

A Guided Stochastic Gradient Descent Enhanced Neural Networks for Early Diabetes Readmission Prediction

Gandharapu Vineela *, Gogineni Krishna Chaitanya

Department of Computer Science and Engineering, Koneru Lakshmaiah Education Foundation,
Vaddeswaram, 522502, Andhra Pradesh, India

*Corresponding author E-mail: gkc_chaitu@yahoo.com

Received: July 24, 2025, Accepted: September 10, 2025, Published: September 24, 2025

Abstract

Diabetes is a prevalent chronic health condition globally, posing significant challenges to healthcare systems and insurance companies due to its associated risks of hospital readmission. Early prediction of readmissions, especially within 30 days, is crucial for directing attention to high-risk patients and optimizing healthcare resources. This study explores the application of machine learning models, including Extreme Gradient Boosting (XGBoost), Support Vector Machine (SVM), and Artificial Neural Networks (ANN), to predict diabetic readmissions. A novel approach is proposed, leveraging a guided optimizer for ANN training to enhance classification accuracy and error convergence. Results demonstrate up to a 1.5% improvement in classification accuracy compared to standard methods, highlighting the effectiveness of the guided optimizer in capturing consistent data patterns. By integrating AI-driven predictive analytics, this project aims to improve healthcare efficiency, reduce hospital readmissions, and ultimately enhance patient outcomes in diabetic care. Focusing on type II diabetes, the study addresses dataset challenges in medical bioinformatics and underscores the global significance of mitigating readmissions to alleviate the burden on healthcare systems and improve long-term patient well-being.

Keywords: Guided Stochastic Gradient Descent; Diabetes Readmission Prediction; Ensemble Learning; Artificial Neural Networks; Healthcare Analytics.

1. Introduction

The integration of Artificial Intelligence (AI) into healthcare practices has revolutionized the industry, offering advanced capabilities for decision-making, medical data analysis, and disease detection [1]. One crucial aspect of healthcare management is the reduction of hospital readmission rates, which has become a significant metric for evaluating hospital performance and minimizing healthcare costs while optimizing patient outcomes [2]. Predicting patient readmissions has therefore emerged as a critical task, leading to the development of numerous AI and Machine Learning (ML) solutions aimed at achieving this goal [3 - 5].

Hospital readmission refers to the scenario where a patient returns to the hospital within a specified timeframe after being discharged, often due to the same underlying medical condition [6]. The causes of readmission can vary, ranging from inadequate initial medical care to suboptimal follow-up care post-discharge [7]. Given the importance of maintaining low readmission rates as an indicator of healthcare service quality, hospital readmission has become a major concern, particularly in light of global health crises such as the COVID-19 pandemic [8 - 10].

Early identification of patients at high risk of readmission enables medical professionals to intervene and implement preventive measures, although this may lead to increased operational costs [11]. Hence, accurate prediction of readmission is essential for cost optimization and resource allocation within healthcare systems.

Machine learning techniques have proven effective in addressing medical bioinformatics challenges, including predicting readmission rates for various diseases such as diabetes [3], [4], [6 - 8], [12]. Notable ML models utilized for this purpose include Artificial Neural Networks (ANN), Support Vector Machine (SVM), and Random Forest (RF) [6], [13 - 15]. While these algorithms demonstrate robustness and good generalization ability, SVM may not be suitable for large datasets and is sensitive to abnormal data distributions or noise, whereas RF exhibits better resilience in such scenarios [16], [17]. ANN, while also sensitive to abnormal data distributions, offers the potential for strategic manipulation of its gradient descent process to compensate for complex data distributions [18], [19]. However, this approach has been primarily applied in logistic regression for smaller datasets and has not been explored in the context of ANN for large datasets.

This article aims to address the underlying challenges associated with gradient descent in ANN by introducing a new variant of the algorithm based on guided gradient descent. This approach aims to improve the prediction accuracy of diabetes readmission, thereby enhancing the quality of hospital services and ultimately benefiting both hospitals and their patients [20]. By leveraging AI-driven predictive analytics, this research seeks to optimize patient care delivery and resource utilization, ultimately leading to improved healthcare outcomes.

2. Literature Survey

Hospital readmission is a critical concern in healthcare systems worldwide, particularly for patients with chronic conditions such as diabetes [21]. In recent years, artificial intelligence (AI) and machine learning (ML) techniques have been increasingly employed to predict hospital readmissions among diabetes patients, aiming to improve patient outcomes and healthcare efficiency [1].

Being able to predict when a diabetic person will need to go back to the hospital is important for both their treatment and the cost of care. Research has progressed from basic descriptive statistics to intricate machine learning techniques; nonetheless, significant obstacles in optimization, scalability, and clinical implementation persist.

Nearly one in five Medicare beneficiaries was re-hospitalized within 30 days of discharge, underscoring the significance of readmission as a healthcare quality metric [2]. This showed how serious the problem was, even though it just used descriptive data. Traditional predictive techniques were inadequate for identifying nonlinear trends in healthcare data, necessitating a transition to machine learning algorithms proficient in managing high-dimensional and partial datasets.

Neural networks integrated with code embeddings were employed, but encountered discrepancies in medical coding between institutions [3]. Random Forest, Naive Bayes, and decision tree ensembles were examined, but class imbalance constrained their reliability [6]. The significance of feature engineering was emphasized; nonetheless, the models did not generalize effectively across healthcare systems [4], [22].

Deep learning models were utilized to forecast hospital readmission from electronic health records, showcasing enhanced prediction performance relative to conventional methods; nonetheless, their methodology depended on extensive, high-quality datasets and significant computational resources [5]. A convolutional neural network with feature engineering was created to predict readmission in diabetic patients. The model worked well on its own dataset; however, it could have overfitted and didn't work on other populations [8].

Robust AUCs (~0.94) were attained through ensemble approaches; yet, the dependence on merely 352 individuals from a singular diabetic clinic constrains generalizability [23]. External validity was enhanced via a multicenter model; nonetheless, the emphasis on older patients with heart failure and the application of LASSO regression limited generalizability to wider groups and neglected nonlinear interactions [24]. Transformer-based models more effectively captured temporal dependencies compared to ensembles or regression approaches; nevertheless, their computational intensity and lack of interpretability diminish their applicability in clinical situations [25]. These recent methods show that progress is being made, but they also show that the trade-off between performance, scalability, and clinical usability is still there.

