Frontiers in Computer Science and Artificial Intelligence
ISSN: 2978-8048
DOI: 10.32996/fcsai
Journal Homepage: www.al-kindipublisher.com/index.php/fcsai

FCSAI
AL-KINDI CENTER FOR RESEARCH AND DEVELOPMENT

| RESEARCH ARTICLE

From Real to Synthetic: An Analytical Machine Learning and Generative AI Approach for Reducing Early Readmissions in U.S. Diabetic Patients

Md. Kamruzzaman — Pompea College of Business, University of New Haven, West Haven, Connecticut, United States; Email: mdkamruzzamandu15@gmail.com ; ORCID: 0009-0005-0671-6397
Sujoy Saha — Pompea College of Business, University of New Haven, West Haven, Connecticut, United States; Email: ssujoy26@gmail.com ; ORCID: 0009-0000-9358-7813
Md Nazmul Alam Bhuiyan — Pompea College of Business, University of New Haven, West Haven, Connecticut, United States; Email: bhuiyan.unh@gmail.com ; ORCID: 0009-0003-2830-3396
Md. Shoeb Siddiki — Pompea College of Business, University of New Haven, West Haven, Connecticut, United States; Email: msidd9@unh.newhaven.edu ; ORCID: 0009-0001-5811-6677
Rabi Sankar Mondal — Pompea College of Business, University of New Haven, West Haven, Connecticut, United States; Email: rabi.s.mondal@gmail.com ; ORCID: 0009-0006-0136-9354
Dr Tufael Ahmed — Department of Biochemistry, Parul University, India; Ibn Sina Diagnostic and Imaging Center, Dhaka, Bangladesh; Email: tofayelahmed083@gmail.com ; ORCID: 0000-0001-5809-4319
Corresponding Author: Md. Kamruzzaman, **E-mail**: mdkamruzzamandu15@gmail.com

| ABSTRACT

US healthcare systems struggle with hospital readmissions, especially within 30 days of release. ML-based prediction of hospital readmission is a significant milestone for healthcare professionals to efficiently manage healthcare quality, premature readmissions, post-discharge follow-ups, incur large costs and identify high-risk patients. Many disadvantages still prohibit these models from being widely adopted and successful. This research proposes a promising solution to the challenge of early hospital readmissions among diabetic patients in the U.S. by combining ML and generative artificial intelligence (AI) techniques. Utilizing the UCI Diabetes 130-US Hospitals Dataset, the study develops predictive models for 30-day readmission risks using conventional Random Forest (RF), K-Nearest Neighbors (KNN), and Gradient Boosting Classifier (GBC)—alongside Conditional Tabular Generative Adversarial Networks (CTGAN) and Variational Autoencoders (VAE) for synthetic data augmentation. These models effectively address the issue of class imbalance, significantly improving model accuracy, fairness, and interpretability. Among all classifiers, RF achieved the highest accuracy (83%), indicating its robustness in handling complex clinical data. The results underscore the value of integrating structured EHRs with advanced AI to enable early interventions, equitable resource distribution, and improved patient outcomes. Future research will incorporate unstructured clinical data and validate models across diverse hospital environments.

| KEYWORDS

Hospital readmission, Machine learning (ML), Generative AI, CTGAN, VAE, Electronic health records (EHR).

| ARTICLE INFORMATION

ACCEPTED: 02 August 2026 **PUBLISHED:** 11 September 2026 **DOI:** 10.32996/fcsai.2026.5.9.22

1. Introduction

A significant challenge for healthcare systems in the United States is the readmission of patients to hospitals, particularly those readmitted within 30 days of their initial discharge. Premature readmissions impose significant financial burdens on hospitals and highlight deficiencies in patient care coordination, post-discharge follow-ups, and overall healthcare quality [1]. The Centres for

Medicare and Medicaid Services (CMS) has included the penalties based on excessive readmission rates, which further highlights the necessity of effective interventions to alleviate the problem [2]. In recent years, machine learning (ML) has been used to solve readmission utilizing the massive amount of electronic health record (EHR) data. ML models predicted readmission risks well, allowing healthcare professionals to efficiently manage resources and identify high-risk patients [3,4]. Nevertheless, these models have several drawbacks that prevent their widespread adoption and success. For instance, real-world healthcare data are skewed because readmission rates are disproportionately low [5]. Numerous ML models cannot be objectively understood, and physicians who struggle to incorporate their data into clinical practice cannot simply justify their results [6]. However, models are not able to achieve detailed information about patients because they do not capture unstructured data that includes clinical notes, discharge summaries, etc. [7]. In addition, Generative artificial intelligence (AI), like Generative adversarial networks (GANs) and Variational Autoencoders (VAEs), can provide promising solutions. The techniques create natural artificial data to equalize the classes, reduce bias, and also improve the fairness and transparency of models [8,9]. Generative AI also helps overcome the constraints of data to create fair and clinically useful predictive models. Even though hospital readmission prediction has improved, it still has certain gaps. First, class imbalance in healthcare datasets has long been a major issue, with previous ML algorithms failing to forecast the outcomes of minority classes like readmitted patients [5]. Second, there is the barrier to clinical adoption because complex models like deep neural networks do not have clear interpretations, which are a must for implementing the model in the clinic [6]. Thirdly, EHR unstructured information is underutilised because of difficulty in natural language processing (NLP) [7]. Lastly, algorithmic bias will negatively affect vulnerable groups such as racial minorities and those with low incomes [10]. This study seeks to address these challenges by combining analytical ML techniques with generative AI approaches and fair to minimise early admissions in the U.S. hospital systems. The main Objective of the study is to establish the integrated predictive framework, including both the existing models of the traditional machine learning classification as Logistic Regression (LR), Random Forest (RF), and XGBoost, and the models of the generative artificial intelligence, namely Conditional Tabular GAN (CTGAN) and Variational Autoencoders (VAE), to predict the 30-day hospital readmission. The dataset neuro uses in this framework was the UCI Diabetes 130-US Hospitals Dataset, a structured dataset with clinically rich information that was used to capture the number of admissions by diabetic patients in various healthcare facilities within the United States. Additionally, another aim is to calculate the resilience of the models with the help of the resilient methodologies, e.g., AUC-ROC, F1 score, Precision-Recall curves, and the SHAP incorporation into the models to better understand them. Additionally, the implications of this research are significant to achieve better outcomes in regard to healthcare and minimize the expenses related to readmission of patients in hospitals [1]. Having well-made and interpretable predictive models, the hospital can be proactive when singularly identifying high-risk patients and taking specific steps like improving post-discharge assistance and individual care plans [11]. Such proactive practice not only decreases the rates of readmission but also increases satisfaction among patients and the quality of healthcare in general. Furthermore, the integration of generative AI methods can solve urgent problems of healthcare data analytics (data scarcity, imbalance, and bias) [9]. The use of synthetic data creation and GANs can introduce new approaches to these concerns and open the way for predictive models to be more equitable and reliable [10]. The proposed models can reduce algorithmic biases and thus AI-based tools will not reinforce institutional racism and other forms of inequality that exist within the healthcare system [8].

