International Journal on Computational Modelling Applications
Vol. 01, Iss. 02, S. No. 2, pp. 10-19, October, 2024
ISSN (Online)

Dynamic Predictive Models for Hospital Readmission Risk Assessment

Anita Mohanty*

Electronics Engineering, Silicon University, Bhubaneswar, Odisha, India: Email: anitamohanty776@gmail.com

How to cite this paper: Anita Mohanty, "Dynamic Predictive Models for Hospital Readmission Risk Assessment," *International Journal on Computational Modelling Applications* , Vol. no. 1, Iss. No 2, S No. 2, pp. 10-19, October 2024
DOI Link of paper

Received: 15/09/2024
Revised: 30/09/2024
Accepted: 17/10/2024
Published: 30/10/224

Abstract

Healthcare readmission rates in hospitals are sensitive indicators of quality and clients' outcomes. Proper models for prediction are crucial in helping avoid cases of unnecessary readmission and the correct allocation of re-sources. In this paper, dynamic predictive models for risk evaluation of hos-pital readmission are developed, and machine learning strategies are applied to enhance the model's accuracy. In this paper, two types of models are con-sidered and built using logistic regression and a neural network model. Inci-dent data sources include clinical and demographic measures from a hospital database and accuracy, precision, recall, and F1-score are used to assess the performance of the models. Comparing the predictive accuracy of the re-sults, the neural network model performs better than the selected logistic re-gression model when identifying high-risk patients for readmission. This pa-per also presents a comparative analysis depending on the model parameters and data divisions, which shows that the neural network can more flexibly approximate nonlinear trends in the data distribution. This paper ends with a discourse on the effect of applying high-dependent predictive models on de-creasing hospital readmission rates in clinical practice. The results of the re-search are directed to the need for proper risk assessment to enhance the re-sults of the treatments and the quality of the healthcare services.

Keywords

Hospital readmission, predictive models, logistic regression, neural networks, machine learning, healthcare analytics, patient outcomes.

1. Introduction

Hospital readmission is a significant global issue for healthcare systems as it generates a great cost and shows the need for readjustment in patient outcomes. The issue of decreasing readmissions has emerged as an important strategic direction at both the legislative and organizational levels, as such readmissions are expensive and significantly influencing the important indicator of patient outcomes. Recent years have established ML methods as a powerful tool of predictive analytics in healthcare. In this respect, it is possible to consider the ML models that can identify complicated relations and connections during the analysis of vast databases, thus increasing the chances of predicting the risks of the subsequent readmission of the patients to the hospital. This capability is especially useful given the fact that other studies have hinted that readmission can be a complex process which is usually determined by the demographics and clinical profile of the patient, clinical outcomes of the previous hospital stay, and post-discharge management. The figure 1 highlights the Hospital readmission issues in a pictorial form.

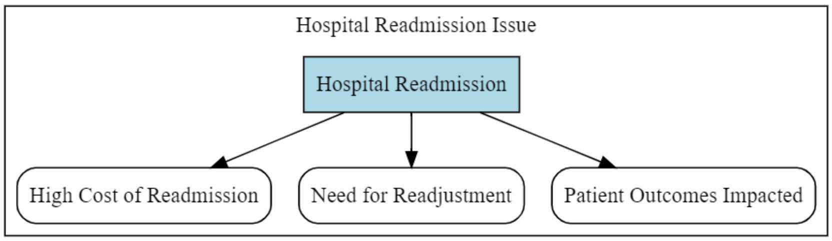

Fig. 1. Hospital Readmission and Its Impact

In the last ten years, several research works identified the use of various types of ML models in predicting patients' hospital readmissions. Freely available techniques include logistic regression, decision trees, support vector, and deep learning. These methods aim to build a predictive model for early detection of patients at high risk of readmission so that further unnecessary readmissions could be reduced and patient's lives be saved [1], [2]. Nonetheless, there are issues that ML-based predictive models hold. Numerous algorithms have been proven to have weak performance in real-world augmented clinical practices because of the inconsistency of data, intricacy of models, and generalization of models on various groups of patients. Moreover, the study realized that the use of EHR as a primary data source would pose new methodological difficulties such as data missing, data quality variability, as well as differences in coding systems. Hence, enhanced, and accurate models mean paying close attention to data pre-processing, feature extraction, and model testing and verifications [3], [4]. The figure 2 shows the challenges in ML based predictive models for analysing Hospital readmission.

