = 0.5118$, and at $\epsilon = 2.0$, all mechanisms converge to the No-DP baseline, Feature-Adaptive DP (0.5952) demonstrates superior performance.

Fig. 8 illustrates these trajectories. The FADP curve is maximized throughout all evaluated budgets, with the exception of the separate $\epsilon = 0.2$, where uniform noise temporarily outperforms due to the compounding noise effect on high-sensitivity features. Practitioners in regulatory environments necessitating low ϵ should meticulously observe this crossover at the highest privacy thresholds.

The mechanistic rationale for the superiority of ADP-

FL is that high-sensitivity clinical features receive stronger perturbation, hence preserving the gradient signal that facilitates the memory of minority classes. Uniform differential privacy applies the same amount of noise to each feature, regardless of its predictive or privacy significance, leading to an overemphasis on informational gradients that adversely affect the imbalanced ventilation prediction task examined in this instance.

D. Literature-Level Comparison with Recent Published Studies

To meet the need of wider validation with the recent literature, Table V presents a comparison between the proposed ADP-FL model and the recently published federated learning studies for privacy, in the range of 2020–2026. The comparison reveals that recent work on healthcare FL has been successful; most of the studies concentrated on imaging data, privacy in blockchain, multimodal fusion, or general healthcare analytics, but did not include FL adaptive privacy budgeting for structured data in an ICU decision support system. The comparison indicates that the proposed ADP-FL model is not intended to simply maximize classification accuracy across different healthcare datasets; instead, its main contribution is to preserve ICU prediction performance while providing feature-sensitive privacy protection. Unlike blockchain-based and multimodal FL approaches, ADP-FL directly addresses the unequal privacy sensitivity of ICU variables by assigning different privacy budgets to different clinical features. This makes the proposed model more suitable for edge-based ICU clinical decision support where laboratory biomarkers, vital signs, and demographic attributes do not carry equal disclosure risk. Therefore, the new information offered by this study is not only improved internal performance over DP baselines, but also evidence that feature-level privacy allocation can preserve ICU prediction utility more effectively than uniform privacy treatment when clinical variables have unequal sensitivity and predictive importance. The performance values of the comparator studies in Table V are taken directly from their respective publications. Since these studies use different datasets, models, prediction tasks, and privacy settings, the comparison is literature-level rather than a direct head-to-head experimental benchmark.

The reported performance values should not be interpreted as a direct ranking of the methods because the studies use different datasets, prediction tasks, models, and privacy configurations. The contribution of ADP-FL is instead demonstrated through its controlled comparison with the privacy baselines under identical experimental conditions, where feature-sensitive privacy allocation preserves predictive utility while providing privacy protection for structured ICU data.

E. Operational Interpretation of Privacy-Utility Findings

For operational deployments, all evidence in favor of ADP-FL is within the range of $\varepsilon \in [0.5, 2.0]$. During this period, accuracy increases by approximately one percentage point or more when compared to uniform noise, while F1 also improves due to higher recall while maintaining a similar AUROC value. The operating margin is comparable with the typical ε of the order of few digits reported in current clinical machine learning

pipelines. If the policy requires strong privacy ($\varepsilon = 0.1$ – 0.2), then the dynamic mechanism may over-perturb high-sensitivity features, which can reduce minority-class detection. In such cases, practitioners are advised to carefully consider the clinical miss-rate tolerance and privacy considerations.

The fact that privacy is pure- ε Laplace allows for simpler privacy accounting for clinical governance and regulatory reporting than (ε, δ) -DP techniques, and ADP-FL is well suited for installations regulated by HIPAA or GDPR, where audit trails require transparency in privacy accounting. The steady improvement of Recall with ADP-FL is important for use in the context of healthcare decision support in ICUs. The clinical cost of missing a mechanical ventilation warning is higher than the cost of a false alarm, and the dynamic control mechanism that keeps recall close to the No-DP baseline and at the same time ensures formal privacy is directly related to this cost imbalance.

