Prediction of Hospital Readmission using Federated Learning

Borjan Sazdov[1], Matea Tashkovska[1], Stefan Krsteski[1], Borche Jovanovski[1],
Stefan Kalabakov[2], Valentin Rakovic[1], Daniel Denkovski[1], and Hristijan Gjoreski[1]

[1] Faculty of Electrical Engineering and Information Technologies, Ss. Cyril and Methodius University in Skopje, Macedonia
[2] Digital Health – Connected Healthcare, Hasso Plattner Institute, University of Potsdam, Potsdam, Germany

***Abstract*—Wearable devices have the ability to generate vast amounts of data that can be put to use in a multitude of applications, particularly in the field of e-health. However, the potential invasion of privacy that comes with utilizing personal data collected by these devices cannot be overlooked. Federated Learning (FL) is a promising solution to this issue that allows models to be trained in a decentralized manner while keeping user data on their own devices. This approach effectively minimizes the risk of privacy breaches and has the potential to be employed in a variety of applications where the protection of user data is of utmost importance. This paper focuses on the use of FL in predicting hospital readmission in 130-US diabetes hospitals for a data set collected over an 9-year period. The results suggest that FL can achieve comparable performance while maintaining privacy and diversity of data. This is an essential aspect of FL, as it enables continuous real-time learning without compromising privacy.**

***Index Terms*—Wearable Technologies, e-Health , Federated Learning, Hospital Readmission**

I. Introduction

Ubiquitous Internet of Things - IoT [1] devices pave the way for the everyday use and application of Wearable Sensor Technology (WST) [2]. WST has shown great potential in the medical field, especially in the areas of remote patient monitoring and telemedicine. With the use of WST, medical professionals can remotely monitor a patient's vital signs and health parameters, such as heart rate, blood pressure, and glucose levels, in real-time. This allows for early detection of any potential health problems, and medical intervention can be provided promptly. In addition, WST can also be used for post-surgery rehabilitation, physical therapy, and monitoring the progress of patients with chronic conditions. WST can provide real-time feedback on a patient's progress, and healthcare professionals can adjust the treatment plan accordingly. Moreover, WST can also be used for fitness and training purposes, as it can track a person's physical activity levels, sleep patterns, and other metrics related to their overall health and wellness. The use of WST in healthcare has the potential to improve patient outcomes, reduce healthcare costs, and enhance the overall quality of care provided to patients. This technology has made it possible to remotely monitor patients' health conditions, and the data collected can be sent to the appropriate doctor for analysis, which can help to identify potential problems and prevent them through early intervention. However, due to the small form factor of these devices, they may not have sufficient processing power or energy to apply complex machine learning models such as Deep Neural Networks (DNNs) [3]. Therefore, the data needs to be forwarded to a central server for processing, which raises security and real-time continuous learning concerns. In this context, Federated Learning (FL) [4] can be the most compatible approach to address these concerns. FL enables the training of ML models without centralizing the data, and instead, the learning takes place locally on the device itself or at the edge of the network. FL ensures data privacy and security while also enabling real-time continuous learning.

Federated Learning (FL) can be a useful approach for addressing the problem of hospital readmission. Hospital readmission [5], occurs when a patient is readmitted to the hospital within a certain period after being discharged. It is a costly and often preventable problem that can lead to poor patient outcomes and increased healthcare costs. FL can be used to train predictive models that identify patients who are at high risk of readmission. These models can be trained using data from multiple hospitals or healthcare systems without sharing patient data or violating privacy regulations. FL allows the participating hospitals to collaborate in training the model while maintaining data privacy. This can also be used to monitor patients' health conditions remotely, which can help to prevent readmission. By using wearable sensor technology to collect patient data, FL can enable real-time monitoring of patients' health conditions and alert healthcare providers when there is a need for intervention. Overall, FL can be a promising approach to address the problem of hospital readmission, as it can enable collaboration among healthcare providers while also protecting patient privacy and improving patient outcomes.

The paper is organized as follows. Section II resents the related work at the intersection of FL and prediction of hospital readmission. Section III describes the characteristics of the dataset we use for the experiments. Section IV presents the data preprocessing, feature engineering, and model development steps. After that Section V the main goal will be towards experimental setup. Finally, Section VII provides a summary of the paper.

II. Related Work

Implementing FL for hospital readmission prediction can provide several benefits, such as improved model accuracy and the ability to train models on larger, more diverse datasets while maintaining data privacy and security. However, it is important to carefully consider the potential ethical and legal implications of sharing sensitive healthcare data and to ensure that appropriate security measures are in place to protect patient privacy.

