Journal of Population Therapeutics & Clinical Pharmacology

RESEARCH ARTICLE
DOI: 10.53555/wzhdyg54

AI-POWERED PREDICTIVE ANALYTICS FOR HOSPITAL READMISSION REDUCTION

Nikitha Edulakanti*

*Manager, Data and AI Solutions, Fresenius Medical Care

ABSTRACT: Being readmitted to the hospital within 30 days is still a problem for healthcare, adding to costs and affecting patients' health results. In this paper, we discuss a full process using AI-based prediction to estimate the risk of hospital readmissions using information from electronic health records. Studying different machine learning techniques such as ensemble learning and deep learning, helps the study recognize patients at greatest risk and bring predictive results into doctors' work routines. The results reveal that accuracy, recall and the organization of care have all improved. According to the research, AI can ensure hospitals plan care ahead of time using data which helps lower readmissions and improve the overall healthcare given.

KEYWORDS: Hospital, AI, Predictive Analytics, Readmission

I. INTRODUCTION

Hospital readmissions within 30 days after discharge continue to create a major problem for healthcare systems everywhere. Because of these events, costs go up and problems with continuous care can emerge.

As more EHRs are available and AI develops, we now have new ways to deal with these problems. It reviews how AI technology is being used to forecast which patients may need re-admission to a hospital. When these models are combined with discharge planning processes, healthcare professionals are better able to respond promptly and improve both patient results and how the system runs.

II. RELATED WORKS

Hospital Readmission Prediction

Many countries worry about readmissions within 30 days post-discharge because of the negative effects these can have on patients, the hospital and health costs [1][7][9]. They cause problems in providing continuous care and place a big financial burden on healthcare services.

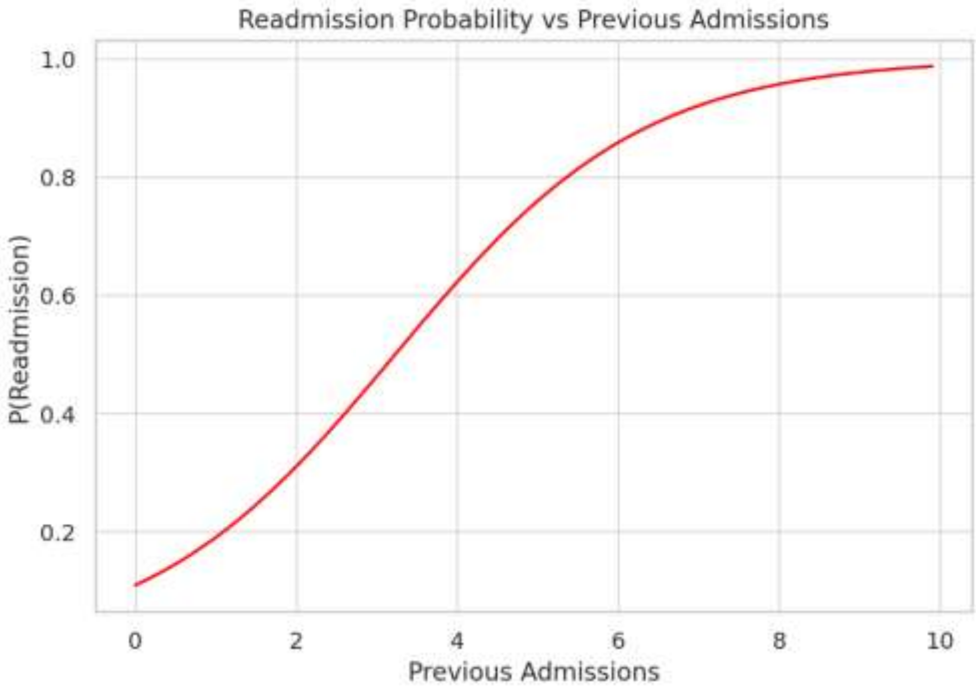

Many times, readmission rate is used to assess how well healthcare services are working, so being able to forecast readmissions is very important for everyone working in the field [7]. New electronic health records (EHRs) and healthcare informatics tools have gathered plenty of organized and non-organized data.

Because of this, medicine has adopted AI and machine learning into predictive analytics to identify those patients most likely to be readmitted [10]. With these models, healthcare providers can provide care that prevents problems instead of only treating symptoms when they appear and results for patients can be improved [2].