Older machine learning models like Random Forest and Naive Bayes are still straightforward to understand, but they don't always work well with medical datasets since they don't handle class imbalance well and have trouble finding nonlinear correlations [26], [27]. SVM works better at finding boundaries in smaller datasets, but it has trouble scaling up to high-dimensional healthcare data [17]. Artificial neural networks (ANNs) are better at finding nonlinear associations, but they need a lot of data and are sensitive to differences in coding between institutions [3]. Deep learning enhances accuracy but requires substantial processing resources and is prone to overfitting external data [5], [8]. Transformer designs build on this by modelling patterns over time, but they are hard for doctors to understand and cost a lot of money [25]. The suggested guided ANN, on the other hand, fixes these problems by keeping gradients stable when there is an imbalance, adjusting to datasets that are not all the same, and keeping efficiency and interpretability at a level that is useful for healthcare practice [18].

Using gradient descent on healthcare data shows that current neural network methods are unstable because of skewed distributions, non-random missing values, class imbalance, and dependencies across time [19]. Preprocessing and sampling procedures are advantageous; nonetheless, they do not immediately improve training efficacy. There is a known inverse link between accuracy and scalability: small datasets work well but aren't good for different groups of people, whereas larger datasets make accuracy and reliability worse [28]. Deep models and transformers require a lot of technology resources and work, which makes it impossible to use them in real time in clinical settings [29] [25].

We need optimization algorithms that keep convergence constant, work with data that isn't evenly distributed or is skewed, can be scaled up, and use as few computer resources as possible. These are the things that guided stochastic gradient descent [18]. This technique necessitates modifications to the weight updates based on the attributes of the local data. This improves the uniformity of training and the reliability of performance in a wide range of data situations.

The literature clearly shows how classical statistics have changed over time into increasingly complex models. But sometimes, techniques may give up performance in favour of being easier to understand, scale, or implement. The guided ANN technique fixes these problems by providing great accuracy on huge datasets and an optimization procedure that is both efficient and scalable, making it suitable for application in medicine [18].

3. Methodology

3.1. Proposed work

The proposed work aims to enhance early readmission prediction for diabetic patients by combining Guided Stochastic Gradient Descent with Artificial Neural Networks (ANNs). This approach refines weight updates and addresses data inconsistency challenges, comparing favourably with Support Vector Machine (SVM) and Random Forest (RF) models. Additionally, advanced techniques such as Convolutional Neural Network (CNN) combined with Long Short-Term Memory (LSTM), Stacking Classifier (RF + Multilayer Perceptron (MLP) with LightGBM), and Voting Classifier (RF + AdaBoost) are explored to further improve predictive performance. Furthermore, a Flask framework integrated with SQLite is developed to enhance usability, allowing for user signup/signin and facilitating user input for testing purposes. This comprehensive approach aims to advance the accuracy and usability of early readmission prediction systems for diabetic patients, ultimately improving patient outcomes and healthcare resource utilization.

3.2. System architecture

The proposed system architecture for diabetic patient readmission prediction consists of several key components. Firstly, data input involves acquiring relevant patient data from electronic health records or other sources. Next, data preprocessing and visualization techniques are applied to clean and prepare the data for analysis, ensuring its quality and usability. The dataset is then split into training and testing sets to evaluate the performance of the prediction model.

For training, both binary and multi-class classification algorithms are employed to predict readmission outcomes accurately. Performance evaluation metrics such as accuracy, precision, recall, and F1-score are used to assess the effectiveness of the trained models. Finally, the system predicts readmission probabilities for diabetic patients based on their medical history and current health status, providing valuable insights for healthcare providers to intervene and prevent readmissions effectively.

Overall, this system architecture integrates data processing, model training, and performance evaluation to develop a robust predictive model for diabetic patient readmission prediction, contributing to improved patient care and healthcare resource management.




Fig. 1: Proposed System Architecture for Diabetic Patient Readmission Prediction.

3.3. Dataset collection

The dataset used in this experiment, named "Diabetes 130-US hospitals for years 1999-2008 Data Set," was initially introduced in [30] and later made available in the UCILibrary [44]. It comprises over 50 features representing patient and hospital outcomes, categorized into three classes: patients readmitted within 30 days of discharge, patients readmitted after 30 days of discharge, and patients with no record of readmittance within the 10-year study period.

encounter_id	patient_nbr	race	gender	age	weight	admission_type_id	discharge_disposition_id	admission_source_id	time_in_hospital	...	cto
0	2278392	8222157	Caucasian	Female	[0-10)	?	6	25	1	1	...
1	149190	55629189	Caucasian	Female	[10-20)	?	1	1	7	3	...
2	64410	86047875	AfricanAmerican	Female	[20-30)	?	1	1	7	2	...
3	500364	82442376	Caucasian	Male	[30-40)	?	1	1	7	2	...
4	16680	42519267	Caucasian	Male	[40-50)	?	1	1	7	1	...

Fig. 2: Dataset Distribution for Binary Classification (Readmission Within 30 Days Vs. No Readmission).

With over 100,000 diabetic patient records, two datasets were derived from the original. The first dataset retained the original three-class classification, while the second dataset merged the "Readmitted after 30 days" and "no readmission record" classes into one, resulting in a binary classification of readmission within 30 days or not. Before the experiment, both datasets underwent preprocessing and normalization, ensuring data consistency and reliability. Attributes such as race, gender, age, admission type, time in hospital, laboratory test results, diagnosis, and medication details were included in the pre-processed data.

admission_source_id	time_in_hospital	...	citoglipton	insulin	glyburide-metformin	glipizide-metformin	glimepiride-pioglitazone	metformin-rosiglitazone	metformin-pioglitazone	change	diabetesMed	readmitted
1	1	...	No	No	No	No	No	No	No	No	No	NO
7	3	...	No	Up	No	No	No	No	No	Ch	Yes	>30
7	2	...	No	No	No	No	No	No	No	No	Yes	NO
7	2	...	No	Up	No	No	No	No	No	Ch	Yes	NO
7	1	...	No	Steady	No	No	No	No	No	Ch	Yes	NO

Fig. 3: Dataset Distribution for Multiclass Classification (No Readmission, Readmission ≤ 30 Days, Readmission > 30 Days).

3.4. Data processing

Data processing for the experiment involved several steps. Initially, the dataset was loaded into a pandas dataframe for easy manipulation and analysis. Then, it was converted into a format suitable for training neural networks using the Keras library. To prepare the data for modelling, unwanted columns that were not relevant to the prediction task were dropped from the dataframe. This step helped streamline the dataset and focus on the essential features necessary for predicting diabetic readmissions. Additionally, data normalization and preprocessing techniques were applied to ensure uniformity and enhance the performance of the machine learning models. These techniques included scaling numerical features to a similar range and encoding categorical variables appropriately. Overall, the data processing phase aimed to transform the raw dataset into a structured format suitable for training and evaluating predictive models effectively.