F. Comparison with Alternative DP Approaches

When a trusted central aggregator or safe compute environment exists, Gaussian-noise DP-SGD may prove to be better suited to deeper models by altering the gradient-level noise based on the dynamic of the perturbation, instead of the scale of the features themselves. Standard smoothness assumptions give rise to an interesting pure- ε central differential privacy analog to objective perturbation in the convex learner case, e.g., logistic regression. Both strategies, however, presuppose a more reliable model than the local-edge danger model presented here, wherein data must be perturbed in situ prior to any uplink transmission.

Within the local-DP framework, the dynamic allocation of Laplace noise, which is responsive to variations in sensitivity, performed as the most suitable option in this study. Recent advancements in differential privacy for medical deep learning [27] suggest that other forms of DP, including generative modeling-based DP and hybrid federated designs, are forthcoming; nevertheless, these approaches have predominantly been evaluated using imaging data rather than structured tabular clinical data. The distinction between local Laplace and central Gaussian DP-SGD in terms of trust model and injection locus is not one of absolute superiority, but rather pertains to the trust model and the acceptability of central training: ADP-FL should be utilized when a trusted curator is not permitted, while Gaussian DP-SGD is likely to yield greater utility when central training is permissible.

For practical deployment feasibility assessment, the computational and communication features of ADP-FL were compared to FL without DP, Uniform DP-FL, and Time-Adaptive DP-FL. Additional overhead of ADP-FL is small as it does feature sensitivity scoring once before training and adds Laplace noise locally for each round of communication. Furthermore, the communication cost is similar to that of standard federated learning, where only updated noisy models are sent, and the original data, feature values, sensitivity scores, and feature-specific privacy budgets are not sent from the edge nodes.

G. Computational Overhead, Scalability, and Generalizability

To evaluate practical deployment feasibility, the computational and communication characteristics of ADP-FL were

TABLE V. LITERATURE-LEVEL COMPARISON WITH RECENT PRIVACY-PRESERVING FL APPROACHES

Reference	Application/Dataset	Privacy Approach	Reported Performance	Difference from Present Study
Proposed ADP-FL	eICU; 113,265 ICU records	Feature-adaptive DP with per-feature Laplace noise	Accuracy 82.76%; F1 = 0.6672 at $\epsilon = 0.5$	Feature-level privacy allocation for structured ICU data.
Fu et al. [3]	MNIST and Fashion-MNIST	Adaptive clipping and time-decaying Gaussian noise with RDP	94.15% accuracy on MNIST and 79.56% on Fashion-MNIST at $\epsilon = 2$	Adapts privacy across clients and rounds rather than individual clinical features.
Wu et al. [4]	CIFAR-10	Adaptive federated learning with differential privacy	81.25% accuracy (IID) and 79.05% (non-IID)	Uses adaptive DP for image classification rather than feature-sensitive ICU data.
Ahmed et al. [7]	COVID-19 chest X-ray images	Adaptive differential privacy federated learning	91.2% accuracy after 80 epochs with early stopping	Evaluated on medical images rather than structured ICU records.

compared with FL without DP, Uniform DP-FL, and Time-Adaptive DP-FL. Since ADP-FL performs feature-sensitivity scoring once before training and applies per-feature Laplace noise locally during each communication round, its additional overhead remains small. Moreover, the communication cost is comparable to standard federated learning because only noisy model updates are transmitted, while raw data, feature values, sensitivity scores, and feature-specific privacy budgets remain within the local edge nodes.

Table VI shows that for $\epsilon = 0.5$, ADP-FL has the highest F1-score and other federated approaches have similar communication costs. The minor extra training cost is compensated by the better privacy-utility trade-off, as the proposed approach maintains performance close to the no-DP baseline. Additional multi-institutional and heterogeneous ICU datasets are needed, however, for further validation to prove the robustness and generalizability of ADP-FL in other real-world healthcare settings.