2. Literature Review

Numerous studies have used classic and group ML to forecast hospital readmission on the basis of clinical data. The logistic regression has a history as a baseline technique because it is easy to interpret, but it has a tendency to fail under nonlinear or high-dimensional problems [12]. RF and XGBoost, in their turn, have been shown to perform better as they can capture the multi-feature interaction and deal with missing values in a fairly efficient manner [13]. These ensemble approaches are gaining popularity in the area of healthcare analytics, especially because of their capability to analyse large amounts of data and make strong predictions. Moreover, deep learning models have expanded the field in terms of their ability to train temporal patterns in the EHRs, particularly in the case of attention-based architecture and Long Short-Term Memory (LSTM) networks [14]. As an example, transformer-based models have been demonstrated to be effective in predicting diabetes patient readmissions by utilizing the sequence of data to increase the accuracy of the predictions [15]. However, their opacity raises concerns for the clinical adoption because clinicians need interpretable results to make decisions [6]. SHAP (SHapley Additive exPlanations) has been recently introduced, which makes these models more interpretable, so clinicians can learn about what characteristics affect the readmission risk the most [16]. SHAP values can be used to provide information about the importance of features, which can help healthcare providers get vital information on aspects that lead to readmission. However, these models find it difficult to deal with class imbalance, where the minority class (readmitted patients) is a very tiny representation of the data. This inconsistency can give way to prejudiced forecasts and a lessening in the reliability of the models, especially in the real world. In order to deal with this task, synthetic data have been generated using generative models. Techniques like CTGAN and VAE have already been applied to balance the minority classes without tampering with either the statistical or clinical rest of EHR data [17]. Moreover, such generative models are able to create realistic (but imbalanced) data and boost the accuracy and recall overhead of predictive models. In scientific works, there are noted instances of the application of CTGAN and VAE to extremely skewed data, such as the UCI Diabetes data set, when there is a plausibly beneficial effect on the results of the created models [5]. Although research in hospital readmission prediction has progressed, however, there are still gaps in the literature that should be addressed. First, few studies have been able to combine generative AI with conventional ML models with the application towards 30-day hospital readmissions due to diabetes-related admissions [18]. Second, the statistical and clinical performance of synthetic data usually cannot be trusted, and the products of the model in practice are extremely uncertain [9]. The relationship between the distributions of synthetic data to those in realistic datasets should be of paramount importance in terms of model fidelity and clinical relevance. Lastly, despite the availability of explainability tools, such as SHAP, their application in assessing the impact of synthetic data augmentation is rather limited, to date [6]. Transparency and interpretability are critical so that the use of synthetic features can be understood and how they are used to make predictions by models. Additionally, the issue of algorithmic sensitivity and applicability, in particular, in hospital systems having different backgrounds, is also unexplored [8]. It is possible that a trained model on one dataset cannot work effectively in a different healthcare institution, emphasising the necessity of a study with a focus on the cross-institutional variability. Moreover, the experimental study rigorously integrates and quantifies the impact of ML classifiers with generative models and interprets the results of predicting the reliability of synthetic features using standard and adjusted scales [19]. The current study will attempt to plug that gap by designing an interlinked framework that takes advantage of both conventional ML and generative AI methods, so that the models derived are accurate, interpretable and fair.

3. Methodology

3.1 Dataset Description

The data in the present study is the UCI Diabetes 130-US Hospitals Dataset, a publicly available structured dataset collected via the UCI Machine Learning Repository [4]. It includes 100,000 inpatient stays of patients who have diabetes and has been gathered in 130 hospitals in the United States over a decade, and it consists of 22 key features.

Table 1: Key Attributes in the UCI Diabetes Dataset

S/No	Key Attribute / Feature	Description
1	Encounter ID	Unique identifier for each hospital admission instance.
2	Patient Number	Anonymised unique identifier for each patient.
3	Race	Self-reported patient race (e.g., Caucasian, African American).
4	Gender	Patient's gender (Male, Female, Unknown/Invalid).
5	Age	Age group in 10-year intervals (e.g., [50-60), [70-80)).
6	Admission Type	Type of admission (e.g., Emergency, Urgent, Elective).
7	Discharge Disposition	Outcome or status at discharge (e.g., Discharged to home, Died).
8	Admission Source	Source of admission (e.g., Physician referral, Emergency room).
9	Time in Hospital	Length of stay during the admission (in days).
10	Number of Lab Procedures	Count of lab tests performed during the stay.
11	Number of Medications	Count of distinct medications prescribed.
12	Number of Diagnoses	Number of distinct ICD-9 diagnoses recorded.
13	Diagnosis Codes (diag_1, diag_2, diag_3)	ICD-9 codes representing primary and secondary diagnoses.
14	Change	Indicates whether there was a change in medication (Yes/No).
15	Diabetes Medication	Indicates if diabetes medication was prescribed (Yes/No).

16	Readmission	Target variable indicating readmission: <30, >30, or NO.
17	A1C Result	Most recent A1C test result (e.g., >7, Normal).
18	Insulin	Insulin use during hospital stay (e.g., Up, Down, No, Steady).
19	Max Glu Serum	Maximum blood glucose level recorded (e.g., >200, Normal).
20	Number of Inpatient/Outpatient/ER Visits	Number of previous hospital encounters across care settings.
21	Weight	Patient's weight group (though many values are missing).
22	Medical Specialty	Specialty of the attending physician (e.g., Cardiology, Endocrinology).

3.2 Data Pre-Processing and Cleaning

To facilitate the data analysis through machine learning and achieve the correct data quality, pre-processing took place in several methodological stages, such as missing values and relevancy, categorical features encoding, and predictive attributes [3].

Step 1: Check for Missing Values: The dataset was analysed regarding missed entries in all the variables.

The missing values in individual columns were identified and counted. On the basis of this evaluation, correct approaches like imputation or removal were implemented. This procedure allowed the completeness of the dataset to be guaranteed and eliminated biases in the data due to missing records [20].

Step 2: Dropping insignificant columns: Unnecessary columns, complete or partial identifiers and low or

no predictor variables were dropped. Filtering to the clinically relevant attributes skewed the data set, minimized noise, and enhanced the speed of the next processing [21].

Step 3: Categorical Variable Encoding: Categorical values were coded into numeric values to facilitate

the compatibility of the machine learning models, and continuous variables were normalised to represent consistency [22]. The targeted variable "readmitted_le" in the dataset reflects three categories of hospital readmission outcomes- a value of 0 would mean patients who were not readmitted at all, a value of 1 would mean patients have been readmitted within 30 days, and a value of 2 would mean that patients have been readmitted after 30 days. Following the data cleaning process, the total dataset consists of 47,813 patients, whereby 27,048 of the patients (approximately 56.6%) fit into the category of patients readmitted after the 30-day mark and 15,366 patients (32.1%) were readmitted before the 30-day mark. The number of patients who were not readmitted altogether was only 5,399 (11.3%). This distribution in Figure 1 shows that there was a natural class imbalance, with non-readmitted patients forming a minority group.