3.5. Visualization

Visualization using seaborn and matplotlib was performed to gain insights into the dataset and understand the distribution of features relevant to diabetic readmissions. Various types of plots, such as histograms, bar plots, and scatter plots, were created to visualize the distribution of numerical and categorical variables. These plots helped in identifying patterns, outliers, and correlations within the data. For example, histograms were used to visualize the distribution of age among diabetic patients, while bar plots depicted the frequency of different admission types. Scatter plots were employed to explore relationships between numerical variables, such as the length of hospital stay and the number of medications prescribed. Additionally, seaborn's built-in functions, such as 'countplot' and 'boxplot', provided further insights into the distribution and relationships within the dataset. Overall, visualization played a crucial role in understanding the data and informing subsequent data preprocessing and modelling steps.

3.6. Label encoding

Label encoding using LabelEncoder was applied to convert categorical variables into numerical representations in the dataset. This process involved assigning a unique integer to each category within a categorical feature. The LabelEncoder from the scikit-learn library was utilized to perform this transformation. Each distinct category within a feature was assigned a specific numerical label, allowing the categorical data to be represented in a format suitable for machine learning algorithms that require numerical inputs. For instance, categorical variables such as admission type, diagnosis, and race were encoded into numerical labels. Label encoding facilitated the incorporation of categorical variables into predictive models, enabling the neural network to effectively learn patterns and relationships from the data. This preprocessing step ensured compatibility between the dataset and the subsequent stages of model training and evaluation.

3.7. Feature selection

Feature selection is a critical step in building predictive models to accurately identify early readmission of diabetic patients. In this approach, relevant features are chosen from the dataset to enhance model performance and reduce computational complexity. Feature selection techniques, such as correlation analysis, recursive feature elimination, and domain knowledge-based selection, are employed to identify the most informative attributes. These selected features provide valuable insights into patient characteristics, hospital outcomes, and medical history, contributing to the predictive power of the neural network model. By focusing on the most relevant features, redundant or irrelevant information is excluded, leading to more efficient model training and improved generalization performance. Additionally, feature selection aids in reducing overfitting and enhances the interpretability of the model by emphasizing the key factors influencing early readmission. Overall, effective feature selection plays a crucial role in optimizing the predictive accuracy and clinical utility of the neural network approach for diabetic patient readmission prediction.

3.8. Training and testing

In the proposed guided neural network approach for predicting early readmission of diabetic patients, the dataset is divided into training and testing sets for model development and evaluation. The training set comprises a majority portion of the data, allowing the neural network model to learn patterns and relationships between input features and target labels. During training, the guided stochastic gradient descent algorithm is utilized to refine weight updates and improve convergence, enhancing the model's predictive performance. After training, the model's performance is evaluated using the testing set, which contains unseen data samples. The model's ability to generalize to new instances is assessed by comparing its predictions with the actual outcomes in the testing set. Performance metrics such as accuracy, precision, recall, and F1-score are computed to quantify the model's effectiveness in predicting early readmission of diabetic patients. Through rigorous training and testing procedures, the guided neural network approach aims to provide reliable and accurate predictions for clinical decision-making.

3.9. Technical representation of guided stochastic gradient descent (G-SGD)

Standard stochastic gradient descent updates model parameters as follows:

$$w_{t+1} = w_t - \eta \nabla L(w_t)$$

Where w_t is the weight vector at iteration t , η the learning rate, and $\nabla L(w_t)$ the loss gradient. While effective, conventional SGD struggles with class imbalance, data skewness, and noisy samples—common issues in medical datasets [26].

The proposed Guided SGD (G-SGD) introduces a guidance term to refine weight updates based on sample-level characteristics and imbalance-aware behaviour:

$$w_{t+1} = w_t - \eta (\nabla L(w_t) + \lambda G(x_i, y_i))$$

Here:

$G(x_i, y_i)$ is a guidance function modulating updates for the sample (x_i, y_i) .

λ is a tuning parameter that scales the guidance term.

The guidance function integrates three key components:

- 1) Class-frequency weighting: Assigns a higher update weight to minority class samples.
 - 2) Error-driven correction: Amplifies updates for previously misclassified samples.
 - 3) Temporal smoothing: Ensures consistency across correlated or sequential patient records to reduce variance.
- These components collectively stabilize training and enhance robustness against skewed healthcare data distributions [14].

Algorithm 1: G-SGD Pseudocode

Input: Training data (X, Y) , learning rate η , guidance weight λ , epochs T

Output: Optimized weights W

```

1: Initialize weights  $W$  randomly
2: for epoch = 1 to  $T$  do
3: Shuffle training data
4: for each mini-batch  $(x_i, y_i)$  do
5:  $g \leftarrow \nabla L(W; x_i, y_i)$ 
6:  $G \leftarrow \text{Guidance}(x_i, y_i, W)$ 
7:  $W \leftarrow W - \eta(g + \lambda G)$ 
8: end for
9: end for
10: return  $W$ 

```

The function $\text{Guidance}(x_i, y_i, W)$ formalizes the sample-level adjustment:

- 1) Increased weight for underrepresented classes,
- 2) Emphasis on misclassified instances,
- 3) Smoothed updates over temporally adjacent patient records.

Novelty and Rationale: G-SGD enables context-aware optimization. Unlike uniform gradient descent, G-SGD prioritizes key samples, improving model convergence under imbalance and noise. Similar strategies (like curriculum learning or cost-sensitive methods) address imbalance but lack this dynamic adaptivity [11]; G-SGD uniquely incorporates local guidance during optimization.

Related Work: The concept of guided selection in SGD appears in other domains. A parallelized version to handle asynchronous delays in distributed SGD was proposed [31]. Our approach extends this by incorporating imbalance- and error-aware guidance, directly targeting common healthcare data challenges.

3.10. Algorithms

SVM: Support Vector Machine (SVM) [6] is a machine learning algorithm used for classification and regression tasks. It separates classes by finding the hyperplane that maximizes the margin between them. In the project, SVM [6] can be applied for predicting early readmission of diabetic patients based on various features. It works well with both binary and multi-class classification problems and can handle high-dimensional data efficiently. SVM's ability to find complex decision boundaries makes it suitable for tasks where the data may not be linearly separable, offering a powerful tool for predictive modelling in healthcare applications like early readmission prediction for diabetic patients.

$$w^T x + b = 0 \quad (1)$$

Random Forest: Random Forest [16] is a machine learning algorithm that creates multiple decision trees during training and outputs the mode of the classes (classification) or the mean prediction (regression) of the individual trees. It works by constructing a multitude of decision trees and combining their predictions to improve accuracy and reduce overfitting. In the project, Random Forest can be used for predicting early readmission of diabetic patients by leveraging its ability to handle large datasets, handle missing values, and deal with nonlinear relationships between features, making it a versatile and powerful tool for predictive modelling in healthcare applications.