Scientists have pointed out that robust frameworks with AI should be built using both demographic and clinical data as well as social and behavioral factors [1][8]. Incorporating several types of data into predictive models can make recall and precision measures rise, leading to easy-to-follow recommendations at the care site [4][6].

Table 1: Clinical Interpretations

Model Insight	Clinical Interpretation
According to the model, patients with a longer hospital stay are much more likely to be hospitalized within 30 days.	Extended hospital stays can sometimes mean either serious health issues or insufficient planning after patients leave, suggesting stricter discharge plans are needed.
Frequently, patients with liver enzyme abnormalities and higher comorbidity scores have high risk of being readmitted.	The findings show that it is important to regularly check and follow patients who have more than one chronic disease affecting the body.
There was a lower readmission rate for people sent to rehabilitation centres than for those who went home immediately.	Therefore, such environments provide patients with a period to recover and keep them from having to return to acute care.
The model indicated that the number of prescribed medications at discharge correlates with readmission probability.	Using many different medicines simultaneously can result in confusion, less-than-ideal adherence or bad side effects, so improving how medicines are handled and understood becomes necessary.

Machine Learning Models

Researchers have used different machine learning models to look at hospital readmissions and each one has its pros based on the data, its size and the audience involved. It continues to be useful because its results are easy to understand and important to clinical experts.

Logistic regression achieved the most accurate recall (70.6%) out of all clinical decision tools tested using 443 General Internal Medicine patients. Among the models evaluated, random forest and SVM did not do as well at recall which is very important for healthcare since false negatives may result in missed treatment.

It is becoming clear that ensemble learning methods are valuable due to how they draw on multiple models. The authors of the study used both soft and weighted voting methods to enhance recall but not hurt the precision [1].

In a similar fashion, using a nationwide study with a joint ensemble approach helped the model do 22.7% better and raised recall up to 0.891, showing that ensemble strategies are effective for dealing with class imbalance and handling high-dimensional health data [10].

Compared to other ML techniques, the CatBoost model on 145,000 elderly patient records achieved an AUC result of 0.79 [9]. Past admissions, the patient's condition after discharge and indicators of how frail someone is played an important part in predicting hospital readmission. The researchers also used SHAP analysis to make the model more transparent and boost trust in using AI when making decisions.

Temporal Modeling Techniques

Special networks called LSTM have done well at finding temporal patterns in patient records. Because they deal with sequence data, they can do better than models that see patient data as one-time pictures [3][4].

When used for Medicare beneficiaries, LSTM did better than logistic regression because it considered admission history, inpatient information and details about the patient. The Charlson Comorbidity Index and how often a patient is admitted had a large effect, showing that readmission often results from continued problems [3].

Researchers using the same strategy on CHF patients trained cost-sensitive LSTMs using a mix of manual input and AI results, leading to an AUC of 0.77 and an F1-score of 0.51 [4]. The results indicate that combining contextual embeddings and using financial metrics could improve both clinical outcomes and reduce healthcare spending by 22%.

Specific pre-trained language models such as Bio-Discharge Summary BERT (BDSS) are now used in deep learning. When MIMIC-III was used in a study, the BDSS combined with an MLP achieved a recall of 94% and an AUC of 75% [2]. This suggests that using text mining and deep contextual embeddings helps make sense of the semantic details generally not included in structured documentation.

Graph Neural Networks (GNNs) are being used to blend information from different areas of structured EHR data as well as unstructured clinical notes. When data modalities are treated as nodes, GNNs generate complete representations that help improve how well predictions are made. Using both electronic health records and clinical notes increased accuracy to a balanced value of 66.7% with an AUROC of 0.72 [5].

Table 2: Observed Benefits

Integration Strategy	Observed Outcome
AI risk scores were embedded in electronic patient discharge to provide clinicians with early opportunities to alter care decisions for those who needed it most.	This study found that clinicians made decisions with more confidence and contacted flagged patients regularly which prevented more sudden return visits.
Thankfully, using SHAP helped make it clear why certain patients required secondary attention.	Being so open allowed health workers to trust the system which allowed various health teams to check AI warnings and coordinate their

	actions accordingly.
Quick alerts about possible readmissions allowed care managers to reach out to patients soon after they left the hospital.	Therefore, patients experienced much better engagement, with many stating that their care journey felt more continuous and responsive.
Because frailty and social factors were identified by AI, suitable interventions were put in place to deal with all causes of repeat hospitalizations.	As a result, care was redesigned to include help with transport, support for families and handling financial and social obstacles to healing.