The study [6] initiates a discussion on the need for FL in healthcare and reviews recent review papers in the field. The study then explains the fundamentals of FL and the motivations behind its use in healthcare. Next, the study presents various applications of FL in healthcare, including recent state-of-the-art techniques across different verticals of healthcare. The study also highlights lessons learned from using FL, open issues that need to be addressed, and challenges that still need to be solved in the field. Finally, the study provides future directions for researchers who are interested in conducting research in this area.

The study [7] employed specific measures to assess the potential association between patient-related factors and hospital readmission risks. Gait speed was assessed using a 10-meter walk test, muscle strength was determined using a hydraulic handgrip dynamometer, and functional status was evaluated using the Functional independence measure - FIM. The study's primary outcome was the cumulative incidence of the first unplanned early rehospitalization within 30 days of hospital discharge for the entire cohort. This outcome measure is a commonly used indicator of hospital readmission risks. To determine the accuracy of the predictors, the study used a Receiver Operator Characteristic (ROC) analysis, which is a statistical method commonly used to evaluate the performance of prediction models. Overall, it seems that the study employed rigorous and standardized measures to assess the potential association between patient-related factors and hospital readmission risks. The use of ROC analysis adds further rigor to the study and helps to determine the predictive accuracy of the measures used.

The objective of the study [8] is to assess the performance of a decision support tool based on a machine learning model for predicting readmission or death within 7 days after ICU discharge. The study aims to evaluate the tool's performance before, during, and after retraining and recalibration in different settings. The study highlights an important issue in the development and deployment of machine learning models in clinical practice. While there has been a surge in the development of ML models for clinical decision support, their external validation and generalizability to new settings have not been adequately assessed. The study's approach of evaluating the tool's performance before, during, and after retraining and recalibration is commendable.

The article [9] reviewed recent studies that utilized Federated Learning (FL) in clinical studies with structured medical data. FL is a machine learning technique that enables multiple institutions or users to collaborate and build a shared model without sharing their data. The review of recent studies provides insights into the potential of FL in clinical studies. FL has been used in various applications, such as predicting readmissions, detecting heart failure, and predicting mortality. The studies have demonstrated that FL can achieve similar performance to centralized machine learning while preserving data privacy and security. However, the article also highlights some challenges and open questions in FL in clinical studies with structured medical data. These include the lack of standardization in FL implementations, the need for privacy-preserving mechanisms, and the difficulty of dealing with heterogeneity across different datasets.

The proposed approach [10] of using federated deep extreme learning entangled with edge computing and a fused weighted deep extreme machine learning methodology for the diagnosis of lung disease sounds promising. It's great to hear that the model achieved an accuracy of 97.2%, which is better than the state-of-the-art published methods. However, as you have mentioned, data privacy is a crucial concern in the healthcare industry. It would be important to ensure that patient data is adequately protected, and any data sharing or processing is done with the patient's consent and in compliance with privacy regulations. It would also be useful to consider the interpretability of the model. Deep learning models can often be considered as "black boxes" due to their complex nature. It would be important to be able to explain the model's predictions to healthcare professionals and patients for better understanding and acceptance.

III. Data

The data set [11] describes a Modified version of "Diabetes 130-US hospitals for years 1999-2008", data set was used at the competition that is organized as part of the Wide-Health2023 Thematic Winter School in Potsdam, Germany. This data set represents 9 years of clinical care at 130 US hospitals and integrated delivery networks. It includes over 50 features representing patient and hospital outcomes. Information was extracted from the database for encounters that satisfied the following criteria.

- It is an inpatient encounter (a hospital admission).
- It is a diabetic encounter, that is, one during which any kind of diabetes was entered to the system as a diagnosis.
- The length of stay was at least 1 day and at most 14 days.
- Laboratory tests were performed during the encounter.
- Medications were administered during the encounter.

The data contains such attributes as patient number, race, gender, age, admission type, time in hospital, medical specialty of admitting physician, number of lab test performed, HbA1c test result, diagnosis, number of medication, diabetic medications, number of outpatient, inpatient, and emergency visits in the year before the hospitalization, etc.

It is important to know if a patient will be readmitted in some hospital. The reason is that you can change the treatment, in order to avoid a readmission. In this database, you have 3 different outputs:

- No readmission.
- A readmission in less than 30 days (this situation is not good, because maybe your treatment was not appropriate).
- A readmission in more than 30 days (this one is not so good as well the last one, however, the reason can be the state of the patient.