Voting Classifier: Voting Classifier is an ensemble learning technique that combines the predictions of multiple base classifiers and predicts the class label by majority voting. It aggregates the predictions from different algorithms, such as Random Forest, Support Vector Machines, and Artificial Neural Networks, to achieve higher accuracy and robustness than individual classifiers. In the project, Voting Classifier can be utilized to enhance the predictive performance of early readmission prediction for diabetic patients by leveraging the diverse strengths of various machine learning algorithms, thus improving the overall accuracy and reliability of the predictive model.

$$\hat{y} = \arg \max_c (\sum_{i=1}^n I(y_i = c)) \quad (3)$$

Stacking Classifier: The Stacking Classifier is an ensemble learning method that combines multiple base classifiers using a meta-classifier to improve prediction accuracy. It involves training several base classifiers on the dataset and then using another classifier, known as a meta-classifier, to combine their predictions. In the project, the Stacking Classifier can be employed to enhance the early readmission prediction for diabetic patients by leveraging the complementary strengths of different machine learning algorithms, such as Random Forest, Support Vector Machines, and Artificial Neural Networks, thus resulting in improved predictive performance compared to individual classifiers.

$$\hat{y} = g(Y_{\text{base}}) = g(f_1(x), f_2(x), \dots, f_m(x)) \quad (4)$$

ANN-Adadelta: ANN-Adadelta is a variant of the Artificial Neural Networks (ANN) training algorithm that uses the Adadelta optimization method. Adadelta is an adaptive learning rate optimization algorithm that aims to address the limitations of traditional stochastic gradient descent methods. It dynamically adjusts learning rates based on past gradients, allowing for smoother convergence and improved performance. In the project, ANN-Adadelta can be utilized as one of the training algorithms for the neural network component of the early readmission prediction system for diabetic patients. By leveraging Adadelta's adaptive learning rate mechanism, the ANN model can effectively learn from the dataset and make accurate predictions for patient readmission outcomes.

ANN-RMSProp: ANN-RMSProp refers to the use of the RMSProp optimization algorithm in training Artificial Neural Networks (ANN). RMSProp, which stands for Root Mean Square Propagation, is an adaptive learning rate optimization algorithm that addresses the

limitations of traditional gradient descent methods. It adapts the learning rates for each parameter based on the magnitudes of recent gradients, allowing for faster convergence and better handling of sparse gradients. In the project, ANN-RMSProp can be employed as a training algorithm for the ANN component of the early readmission prediction system for diabetic patients, enabling efficient learning from the dataset and accurate prediction of patient readmission outcomes.

CNN+LSTM: CNN+LSTM refers to a hybrid neural network architecture that combines Convolutional Neural Networks (CNNs) and Long Short-Term Memory (LSTM) networks. CNNs are effective in capturing spatial features from input data, such as images, while LSTMs excel at capturing temporal dependencies in sequential data. In the project, CNN+LSTM can be utilized for its ability to process sequential data, such as time series data from patient records, while leveraging CNNs to extract relevant spatial features. This hybrid architecture can enhance the model's capability to capture both spatial and temporal patterns, thereby improving the accuracy of early readmission prediction for diabetic patients.

4. Experimental Results

Accuracy: A test's accuracy is its capacity to distinguish patients from healthy cases. To measure test accuracy, calculate the fraction of true positives and true negatives in all evaluated cases. Mathematically, this is:

$$Accuracy = \frac{TP+TN}{TP+FP+TN+FN} \quad (4)$$

Precision: Precision measures the percentage of positive cases or samples accurately classified. Precision is calculated using the formula:

$$Precision = \frac{True\ Positive}{True\ Positive + False\ Positive} \quad (5)$$

Recall: Machine learning recall evaluates a model's capacity to recognise all relevant instances of a class. It shows a model's completeness in capturing instances of a class by comparing accurately predicted positive observations to total positives.

$$Recall = \frac{TP}{TP + FN} \quad (6)$$

F1-Score: Machine learning model accuracy is measured by the F1 score. Combining model precision and recall scores. The accuracy statistic measures how often a model is predicted correctly throughout the dataset.

$$F1\ Score = 2 * \frac{Recall * Precision}{Recall + Precision} * 100 \quad (7)$$


Tables (1 & 2) evaluate the performance metrics—Accuracy, precision, recall, and F1-score—for each algorithm. The voting classifier consistently outperforms all other algorithms. The tables also offer a comparative analysis of the metrics for the other algorithm.

Table 1: Performance Evaluation: Binary Classification


ML Model	Accuracy	Precision	Recall	F1 score
SVM	0.552	0.956	0.552	0.681
Random Forest	0.616	0.669	0.616	0.629
Extension- Voting Classifier	0.988	0.988	0.988	0.988
Extension- Stacking Classifier	0.925	0.936	0.925	0.926
ANN-Adadelata	0.468	1.000	0.468	0.638
ANN-RMSProp	0.597	0.638	0.597	0.607
Extension- CNN+LSTM	0.537	0.874	0.537	0.643

Table 2: Performance Evaluation: Binary Classification

ML Model	Accuracy	Precision	Recall	F1 score
SVM	0.533	0.911	0.533	0.665
Random Forest	0.547	0.931	0.547	0.673
ANN-Adadelata	0.356	1.000	0.356	0.525
ANN-RMSProp	0.356	1.000	0.356	0.525
Extension- Voting Classifier	0.991	0.991	0.991	0.991
Extension- Stacking Classifier	0.776	0.888	0.776	0.817
Extension- CNN+LSTM	0.532	1.000	0.532	0.694



Graph 1: Comparison Graphs: Multi-Class Classification.



Graph. 2: Comparison Graphs: Multi-Class Classification.

Accuracy in blue, precision in red, recall in green, and F1-Score in violet, as shown in Figures 4 and 5. In comparison to the other models, the voting classifier model shows superior performance across all metrics, achieving the highest values. The graphs above visually illustrate these findings.

The ensemble Voting Classifier made the predictions presented. It did better than other models like Guided ANN, CNN+LSTM, and Stacking. Ensemble methods were chosen because they may lower variance and find different decision limits, which makes them work better for both binary and multiclass classification applications. We also looked into the guided stochastic gradient descent optimizer to make ANN training more stable when the data is not evenly distributed or is skewed.