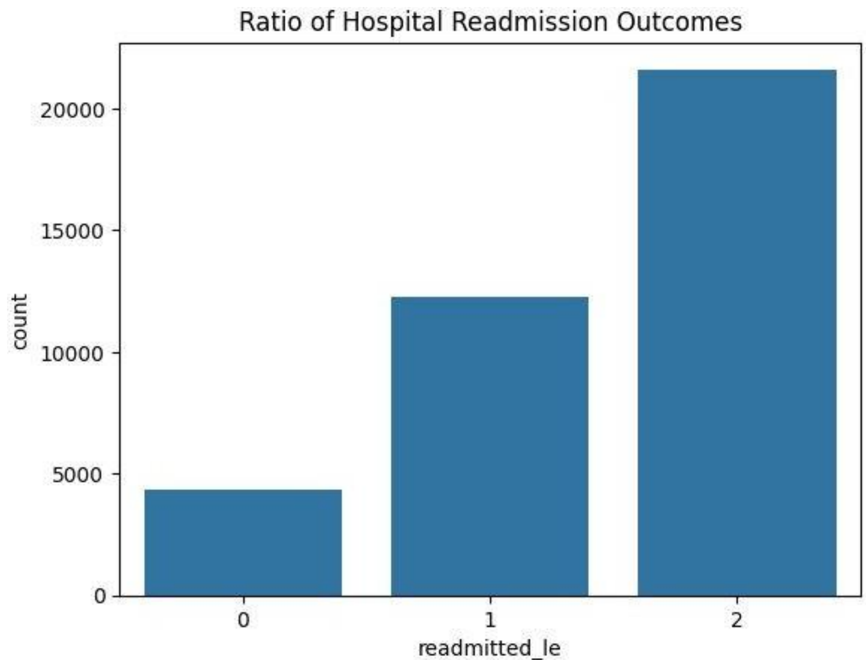

Figure 1: Ratio of Hospital readmission outcomes

3.2.1 Feature Selection

The feature selection algorithm that is employed in this study is the Extra Trees classifier, which generates many decision trees without any pruning of the decision trees and pools the predictions of the decision trees with each other to make the final decision based on the majority vote [20]. K features are randomly assigned to each tree to select the best feature. The feature selection can be observed in the Extra Trees classifier, which is shown in Figure 2.

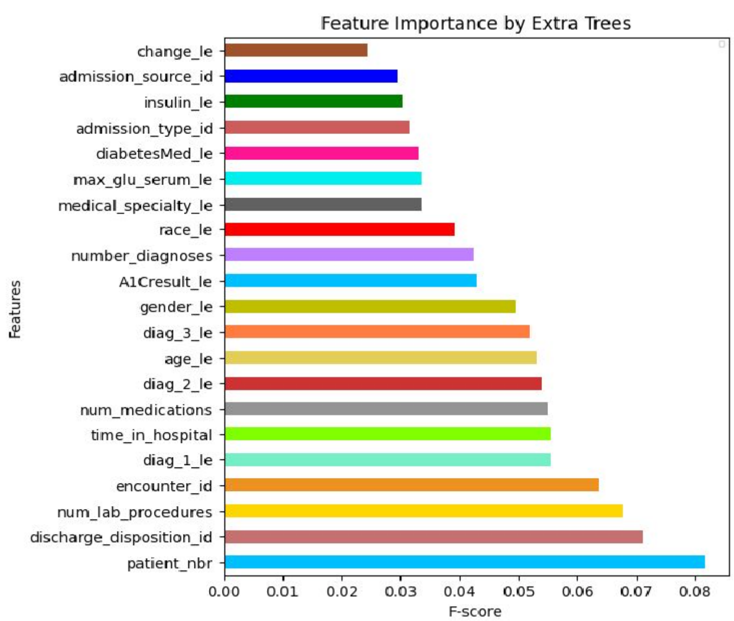

Figure 2: Significant features based on scores

The Pearson correlation was considered between the attribute and the class with the view of establishing the value of the attribute. A weighted average can be used to obtain an overall correlation of a nominal property [22]. A general chart of features with correlation basis of feature selection is shown in Figure 3.