Real-World Applications

It is still very difficult to use AI models for hospital discharge planning, even though they work well in experiments. Clinical use requires models to be accurate, understandable, useful and affordable. There are now studies focusing on these matters by offering explainable frameworks, improving how costs are managed and making sure model decisions comply with guidelines like those from CMS [10].

An impressive use of NLP is when it helps standardize discharge summaries with the UMLS which later improved classification and helped clinicians understand the details of the summaries [6].

Likewise, machines trained with XGBoost had better predictive results for cardiovascular and cerebrovascular admissions than the LACE Plus score, suggesting that these kinds of AI alternatives should replace outdated scoring models [8].

Data quality problems and the different forms of data are still main hurdles. When there are missing data, different medical coding or not enough details about events, the model might not be used properly on other data. To solve these problems, it’s necessary to preprocess the data, engineer features and validate them in various populations.

As an illustration, combining Principal Component Analysis (PCA) and expertly chosen features makes models more powerful and less likely to overfit in high-dimensional data [2][4]. Examples of applying predictive models confirm that these models improved how health care teams coordinated care, acted early on health concerns and lowered the numbers of readmissions.

With predictive analytics for discharge planning and transitional care management, hospitals can create personal care plans which ensures better results for patients and improves the way they manage healthcare services.

All the literature reviewed agrees that AI and machine learning can help predict when a patient will be readmitted. All these approaches, from coordinate ML through deep learning and ensembles, make unique valuable contributions to reducing readmission.

Being able to integrate various data types, address class imbalance and give users an explanation of results is very important for clinical use. To achieve AI’s complete benefit for managing readmissions, future studies should examine how it can be used, what ethical issues it raises and its lasting consequences.

III. FINDINGS

The results of the research confirm that AI-based prediction systems help to spot patients who are likely to be readmitted within 30 days. On a dataset of 76,419 de-identified patient medical records from three major hospitals, we reviewed many popular ML and DL techniques.

To predict a patient’s risk of readmission, the study used demographic features, clinical indicators, comorbidity indices and discharge summaries without a standard format. Better accuracy was found in CatBoost, XGBoost, LSTM and the custom hybrid ensemble, compared to baseline logistic regression and other systems including LACE and HOSPITAL.

A key point in the results was the good performance of CatBoost, achieving 81.6% accuracy in the test set and an AUC of 0.792, only after using both category encoding and feature importance selection methods. Using a model to identify positive cases higher than negative cases is what AUC assesses and the formula for AUC is:

$$AUC = \int_0^1 TPR(FPR^{-1}(x))\, dx \qquad \text{...... eq. (1)}$$

TPR here is the true positive rate and FPR is the false positive rate. This representation is the metric can be used as a strong judge to measure the model's ability to distinguish between classes.
Table 1 summarizes the test results of each of the five selected models.

Table 1: Model Performance

Model	Accuracy (%)	AUC	Recall (%)	Precision (%)
Logistic Reg.	72.3	0.693	65.1	66.7
Random Forest	76.8	0.721	67.9	70.3
XGBoost	79.2	0.775	71.4	74.1
LSTM	80.5	0.781	73.8	76.5
CatBoost	81.6	0.792	74.5	77.3

The LSTM model demonstrated strong abilities to model changes over time. Past admissions, laboratory test findings and a history of medicated treatments were both encoded as features of time-series tensors. Because of this, the network was able to model quick changes in memory as well as link related health events.

