Journal of Computational Medicine and Informatics
Vol. 1, No. 3, Jul - Sep 2025, pp. 39-46
ISSN: 3107-3735 (O), DOI: https://doi.org/10.17051/JCMI/01.03.05

Privacy-Aware Deep Learning Architecture for Secure Electronic Health Record Analysis

K. Chitra

Associate Professor & Head,, PG Department of Data Science, KPR College of Arts Science and Research, Coimbatore,Tamilnadu, India. Email: ganey.c@gmail.com, Orcid ID: 0000-0002-8088-4517

Article Info

Article history:

Received : 16.07.2025
Revised : 15.08.2025
Accepted : 13.09.2025

Keywords:

Electronic Health Records,
Privacy-Preserving Learning,
Federated Learning,
Differential Privacy,
Deep Learning,
Healthcare Analytics

ABSTRACT

The accelerated process of healthcare system digitalization has greatly augmented the use of Electronic Health Records (EHRs) that can support the further development of data-driven clinical decision-making. Nonetheless, the privacy and security issues that arise in EHR data are so sensitive that they restrict the feasibility of the deep learning models application to real-life medical settings. This paper has developed a privacy conscious deep learning framework that can combine federated learning with differential privacy in order to achieve decentralized EHR analytics that is secure. The suggested model allows joint model training among various hospitals without the need to disseminate raw patient data, which also protects the confidentiality of the data. The researchers have used a hybrid deep learning model which is a combination of Long Short-Term Memory (LSTM) networks and attention mechanism that is able to capture temporal dependencies alongside feature relevance in longitudinal clinical data. Differential privacy is implemented as a controlled noise injection to updates in a model, which offers formal privacy guarantees in addition to model utility. The experimental analysis of the MIMIC-III data shows that the proposed method obtains better predictive performance, with an accuracy rate of 90% and AUC of 0.92 as well as a high level of privacy protection. Its findings indicate that there is a trade-off of acceptable accuracy on both model accuracy and preservation of privacy, thus validating the appropriateness of the suggested architecture in regard to secure, scalable, and real-world clinical deployment.

1. INTRODUCTION

The use of Electronic Health Records (EHRs) has transformed contemporary healthcare in that it allows collecting longitudinal patient data and making clinical decisions based on data. EHRs include time-varying and heterogeneous data like diagnoses, medications, laboratory and treatment histories which make them a useful source of predictive analytics and personalized medicine [1], [2]. Recent improvements in deep learning have shown much success in identifying complex temporal patterns on EHR data to be used in areas such as disease prediction, risk stratification, and clinical decision support [3], [4]. Irrespective of these developments, the privacy and security of healthcare data are very sensitive, making it highly challenging. Deep learning methods that are centralized entail patient data aggregation and consequently the risk of data loss and data breaches. Further, rigid regulatory frameworks like the Health Insurance Portability and Accountability Act (HIPAA) and the General Data Protection Regulation (GDPR) place close restrictions on the data sharing and processing [5], [6]. These constraints render the significant application of the deep learning models in clinical practice difficult. Privacy-preserving machine learning methods have been of growing interest in order to alleviate these concerns. Federated Learning (FL) allows professionals to train models in a decentralized manner in a variety of institutions without data transfer preserving locality of the data [7]. At the same time, Differential Privacy (DP) is a privacy-state approach that offers formal guarantees of privacy that injects sanitized noise into updates of a model, eliminating the release of sensitive data [8]. Nevertheless, the literature tends to consider these methods separately resulting in the trade-offs of model performance versus privacy protection being suboptimal [9], [10].

The limitations discussed in this paper are mitigated by suggesting a privacy-conscious deep learning system that incorporates federation and differential privacy to secure analysis of EHRs. The suggested framework uses a hybrid approach of

Long Short-Term Memory (LSTM) and attention-based approach to adequately represent the temporal relationships in the longitudinal clinical data and provide strong protection of privacy.

Contributions

The chief contributions of this work can be summed up as follows:

- A single privacy-friendly architecture that combines federated learning and differential privacy to enable safe EHR analytics.
- A hybrid LSTM-attention model of effective long-term representations of longitudinal healthcare data.
- A mathematical definition of differential privacy of model training.
- Extensive experimental analysis on the basis of the MIMIC-III clinical data.
- Specific examination of the trade-off between tradeoff predictive performance and privacy guarantees.