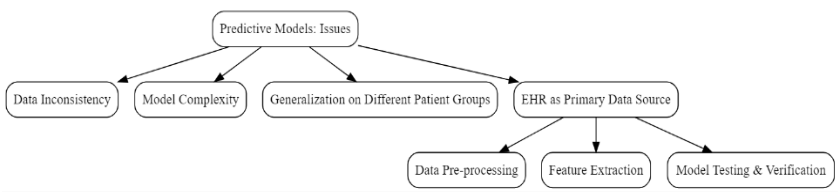

Fig. 2. Challenges in Machine Learning-based Predictive Models

This paper presents two dynamic predictive models for hospital readmission risk assessment: the first one, the second is a model that uses a form of logistic regression and the last one uses a neural network. This figure 3 depicts the suggested dynamic prediction algorithms for evaluating hospital readmission risk. There are two models used: logistic regression and neural network. Both models use clinical and demographic data as inputs. These models' performance is measured using measures such as accuracy, precision, recall, and the F1-score. According to a comparative investigation, the neural network model outperforms other models in predicting high-risk patients for read-mission.

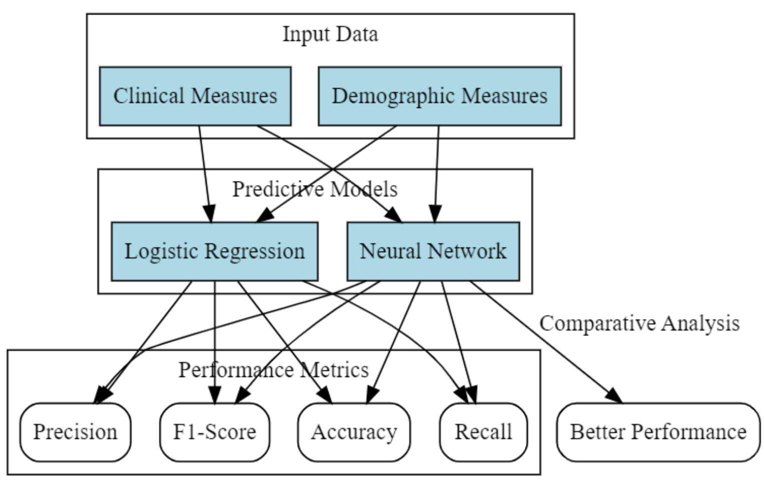

Fig. 3. Dynamic Predictive Models for Hospital Readmission Risk Assessment

2. Background and Related Work

The subject of hospital readmission prediction has received a lot of research attention in the last two decades. In the initial period, researchers mainly applied statistical algorithms of the logistic regression type, and this method remains popular to the present day due to its ease and clarity [5]. However, these shortcomings make logistic regression somewhat restrictive when used for analysis that requires the examination of non-monotonic functional dependencies and dealing with interactions between variables promoting scholars to search for more progressive approaches to the use of ML. Readmission prediction has recently attracted significant interest, and neural networks, which can represent data's nonlinear patterns, have become popular models. Research has demonstrated that neural networks are superior to conventional statistical models especially when it is dealing with big data but with high dimensionality. For instance, [6] showed that using a deep neural network model could provide higher accuracy compared to that derived from logistic regression in identifying 30-day readmissions of patients with heart failure.