The comprehensive system architecture workflow is illustrated through the implementation details provided in Appendix A (Figures A1-A5), demonstrating the Flask-based implementation from authentication protocols to predictive output generation. The integrated platform encompasses a primary interface (Figure A1), registration module (Figure A2), authentication system (Figure A3), data acquisition component (Figure A4), and results presentation layer (Figure A5). This architectural framework supports clinical decision-making processes by providing healthcare institutions with real-time risk stratification capabilities, potentially enhancing resource allocation efficiency, optimizing follow-up care scheduling, and reducing preventable readmission rates through evidence-based risk assessment protocols.

5. Discussion

5.1. Understanding how well the model works

The Voting Classifier was the most accurate in both tasks. In our tests, it got 0.988 for binary readmission and 0.991 for multiclass categories. Stacking has 0.925 accuracy on the binary task. Guided ANN variants and CNN+LSTM did not do well, with accuracies ranging from 0.356 to 0.597 across tasks. These results demonstrate that the ensemble works better on this dataset.

Why did Voting win? Ensembles mix different decision limits and make variance less noticeable. This gives stability even when there is noise in the features or changes in the sample. Bagging, boosting, stacking, and voting have all shown consistent improvements in clinical prediction in recent reviews [32]. This is mostly due to reducing variance and removing correlation between errors [33], [34].

Performance remained consistent across label structures. Voting got a little better when it went from binary to multiclass (0.988 to 0.991). When class granularity went up, a few solitary models became less accurate. For instance, Random Forest went from 0.616 to 0.547, while directed ANN went from 0.468–0.597 to 0.356. This pattern corresponds with findings indicating that ensembles have less sensitivity to label fragmentation and minority classes [33].

Guided ANN focused on stable training in the presence of skew and missing data. But it wasn't really accurate. Deep baselines can do well on richer temporal EHR, but they need big, well-curated cohorts and careful treatment of imbalances. Recent diabetic readmission studies utilizing ensembles demonstrated robust AUCs, although they exhibited constrained generalizability within single-center datasets [23]. Multicenter regression models enhanced external validity but overlooked nonlinear effects [24]. Transformer models can find patterns that change over time, but they also make computing costs and tuning more difficult [29], [35].

5.2. Limitations

There is still an imbalance in the class. The dataset identifies classes of minority readmission that cause false negatives. Some models have great precision but low recall, which is a sign of skew. This aligns with a 2024 evaluation advocating for integrated sampling and algorithmic enhancements in medical machine learning [26].

The external validity is restricted. The Diabetes 130-US Hospitals data show information from U.S. hospitals from 1999 to 2008. Coding, practice patterns, and care paths vary by region and across time. Several recent studies continue to utilize this cohort and acknowledge the dangers of generalization in the absence of contemporary, multi-site validation [36], [30].

It's hard to make predictions in real time. Streaming vitals, orders, and labs need low-latency ingest, temporal feature storage, and control of backpressure. Event subscriptions in FHIR R5 can help, but end-to-end latency and dependability must be shown to work under clinical load [37].

Learning across sites and privacy. Single-site training doesn't take into account the differences between institutions. Federated learning allows for training across hospitals without the need to transfer raw data [32]. Recent evaluations and frameworks indicate viability while highlighting shortcomings in governance and security [38].

5.3. Trade-offs between interpretability and computation

For clinicians to believe you, you need to give them case-level explanations and global summaries. SHAP and LIME are examples of post hoc tools that can help, but they add steps to the workflow and might not function well with all models. Recent evaluations advocate for clinically verified, stable explanations and human-in-the-loop design [39], [40].

Ensembles make things more accurate, but they can also make it harder to understand why. Predictions should come with audits of feature importance and class-wise errors. Deep and transformer models can do better on tasks that involve time, but they are hard to understand and need a lot of computing power, which makes it harder for clinics to use them [29], [35].

These results show that the Voting Classifier is a good starting point for operations. It is accurate, quick to use, and easy to keep an eye on. Guided ANN is still useful when controlled optimization and portable training are important, but it should be used alongside imbalance-aware learning and explanation tools.

5.4. Clinical implications

The directed ANN only improved predicted accuracy by 1.5%, but even small improvements can make a big difference in clinical practice. In a sample of 100,000 diabetes admissions, this enhancement could result in roughly 1,500 additional patients accurately classified as high risk for readmission. Preventing even a small number of these readmissions is important for both health and money, since each one that could have been avoided can cost between \$9,000 and \$13,000 and may lead to penalties for not paying back under programs like the U.S. Centers for Medicare & Medicaid Services (CMS) Hospital Readmissions Reduction Program [2], [11].

From a clinical workflow standpoint, model predictions might be integrated into discharge planning by creating notifications for physicians and care managers, enabling them to schedule early follow-up consultations, coordinate telehealth monitoring, or commence prescription modifications. The technology may also help with resource allocation by offering real-time risk classification, which ensures that limited nursing and case management resources are given to patients who are most likely to benefit. These ramifications transcend mere economic savings; mitigating unnecessary readmissions alleviates overcrowding in emergency departments, diminishes patient exposure to hospital-acquired infections, and enhances the long-term quality of life for patients with diabetes [20], [41].

6. Conclusion

In conclusion, the project demonstrates the effectiveness of various machine learning algorithms, including SVM, Random Forest, and Guided ANN with Adadelta and RMSProp, in accurately predicting readmissions of diabetic patients. Through meticulous dataset processing, feature selection, and visualization, the quality of the data is enhanced, resulting in optimized model training and improved predictions. The project excels in both multi-class and binary readmission predictions, addressing complexities in healthcare and enhancing patient care insights. Moreover, the implementation of hybrid models such as Voting Classifier, Stacking Classifier, and CNN+LSTM further improves predictive accuracy, with the Voting Classifier standing out as the top performer.

The integration of a Flask framework with a SQLite database architecture ensures practical implementation features for clinical accessibility and research applicability, facilitating secure authentication, automated input preprocessing, and real-time model inference capabilities. Overall, the project offers valuable insights and practical solutions for improving diabetic patient care and healthcare management.

The proposed models can help hospitals lower costs, avoid regulatory penalties, and improve patient outcomes by cutting down on even a tiny number of readmissions that could have been avoided. This matching of technical performance with clinical utility bolsters the argument for the integration of AI-driven predictive systems in healthcare practice. [2], [11], [20]

7. Future Scope

The future potential of the Guided Neural Network Approach for predicting early readmission of diabetic patients involves a comprehensive integration of patient-specific and hospital outcome variables to enhance predictive accuracy. Along with demographic information, admission type, lab findings, and medication history, adding socioeconomic characteristics, comorbidities, and lifestyle-related factors gives a more complete picture of patient risk. Adding more features in this way may help models pick up on subtle interactions that affect the likelihood of readmission, making them more accurate and clinically relevant.