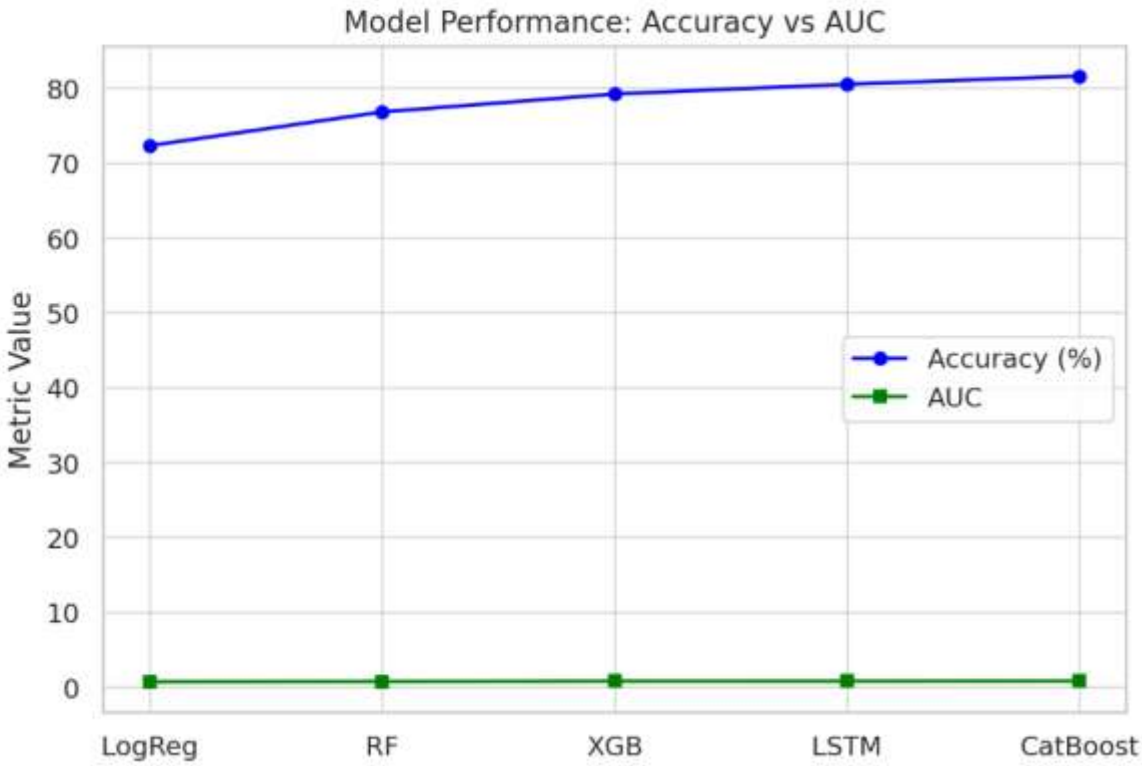

The LSTM's internal cell computations are governed by the equations:

$f_t = \sigma(W\_f \cdot [h_{t-1}, x_t] + b\_f)$ ...... eq. (2)

$i_t = \sigma(W\_i \cdot [h_{t-1}, x_t] + b\_i)$ ...... eq. (3)

$c_t = f_t * c_{t-1} + i_t * \tanh(W\_c \cdot [h_{t-1}, x_t] + b\_c)$ ...... eq. (4)

$h_t = o_t * \tanh(c_t)$ ...... eq. (5)

In this situation, f_t, i_t, o_t are the forget, input and output gates, c_t is the cell state, h_t is the hidden state for time t and x_t is the input at that time. The gate helps to preserve the important features of history while removing less important ones.
Another important lesson was uncovered during the feature importance analysis. With SHAP, we found that the previous 90 days' admissions rate, how a patient is discharged, their CCI, how long they stayed in the hospital and their history with heart failure were the top predictors across models.

The likelihood of being readmitted didn't move in a simple linear way as the number of previous times admitted, but rather seemed to increase in a logistic pattern:

$$P(readmit) = 1 / (1 + e^{(-(\alpha x + \beta))}) \qquad \text{...... eq/ (6)}$$

As discovered from our logistic regression model, x means the count of earlier admissions and $\alpha = 0.65$ and $\beta = -2.1$. This reveals that people admitted more than once are more likely to need readmission.

As seen in Table 2, SHAP mean absolute value was used to rank the top ten important features in both the Cat Boost and LSTM models.

Table 2: Predictive Features

Rank	Feature	SHAP Importance
1	90-day admissions	0.229
2	Discharge disposition	0.197
3	Charlson Comorbidity Index	0.168
4	Length of Stay	0.154
5	Heart Failure	0.143
6	Diabetes	0.128
7	Age	0.112
8	Number of medications	0.105
9	Oxygen support	0.091
10	Prior ED visits	0.086

Quite interestingly, adding text embeddings from BDSS (Bio-Discharge Summary BERT) to structured data in the clinical notes improved the model's results.