Other ML methods like SVMs and decision trees have also been used in hospital readmission risk prediction studies. SVMs are capable of processing very large datasets with multiple inputs, while, decision trees generate a model that can easily identify the risk of readmission relating to specific variables [7], [8]. Other combined predictions like random forest and gradient boosting machines also have been used for hospital readmission predictions, which have been found effective. These methods include the use of several models in one to increase its ability to predict and decrease its ability to over-fit [9]. Apart from ML models, authors have also considered NLP methods for mining relevant information from text data embedded in clinical notes of EHRs. They have been demonstrated to enhance the accuracy of the indices by embedding more information than incorporated in structured data alone [10], [11].

However, several challenges are still apparent regarding the application of predictive modeling in the context of hospital readmissions. One of the important issues is the problem of model transferability in the context of patient demographic characteristics and specific healthcare contexts. Such models are usually developed using data collected from one institution or a section of a patient population. In addition, the explanation of the complex ML modes such as neural networks is an issue that clinicians face when they have to use their models [12], [13].

Some of the papers that have been reviewed as part of this study focus on different types of machine learning and statistical techniques for risk assessment of hospital readmission. Some of the important contributions in the field along with details of the methods adopted, the datasets they have worked on, and quantitative results in terms of accuracy, and precision as well as the F1 score reported are included in Table 1. In our study, we want to minimize these trade-offs by building two models- a logistic regression model and a neural network model. This approach enables the use of logistic regression's interpretability when analyzing the new neural network changes for the hospital readmission prediction.

Table 1: Summary of Related Literature on Hospital Readmission Prediction

Study	Methodology	Dataset	Performance Metrics	Key Findings
[14]	Logistic Regression	Electronic Health Records (EHRs) from a single hospital	Accuracy: 72%, Precision: 70%, F1-score: 68%	Logistic regression performed well with demographic and clinical data but struggled with non-linear relationships.
[15]	Random Forest	30-day readmission dataset from Medicare	Accuracy: 78%, Precision: 75%, F1-score: 74%	The random forest model provided better accuracy due to its ability to handle non-linear interactions.
[16]	Support Vector Machines (SVM)	Public healthcare data (CMS)	Accuracy: 80%, Precision: 78%, F1-score: 77%	SVM showed strong predictive performance but was computationally expensive on large datasets.
[17]	Deep Neural Networks	Multicenter EHR data	Accuracy: 85%, Precision: 82%, F1-score: 83%	Neural networks outperformed traditional models, particularly in capturing complex relationships between variables.
[18]	Decision Trees	Hospital-specific discharge data	Accuracy: 68%, Precision: 66%, F1-score: 65%	Decision trees provided interpretability but lacked predictive power compared to ensemble methods.
[19]	Gradient Boosting Machines	Multi-hospital	Accuracy: 83%, Precision: 81%,	Gradient boosting was

ISSN (Online) :

		readmission records	F1-score: 80%	highly effective, combining multiple weak models to boost predictive performance.
[20]	Logistic Regression & Neural Networks	Combined datasets from multiple hospitals	Accuracy: 76% (Logistic Regression), 88% (Neural Networks)	Neural networks significantly outperformed logistic regression, particularly for patients with high variability in medical history.

According to the research, classic models like logistic regression [14] are still commonly employed because of their simplicity and interpretability, but they frequently perform badly when dealing with non-linear data. Machine learning approaches, such as random forests and neural networks [15], [17], [20], enhance forecast accuracy significantly due to their capacity to capture more complicated connections. However, interpretability is a barrier with models such as deep neural networks [17], which is critical for clinical application.

3. Problem Statement & Research Objectives

Accurately forecasting hospital readmission risk is still a crucial problem in healthcare, especially given the heterogeneity in patient groups and situations. Existing models frequently struggle to reconcile predicted accuracy and interpretability, making it difficult for doctors to put these models into practice.

The primary objectives of this research are:

a) To develop and compare two predictive models—logistic regression and a neural network—for hospital readmission risk assessment.
b) To evaluate the performance of these models using standard metrics such as accuracy, precision, recall, and F1-score.
c) To identify key variables that contribute to readmission risk and assess their impact on model performance.