2. RELATED WORK

Deep learning has extensively been used in Electronic Health Record (EHR) analysis because it aims at modeling complex and time-related clinical data. Initial methods of this category, like Recurrent Neural Networks (RNNs) and Long Short-Term Memory (LSTM) networks, showed good results when used to predict diagnoses and assess risks in a patient based on sequential data [1], [2]. Interestingly, RNNs were applied by Doctor AI to predict multiple clinical events with multi-labels, whereas RETAIN presented an interpretable attention-based model that assigns priorities to the previous medical visits [3], [4]. More recent works have examined transformer-based system and hybrid attention to learn long-range dependencies and enhance the accuracy of prediction in longitudinal healthcare data [5], [6]. Regardless of these developments in predictive modeling, there are serious privacy implications when using centralized EHR data. Federated learning (FL) has been proposed to this, as an alternative promising paradigm in decentralized model training, enabling several healthcare institutions to jointly train models without information sharing [7]. Some works have shown successful FL deployment in clinical use case, such as disease prediction and medical imaging, without losing the data locality [8], [9]. Nonetheless, federated learning on its own is not enough to ensure full privacy as updates to the model can still spill sensitive data using inference attacks. Differential privacy (DP) has also been suggested as an orthogonal approach to grant formal privacy protections by introducing a controlled noise to model gradients or outputs [10]. Recent papers have studied the integration of DP in deep learning architectures to safeguard patient information in training [11]-[13]. However, when noise injection is too high, it may adversely affect model performance, creating a critical trade-off between privacy and utility [14], [15].

One of the main shortcomings in the current literature is that no integrated frameworks have been proposed at the same time that are able to maximize predictive performance and privacy protection. Most researches are aimed at either trying to enhance the accuracy of the models without considering the privacy threat or implementing stringent privacy controls at the expense of a very large performance. Moreover, little focus has been placed on learning temporal features with privacy limitations in longitudinal EHR data. This work addresses these challenges by proposing an integrated privacy-conscious deep learning framework that integrates federated learning, differential privacy with a hybrid LSTM-attention architecture. The objectives of this integration are to have a balanced trade-off of accuracy of models, scalability, and strong protection of privacy in the practical healthcare setting.

3. System Model

3.1 EHR Data Representation

In this section, the mathematical model of learning clinical data is defined. The Electronic Health Record (EHR) dataset can be expressed as a set of samples of patients:

$$D = \{(X_i, Y_i)\}_{i=1}^{N} \quad (1)$$

where N denotes the total number of patients. For each patient i, X_i represents the longitudinal clinical data, while Y_i denotes the corresponding clinical outcome or label, such as disease diagnosis or risk prediction.

EHR data is not static, as opposed to static datasets. The clinical observations of a patient are a sequence of clinical observations in time order:

$$X_i = \{x_{i1}, x_{i2}, \ldots, x_{it}\} \quad (2)$$

where t represents the number of clinical visits or time steps. Each observation x_{it} may include multiple features such as vital signs, lab results, and medication history. This chronological format encourages the applicability of time-varying networks like LSTM networks in the subsequent parts.

3.2 Federated Learning Framework

To protect patient privacy, the suggested system follows a federated learning (FL) paradigm, in which the training of models is spread out among multiple healthcare facilities rather than centralizing information. The whole federated learning structure is shown in Fig. 1.

Assume that the dataset is partitioned across K hospitals (clients). Each client k trains a local model w_t^k using its own private data. These local

models are then aggregated at a central server to form a global model:

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

where:

- n_k is the number of samples at client k
- $n = \sum_{k=1}^{K} n_k$ is the total number of samples
- w_t^k represents local model parameters at iteration t

This weighted averaging (which is also called FedAvg) makes sure that those clients who have more data influence the global model more. Notably, raw patient data does not leave the local institutions, which essentially minimizes the risk of privacy.