3.3 Model Development

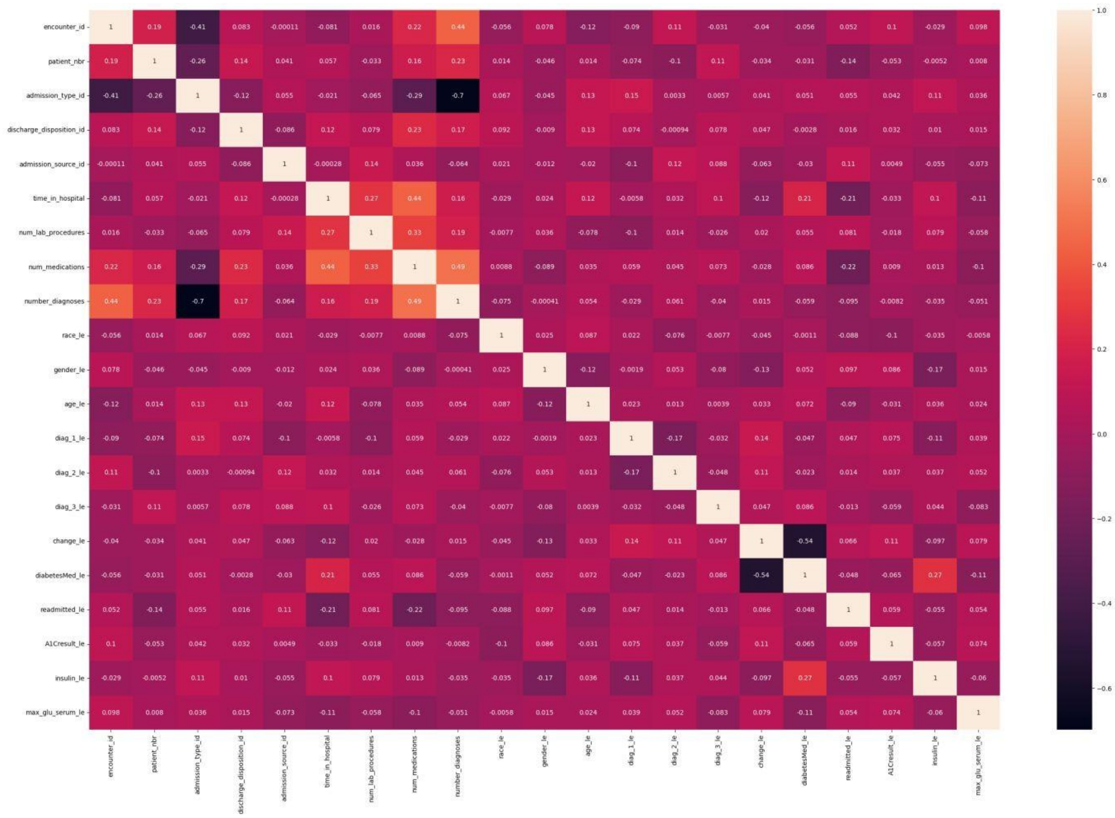

Figure 3: Correlation analysis of key features

The objective of the study was to address the issue of patient diabetes hospitalisation more precisely, predicting readmissions and improving the degree of satisfaction and logicality by integrating generative AI with traditional machine learning approaches. The specific algorithms used were KNN, RF and Gradient Boosting Classifier (GBC). They were selected due to the capacity to work with structured healthcare data and model interesting feature relationships [13]. KNN makes similarity-based classification, which relies on the nearest neighbors in the feature space. Although it is practical with small datasets, with large amounts, it can be computationally costly. KNN was used to benchmark more complex methods, such as augmenting synthetic data [23]. RF is an ensemble technique which trains several decision trees and then consolidates their decision on testing. It can resist overfitting, it can work with missing values, and it can deal with high-dimensional data, which is why it can be successfully applied in healthcare [20]. RF also gave feature importance measurements, which improved the interpretability by establishing important factors affecting readmission risks. GBC successively corrects errors by replicating weak learners, whose results are successful in skewed data, such as that around hospital readmissions [5]. It performed better than other models in prediction and the issue of class imbalance. Each of the models was compared based on typical metrics, accuracy, precision, recall, and F1, to evaluate its strengths and weaknesses regarding the ability to predict a readmission.

3.4 Generative AI Integration

Applying generative AI models to improve model generalizability by addressing the critical issue of class imbalance and preventing an accurate prediction of minority classes, such as early readmissions, as a limiting element, the proposed study uses both CTGAN and VAE. CTGAN compiles synthetic data that looks real, given labels of the classes, which effectively oversamples minority groups (e.g., readmitted patients). This method overcomes the problem of performance degradation induced by unequal datasets and high standards of statistical and visual validation are confident in clinical relevance [24, 25]. Additionally, VAE reproduces patient records directly in a latent space with their original structure and content, producing synthetic samples. The models are made more robust and reliable by doing this to include VAE-generated data in training, resolving the issue of the shortage of healthcare-related data without suffering a patient privacy reduction [26]. The combination of CTGAN and VAE evenly balances the data, reduces algorithm biases, and evolves models for predicting in line with ethical norms in health care [8].

3.5 Evaluation Metrics

The performance of this study was quantified using such performance indicators as accuracy, precision, recall and F1 score of the classification models [1]. The number of correctly predicted cases, divided by the number of total cases, gives accuracy. Precision is the fraction of true positives out of the total positives that are anticipated and indicates the quality of the model to avoid false positives [6]. Recall is how much interest rate the model has been attaining, represented as a ratio between true positives and actual positives. F1 is a harmonic mean of the precision and recall measures, where they are balanced against one another [5].

4. Findings

4.1 Baseline Performance

The baseline performance was assessed using RF, KNN, and GBC. These models were trained and tested (80-20) on the preprocessed dataset to establish a foundation for comparison with enhanced models incorporating synthetic data generation.

4.1.1 RF

The Random Forest model showed a great performance in all three classes of readmission. As demonstrated in Table 2 and Figure 4, the classifier realised precisions, recalls, and F1-scores of about 0.84-0.85, resulting in a general accuracy of 83%. Both Macro and weighted average were 0.83, indicating the equality of the performance among the classes. Although it works, RF is believed to yield better results in conjunction with synthetic data in order to enhance generalizability, especially to minority populations [5].

Table 2: Depicts RF Classification Report

Classification Report is:

	precision	recall	f1-score	support
0	0.84	0.85	0.85	3149
1	0.82	0.80	0.81	3261
2	0.84	0.85	0.85	3153
accuracy			0.83	9563
macro avg	0.83	0.83	0.83	9563
weighted avg	0.83	0.83	0.83	9563

1) *4.1.2 KNN*

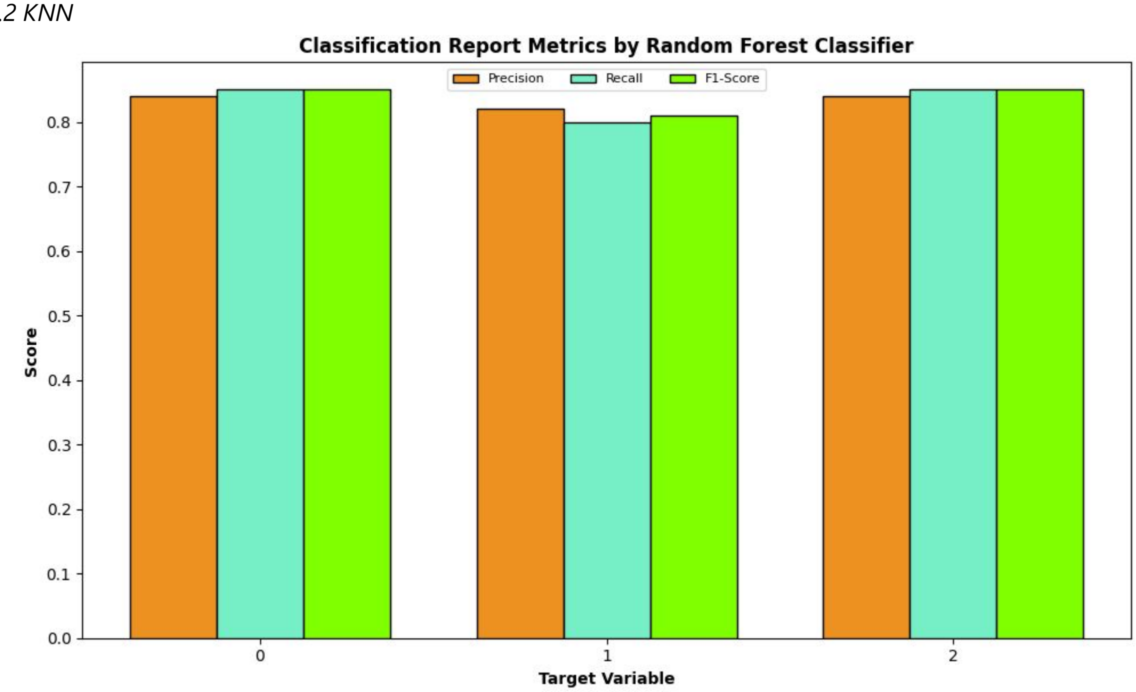

Figure 4: Performance metrics of RF

The results of the KNN model in terms of classification are listed in Table 3 and Figure 5, which offers an in-depth overview of the classification with their percentages. The following classification report is the performance of a binary classification model where three classes are included, i.e., 0, 1, and 2, whereas there are around 9,563 samples of classes 0, 1 and 2. The rest of the classes have similar effectiveness of the model, with the precision and recall scores of approximately 1.00-0.52, hence producing f1 scores of 0.68, 0,73 and 0,74, respectively. The overall model accuracy is 0.73, and this same value is reflected as a weighted average; however, this does not have equal priority to each class of the macro average.