IV. METHODS

Hospital readmission prediction is a critical task in the healthcare domain as it can help reduce the burden on healthcare facilities and improve the quality of care provided to patients. In this study, we investigated the use of federated learning, a privacy-preserving machine learning approach, to develop a hospital readmission prediction model. In this section we present the data preprocessing, feature engineering, and model development steps that were employed to achieve our goal.

A. Preprocessing

Data preprocessing is an essential step in any machine learning project. It involves cleaning, transforming, and preparing the raw data before it can be fed into a machine learning model. In this step, data is refined and made ready for analysis to ensure accurate and reliable results. The first and foremost task in data preprocessing is data cleaning, which involves handling missing values, dealing with duplicate data, and removing irrelevant data. The dataset contained missing values in both of its categorical and numerical features. To overcome this problem, as a solution for the numerical features, we employed K-nearest neighbor imputation with 5 neighbors weighted by distance. On the other hand, the missing values in the categorical features where treated as a separate category named "Other". The next step is data transformation, which involves converting the data into a suitable format that can be processed by a machine learning algorithm. To achieve this purpose, the continuous variables were standardized with standard scaler. This ensures that these variables are on the same scales.

B. Feature Engineering

Feature engineering process that involves creating new features from existing data, can help in capturing more relevant information and improving the performance of the machine learning model. As shown in previous research [12], [13] two new features, service utilization, and medication count, proved to be of high importance in the model performance:

- Service utilization: represents the sum of the number of times a patient has used the hospital's services. It is calculated this feature by summing the number of inpatient visits, outpatient visits, and emergency visits.
- Medications count: represents the number of medication changes that occurred when the patient was admitted.

C. Feature Selection

To reduce the dimensionality of the feature space and remove redundant or irrelevant features, we performed feature selection. We dropped the columns with a higher percentage of missing values, such as "Weight," "Medical specialty," "Enocounter ID," and "Payer code." After an analysis on the impact of the medication features, it was concluded that the following medications were not informative and were removed: 'acetohexamide', 'tolbutamide', 'troglitazone', 'tolazamide', 'examide', 'citoglipton', 'glipizide-metformin', 'glimepiride-pioglitazone', 'metformin-rosiglitazone', 'metformin-pioglitazone'.

D. Model

The chosen model for this study is a neural network, a powerful and versatile tool for classification tasks. The specific architecture used in this study consisted of an input layer, two hidden layers with 128 and 64 neurons, respectively, followed by a dropout layer with a 0.1 dropout rate to prevent overfitting. The choice of the number of layers and neurons was based on empirical experimentation, balancing model complexity and performance. The input layer consisted of the preprocessed features, while the output layer consisted of a single neuron representing the predicted probability of readmission. The activation function used in the hidden layers was ReLU (Rectified Linear Unit), which is a widely used activation function in deep learning due to its simplicity and computational efficiency. The model was trained over 20 rounds, with each round consisting of 10 epochs per client. The server learning rate was set to 1.0, and the client learning rate was set to 0.1. The batch size was 128, meaning that the model was trained on 128 samples at a time. The choice of hyperparameters was based on empirical experimentation, balancing model performance and convergence speed.

V. EXPERIMENTAL SETUP

In order to assess the efficacy of our developed model, we partitioned the dataset into three subsets: a training set consisting of 60% of the data, a validation set consisting of 20%, and a testing set consisting of the remaining 20%. Our objective was to accurately predict the probability of a patient being readmitted to the hospital. To achieve this, our model was designed to output a floating-point number between 0 and 1 for each encounter, which represented the prediction certainty. To evaluate the performance of the model, we employed the Area Under the Receiver Operating Characteristic (AUROC) curve. AUROC is a commonly used classification metric that evaluates the trade-off between the true positive rate (TPR) and the false positive rate (FPR). This metric is especially relevant for probabilistic classifiers, where the model's output is a probability distribution over a set of classes.S.

VI. RESULTS AND DISCUSSION

In order to obtain optimal results on the test set, the model was meticulously fine-tuned with the aid of the validation set. The calibration of our model was then compared to a perfectly calibrated model, as illustrated in Figure 1. The x-axis represents the mean predicted probability while the y-axis represents the fraction of positives. The image clearly demonstrates that our model is well calibrated.