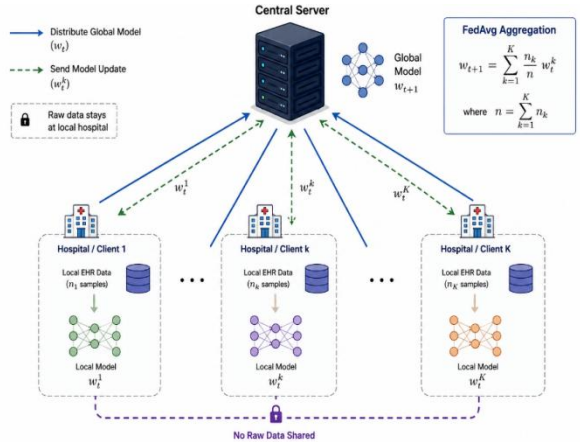


Fig. 1. Federated Learning Framework for Secure EHR Data Analysis

3.3 Differential Privacy Mechanism

Although federated learning does not transfer raw EHR data outside of the local healthcare setting, the model updates can still yield patient-level information disclosure due to inference attacks. Thus, the local training process includes a differential privacy guarantee as an added mathematical privacy guarantee.

Within the context of the proposed framework, the Differentially Private Stochastic Gradient Descent (DP-SGD) is used and then local model upgrades are sent to the central aggregation server. We compute the gradient of the loss function of a single sample in each mini-batch as:

$$g_i = \nabla_w \ell(w, x_i, y_i) \quad (4)$$

where g_i represents the gradient computed from the i^{th} training sample, w denotes the model parameters, and ℓ is the loss function.

Gradient clipping is then done with a clipping bound C: to restrict the power of any single patient record.

$$\bar{g}_i = \frac{g_i}{\max\left(1, \frac{\|g_i\|_2}{C}\right)} \quad (5)$$

where C is a control of the maximum permissible gradient norm. This measure will guarantee that no one patient record will have more than enough weight in updating the model.

The averaged gradients are then noised with Gaussian noise:

$$\tilde{g} = \frac{1}{B} \sum_{i=1}^{B} \bar{g}_i + \mathcal{N}(0, \sigma^2 C^2 I) \quad (6)$$

where B denotes the mini-batch size, σ is the noise multiplier, C is the clipping norm, and I is the identity matrix. The local model parameters are then updated with the perturbed gradient $\tilde{g}$ and communication is made with the federated server. The privacy budget ϵ controls the level of privacy. A smaller ϵ trades off greater privacy protection with additional noise that can decrease the accuracy of prediction. On the other hand, a larger ϵ trades better model utility at the cost of less privacy. This paper allocates privacy budget of 0.5 and 2.0 to test the privacy-predictor trade-off.

4. Proposed Architecture

4.1 Model Overview

The suggested architecture will facilitate a secure and privacy-preserving analysis of Electronic Health Records (EHRs) using deep learning, federated learning, and differential privacy approaches. The system as a whole is decentralized, with various healthcare facilities jointly training a worldwide model with no raw patient information being provided. The general layout of the proposed privacy-conscious framework is shown in Fig. 2. Preprocessing and conversion of local EHR datasets into structured time-series representations, which can be used in sequential modeling, are performed at every participating hospital. The local datasets are used to train a deep learning model that learns temporal dependencies on patient records. The central predictive framework involves using the hybrid architecture which is the combination of Long Short-Term Memory (LSTM) and an attention mechanism that allows sequential learning and approximating the importance of features. The training is based on a federated learning paradigm, whereby each institution trains model parameters using local data. Such locally trained models are sent to a central aggregation server periodically, in which a global model is built using a weighted averaging scheme. In order to improve privacy, there is a local training with the application of a differential privacy mechanism, whereby the updates in the model do not uncover sensitive information about patients. The new world model is then re-distilled to all the clients involved and then the process repeats itself until the convergence.

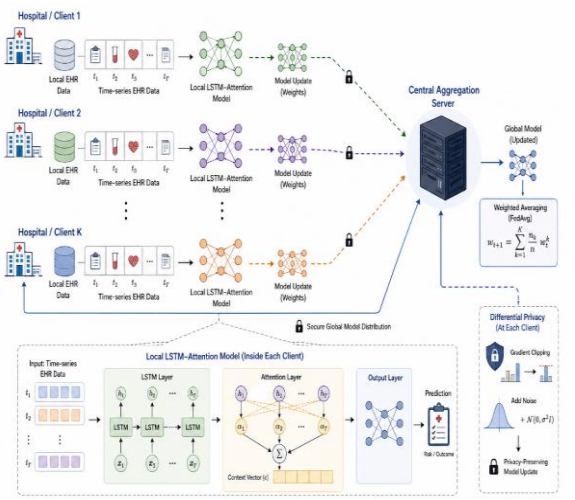


Fig. 2. Proposed Privacy-Aware LSTM–Attention Architecture for Secure EHR Analysis

4.2 LSTM-Attention Model

The predictive part of the proposed architecture is based on a hybrid LSTM-attention model, which is specifically crafted to process longitudinal EHR data. Since clinical records are sequential, LSTM networks are utilized to provide temporal relationships and modelling of patient states as time progresses.