TABLE VI. PRACTICAL OVERHEAD AND UTILITY COMPARISON OF FEDERATED PRIVACY STRATEGIES

Method	Clients	Rounds	Relative Training Overhead	Relative Communication Cost	F1 at $\epsilon = 0.5$
FL without DP	5	100	1.00×	1.00×	0.6668
Uniform DP-FL	5	100	1.01×	1.00×	0.2171
Time-Adaptive DP-FL	5	100	1.03×	1.00×	0.4642
Proposed ADP-FL	5	100	1.02×	1.00×	0.6672

V. CONCLUSION

This research presented and evaluated an adaptive differential privacy approach (ADP-FL) for federated learning in edge-based clinical decision assistance within the ICU. ADP-FL exhibits a markedly superior privacy-utility trade-off compared to Uniform DP and Time-Adaptive DP by dynamically allocating per-feature Laplace noise in accordance with clinically motivated sensitivity scores over the majority of privacy budgets. The novelty of this work is that it allocates feature-level privacy budgets based on clinical relevance, statistical influence, and re-identification risk, instead of having a uniform budget for all the ICU variables. ADP-FL demonstrates strong performance on the eICU database, achieving a test accuracy of 82.7%, closely aligning with the non-private baseline of 83.27%, and an F1-score of 0.6672 at $\epsilon = 0.5$, nearly indistinguishable from the No-DP performance. These results demonstrate that the distribution of noise among the features is as significant as the quantity of noise applied. ADP-FL provides a governance-light, operational privacy-preserving edge Clinical Decision Support System (CDSS) implementation: pure- ϵ composition is audit-transparent, local noise

injection requires no curator, and secure aggregation limits server involvement. The results showed that feature-sensitive privacy allocation is a better approach than uniform noise injection for ICU-based edge deployments, especially when clinical variables have varying disclosure risk and predictive value. The mechanism will grow in three ways in the future: 1) testing more complex classifiers while keeping the same privacy accounting; 2) creating or scheduling sensitivity weights instead of using fixed weights; and 3) combining privacy-budget accounting with edge resource limitations (latency, energy) to allow for a larger adaptive scheduling framework, which is at the heart of ADP-FL's privacy feature. Furthermore, to ensure the generalizability of the findings, further investigation is necessary to evaluate ADP-FL on more multi-institutional critical care unit datasets and across a broader spectrum of non-IID federated contexts. Future research may investigate if metaheuristic optimization methods like genetic algorithms and ant colony optimization can assist sensitivity weight scheduling.

ACKNOWLEDGMENT

The authors would like to thank to College of Computing and Information Technology, Shaqra University, Shaqra, Kingdom of Saudi Arabia for providing the funding to the student do research work in UTM.

REFERENCES

- [1] G. Andrew, O. Thakkar, B. McMahan, and S. Ramaswamy, "Differentially private learning with adaptive clipping," in *Advances in Neural Information Processing Systems*, vol. 34, pp. 17455–17466, 2021.
- [2] I. Mironov, "Rényi differential privacy," in *Proc. 2017 IEEE 30th Computer Security Foundations Symposium (CSF)*, pp. 263–275, 2017, doi: 10.1109/CSF.2017.11.
- [3] J. Fu, Z. Chen, and X. Han, "Adap DP-FL: Differentially private federated learning with adaptive noise," in *Proc. 2022 IEEE International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom)*, pp. 656–663, 2022, doi: 10.1109/TrustCom56396.2022.00094.
- [4] X. Wu, Y. Zhang, M. Shi, P. Li, R. Li, and N. N. Xiong, "An adaptive federated learning scheme with differential privacy preserving," *Future Generation Computer Systems*, vol. 127, pp. 362–372, 2022, doi: 10.1016/j.future.2021.09.015.
- [5] H. Liu, C. Li, B. Liu, P. Wang, S. Ge, and W. Wang, "Differentially private learning with grouped gradient clipping," in *Proc. 3rd ACM International Conference on Multimedia in Asia (MMAsia)*, Art. no. 14, pp. 1–7, 2021, doi: 10.1145/3469877.3490594.
- [6] S. Saifullah, D. Mercier, A. Lucieri, A. Dengel, and S. Ahmed, "The privacy-explainability trade-off: Unraveling the impacts of differential privacy and federated learning on attribution methods," *Frontiers in Artificial Intelligence*, vol. 7, Art. no. 1236947, 2024, doi: 10.3389/frai.2024.1236947.