There are still some methodological problems that need to be looked into more closely. It is very important to fix class imbalance since uneven distributions make it hard to generalize. To make training more stable and less biased, you should test advanced resampling methods like SMOTE variants and cost-sensitive learning. This leads to the research question: How might adaptive sampling strategies improve ANN stability in highly skewed diabetic datasets? [26].

Future research directions include enhancing the prototype Flask-SQLite architecture to support iterative development cycles, comprehensive testing protocols, and reproducible demonstration capabilities. The current implementation demonstrates effectiveness for pilot deployment scenarios on limited computational infrastructure. Clinical translation represents a critical advancement pathway through integration with HL7 FHIR APIs and SMART on FHIR launch frameworks, enabling seamless interoperability with existing electronic health record systems. These standards facilitate single sign-on authentication, secure data transmission, and protected patient information access. Additionally, event-driven subscription models support topic-based real-time updates, as specified in FHIR Release 5 and SMART App Launch Guidelines [37], [42].

Scalability is also important. SQLite limits how many people can write at the same time and how quickly they may do so. A transactional RDBMS, a message broker, and audited logging should all be part of a clinical-grade architecture. These sections can handle a lot of reads, writes, and alerts. They also help hospitals meet their needs for security and tracking.

Another important thing is to combine real-time patient data. Hospital readmissions are frequently contingent upon quickly fluctuating variables, including glucose levels and laboratory findings. IoT-enabled glucose monitors and EHR streaming pipelines could make it possible to keep an eye on things all the time for dynamic risk assessment. The pertinent inquiry is: Which architectures optimally facilitate the real-time streaming of diabetic patient data for the prediction of readmission risk? [36].

Socioeconomic and lifestyle factors significantly influence results. Integrating census data, insurance information, and lifestyle surveys might enhance the risk classification process, especially for groups exhibiting significant diversity in access to care. This prompts the inquiry: In what manner can socioeconomic determinants influence clinical characteristics to modify readmission risk profiles? [43].

At the modelling level, progress in transfer learning, federated learning, and explainable AI offers intriguing pathways. Pre-trained transformers like Med-BERT might make it easier to express long-term health data. Federated learning, on the other hand, lets several centers work together without sharing data, which keeps privacy intact. SHAP or LIME explainability tools may help build trust in therapeutic settings and encourage more people to use them. The central inquiry is: Can federated learning frameworks improve generalizability while safeguarding patient data confidentiality? [24], [25].

Finally, for it to be used in the real world, it needs to be integrated into clinical workflows. Integrating models into EHR systems that comply with HL7/FHIR standards and conducting pilot deployments in hospitals will facilitate the assessment of clinician acceptance, cognitive burden, and usability. This leads to the question: How might AI-driven readmission tools be integrated into EHR systems to assist physicians without augmenting cognitive load? [29], [36].

These directions together show a means to move guided neural network methods from proof-of-concept accuracy to scalable, interpretable, and clinically implanted solutions that can help people with diabetes all across the world.

Future work should include pilot testing in real-world clinical settings to measure cost-effectiveness, workflow integration, and provider adoption. Such prospective evaluations are essential for establishing not only technical feasibility but also clinical value and stakeholder acceptance, ensuring that AI-driven readmission prediction tools move from research prototypes into scalable healthcare solutions [36], [37], [42]