Table 3: Displays KNN Classification Report

Classification Report is:

	precision	recall	f1-score	support
0	0.52	1.00	0.68	1048
1	0.63	0.88	0.73	3090
2	1.00	0.59	0.74	5425
accuracy			0.73	9563
macro avg	0.71	0.82	0.72	9563
weighted avg	0.83	0.73	0.73	9563

4.1.3 GBC

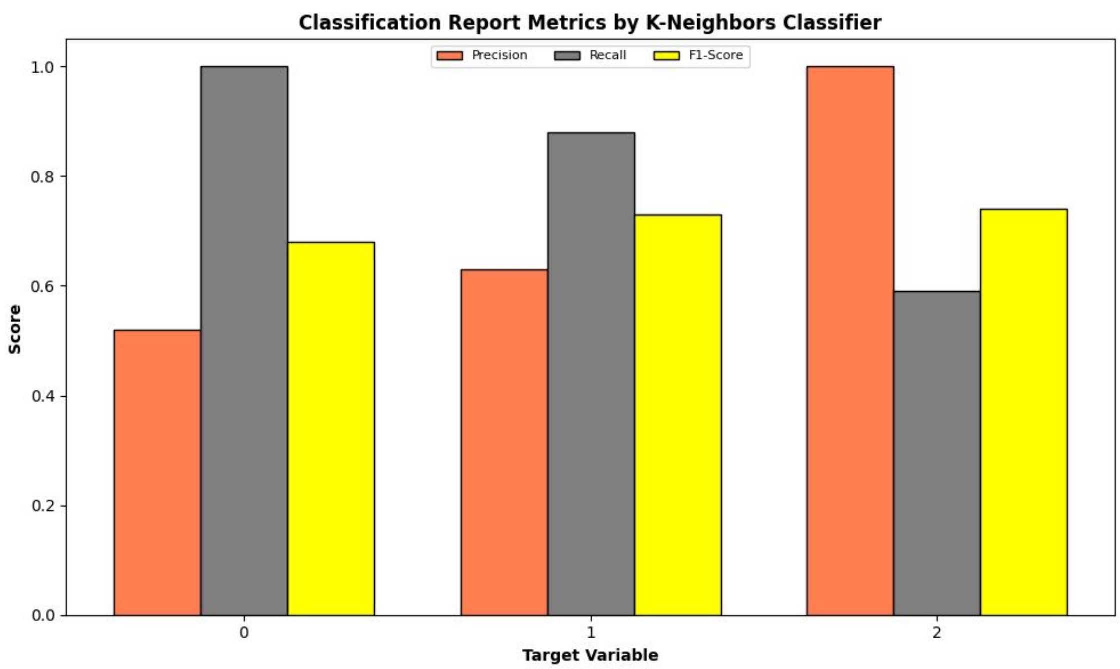

Figure 5: Performance metrics of KNN

The training report of the Gradient Boosting Classifier, along with a detailed classification report, can be seen in Table 4 and Figure 6, respectively. In the classification report, this model is getting terrible accuracy, recall, and F1 score on classes 0 and 1, whereas on class 2, it is getting the same as the precision, recall and F1 score of 0.62, 0.91 and 0.74. The macro average does not give a global accuracy. Weighted average provides 0.60 as an evaluation of precision, recall due to taking into consideration the imbalance of classes. Such findings suggest that, probably, the model is performing well in most of the classes; however, it is performing dismally when it comes to the minority class, which could be one of the areas where it could be improved.

Table 4: Presents GBC Classification Report

Classification Report is:

	precision	recall	f1-score	support
0	0.70	0.01	0.02	1543
1	0.53	0.26	0.35	4602
2	0.62	0.91	0.74	8199
accuracy			0.60	14344
macro avg	0.62	0.39	0.37	14344
weighted avg	0.60	0.60	0.53	14344

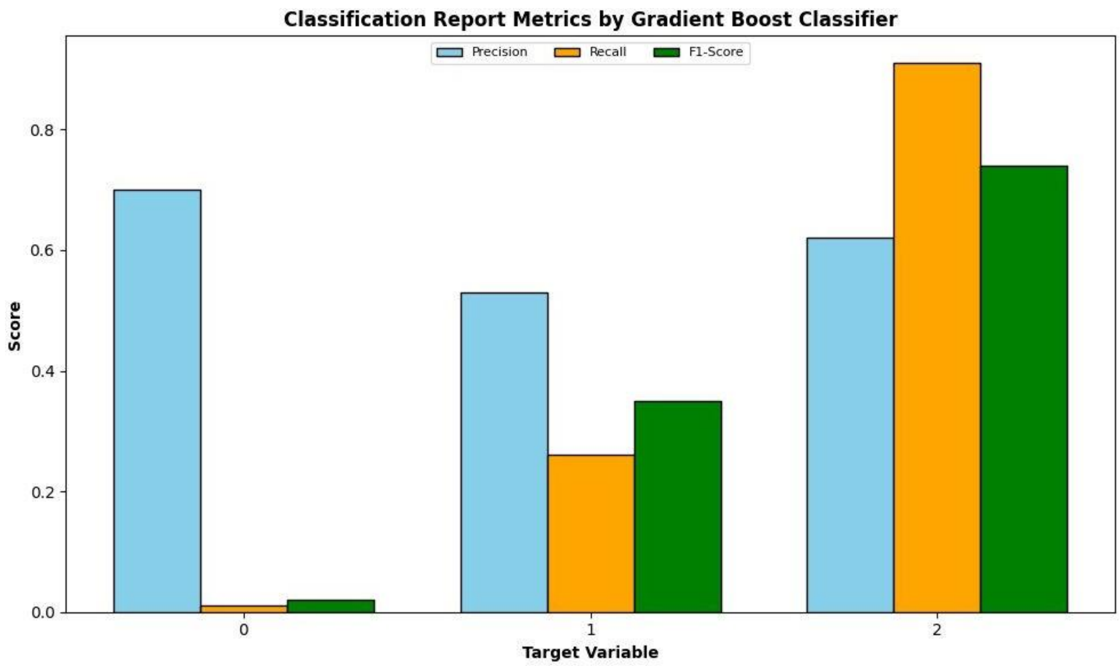

Figure 6: Performance metrics of GBC

The overall findings of the experiment in terms of accuracy, precision, recall, and f1-score are shown in Table 5, and the overall accuracy is given in Figure 7. The receiver operating characteristic (ROC) curve is displayed in Figure 8, where RF provide better performance.

Table 5: Overall accuracy and performance metrics

Model	Outcome	Precision	Recall	F1-score	Accuracy (%)
RF	0	0.84	0.85	0.85	83.0
RF	1	0.82	0.80	0.81	
RF	2	0.84	0.85	0.85	
KNN	0	0.52	1.00	0.68	73.0
KNN	1	0.63	0.88	0.73	
KNN	2	1.00	0.59	0.74	
GBC	0	0.70	0.01	0.02	60.0
GBC	1	0.53	0.26	0.35	
GBC	2	0.62	0.91	0.74	