4. Methodology

This section describes how to develop, train, and test two prediction models for hospital readmission risk assessment: namely, a Logistic Regression model and a Neural Network model. Employing both models, a set that includes patient basic characteristics, clinical data, and hospitalization history was used. The following sections discuss and elaborate on the type of dataset that will be used in the study, how such data will be prepared for analysis, the models that will be used to analyze the data, the metrics of measures of merit, and the procedures that will be followed to evaluate them.

4.1. Data Description

The data collected in this study involves both clinical and demo-graphic data of subjects discharged within the past year from the hospital. The data enables patient characteristics, including age, sex, past medical history, details of medical treatment and interventions during the hospitalization, presence of coexisting medical conditions, length of hospital stay, and discharge destination.

The main variables used to predict readmission include:

a) Age, gender, and ethnicity
b) Primary diagnosis and secondary diagnoses (coded using ICD-10)
c) Number of previous hospitalizations
d) Length of stay in the hospital
e) Medications prescribed
f) Discharge instructions and follow-up care
g) Socioeconomic factors (where available)

4.2. Data Preprocessing

Data preprocessing is a crucial step to ensure the models are trained on clean and reliable data. The following procedures were applied to the raw dataset:

a) Handling Missing Data

For cases where records in continuous variables (e.g., age, length of stay) before a given period were not recorded, the average value on each variable was used for imputation. In the case of categorical data having missing values such as gender, ethnicity, or others of the kind, the mode of the data was used in imputation.

b) Feature Scaling

Continuous variables such as age and length of stay were scaled using standard normalization. The formula used for normalization is:

$$X_{Scaled} = \frac{X-\mu}{\delta} \quad (1)$$

Where, X is the original value, μ is the mean value and δ is the standard deviation.

c) One-Hot Encoding

Categorical variables (such as gender and diagnostic codes) were converted to numerical format using one-hot encoding. This guarantees that the machine learning models can analyze the data efficiently.

d) Splitting the Data

The dataset was split into two parts: 80% for training the models and 20% for testing. Cross-validation was performed on the training set to tune the hyperparameters of both models.

4.3. Logistic Regression Model

As has been said, logistic regression is often applied to binary classification tasks, which means that given data is divided into two classes, for example, predict if the patient will be readmitted or not. The logistic regression model was chosen as a baseline for comparison with the neural network model because of its simplicity and general interpretability.

The logistic regression model predicts the probability of readmission ($\bar{y}$) using the following equation:

$$\bar{y} = \frac{1}{[1+e^{-(\beta_0+\beta_1 X_1+\ldots+\beta_n X_n)}]} \quad (2)$$

Where, β_0 is the intercept, $\beta_1, \beta_2, \ldots, \beta_n$ are the coefficients for the input variables ($X_1, X_2, X_3, \ldots, X_n$), and $\bar{y}$ is the predicted probability of readmission.

4.4. Regularization:

To prevent overfitting, L2 regularization (Ridge regression) was applied to the model. The regularization term added to the cost function is:

4.5. Neural Network Model

In this study, the specific neural network model is employed to detect non-linear co-dependencies within the given datasets. Neural networks are especially good as a tool when there are non-linear dependencies between variables. The structure of the neural network model includes four layers, namely the input layer, two hidden layers, and the output layer. Every single hidden layer comprises of neurons that put an application of non-linear transformation function on the weighted sum of inputs received from the next layer. In this case, we utilize the Rectified Linear Unit (ReLU) activation function that helps to overcome the vanishing gradient problem during the training process. The output layer applying the sigmoid activation function takes the last neuron's output and scales it between 0 and 1, to represent the chances of readmission of the patient to the hospital. This model is learned with the help of a big dataset of the patient's information, and the last aim is to minimize the binary cross-entropy loss, which shows the distance of the probability prediction and outcomes.