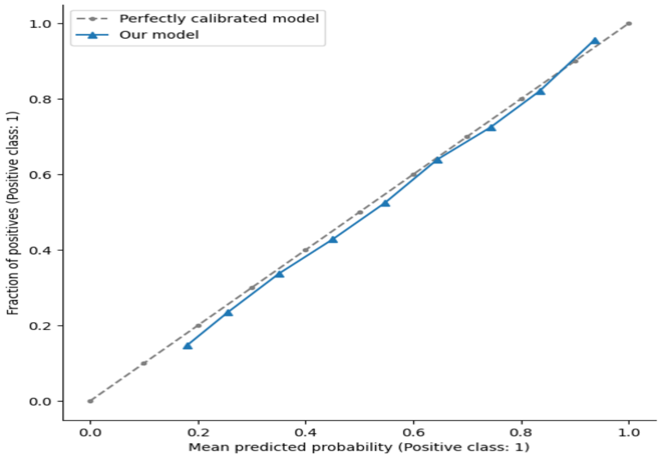

Fig. 1. Mean predicted probability.

Subsequently, the performance of the model was assessed on the test set, where it achieved an AUROC score of 0.69823. The relationship between the false positive and true positive rate, as well as the area under the curve, can be visualized in Figure 2. The attained score serves as an indicator that our model has the ability to predict the probability of readmission with a certain degree of accuracy, making it a promising approach to predicting readmission rates while ensuring the confidentiality and privacy of the patient's sensitive medical information.

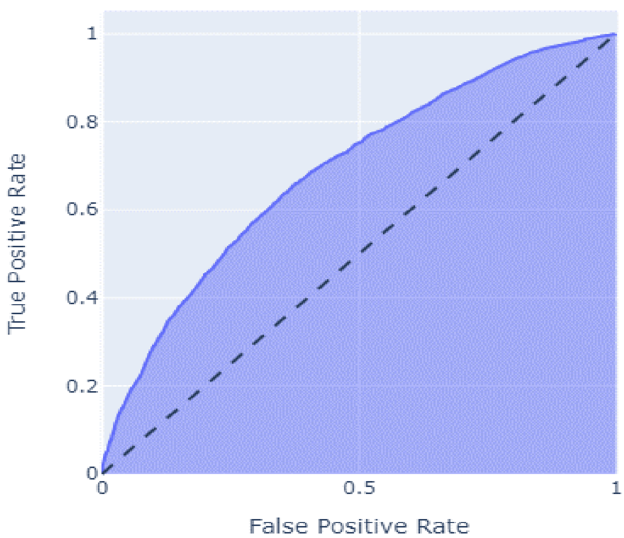

Fig. 2. Area Under the Receiver Operating Characteristic (AUROC) of the model.

The dataset comprising of four prominent ethnicities: Caucasian, African American, Asian and Hispanic. Other ethnicities were grouped under the class 'Other'. Figure 3 presents the AUROC score of our model by ethnicity. The model demonstrated the highest AUROC score for the Asian ethnicity and the lowest for 'Other'. However, the variation in the AUROC score by ethnicity is not substantial enough to draw conclusions about the model's dependency on the patient's ethnicity. Therefore, our model can be considered reliable and unbiased in predicting the likelihood of readmission across all ethnicities.

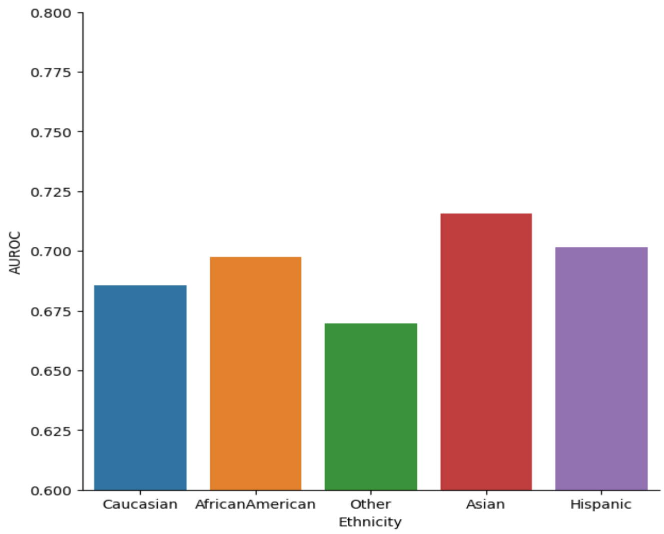

Fig. 3. AUROC score by ethnicity.