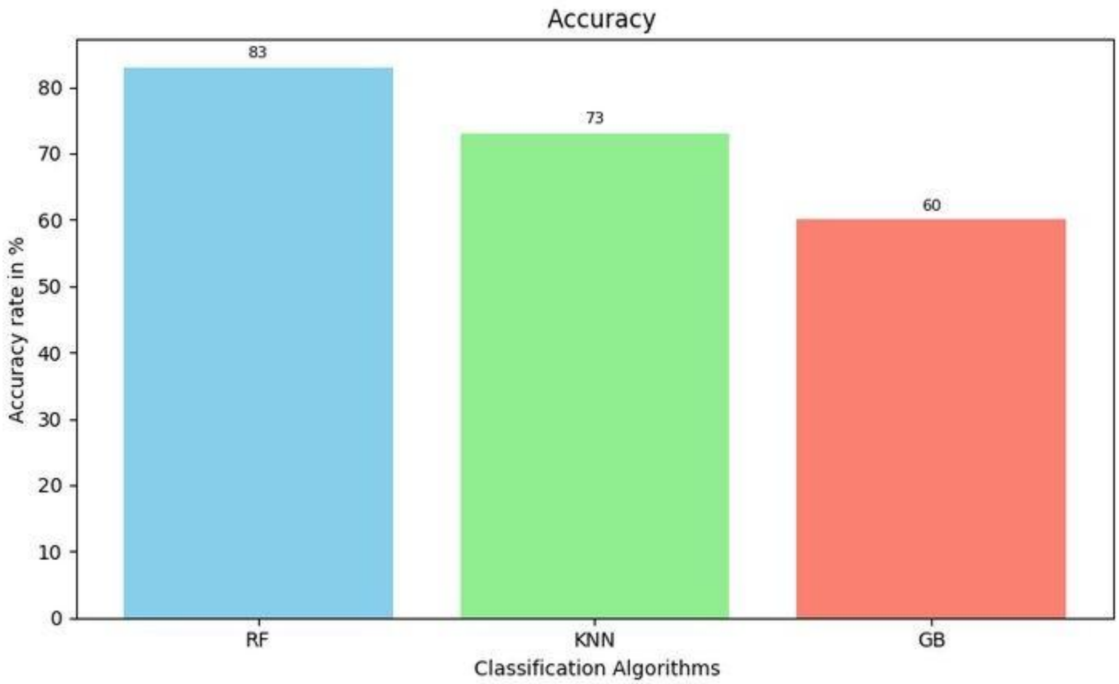

Figure 7: Overall Accuracy of the models

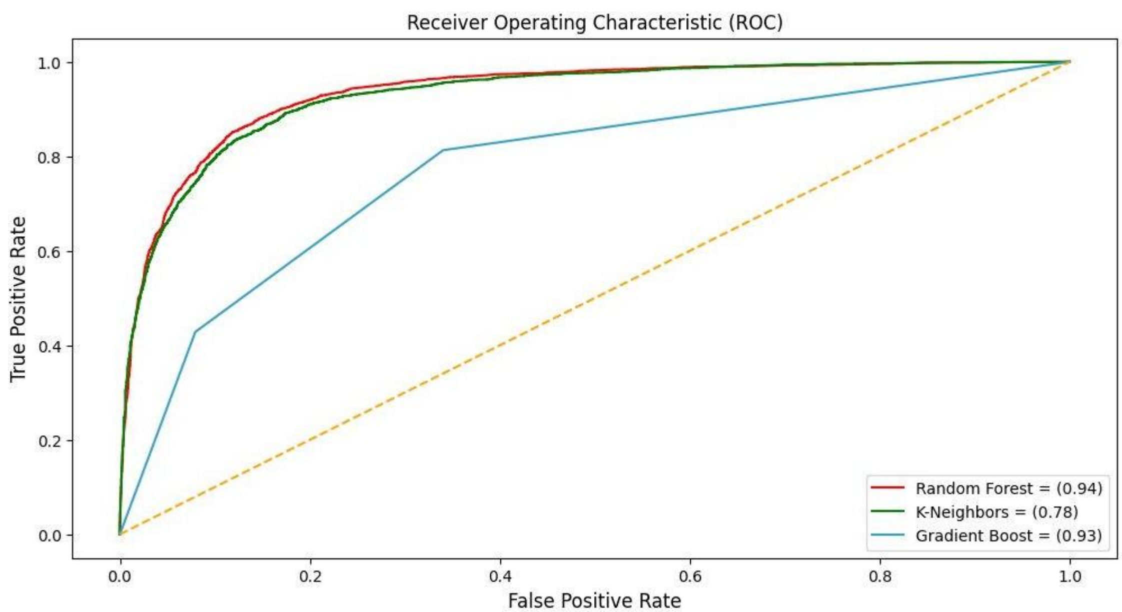

Figure 8: Receiver operating characteristic (ROC) curve

4.2 CTGAN and VAE model 4.2.1 CTGAN: Generating Realistic Synthetic Data The CTGAN was uniquely set up to create synthetic tabular data that is derived under the condition of the class labels. It is especially better when applied to healthcare data, when the population size of minority classes (e.g., patients readmitted to the hospital) is usually too small to train a good machine learning model [5]. CTGAN uses the adversarial training of a generator and a discriminator. The generator is taught to produce the artificial examples which tend to resemble the statistical peculiarities of the initial data, and the discriminator is taught to decide on the validity of examples. CTGAN enforces the generative process to be contingent on the class labels to make the synthetically generated data representative of the minority classes [27]. To be able to balance the dataset, CTGAN was applied to oversample the minority class (readmitted patients) in this study. 4.2.2 VAE: Latent Space-Based Reconstruction The VAE-based synthetic samples were specifically used to enhance the minority classes so that the problem of class imbalance was solved. These synthetic data sets were proven correct depending on the same statistical and visual examinations as the CTGAN ones [9]. The combination of CTGAN and VAE caused the class imbalance problem in the dataset to decrease drastically, as observed in Figure 9, which is a comparison of the distribution of a selected feature density distribution in the original dataset against those of the CTGAN- and VAE-generated synthetic datasets.