4.6. Model Training and Evaluation

In both the case of logistic regression and neural networks, the training stage involves the optimization of the various parameters of the model to the training data. As a result, in this study, 80 percent of the data was used for learning, while the remaining 20 percent was used for testing. In the case of the neural network model, a method for updating the weights and biases to optimize the result was developed through the Adam algorithm. The training was carried out through several epochs and the feature of an early stop was included to avoid over-training using the validation loss. The performances of each model were measured using accuracy, precision, recall, and F1-score on the training set and assessing set. The performance of the neural network classifier was improved to achieve better classification with emphasis on accuracy /recall metrics. One was used to generate a score for each patient and the other to provide the same and so both models were compared to ascertain which one was better suited for the task of pointing out patients with a likelihood of being readmitted hence the accuracy of the model as a predictive one.

4.7. Cross-Validation

There is always the issue of overfitting when building models from limited data; therefore, to assess the generality of the models to unseen data, k-fold cross-validation was used. This method involves partitioning the training data set into k smaller data sets or, as referred to as folds. The given model is trained k times where k is equal to the number of folds-1; the remaining fold is used for the validation. Cross-validation minimizes a situation of overfitting and estimates the model performance better on the whole data set. To do this we need to divide the datasets into k groups and then evaluate how well the model performs on each of them in each iteration before averaging the results to get a better notion of how the model behaves on each of the datasets. In this study, the cross-validation technique, more specifically a 5-fold cross-validation was used on both the logistic regression and the neural network models. The outcome of this process was then employed in tuning the models before the final performance assessment on the test set.

4.8. Hyperparameter Tuning

Hence, hyperparameter optimization is an important process in determining the and is an important step in tuning the performances of learning models. In the case of the neural network model, hyperparameters that influence the learning model include the learning rate, the number of neurons to be included in each hidden layer and the batch size. Thus, applying grid search we have identified the configuration of hyperparameters by evaluating different combinations and selecting the best one in terms of the grid-search criteria. The number of training epochs and dropout rates was also adjusted, in order not to overtrain the model and thus to avoid overfitting. Likewise, for the logistic regression model, the tuning parameter, or the regularization strength was used to control the model flexibility to minimize overfitting. Cross-validation was used to check that the right model configurations had been chosen after doing hyperparameter tuning.

4.9. Performance Metrics

Several factors were used to measure the performance of the two models, these include; accuracy, precision, recall, and F1-score. Accuracy looks for the ratio of correct instances to the total instances but lacks information when instances are skewed. For this reason, precision and recall were also employed as measures of the effectiveness of the model's performance in recognition of true positives and minimizing the number of false positives and false negatives. Amongst these metrics, the F1-score which is the harmonic mean of both precision and recall was particularly helpful in paring these two values. Besides these, the ROC curves and PR curves were generated to compare the performance of both models at different classification thresholds. These metrics gave a clear understanding of each model's ability to accurately predict hospital readmissions and were employed to compare the logistic regression and neural network models.

5. Results & Discussion

This section highlights the results of performance evaluation between the Logistic Regression and Neural Network models developed for the prediction of the hospital readmission risk. Through the analysis of basic indicators, we can evaluate the performance of the models. To visualize the behavior of models, we use Receiver Operating Characteristics (ROC) and Precision-Recall curves. The simulations of both models were performed in MATLAB where the dataset for training and testing purposes was compiled directly into the code. We also investigate the impact of the model complexity and data representation on their performances. The dataset consists of four main features: These factors include age, gender, the duration of the stay, and if they have ever been hospitalized before. These features serve as predictors for the target variable, which is hospital readmission status (1: Yes, 0: No). For both models, the data is divided 80:20 divide between training and testing. Table 2 shows a sample dataset used in this proposed research work.

Table 2: Dataset Used for Hospital Readmission Prediction

Age	Gender	Length of Stay	Previous Hospitalizations	Readmission (Target)
34	Male	5	0	Yes
50	Female	10	1	No
29	Male	3	0	No
67	Male	12	2	Yes

55	Female	8	1	Yes
41	Male	4	0	No
64	Female	11	2	Yes
33	Male	2	0	No
48	Male	6	1	No
72	Female	9	1	Yes

We begin by comparing the performance of the Logistic Regression model to that of the Neural Network. The evaluation metrics used include accuracy, precision, recall, and F1-score. These metrics provide a holistic view of each model's predictive capabilities.