- [7] M. W. Rahman, M. R. Khan, and M. Nijim, "Privacy-preserving deep learning for disease diagnosis in medical imaging: A systematic review," *IEEE Access*, vol. 14, pp. 7248–7269, 2026, doi: 10.1109/ACCESS.2025.3647561.
- [8] S. Pati, S. Kumar, A. Varma, B. Edwards, C. Lu, L. Qu, *et al.*, "Privacy preservation for federated learning in health care," *Patterns*, vol. 5, no. 7, Art. no. 100974, 2024, doi: 10.1016/j.patter.2024.100974.
- [9] M. Mohammadi, M. Vejdanihemmat, M. Lotfinia, M. Rusu, D. Truhn, A. Maier, and S. Tayebi Arasteh, "Differential privacy for medical deep learning: Methods, tradeoffs, and deployment implications," *npj Digital Medicine*, vol. 9, Art. no. 93, 2026, doi: 10.1038/s41746-025-02280-z.
- [10] N. T. Priya, V. S. Kumari, and K. J. Peter, "Evaluating the role of RAPPOR in strengthening data privacy within cybersecurity practices," in *Proc. 2024 International Conference on Communication, Computing, Smart Materials and Devices (ICCCSMD)*, pp. 1–5, 2024, doi: 10.1109/ICCCSMD63546.2024.11015118.
- [11] Z. Qiao, P. Wang, P. Wang, Z. Ning, Y. Fu, Y. Du, *et al.*, "A dual-channel semi-supervised learning framework on graphs via knowledge transfer and meta-learning," *ACM Transactions on the Web*, vol. 18, no. 2, Art. no. 18, pp. 1–26, 2024, doi: 10.1145/3577033.
- [12] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. Agüera y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Proc. 20th International Conference on Artificial Intelligence and Statistics (AISTATS), Proceedings of Machine Learning Research*, vol. 54, pp. 1273–1282, 2017.
- [13] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang, "Deep learning with differential privacy," in *Proc. 2016 ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pp. 308–318, 2016, doi: 10.1145/2976749.2978318.
- [14] A. A. Baylor, J. Li, I. A. Yang, and M. Varnfield, "Designing clinical decision support systems (CDSS)—A user-centered lens of the design characteristics, challenges, and implications: Systematic review," *Journal of Medical Internet Research*, vol. 27, Art. no. e63733, 2025, doi: 10.2196/63733.
- [15] S. Zhan, L. Huang, G. Luo, S. Zheng, Z. Gao, and H.-C. Chao, "A review on federated learning architectures for privacy-preserving AI: Lightweight and secure cloud-edge-end collaboration," *Electronics*, vol. 14, no. 13, Art. no. 2512, 2025, doi: 10.3390/electronics14132512.
- [16] M. K. Quan, P. N. Pathirana, M. Wijayasundara, S. Setunge, D. C. Nguyen, C. G. Brinton, *et al.*, "Federated learning for cyber physical systems: A comprehensive survey," *IEEE Communications Surveys & Tutorials*, 2025, doi: 10.1109/COMST.2025.3570288.
- [17] X. Xu, Y. Li, L. Lu, and T. Wang, "ADPFL: Adaptive differential privacy-enhanced federated learning," *IEEE Internet of Things Journal*, vol. 12, no. 23, pp. 50682–50692, 2025, doi: 10.1109/IJOT.2025.3611410.
- [18] G. Andrew, O. Thakkar, B. McMahan, and S. Ramaswamy, "Differentially private learning with adaptive clipping," in *Advances in Neural Information Processing Systems*, vol. 34, pp. 17455–17466, 2021.
- [19] J. Fu, Z. Chen, and X. Han, "Adap DP-FL: Differentially private federated learning with adaptive noise," in *Proc. IEEE International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom)*, 2022, doi: 10.1109/TrustCom56396.2022.00094.