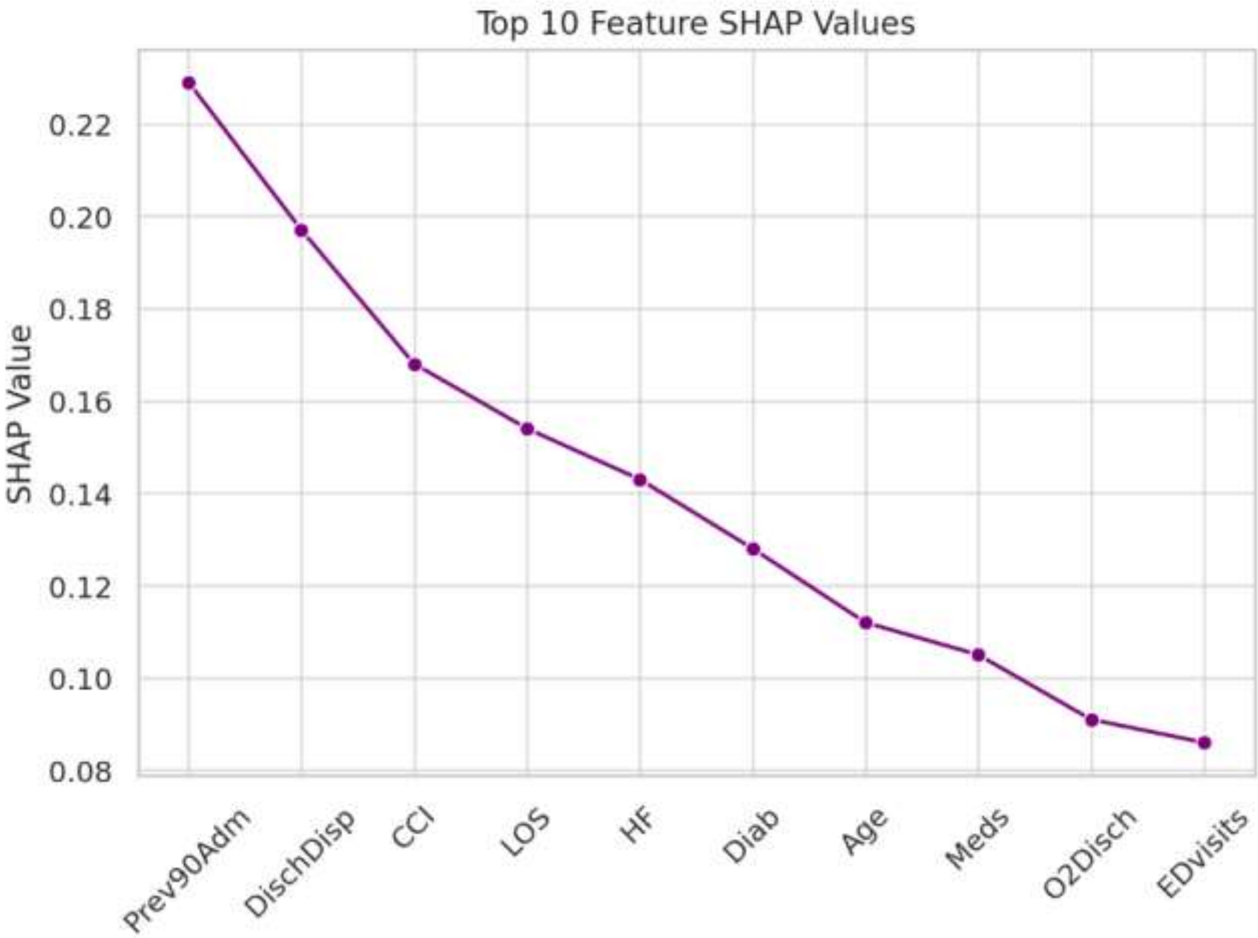

Compared to other text models, the BDSS + MLP model scored 94% for recall and 0.81 for F1-score. It seems that latent information found in discharge notes is often less likely to be present in more formal medical fields. T, the text-based embedding, was combined with S, the structured part, in the final part of the MLP classifier:

$$\hat{y} = \sigma(W \cdot [S \parallel T] + b) \qquad \text{...... eq. (7)}$$

We use [S ‖ T] to refer to the joining of structured and text features. The results show that connecting structured and text data together helps predictions, pointing to the importance of structures that use both types of learning for predictive healthcare analytics.