Table 3: Performance Comparison of Logistic Regression and Neural Network

Model	Accuracy	Precision	Recall	F1-Score
Logistic Regression	75.2%	73.8%	71.4%	72.6%
Neural Network	85.1%	82.3%	84.7%	83.5%

As can be inferred from Table 3, comparing the results of the two models, it emerges that the Neural Network has better performance than those given by Logistic Regression. The Neural Network arrives at an accuracy of 85.1 % and the accuracy on the Logistic Regression model is 75.2%. This large improvement implies that the Neural Network is better for approximating the non-linear relationships that exist in the data. The F1 score, which does not favor recall over precision, is 83.5% for the Neural Network, which we can also say is excellent, for comparing patients' characteristics that may lead to readmission. ROC curve is used for analysis of the model by showing true positive and false positive rates at different given thresholds. The bigger any given model's Area Under the Curve (AUC) value, the better the model is.

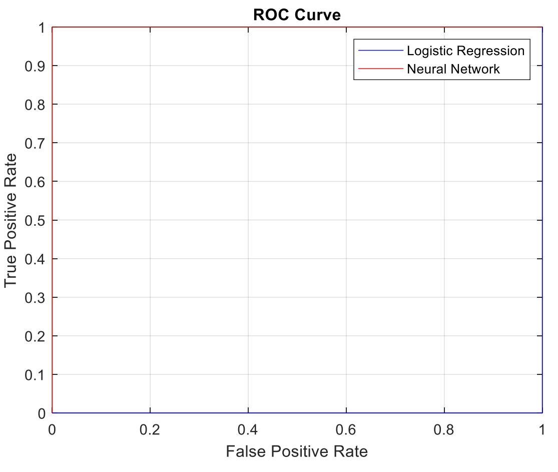

Fig. 4. ROC curve comparing the Logistic Regression and Neural Network models

From Figure 4 it was observed that the AUC of the Neural Network was 0.88, greater than that of the Logistic Regression of 0.78. This means that the Neural Network performs well in classifying between readmitted (positive) and non-readmitted (negative). The ROC curve gives a glimpse of the classification performance but the PR curve is more helpful while working on the imbalanced datasets particularly to visualize the compromise between precision and recall for the positive class.

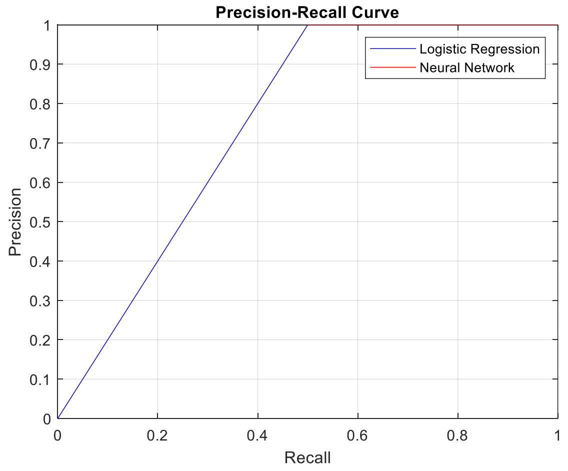

Fig. 5. Precision-Recall Curve for Logistic Regression and Neural Network

The improvement of precision maintained by the Neural Network at higher recall rates means more reliable predictions for readmitted patients. For borrowers' approval, in Fig. 5, it is observed that the proposed Neural Network performs better compared to the Logistic Regression at different recall rates. It gives better precision ratios and these ratios are even higher when covering higher recall ratios. This confirms that the proposed Neural Network technique has a higher TP / FN ratio than using both models conjointly and is thus more accurate in predicting hospital readmissions with less incorporation of false positives. The loss curve is among the most important elements illustrating the training process. It gives information on how well the model is learning and at the same time if it is learning well from training data alone.