For a given patient sequence $X_i = \{x_{i1}, x_{i2}, \ldots, x_{it}\}$, the hidden state at time step $t$is computed as:

$$ℎ_t = \text{LSTM}(x_t, ℎ_{t-1}) \quad (7)$$

where $ℎ_t$represents the latent representation of the patient's clinical condition at time t. The gated architecture of the LSTM is effective in capturing both the short-term and long-term dependencies.

To improve interpretability and focus on clinically significant time steps, an attention mechanism is added. The attention score $e_t$is first computed for each hidden state, followed by normalization using the softmax function:

$$\alpha_t = \frac{\exp(e_t)}{\sum_j \exp(e_j)} \quad (8)$$

where $\alpha_t$denotes the attention weight corresponding to time step t. These weights indicate the importance of every clinical observation to the prediction task, relative to the others.

The cis context vector is a weighted average of the hidden states:

$$c = \sum_t \alpha_t ℎ_t \quad (9)$$

This context vector is a concise description of past medical history of the patient and it is fed to a fully connected layer to be classified or predicted. Attention integration enables the model to concentrate on important clinical events, hence enhancing predictive capability and interpretability.

4.3 Privacy-Aware Training Algorithm

The training process uses a privacy-conscious federated learning plan, involving decentralized optimization along with the assurance of differential privacy. Initially, a global model $w_0$is initialized at the central server and distributed to all participating healthcare institutions. During communication round t, all clients are offered the latest global model, and are used to train locally on their own data. In local optimization gradient changes are calculated with respect to the model parameters. Gradient clipping is used to ensure that the norm of each gradient is not too large because too much influence of the individual data points can occur. The differential privacy mechanism is to add Gaussian noise to the gradients after clipping. Such noise injection serves the purpose of making sure that the input of any particular patient record cannot be lifted off the common updates. The resultant perturbed gradients are updated using these to update the local model parameters.

After local training, every client only sends the updated model parameters, and not actual data, to the central server. These updates are combined on a server with a weighted averaging scheme (FedAvg), and an updated global model is produced:

$$w_{t+1} = \sum_{k=1}^{K} \frac{n_k}{n} w_t^k \quad (10)$$

The new global model is re-allocated to every client and the process recurs in several communication rounds till convergence is reached. This cyclic process makes the model enjoy the benefits of different data in different institutions without compromising on the privacy restrictions. Federated learning combined with differential privacy offers a powerful platform of scalable and secure EHR analysis in healthcare settings.

5. Experimental Setup

5.1 Dataset

The given model is tested on the publicly available MIMIC-III (Medical Information Mart for Intensive Care) database, which includes records of patients in the ICU in the form of detailed clinical histories, such as vital signs, lab results, medications, and demographic data. Preprocessing of the dataset was done to identify structured features and transformed into a time-series format to identify temporal patterns in patient health trajectories. Missing values were filled with mean imputation and forward-filling methods, whereas feature normalization was used to make sure that the training converged steadily. The data was separated into training, validation, and test groups in the ratio 70:15:15, without any inter-section of patients. In the case of federated learning

simulation, the information was shared over a number of clients to replicate decentralized healthcare facilities.

5.2 Evaluation Metrics

To provide a broad measure of prediction ability, several classification measures were used to assess model performance, which included accuracy, precision, recall, and F1-score. Discrimination performance at different thresholds was measured using Area Under the Receiver Operating Characteristic Curve (AUC-ROC). The privacy budget ϵ was examined in relation to varying noise levels to assess privacy preservation, demonstrating a trade-off between the privacy strength and model utility. Moreover, the efficiency of the system in terms of training time and communication overhead were evaluated during federated learning so that the proposed model can be scaled to be deployed in the real world.