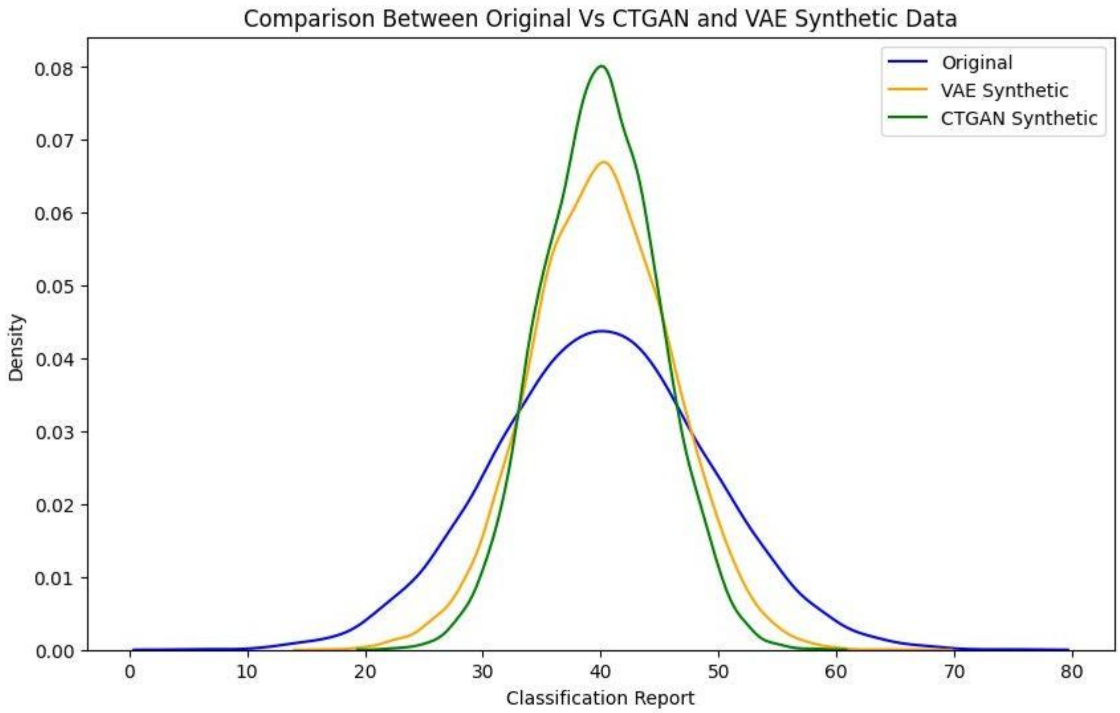

Figure 9: Comparison Between Original, CTGAN and VAE Synthetic Data

4.4 Comparative Analysis The three machine learning models, RF, KNN, and GBC, were subjected to a comparative assessment in terms of their ability to identify whether or not the patient will be readmitted in 30 days as a diabetic. The accuracy of RF was found to be the best, with an accuracy of 83%. RF had good precision and recall values, especially on class 2 (2 patients not read), based on the accuracy of the classification. The ensemble models, such as the RF, were resilient when dealing with multifaceted clinical data [1]. KNN displayed a middle score of 73% accuracy, typical of its sensitivity to scaling of the features and the local concentration of the data [3]. GBC, though equipped with an advanced technique to boost, also performed poorly with an accuracy of 60% because the algorithm is being overfit to the majority class, in light of how class imbalance affects boosting algorithms. The discrepancy between them points to the fact that model selection and fine- tuning are critical when working with imbalanced clinical data. The results support ensemble models, which achieve the best superiority when tuned properly, at preserving the intricate patterns in structured info in EHR [28, 29].

5. Discussion

5.1 Implications of the Findings

This paper indicates that a combination of generative AI and traditional ML classifiers leads to a significant boost in the hospital readmission predictive accuracy within 30 days in diabetic patients. In particular, specific algorithms, including CTGAN and VAE, solved the problem of imbalance between classes, and allowed the reflecting models to represent outcomes in minorities without compromising the precision. The modified Random Forest classifier (enhanced with synthetic data) was always the most efficient one, highlighting the benefit of the ensemble approaches in the field of structured healthcare analytics. Additionally, the results are contributions to predictive modeling studies in two aspects. To begin with, they show how integrating generative models and explainable ML methods generates analogies that are accurate and explainable. This overcomes, long-standing, critiques of deep learning models in healthcare, which, in most cases, prioritize performance at the expense of interpretability [6]. Second, the findings are relevant in the growing body of algorithmic fairness research that demonstrates that synthetic amplification can rehabilitate biases found in biased clinical data [30]. Furthermore, the framework offers practical hospital facilitation in the identification of high-risk patients in advance of discharge, in a clinical sense. Minority cases that are better recalled promote prompt measures like increased post-discharge monitoring, prescription of drugs or timely follow-ups. To healthcare administrators, decreased readmission rates are a direct savings of money, which comes in the form of CMS penalty reduction [2]. To policymakers and AI establishments, the study demonstrates how it is possible to incorporate fairness-sensitive generative models into clinical decision-support systems to ensure that technological innovation is matched by ethical healthcare service provision [11].

5.2 Limitations

Despite all these promising contributions, this study has a number of limitations. First of all, it is based on the secondary information contained in the UCI Diabetes 130-US Hospital data [4], and it does not represent the current changes in clinical practice, treatment guidelines, and EHR system advances. Additionally, the data point does not contain any unstructured data like physician notes, lab reports, and socio-economic data that would bring a notable increase in predictive performance [19]. However, their risk is that generative models such as CTGAN and VAE are potentially subject to the synthetic bias in that synthetic data distributions do not show statistical validations [31]. Additionally, the models designed here have been trained and tested in a retrospective, static environment and deployed in real-time hospital systems, hence limiting their further generalizability in a live setting, in which they would need to be validated prospectively [32].

5.3 Future Directions

Future studies would be needed to verify the proposed model with big, more heterogeneous, and current datasets across multiple healthcare facilities in order to improve robustness and generalizability. A new trend is discharge summaries, lab notes, and prescriptions that are not structured and better integrated using NLP infrastructures [32, 33]. Additionally, privacy-preserving federated learning systems and differential privacy generative models would allow hospitals to work on predictive models without breaching patient confidentiality. Lastly, the practical application of these AI models on the processes of clinical work and patient outcomes and satisfaction should be evaluated in controlled clinical trials.

6. Conclusion

In this research, a hybrid system was proposed that integrates traditional ML methods with the generative AI approach to determine typical readmission based on the UCI Diabetes 130-US Hospitals data 30 days following the patient's discharge. Incorporating both the CTGAN and VAE in augmenting synthetic data helped the framework to rectify the problem of class imbalance, improve predictive ability, and enhance the validity of the minority classes detection. Random Forest outperformed the other tested classifiers in the case of accuracy, precision, recall, and F1-score. In addition, this study has shown that the interpretable and fairness-aware AI models are able to deliver actionable insights to the clinician and healthcare administrator. The framework promotes early interventions, helps reduce avoidable readmissions and eases financial and operational pressures on hospitals by identifying high-risk patients before they are discharged. Moreover, the results show the disruptive capabilities generative AI can have in clinical predictive modeling. Future studies are required to confirm this framework in a variety of healthcare contexts, with unstructured EHR data, and by assessing practical implementation using clinical trials. Together, these will help to further the progress of clear, fair, and clinical decision-support AI-based systems.

Funding Statement

No Funding is available for this research

Conflicts of Interest

The author(s) have no conflict of interest to declare

Ethical Statement

This study did not need any ethical approval, as the UCI Diabetes 130-US Hospitals Dataset, a publicly available structured dataset, was used for this study.