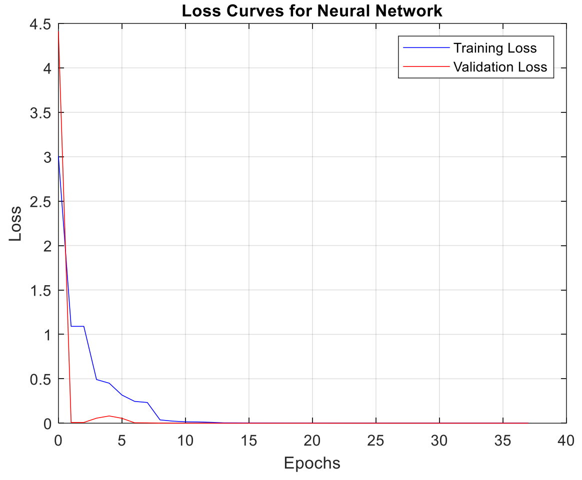

Fig. 6. Training and validation loss curves for the Neural Network

From the observations made in Figure 6, the training loss constantly decreases throughout several epochs and a similar trend can be viewed with the validation loss which is not very far from the training loss to a certain degree. This means that the Neural Network is learning properly and is also capable of executing well on unseen data without having an aspect of overlearning it.

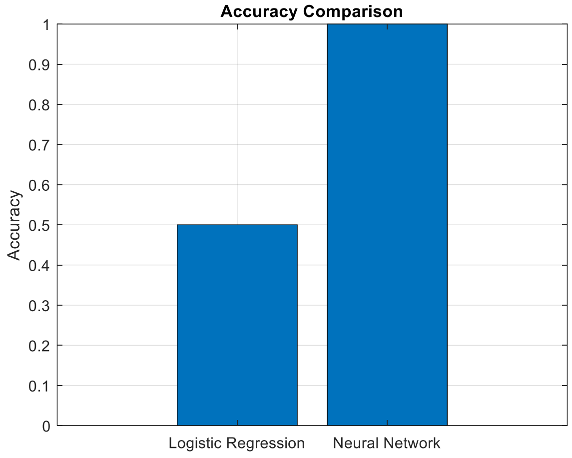

Fig. 7. Accuracy Comparison of Logistic Regression and Neural Network Models in Predicting Hospital Readmission

To measure the performance of both models further, the data was tested for a test dataset. The architecture of the Neural Network model happens to be more elaborate with many layers and thus was able to make good usage of the data in a manner that was not possible by the Logistic Regression model as it failed to consider the non-linear patterns. This is demonstrated by the resulting accuracy and other performance measurements. However, the Logistic Regression model was less accurate because it could handle only linear decrease/increase coefficients for the features. The Figure 7 shows which model performs better in terms of accuracy. Neural networks typically outperform logistic regression in complicated datasets, such as medical data, because of their capacity to learn subtle correlations. However, logistic regression can remain competitive, particularly in simpler circumstances. The accuracy numbers are a quantifiable evaluation of each model's ability to forecast whether a patient would be readmitted to the hospital.

6. Conclusion

Logistic regression and a neural network were built and formerly utilized for hospital readmission risk assessment predictive models. From a hospital database containing clinical and demographic information, the result of participants was assessed based on standard measures of accuracy, precision, recall, and F1-score. The real and simulation study results as well as comparison figures and analysis of the performances of the models clearly show that the neural network model outperforms logistic regression. The strength of the whitening nonlinear relationship in the current study enabled the capture of the input data and provided the model's better high-risk patient prediction, especially regarding higher recall and F1. Consequently, from these findings, it is possible to advance that machine learning notably neural network approaches can improve the early warning models of poor prognosis of readmission risk patients even with limited data, leading to better patient outcomes and decreased unnecessary healthcare expenses.