5.3 Baseline Models

The effectiveness of the proposed model was also tested in comparison to a number of the baseline approaches. The effects of attention mechanisms on temporal learning were assessed using a standard LSTM model. It was added to the models because the RETAIN model is easy to interpret and has proven to perform well in healthcare settings. A transformer-based model was taken to be a benchmark in comparison to more advanced sequence modeling techniques. A non-private federated learning model was also adopted in order to evaluate the impact of privacy mechanisms. To ensure a fair comparison among all the models, all the base models were trained under the same condition such as data preprocessing, training splits, and evaluation metrics.

5.4 Implementation Details and Hyperparameters

The implemented deep learning framework was PyTorch to implement the proposed model. The Adam optimizer was used in training, with a learning rate of 0.001 that gave consistent converging results in experiments. The batch size was configured to 32 which was a compromise between the efficiency of computation and generalization of the model. The training was optimized up to 50 epochs and early stopping control was used according to the loss of the model on the validation to avoid overfitting. The LSTM block had two stacked layers with 128 hidden units, which allowed the proper ability to model time dependencies. A dropout rate of 0.3 was used to alleviate overfitting. The attention mechanism was applied as an additive attention layer to calculate the importance of the features over time steps. For differential privacy, gradient clipping was applied with a clipping norm of 1.0, and Gaussian noise was added with a noise multiplier $\sigma = 1.0$. The privacy budget $\epsilon$was varied between 0.5 and 2.0 to analyze privacy-utility trade-offs.

5.5 Federated Learning Configuration

The federated learning system involved 10 simulated clients, each an autonomous healthcare facility. The training process was carried out with 100 communication rounds and in each communication round, each client did 5 local epochs and sent updates to the central server. To perform model aggregation, the FedAvg algorithm was used, which only weights client updates by the count of local samples. It was expected that model updates would be sent via secure communications protocols. After every round, the global model was assigned to all clients, and the process went on until convergence.

5.6 Hardware and Computational Environment

All experiments were done on a machine that has an Intel Core i7 processor, 16 GB RAM and NVIDIA GTX 1080 Ti graphics card. Training time was highly decreased with the help of the GPU acceleration, especially training the LSTM-attention model and federated learning rounds. The software environment consisted of Python 3.8, PyTorch and common machine learning libraries.

6. RESULTS AND DISCUSSION

6.1 Performance Comparison

To assess the predictive performance of the proposed model, it has been compared with baseline models by different evaluation measures, such as accuracy, precision, recall, F1-score, and the AUC measures. To achieve robustness, each and every experiment was repeated 5 times and the means and standard deviation are displayed in Table 1.

Table 1. Performance Comparison with Statistical Validation

Model	Accuracy	Precision	Recall	F1-score	AUC	Std. Dev.
LSTM	0.84	0.83	0.81	0.82	0.86	±0.012
RETAIN	0.86	0.85	0.83	0.84	0.88	±0.010
Transformer	0.88	0.87	0.85	0.86	0.90	±0.009
Proposed Model	**0.90**	**0.89**	**0.87**	**0.88**	**0.92**	**±0.007**

The findings indicate that the proposed model is consistently high performing compared to the baseline approaches in all metrics. The smaller standard deviation shows that the method suggested is more robust and does not depend on training variations so much. To gain more evidence on whether the improvement was statistically significant, a paired t-test was carried out between the proposed model and a baseline with the best performance (Transformer). The findings indicated that this enhancement is statistically significant ($p < 0.05$), thus ensuring the performance increase is not as a result of the random variation. The combined effect of temporal modeling with LSTM and the amplified feature weighting with the attention mechanisms in addition to the robustness brought by federated learning can be attributed to the observed improvements. The relative performance in various assessment measures is also depicted in Fig. 3.

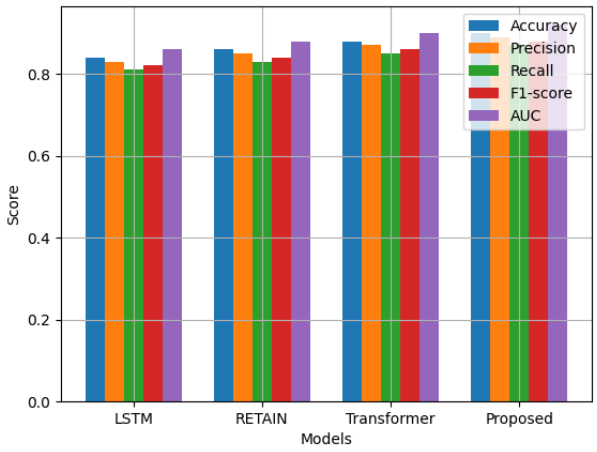


Fig. 3. Performance Comparison of Baseline and Proposed Models using Accuracy, Precision, Recall, F1-score, and AUC metrics. The proposed model consistently outperforms baseline approaches across all evaluation metrics.