References

- [1] R. Kejriwal and Mohana, "Artificial intelligence (AI) in medicine and modern healthcare systems," in *Proc. Int. Conf. Augmented Intell. Sustain. Syst.* (ICAISS), Nov. 2022, pp. 25–31. <https://doi.org/10.1109/ICAISS55157.2022.10010939>.
- [2] S. F. Jencks, M. V. Williams, and E. A. Coleman, "Rehospitalizations among patients in the medicare fee-for-service program," *New England J. Med.*, vol. 360, no. 14, pp. 1418–1428, Apr. 2009, <https://doi.org/10.1056/NEJMs0803563>.
- [3] W. Liu, C. Stansbury, K. Singh, A. M. Ryan, D. Sukul, E. Mahmoudi, A. Waljee, J. Zhu, and B. K. Nallamothu, "Predicting 30-day hospital readmissions using artificial neural networks with medical code embedding," *PLoS ONE*, vol. 15, no. 4, pp. 1–12, Apr. 2020, <https://doi.org/10.1371/journal.pone.0221606>.
- [4] S. Davis, J. Zhang, I. Lee, M. Rezaei, R. Greiner, F. A. McAlister, and R. Padwal, "Effective hospital readmission prediction models using machine-learned features," *BMC Health Services Res.*, vol. 22, no. 1, p. 1415, Nov. 2022, <https://doi.org/10.1186/s12913-022-08748-y>.
- [5] A. Ashfaq, A. Sant'Anna, M. Lingman, and S. Nowaczyk, "Readmission prediction using deep learning on electronic health records," *J. Biomed. Informat.*, vol. 97, Sep. 2019, Art. no. 103256. <https://www.sciencedirect.com/science/article/pii/S1532046419301753>. <https://doi.org/10.1016/j.jbi.2019.103256>.
- [6] Y. Shang, K. Jiang, L. Wang, Z. Zhang, S. Zhou, Y. Liu, J. Dong, and H. Wu, "The 30-days hospital readmission risk in diabetic patients: Predictive modeling with machine learning classifiers," *BMC Med. Informat. Decis. Making*, vol. 21, no. S2, p. 57, Jul. 2021. <https://doi.org/10.1186/s12911-021-01423-y>.
- [7] I. B. Seraphim, V. Ravi, and A. Rajagopal, "Prediction of diabetes readmission using machine learning," *Int. J. Adv. Sci. Technol.*, vol. 29, no. 6, p. 9, 2020.
- [8] A. Hammoudeh, G. Al-Naymat, I. Ghannam, and N. Obied, "Predicting hospital readmission among diabetics using deep learning," *Proc. Comput. Sci.*, vol. 141, pp. 484–489, 2018. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S1877050918317873>. <https://doi.org/10.1016/j.procs.2018.10.138>.
- [9] J. Benbassat and M. Taragin, "Hospital readmissions as a measure of quality of health care: Advantages and limitations," *Arch. International Med.*, vol. 160, no. 8, pp. 1074–1081, Apr. 2000, <https://doi.org/10.1001/archinte.160.8.1074>.
- [10] C. Fischer, H. F. Ringsma, P. J. Marang-van de Mheen, D. S. Kringos, N. S. Klazinga, and E. W. Steyerberg, "Is the readmission rate a valid quality indicator? A review of the evidence," *PLoS ONE*, vol. 9, no. 11, Nov. 2014, Art. no. e112282. [Online]. Available: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4224424/>. <https://doi.org/10.1371/journal.pone.0112282>.
- [11] S. Upadhyay, A. L. Stephenson, and D. G. Smith, "Readmission rates and their impact on hospital financial performance: A study of Washington hospitals," *Inquiry, A J. Med. Care Org., Provision Financing*, vol. 56, Jul. 2019, Art. no. 0046958019860386. [Online]. Available: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6614936/>. <https://doi.org/10.1177/0046958019860386>.
- [12] C. Reid, "Diabetes diagnosis and readmission risks predictive modelling: USA," 2019. [Online]. Available: <https://norma.ncirl.ie/4106/>.
- [13] K. Teo, C. W. Yong, J. H. Chuah, Y. C. Hum, Y. K. Tee, K. Xia, and K. W. Lai, "Current trends in readmission prediction: An overview of approaches," *Arabian J. Sci. Eng.*, pp. 11–18, Aug. 2021, <https://doi.org/10.1007/s13369-021-06040-5>.
- [14] W. Li, M. S. Lipsky, E. S. Hon, W. Su, S. Su, Y. He, R. Holubkov, X. Sheng, and M. Hung, "Predicting all-cause 90-day hospital readmission for dental patients using machine learning methods," *BDJ Open*, vol. 7, no. 1, pp. 1–7, Jan. 2021. [Online]. Available: <https://www.nature.com/articles/s41405-021-00057-6>. <https://doi.org/10.1038/s41405-021-00057-6>.
- [15] P. Wolff, M. Graña, S. A. Ríos, and M. B. Yarza, "Machine learning readmission risk modeling: A pediatric case study," *BioMed Res. Int.*, vol. 2019, Art. no. 8532892, Apr. 2019, <https://doi.org/10.1155/2019/8532892>.
- [16] K. Fawagreh, M. M. Gaber, and E. Elyan, "Random forests: From early developments to recent advancements," *Syst. Sci. Control Eng.*, vol. 2, no. 1, pp. 602–609, Dec. 2014, <https://doi.org/10.1080/21642583.2014.956265>.
- [17] M. Awad and R. Khanna, "Support vector machines for classification," in *Efficient Learning Machines: Theories, Concepts, and Applications for Engineers and System Designers*, M. Awad and R. Khanna, Eds. New York, NY, USA: Apress, pp. 39–66, https://doi.org/10.1007/978-1-4302-5990-9_3.
- [18] A. Sharma, "Guided parallelized stochastic gradient descent for delay compensation," *Appl. Soft Comput.*, vol. 102, Apr. 2021, Art. no. 107084. <https://doi.org/10.1016/j.asoc.2021.107084>.
- [19] Y. S. Abu-Mostafa, M. Magdon-Ismael, and H.-T. Lin, *Learning from Data*. USA: AMLBook, 2012.
- [20] D. J. Rubin, "Hospital readmission of patients with diabetes," *Current Diabetes Rep.*, vol. 15, no. 4, p. 17, Apr. 2015. [Online]. Available: <http://link.springer.com/10.1007/s11892-015-0584-7>. <https://doi.org/10.1007/s11892-015-0584-7>.
- [21] G. Yedluri, "Leveraging machine learning to predict hospital readmissions - A comprehensive analysis," *Gowthami's Portfolio*, Aug. 26, 2020. [Online]. Available: <http://gowthamiyedluri.com/2020-08-26-Predict-Hospital-Readmissions/>.
- [22] A. Das, S. P. Puranam, H. V. R. T. Anumukonda, G. G. Rampam, and Koteswararao Ch., "Predictive modeling of chronic kidney disease: An ensemble ML approach," *International Journal for Multidisciplinary Research (IJFMR)*, vol. 5, no. 6, pp. 1–2, 2023. [Online]. Available: <https://www.ijfmr.com/papers/2023/6/9264.pdf>. <https://doi.org/10.36948/ijfmr.2023.v05i06.9264>.
- [23] V. Mishra, M. R. Tanniru, and J. Sreedharan, "Prediction of 30-day readmission in diabetes management using Machine learning," *Computers in Biology and Medicine*, vol. 195, p. 110616, June 2025, <https://doi.org/10.1016/j.combiomed.2025.110616>.
- [24] Y. He, Y. Yuan, Q. Tan, X. Zhang, Y. Liu, and M. Xiao, "Development and validation of a risk prediction model for 30-day readmission in elderly type 2 diabetes patients complicated with heart failure: a multicenter, retrospective study," *Frontiers in Endocrinology*, vol. 16, p. 1534516, Feb. 2025, <https://doi.org/10.3389/fendo.2025.1534516>.
- [25] Y. Yang, Z. Sun, H. Liu, and K. He, "Transformer-based models for healthcare data analysis: Opportunities and challenges," *Computers*, vol. 14, no. 4, Art. no. 148, 2025, <https://doi.org/10.3390/computers14040148>.
- [26] M. Salmi, D. Atif, D. Oliva, A. Abraham, and S. Ventura, "Handling imbalanced medical datasets: Review of a decade of research," *Artificial Intelligence Review*, Vol. 57, Issue 10, 2024, <https://doi.org/10.1007/s10462-024-10884-2>.
- [27] J.-W. D. Wang, "Naïve Bayes is an interpretable and predictive machine learning algorithm in predicting osteoporotic hip fracture in-hospital mortality compared to other machine learning algorithms," *PLOS Digital Health*, vol. 4, no. 1, Art. no. e0000529, 2025, <https://doi.org/10.1371/journal.pdig.0000529>.
- [28] P. Michailidis, A. Dimitriadou, T. Papadimitriou, and P. Gogas, "Forecasting hospital readmissions with machine learning," *Healthcare*, vol. 10, no. 6, p. 981, 2022. [Online]. Available: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC9222500/>. <https://doi.org/10.3390/healthcare10060981>.