References

1. Davis S, Zhang J, Lee I, Rezaei M, Greiner R, McAlister FA, et al. Effective hospital readmission prediction models using machine-learned features. BMC Health Serv Res. 2022;22(1):1415. doi:10.1186/s12913-022-08748-y
2. Shang Y, Jiang K, Wang L, Zhang Z, Zhou S, Liu Y, et al. The 30-days hospital readmission risk in diabetic patients: Predictive modeling with machine learning classifiers. BMC Med Inform Decis Mak. 2021;21(Suppl 2):57. doi:10.1186/s12911-021-01423-y
3. Baek H, Cho M, Kim S, Hwang H, Song M, Yoo S. Analysis of length of hospital stay using electronic health records: A statistical and data mining approach. PLoS One. 2018;13(4):e0195901. doi:10.1371/journal.pone.0195901
4. Clore J, Cios K, DeShazo J, Strack B. Diabetes 130-US Hospitals for Years 1999–2008 [dataset]. UCI Machine Learning Repository; 2014. doi:10.24432/C5230J
5. Wang AX, Le VT, Trung HN, Nguyen BP. Addressing imbalance in health data: Synthetic minority oversampling using deep learning. Comput Biol Med. 2025;188:109830. doi:10.1016/j.compbiomed.2025.109830
6. Mohanty SD, Lekan D, McCoy TP, Jenkins M, Manda P. Machine learning for predicting readmission risk among the frail: Explainable AI for healthcare. Patterns. 2022;3(1):100395. doi:10.1016/j.patter.2021.100395
7. Fierro C, Pérez J, Mora J. Predicting unplanned readmissions with highly unstructured data. arXiv. 2020. doi:10.48550/arXiv.2003.11622

8. Wang HE, Weiner JP, Saria S, Kharrazi H. Evaluating algorithmic bias in 30-day hospital readmission models: Retrospective analysis. J Med Internet Res. 2024;26:e47125. doi:10.2196/47125
9. Lan L, You L, Zhang Z, Fan Z, Zhao W, Zeng N, et al. Generative adversarial networks and their applications in biomedical informatics. Front Public Health. 2020;8:164. doi:10.3389/fpubh.2020.00164
10. Smolyak D, Bjarnadóttir MV, Crowley K, Agarwal R. Large language models and synthetic health data: Progress and prospects. JAMIA Open. 2024;7(4):ooae114. doi:10.1093/jamiaopen/ooae114
11. Ahmed Z, Mohamed K, Zeeshan S, Dong X. Artificial intelligence with multi-functional machine learning platform development for better healthcare and precision medicine. Database. 2020;2020:baaa010. doi:10.1093/database/baaa010
12. Badawy M, Ramadan N, Hefny HA. Healthcare predictive analytics using machine learning and deep learning techniques: a survey. Journal of Electrical Systems and Information Technology. 2023 Aug 29;10(1):40.
13. Al Alawi AM, Al Abdali M, Al Mezeini AZA, Al Rawahia T, Al Amri E, Al Salmani M, et al. Development and evaluation of a machine learning model for predicting 30-day readmission in general internal medicine. Computers. 2025;14(5):177. doi:10.3390/computers14050177
14. Chiu CC, Wu CM, Chien TN, Kao LJ, Li C. Predicting ICU readmission from electronic health records via BERTopic with long short-term memory network approach. J Clin Med. 2024;13(18):5503. doi:10.3390/jcm13185503
15. García-Mosquera J, Villa-Monedero M, Gil-Martín M, San-Segundo R. Transformer-based prediction of hospital readmissions for diabetes patients. Electronics. 2025;14(1):174. doi:10.3390/electronics14010174
16. Emi-Johnson OG, Nkrumah KJ. Predicting 30-day hospital readmission in patients with diabetes using machine learning on electronic health record data. Cureus. 2025;17(4):e82437. doi:10.7759/cureus.82437
17. Bietsch D, Stahlbock R, Voß S. Synthetic data as a proxy for real-world electronic health records in the patient length of stay prediction. Sustainability. 2023;15(18):13690. doi:10.3390/su151813690
18. Cara R, Dhoska K, Cara F, Kayusi F, Juma L. An artificial intelligence model for predicting hospital readmission using electronic health records data. Mesopotamian J Artif Intell Healthc. 2025;2025:116–23. doi:10.58496/MJAIH/2025/012
19. Samani R, Dehghani M, Shahrokh F. Enhancing Readmission Prediction with Deep Learning: Extracting Biomedical Concepts from Clinical Texts. arXiv preprint arXiv:2403.09722. 2024 Mar 12.
20. Suchi TA, Rabbi MA, Layek MA. Effective feature selection and soft voting classifier-based diabetes detection using machine learning approaches. In: Proc Int Conf Next-Generation Computing, IoT and Machine Learning (NCIM). IEEE; 2023. p.1–7. doi:10.1109/NCIM59001.2023.10212616
21. Estiri H, Strasser ZH, Klann JG, Naseri P, Wagholikar KB, Murphy SN. Predicting COVID-19 mortality with electronic medical records. NPJ digital medicine. 2021 Feb 4;4(1):15. doi: /10.1038/s41746-021-00383-x
22. Saxena R, Sharma SK, Gupta M, Sampada GC. A novel approach for feature selection and classification of diabetes mellitus: machine learning methods. Comput Intell Neurosci. 2022;2022:3820360. doi:10.1155/2022/3820360
23. Araf I, Idri A, Chairi I. Cost-sensitive learning for imbalanced medical data: a review. Artif Intell Rev. 2024;57:80. doi:10.1007/s10462-023-10652-8
24. Parise O, Kronenberger R, Parise G, de Asmundis C, Gelsomino S, La Meir M. CTGAN-driven synthetic data generation: A multidisciplinary, expert-guided approach. Comput Methods Programs Biomed. 2025;259:108523. doi:10.1016/j.cmpb.2024.108523
25. Budu E. Evaluation of synthetic electronic health records: Metrics and methodologies. Artificial Intelligence in Medicine. 2024. doi: 10.1016/j.artmed.2024.102450
26. Li J, Cairns BJ, Li J, et al. Generating synthetic mixed-type longitudinal electronic health records for artificial intelligent applications. NPJ Digit Med. 2023;6:98. doi:10.1038/s41746-023-00834-7
27. Nik AH, Riegler MA, Halvorsen P, Storås AM. Generation of synthetic tabular healthcare data using generative adversarial networks. In International Conference on Multimedia Modeling 2023 Jan 9 (pp. 434-446). Cham: Springer International Publishing.
28. Zarghani A. Comparative analysis of LSTM neural networks and traditional machine learning models for predicting diabetes patient readmission. arXiv. 2024. doi:10.48550/arXiv.2406.19980
29. Alabdulhafith M, Saleh H, Elmannai H, Ali ZH, El-Sappagh S, Hu JW, et al. A clinical decision support system for edge/cloud ICU readmission model based on particle swarm optimization, ensemble machine learning, and explainable artificial intelligence. IEEE Access. 2023;11:100604–21. doi:10.1109/ACCESS.2023.10241995
30. Hu, T.-L., Chao, C.-M., Wu, C.-C., Chien, T.-N., & Li, C. (2024). Machine learning-based predictions of mortality and 30-day post-discharge readmission among type 2 diabetes ICU patients. Applied Sciences, 14(18), 8443. doi: 10.3390/app14188443
31. Mishra V, Tanniru MR, Sreedharan J. Prediction of 30-day readmission in diabetes management using Machine learning. Computers in Biology and Medicine. 2025 Sep 1;195:110616. doi: 10.1016/S0010-4825(25)00096-79
32. Liu VB, Sue LY, Wu Y. Comparison of machine learning models for predicting 30-day readmission rates for patients with diabetes. J Med Artif Intell. 2024;7:23. doi:10.21037/jmai-24-70
33. El-Kenawy E.-S., M. Zaki A. M., Eid, M. M. AdaBoost-based classification for bone sarcoma outcome prediction: A comparative machine learning approach. Mesopotamian Journal of Artificial Intelligence in Healthcare.