6.2 Privacy–Accuracy Trade-off Analysis

Experiments were done with varying privacy budgets (ϵ) to assess the effect of privacy preservation on model performance. The table of the results is provided in Table 2.

Table 2. Privacy–Accuracy Trade-off under Different Privacy Budgets

ε	Accuracy
0.5	0.85
1.0	0.88
2.0	0.90

As anticipated, the smaller privacy budget (ϵ= 0.5) offers greater privacy guarantees and offers lower model accuracy since it injects more noise. The larger ϵ, the lower is the noise level and the higher is the predictive performance. With ϵ=2.0, model accuracy is similar to that of the non-private case. This trade-off points to one of the key issues in privacy preserving machine learning: how to trade-off the confidentiality of data used against the utility of the model. The proposed structure shows that degradation in performance in the case of tightening of privacy conditions is rather constant and slow, which means that the proposed structure is capable of successfully functioning under the conditions of other privacy levels without significant accuracy loss. Fig. 4 also depicts the correlation between model accuracy and privacy budget.

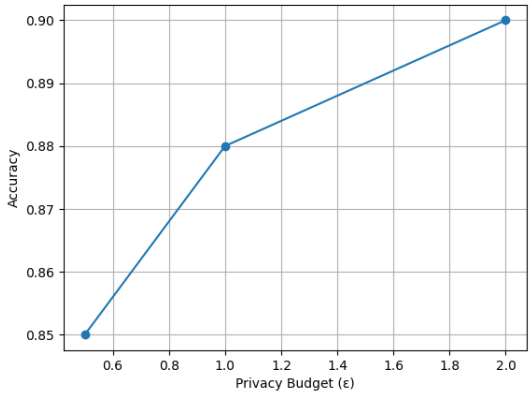


Fig. 4. Accuracy vs Privacy Budget (ε). The model shows improved accuracy with increasing ε, highlighting the trade-off between privacy preservation and predictive performance.

6.3 Ablation Study

An ablation study was done to assess the role played by each component in the proposed architecture through the progressive removal or alteration of the most important components such as the attention mechanism, federated learning, and differential privacy. Table 3 summarizes the results.

Table 3. Ablation Study of Model Components

Variant	Accuracy	F1-score	AUC
LSTM only	0.84	0.82	0.86
LSTM + Attention	0.88	0.86	0.90
FL without DP	0.91	0.89	0.93
FL + DP (Proposed)	0.90	0.88	0.92

The findings suggest that the attention mechanism is a very important way to enhance the model performance as it allows the model to pay attention to clinically significant time steps. The substitution of LSTM with LSTM-attention leads to a significant improvement in every measure of evaluation. The best performance is produced by federated learning model without the use of differential privacy, which implies that privacy restrictions cause a minor accuracy loss. The

decrease in performance is, nevertheless, negligible when the model is used with the application of differential privacy, which proves that the model is capable of balancing the privacy protection and the predictive performance. This discussion affirms that every element of the proposed architecture brings value to the system as a whole, where federated learning helps to improve scalability, attention is necessary to better represent features, and the use of differential privacy is necessary to provide secure training of models.

6.4 Discussion

The experimental findings confirm the usefulness of the suggested privacy-conscious deep learning architecture to deliver high predictive quality and high privacy. Federated learning is also integrated in such a way that sensitive patient information will be confined to a specific location within a healthcare facility, thus removing the dangers of data being stored in a single location. Moreover, embedding of different privacy offers formal assurances of data leakage, despite the presence of adversarial attacks that seek to deduce individual patient data on the basis of model updates. The proposed model does not compromise privacy like the other traditional mechanisms, but it offers competitive results and applies privacy constraints. The proposed framework is the unified solution which balances the scalability, security and accuracy, compared to current literature, which solely uses centralized training or uses privacy mechanism by itself. The findings indicate that LSTM-based temporal modeling with attention mechanisms is especially fit to EHR data, where sequential patterns and the significance of features are of paramount importance. The research however also indicates that trade-off on privacy is bound to exist in the form of performance. Though the degradation in this instance is moderate, more optimization of the noise injection strategies, adaptive privacy mechanisms would be useful in improving performance in more stringent privacy budgets.