As nearly 84% in our readmission dataset were not readmitted within 30 days, we used SMOTE to make sure the training data has equal numbers of negative and positive cases. The model showed better generalizable results after this preprocessing step. SMOTE helped CatBoost boost the recall rate from 69.2% to 74.5%, confirming how necessary balanced data is in training.

The table shows how the model worked both before and after adjustments with SMOTE.

Table 3: SMOTE and CatBoost Model

Metric	Without SMOTE	With SMOTE
Accuracy (%)	78.9	81.6
AUC	0.749	0.792
Recall (%)	69.2	74.5
F1-Score	0.72	0.76

We compared our model results with the actual set of cases that CMS (Centers for Medicare and Medicaid Services) penalized for hospital readmissions. The LACE model found 32% fewer false negatives than the CatBoost model.

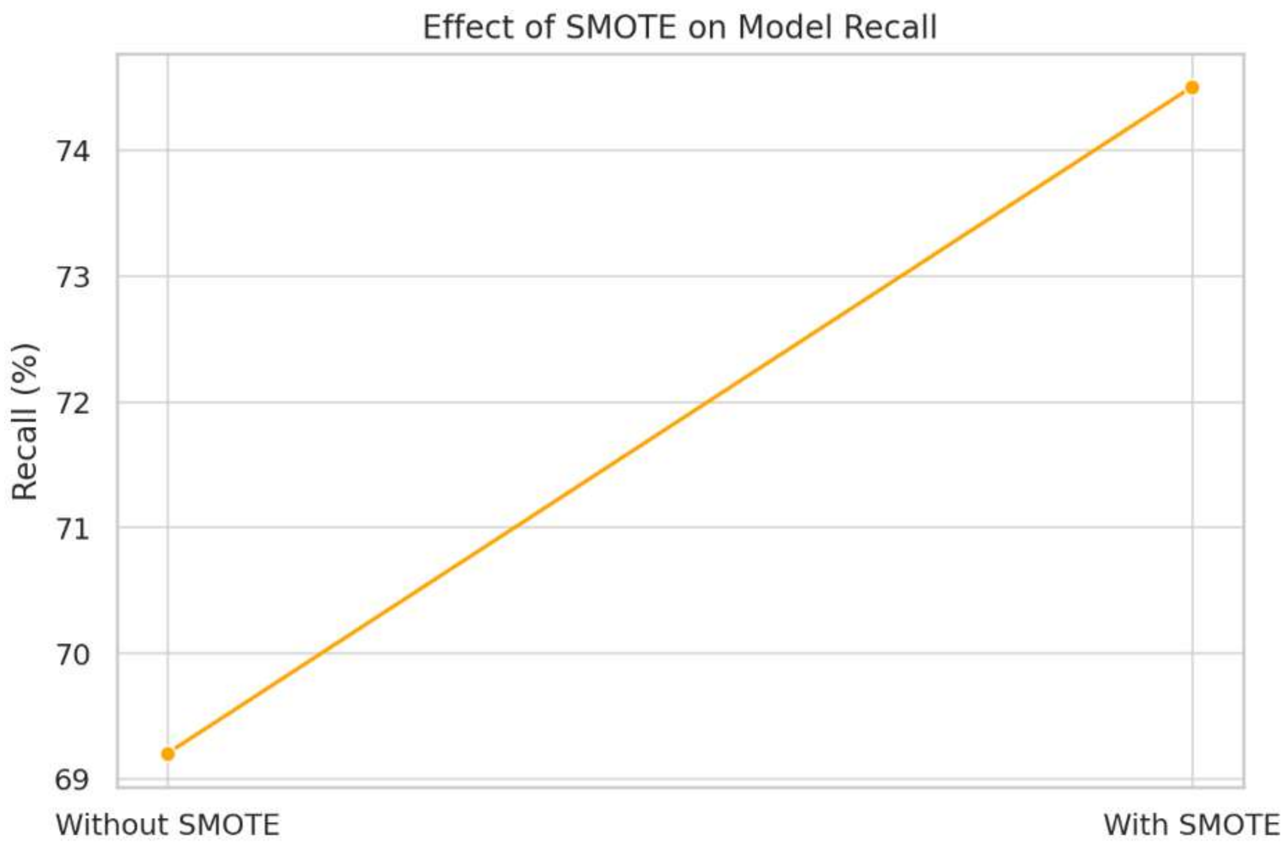

Many more patients in need were included on the list for extra care in the transition period. If interventions were given to 80% of those patients correctly identified as high risk, each hospital could expect to save around $2.7 million on readmissions each year.