- [20] X. Wu, Y. Zhang, M. Shi, P. Li, R. Li, and N. N. Xiong, "An adaptive federated learning scheme with differential privacy preserving," *Future Generation Computer Systems*, vol. 127, pp. 362–372, 2022, doi: 10.1016/j.future.2021.09.015.
- [21] M. Adnan, S. Kalra, J. C. Cresswell, G. W. Taylor, and H. R. Tizhoosh, "Federated learning and differential privacy for medical image analysis," *Scientific Reports*, vol. 12, no. 1, Art. no. 1953, 2022, doi: 10.1038/s41598-022-05539-7.
- [22] J. Geiping, H. Bauermeister, H. Dröge, and M. Möller, "Inverting gradients—How easy is it to break privacy in federated learning?" in *Advances in Neural Information Processing Systems*, vol. 33, pp. 16937–16947, 2020.
- [23] M. Almanian, A. Zainal, F. A. Ghaleb, A. Alnawasrah, and M. Al Qerom, "Adaptive intrusion detection system with ensemble classifiers for handling imbalanced datasets and dynamic network traffic," *Journal of Robotics and Control*, 2025, doi: 10.18196/jrc.v6i1.23648.
- [24] M. Almanian, A. Zainal, F. A. Ghaleb, A. Alnawasrah, and M. Al Qerom, "Two-level feature selection for enhanced accuracy and reduced computational overhead in intrusion detection systems using rough set theory and binary particle swarm optimization," *Journal of Robotics and Control*, vol. 6, no. 1, pp. 262–271, 2025, doi: 10.18196/jrc.v6i1.23649.
- [25] A. S. Khattak, A. B. Zain, R. B. Hassan, F. Nazar, M. Haris, and B. A. Ahmed, "Hand gesture recognition with deep residual network using sEMG signal," *Biomedical Engineering/Biomedizinische Technik*, vol. 69, no. 3, pp. 275–291, 2024, doi: 10.1515/bmt-2023-0208.
- [26] S. Pati, S. Kumar, A. Varma, B. Edwards, C. Lu, Q. Le, J. J. Wang, A. Lakshminarayanan, S. H. Wang, M. J. Sheller, and K. Chang, "Privacy preservation for federated learning in health care," *Patterns*, vol. 5, no. 7, Art. no. 100974, 2024, doi: 10.1016/j.patter.2024.100974.
- [27] M. Mohammadi, M. Vejdanihemmat, M. Lotfinia, M. Rusu, D. Truhn, A. Maier, and S. T. Arasteh, "Differential privacy for medical deep learning: Methods, tradeoffs, and deployment implications," *npj Digital Medicine*, vol. 9, Art. no. 93, 2026, doi: 10.1038/s41746-025-02280-z.
- [28] K. H. Sayed, M. J. Cruyff, and P. G. M. van der Heijden, "Modeling evasive response bias in randomized response: Cheater detection versus self-protective no-saying," *Psychometrika*, vol. 89, no. 4, pp. 1261–1279, 2024, doi: 10.1007/s11336-024-10000-x.
- [29] Z. Qiao, P. Wang, P. Wang, Z. Ning, Y. Fu, Y. Du, Y. Zhou, J. Huang, X.-S. Hua, and H. Xiong, "A dual-channel semi-supervised learning framework on graphs via knowledge transfer and meta-learning," *ACM Transactions on the Web*, vol. 18, no. 2, pp. 1–26, 2024, doi: 10.1145/3577033.
- [30] W. Zhang, L. Wu, and S. Zhang, "Clinical phenotype of ARDS based on K-means cluster analysis: A study from the eICU database," *Heliyon*, vol. 10, no. 20, Art. no. e39198, 2024, doi: 10.1016/j.heliyon.2024.e39198.
- [31] S. M. Orthi, M. H. Rahman, K. B. Siddiqua, M. Uddin, S. Hossain, A. Al Mamun, and M. N. Khan, "Federated learning with privacy-preserving big data analytics for distributed healthcare systems," *Journal of Computer Science and Technology Studies*, vol. 7, no. 8, pp. 269–281, 2025, doi: 10.32996/jcsts.2025.7.8.31.