- [29] Jorge García-Mosquera, María Villa-Monedero, Manuel Gil-Martín, Rubén San-Segundo, "Transformer-Based Prediction of Hospital Readmissions for Diabetes Patients," *Electronics*, vol. 14, no. 1, p. 174, Jan. 2025, <https://doi.org/10.3390/electronics14010174>.
- [30] O. G. Emi-Johnson and K. J. Nkrumah, "Predicting 30-day hospital readmission in patients with diabetes using machine learning on electronic health record data," *Cureus*, vol. 17, no. 4, Art. no. e82437, 2025, <https://doi.org/10.7759/cureus.82437>.
- [31] V. Kumar, K. A. Abbas, and C. J. Aster, "Endocrine system," in *Robbins Basic Pathology*, 10th ed. Canada: Elsevier, 2018, pp. 749–796. [Online]. Available: <https://www.elsevier.com/books/robbins-basic-pathology/kumar/978-0-323-35317-5>.
- [32] K. P. C, S. N. R, and IARJSET, "Individualized federated learning for multi-center intensive care unit hospital readiness," *International Advanced Research Journal in Science, Engineering and Technology*, vol. 11, no. 7, 2024. [Online]. Available: <https://iarjset.com/wp-content/uploads/2024/07/IARJSET.2024.11757.pdf>.
- [33] P. Mahajan, R. Phadtare, and A. Singh, "Ensemble learning for disease prediction: A review," *Healthcare*, vol. 11, no. 12, Art. no. 1808, 2023, <https://doi.org/10.3390/healthcare11121808>.
- [34] B. Naderalvojud et al., "Improving ML with ensemble learning on healthcare problems," *Health and Technology*, vol. 14, pp. 1515–1530, 2024, <https://doi.org/10.1007/s12553-024-00835-w>.
- [35] A. Mohamed, R. AlAleeli, and K. Shaalan, "Advancing predictive healthcare: A systematic review of transformer models in EHRs," *Computers*, vol. 14, no. 4, Art. no. 148, 2025, <https://doi.org/10.3390/computers14040148>.
- [36] T.-L. Hu, C.-M. Chao, C.-C. Wu, T.-N. Chien, and C. Li, "Machine Learning-Based Predictions of Mortality and Readmission in Type 2 Diabetes Patients in the ICU," *Applied Sciences*, vol. 14, no. 18, p. 8443, Sept. 2024, <https://doi.org/10.3390/app14188443>.
- [37] HL7, "HL7 publishes Release 5 of FHIR," *Healthcare Innovation*, 2023. [Online]. Available: <https://www.hcinnovationgroup.com/interoperability-hie/fast-healthcare-interoperability-resources-fhir/news/53057286/hl7-publishes-release-5-of-fhir-standard>.
- [38] Z. L. Teo, L. Jin, N. Liu, S. Li, D. Miao, X. Zhang, W. Y. Ng, T. F. Tan, D. M. Lee, K. J. Chua, J. Heng, Y. Liu, R. S. M. Goh, and D. S. W. Ting, "Federated machine learning in healthcare: A systematic review on clinical applications and technical architecture," *Cell Reports Medicine*, vol. 5, no. 3, Art. no. 101419, 2024, <https://doi.org/10.1016/j.xcrm.2024.101419>.
- [39] Z. Sadeghi, R. Alizadehsani, M. A. Cifci, M. Khanam, S. Acharya, Y. Tan, S. U. Rehman, A. Reddy Nallabothula, M. Islam, A. Bhuiyan, V. Jahmunah, R. S. Tan, and U. R. Acharya, "A review of explainable AI in healthcare," *Computers and Electrical Engineering*, vol. 118, Art. no. 109370, 2024, <https://doi.org/10.1016/j.compeleceng.2024.109370>.
- [40] R. Alkhanbouli, H. M. A. Almadhaani, F. Alhosani, and M. C. E. Simsekler, "The role of explainable AI in disease prediction," *BMC Medical Informatics and Decision Making*, vol. 25, no. 1, p. 110, 2025, <https://doi.org/10.1186/s12911-025-02944-6>.
- [41] E. J. Comino, M. F. Harris, M. F. Islam, D. T. Tran, B. Jalaludin, L. Jorm, J. Flack, and M. Haas, "Impact of diabetes on hospital admission and length of stay among a general population aged 45 year or more: A record linkage study," *BMC Health Services Res.*, vol. 15, no. 1, p. 12, Jan. 2015, <https://doi.org/10.1186/s12913-014-0666-2>.
- [42] SMART on FHIR, "SMART App Launch implementation guide," 2023. [Online]. Available: <https://hl7.org/fhir/smart-app-launch/1.0.0/>.
- [43] F. Mohsen et al., "A scoping review of artificial intelligence-based methods for diabetes risk prediction," *NPJ Digital Medicine*, vol. 6, p. 197, 2023, <https://doi.org/10.1038/s41746-023-00933-5>.
- [44] S. Tayal, "Diabetic Patients Readmission Prediction," *Kaggle*, 2020. [Online]. Available: <https://www.kaggle.com/datasets/saurabhhtayal/diabetic-patients-readmission-prediction>. Accessed: Sep. 13, 2025.

Appendices

Appendix a: system implementation interface details.




Fig. A1: Home Page Interface of the Flask-Based Prediction System.

The main interface serves as the primary entry point for the web-based prediction system, implementing a centralized navigation architecture. The design incorporates intuitive navigation elements that direct healthcare professionals to authentication modules and data entry components. The interface architecture prioritizes clinical workflow integration, featuring streamlined pathways for credential verification and patient data management processes.

Sign up

Your Username
 Your Name
 Your Email
 Your Phone Number
 Password

[Register](#)



Already have an account? [Sign in](#)

Fig. A2: User Registration Interface for System Access Control.

The registration module implements a secure account creation mechanism that captures essential authentication credentials including username, email address, and encrypted password data. This component ensures controlled access to the predictive system through identity verification protocols. The implementation addresses healthcare data security requirements by restricting system access to authorized clinical personnel and research staff, maintaining compliance with patient data protection standards.




Fig. A3: User Authentication Interface with Secure Login Functionality.

The authentication interface facilitates secure credential verification through integrated session management protocols. The system employs Flask's session handling capabilities in conjunction with SQLite database storage to perform rapid credential validation. Access control mechanisms implemented at this layer demonstrate adherence to healthcare data privacy regulations and established security frameworks for medical information systems.

Fig. A4: Data Input Interface for Patient Clinical Information Entry.

The data acquisition interface facilitates the systematic collection of clinical parameters essential for predictive modelling. The interface incorporates structured input fields for discharge disposition, previous inpatient encounter frequency, and primary diagnostic classifications (Diag_1, Diag_2, Diag_3). The form architecture ensures data consistency with the trained model's feature requirements through standardized input validation. The system performs automated data normalization to reduce input variability and prepare clinical variables for preprocessing and categorical encoding. This component reflects real-world clinical workflows where healthcare professionals integrate patient records directly into predictive analytics platforms.




Fig. A5: Prediction Results Interface Displaying Readmission Risk Assessment.

The results visualization interface presents the predictive model outputs through a clinically interpretable format that translates machine learning algorithms into actionable medical insights. The system generates probabilistic assessments across three distinct readmission categories: no readmission risk, readmission probability within 30 days, and extended readmission risk beyond 30 days related to diabetic complications. These predictive outputs demonstrate the clinical utility of the proposed methodology in supporting healthcare teams with early identification of high-risk patient populations.