VII. Conclusions

Federated learning is a machine learning approach that allows multiple parties to collaborate and train a model without sharing their data with each other. Instead, the data remains on the local devices, and only the model updates are shared between them. This approach can be particularly useful in the healthcare domain, where patient privacy is a critical concern. Predicting hospital readmission is an important task in healthcare, as it can help hospitals identify patients who are at risk of being readmitted after discharge. This can lead to better care coordination and better patient outcomes. The 130-US Diabetes Hospital Database for the years 1999-2008 in Health Research used in this paper contains data on patients with diabetes admitted to 130 US hospitals. By incorporating federated learning into our model, multiple hospitals can collaborate to develop a hospital readmission prediction model without sharing their patient data with each other. Each hospital can train a local model using its own data, and then model updates can be aggregated to create a global model that can be used for predictions. This approach can help improve model accuracy by incorporating data from multiple hospitals while still maintaining patient privacy.

Acknowledgment

This work has been supported by the WideHealth project - European Union's Horizon 2020 research and innovation programme under grant agreement No. 952279.

REFERENCES

[1] Khurshid, Kiran, et al. "An In-Depth Survey Demystifying the Internet of Things (IoT) in the Construction Industry: Unfolding New Dimensions." Sustainability 15.2 (2023): 1275.

[2] Brun, Romana, et al. "Wearable sensors for prediction of intraamniotic infection in women with preterm premature rupture of membranes: a prospective proof of principle study." Archives of Gynecology and Obstetrics (2022): 1-10.

[3] Nazir, Sajid, Diane M. Dickson, and Muhammad Usman Akram. "Survey of explainable artificial intelligence techniques for biomedical imaging with deep neural networks." Computers in Biology and Medicine (2023): 106668.

[4] Ouyang, Liwei, et al. "Artificial identification: a novel privacy framework for federated learning based on blockchain." IEEE Transactions on Computational Social Systems (2023).

[5] Zhou, Huaqiong, et al. "Risk factors associated with 30-day unplanned hospital readmission for patients with mental illness." International Journal of Mental Health Nursing 32.1 (2023): 30-53.

[6] K. Dasaradharami Reddy, Thippa Reddy Gadekallu, "A Comprehensive Survey on Federated Learning Techniques for Healthcare Informatics", Computational Intelligence and Neuroscience, vol. 2023, Article ID 8393990, 19 pages, 2023. https://doi.org/10.1155/2023/8393990.

[7] de Melo, Thiago Araújo, Fernando Silva Guimarães, and José Roberto Lapa E. Silva. "Gait speed, muscle strength, and functional status predict 30-day unplanned hospital readmission of critical care survivors." Physiotherapy Research International (2023): e2001.

[8] de Hond, Anne AH, et al. "Predicting Readmission or Death After Discharge From the ICU: External Validation and Retraining of a Machine Learning Model." Critical Care Medicine 51.2 (2023): 291-300.

[9] Oh, Wonsuk, and Girish N. Nadkarni. "Federated Learning in Health care Using Structured Medical Data." Advances in Kidney Disease and Health 30.1 (2023): 4-16.

[10] Sagheer Abbas, Ghassan F. Issa, Areej Fatima, Tahir Abbas, Taher M. Ghazal, Munir Ahmad, Chan Yeob Yeun, Muhammad Adnan Khan, "Fused Weighted Federated Deep Extreme Machine Learning Based on Intelligent Lung Cancer Disease Prediction Model for Healthcare 5.0", International Journal of Intelligent Systems, vol. 2023, Article ID 2599161, 14 pages, 2023. https://doi.org/10.1155/2023/2599161

[11] Beata Strack, Jonathan P. DeShazo, Chris Gennings, Juan L. Olmo, Sebastian Ventura, Krzysztof J. Cios, and John N. Clore, "Impact of HbA1c Measurement on Hospital Readmission Rates: Analysis of 70,000 Clinical Database Patient Records," BioMed Research International, vol. 2014, Article ID 781670, 11 pages, 2014.

[12] Jauhari, Sarthak Shukla, Shikhar Tripathi, Surya. (2021). EmbPred30: Assessing 30-Days Readmission for Diabetic Patients Using Categorical Embeddings. 10.1007/978-981-15-5345-5-7.

[13] Goudjerkan, T., and Jayabalan, M. Predicting 30-day hospital readmission for diabetes patients using multilayer perceptron. International Journal of Advanced Computer Science and Applications 10 (03 2019), 268–275.