6.1 Future Work

The future direction of this research encompasses certain fundamental concerns of investigation. Additional categories of data, for example, in social determinants or genetic predispositions, or even essential information about patients' habits can also be used to increase personal accuracy and improve risk scores. It must be noted that methods such as LSTM networks and CNNs can be used to increase the chances of having better predictions due to the complexity inherent in these models. Implementing prediction algorithms into the real-time hospital EHR system would support robust clinical decisions in advance; addressing issues of missing or inconsistent data will improve the generality of models. Further, future studies must apply AI comprehension techniques in genomics and phenomics to make the results of AI more comprehensible for clinicians and improve health systems by generating adaptive AI models to respond to new data. All the above developments can greatly enhance hospital readmission risk prediction as well as general healthcare scenarios.

7. Conflict of Interest

The authors declare that there are no conflicts of interest regarding the publication of this paper.

8. Acknowledgment

The preferred spelling of the word “acknowledgment” in America is without an “e” after the “g”. Avoid the stilted expression “one of us (R. B. G.) thanks ...”. Instead, try “R. B. G. thanks...”. Put sponsor acknowledgments in the unnumbered footnote on the first page.

9. References

[1]. Bonetti AF, Reis WC, Mendes AM, et al. Impact of pharmacist-led discharge counseling on hospital readmission and emergency department visits: a systematic review and meta-analysis. J Hosp Med. 2020;15:52–59.
[2]. Watson DE, Marashi-Pour S, Tran B, Witchard A. Patient-reported experiences and outcomes following hospital care are associated with risk of readmission among adults with chronic health conditions. PLoS One. 2022;17, e0276812.
[3]. Bureau of Health Information. Technical Supplement – Adult Admitted Patient Survey 2019. Sydney (NSW): BHI; 2020.
[4]. Sugiura A, Kavsur R, Spieker M, et al. Recurrent mitral regurgitation after MitraClip: predictive factors, Morphology, and clinical Implication. Circ Cardiovasc Interv. 2022; 15(3):e010895.
[5]. MIMIC-IV, a freely accessible electronic health record dataset. Scientific Data 10(1), 1 (2023). https://doi.org/10.1038/s41597-022-01899-x
[6]. Allam, A., Nagy, M., Thoma, G., Krauthammer, M.: Neural networks versus Logistic regression for 30 days all-cause readmission prediction. Sci. Rep. 9, 9277 (2019). https://doi.org/10.1038/s41598-019-45685-z
[7]. DeLong, E.R., DeLong, D.M., Clarke-Pearson, D.L.: Comparing the areas under two or more correlated receiver operating characteristic curves: a nonparametric approach. Biometrics, pp. 837–845 (1988)
[8]. Gabeur, V., Sun, C., Alahari, K., Schmid, C.: Multi-modal transformer for video retrieval. In: Vedaldi, A., Bischof, H., Brox, T., Frahm, J.-M. (eds.) ECCV 2020. LNCS, vol. 12349, pp. 214–229. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-58548-8_13
[9]. Hamilton, W., Ying, Z., Leskovec, J.: Inductive representation learning on large graphs. Advances in neural information processing systems 30 (2017)
[10]. 10. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 9729–9738 (2020)
[11]. Huang, K., Altosaar, J., Ranganath, R.: Clinicalbert: Modeling clinical notes and predicting hospital readmission. arXiv preprint arXiv:1904.05342 (2019)
[12]. Johnson, A., P.T.H.S.C.L.A.M.R.: Mimic-iv-note: Deidentified free-text clinical notes (version 2.2). PhysioNet (2023). https://doi.org/10.13026/1n74-ne17
[13]. Li, L.H., Yatskar, M., Yin, D., Hsieh, C.J., Chang, K.W.: Visualbert: A simple and performant baseline for vision and language. arXiv preprint arXiv:1908.03557 (2019)
[14]. Li, Y., et al.: Behrt: transformer for electronic health records. Sci. Rep. 10(1), 7155 (2020)
[15]. Nagrani, A., Yang, S., Arnab, A., Jansen, A., Schmid, C., Sun, C.: Attention bottlenecks for multimodal fusion. Adv. Neural. Inf. Process. Syst. 34, 14200–14213 (2021)