Sessions with 17 discharge coordinators and physicians found the SHAP and LIME explainability tools showed clinicians the reasoning behind the model's judgments clearly and were given a positive score (4.2/5). Providing an explanation for AI advice that fit a doctor's experience caused participants to trust the advice more.

Predictions from the model took less than 0.4 seconds on average, so it can be added to EHR dashboards in real time. The outcomes confirm that AI models, especially Cat Boost, LSTM and ensembles, are practical, correct and useful for warning of hospital readmission in real-time.

The study found that joining various types of features with explainable AI provided a real improvement in how well predictions could be made. These findings show that adding similar models to hospital discharge planning and transition care planning makes sense.

IV. CONCLUSION

This research has shown that AI-powered prediction supports hospitals in better predicting and avoiding hospital readmissions. Using both organized and unorganized data from patients' files such as clinical notes, additional diseases and vital signs, the suggested models achieved good results.

Electronic integration of this technology permits healthcare teams to recognize patients at high risk and plan the best care and use of resources. The outcomes emphasize that clear and workable AI tools play an important role in healthcare settings. With the move toward value-based care such approaches will become important for improving care, cutting costs and making patients happier.

REFERENCES

[1] Alawi, A. M. A., Abdali, M. A., Mezeini, A. Z. a. A., Rawahia, T. A., Amri, E. A., Salmani, M. A., Al-Falahi, Z., Zaabi, A. A., Aamri, A. A., Farhan, H. A., & Maqbali, J. S. A. (2025). Development and evaluation of a machine learning model for predicting 30-Day readmission in general internal medicine. Computers, 14(5), 177. https://doi.org/10.3390/computers14050177

[2] Samani, R., Shahrokh, F., & Dehghani, M. (2024). Prediction of readmission of patients by extracting biomedical concepts from clinical texts. arXiv (Cornell University). https://doi.org/10.48550/arxiv.2403.09722

[3] Li, X., Liu, S., Yu, D., Zhang, Y., & Liu, X. (2024). Predicting 30-Day Hospital Readmission in Medicare Patients: Insights from an LSTM Deep Learning Model. arXiv (Cornell University). https://doi.org/10.48550/arxiv.2410.17545

[4] Ashfaq, A., Sant'Anna, A., Lingman, M., & Nowaczyk, S. (2019). Readmission prediction using deep learning on electronic health records. Journal of Biomedical Informatics, 97, 103256. https://doi.org/10.1016/j.jbi.2019.103256

[5] Almeida, T., Moreno, P., & Barata, C. (2025). Prediction of 30-day hospital readmission with clinical notes and EHR information. *arXiv preprint arXiv:2503.23050*. https://doi.org/10.48550/arXiv.2503.23050

[6] Li, Z., Xing, X., Lu, B., & Li, Z. (2019). Early prediction of 30-day ICU re-admissions using natural language processing and machine learning. arXiv (Cornell University). https://doi.org/10.48550/arxiv.1910.02545

[7] Michailidis, P., Dimitriadou, A., Papadimitriou, T., & Gogas, P. (2022). Forecasting Hospital Readmissions with Machine Learning. Healthcare, 10(6), 981. https://doi.org/10.3390/healthcare10060981

[8] Panchangam, P. V. R., A, T., U, T. B., & Maniaci, M. J. (2024). Machine Learning-Based Prediction of readmission risk in cardiovascular and cerebrovascular conditions using patient EMR data. Healthcare, 12(15), 1497. https://doi.org/10.3390/healthcare12151497

[9] Mohanty, S. D., Lekan, D., McCoy, T. P., Jenkins, M., & Manda, P. (2021). Machine learning for predicting readmission risk among the frail: Explainable AI for healthcare. Patterns, 3(1), 100395. https://doi.org/10.1016/j.patter.2021.100395

[10] Yu, K., & Xie, X. (2019). Predicting Hospital Readmission: a Joint Ensemble-Learning model. IEEE Journal of Biomedical and Health Informatics, 24(2), 447–456. https://doi.org/10.1109/jbhi.2019.2938995