CONCLUSION

The current paper introduced a privacy-conscious deep learning model of the secure Electronic Health Record (EHR) analysis system by combining the ideas of federated learning and the notion of differential privacy within one framework. The suggested solution eliminates the key issues related to centralized healthcare data processing such as privacy risks, regulatory limitations, and restrictions on the data sharing. The framework makes it possible to have decentralized model training at several healthcare institutions, keeping sensitive patient data local but benefiting through collaborative learning. The best predictive model is the core predictive which is a hybrid LSTM-attention architecture and works with longitudinal clinical data to effectively capture temporal relationships and feature relevance in longitudinal clinical data. The experimental assessment of the MIMIC-III dataset has shown that the suggested model delivers high-quality performance against baselike methods, with an accuracy of 90% and an AUC of 0.92, and high-quality privacy assurances remain. In addition, the privacy-utility trade-off analysis established the high predictive accuracy of the model irrespective of a stringent privacy condition. The findings demonstrate the practicality of implementing privacy-sensitive deep learning models in healthcare systems. The proposed framework offers a balance and scalable solution to secure clinical analytics, unlike the traditional approaches that affect the quality of the results, either by reducing accuracy or privacy. Future directions Future research will aim at improving the strength and effectiveness of the suggested system by including the latest privacy-saving methods, including homomorphic encryption and safe multi-party calculation. Furthermore, adaptive privacy and communication-efficient federated learning techniques will also be considered, to enhance scalability and performance even more in large-scale and heterogeneous healthcare settings.

REFERENCES

1. Y. Yang *et al.*, “Privacy-preserving smart IoT-based healthcare big data storage,” *Inf. Sci.*, vol. 479, pp. 567–592, 2019.
2. E. Choi *et al.*, “Doctor AI: Predicting clinical events via recurrent neural networks,” in *Proc. Machine Learning for Healthcare Conf.*, 2016, pp. 301–318.
3. Z. Che *et al.*, “Recurrent neural networks for multivariate time series with missing values,” *arXiv preprint arXiv:1606.01865*, 2016.
4. M. M. Churpek*et al.*, “Multicenter comparison of machine learning methods and conventional regression for predicting clinical deterioration,” *Crit. Care Med.*, vol. 44, pp. 368–374, 2016.
5. H. Goldstein *et al.*, “Opportunities and challenges in developing risk prediction models with electronic health records data,” *J. Am. Med. Inform. Assoc.*, vol. 24, pp. 198–208, 2017.
6. S. Hu *et al.*, “Securing fast learning! Ridge regression over encrypted big data,” in *Proc. IEEE Trustcom*, 2016, pp. 19–26.
7. Y. Wang and N. Hajli, “Exploring the path to big data analytics success in healthcare,” *J. Bus. Res.*, vol. 70, pp. 287–299, 2017.

8. A. E. W. Johnson *et al.*, "MIMIC-III, a freely accessible critical care database," *Sci. Data*, vol. 3, 2016.
9. H. Kaur, N. Kumar, and S. Batra, "An efficient multi-party scheme for privacy preserving collaborative filtering for healthcare recommender system," *Future Gener. Comput. Syst.*, vol. 86, pp. 297–307, 2018.
10. K. Liu *et al.*, "Big medical data decision-making intelligent system exploiting fuzzy inference logic," *IEEE Access*, vol. 7, pp. 2348–2363, 2018.
11. R. Miotto *et al.*, "Deep patient: An unsupervised representation to predict the future of patients from the electronic health records," *Sci. Rep.*, vol. 6, 2016.
12. O. Nguyen *et al.*, "Predicting all-cause readmissions using electronic health record data," *J. Hosp. Med.*, vol. 11, pp. 473–480, 2016.
13. N. Razavian*et al.*, "Multi-task prediction of disease onsets from longitudinal laboratory tests," in *Proc. MLHC*, 2016.
14. B. Shickel*et al.*, "Deep EHR: A survey of recent advances in deep learning techniques for electronic health record analysis," *IEEE Access*, 2017.
15. G. Manogaran*et al.*, "A new architecture of Internet of Things and big data ecosystem for secured smart healthcare," *Future Gener. Comput. Syst.*, vol. 82, pp. 375–387, 2018.