IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 7, October 2026
Impact Factor: 8.2

# Explainable Random Forest Prediction of 30-Day Hospital Readmission

[1]Ms.Vijeta Dinesh Mandlik and [2]Ms. Chikhale Shivani Subodh
[1]Student Jaihind Institute of Management and Research, Junnar, India
[2]Assistatmt Professor, Jaihind Institute of Management and Research, Junnar, India

**Abstract**: *Early return to hospital after discharge is a widely used signal of care quality and a major driver of avoidable cost. This paper reports an explainable machine-learning study that estimates whether a diabetic patient will be readmitted within 30 days, using the public Diabetes 130-US Hospitals data. Keeping only the first encounter of every patient left 69,987 records, of which 5,143 (7.35%) were early readmissions. A single Random Forest classifier with 150 trees and balanced class weights was fitted on 55,989 records and checked on 13,998 unseen ones. It reached 85.1% accuracy, yet precision was 0.174, recall 0.276, F1-score 0.214 and ROC-AUC 0.665, which shows how misleading accuracy becomes under heavy imbalance. Permutation importance ranked the count of earlier inpatient stays first, followed by discharge to skilled nursing and rehabilitation facilities. A 20-feature subset kept almost all of the performance, and a threshold study exposed the precision–recall trade-off that a hospital must manage. The model is best seen as a decision-support aid that complements, and never replaces, clinical judgement.*

**Keywords**: Hospital readmission, Random Forest, Explainable AI, Permutation importance, Diabetes, Class imbalance, Healthcare analytics

## I. INTRODUCTION

A readmission happens when a patient who has left hospital is admitted again. When the gap is short, the return often points to unfinished recovery, trouble with medicines, weak follow-up or problems at home. For this reason the 30-day readmission rate is treated as a quality indicator, and a high share of avoidable returns burdens clinical staff, uses up hospital capacity, raises treatment cost and worsens the patient's experience [1].

Hospitals now store large volumes of electronic health records covering demographics, diagnoses, laboratory work, drugs, length of stay, earlier admissions and discharge destination. Machine learning can mine such records to move from describing past events to estimating the likelihood of a future one. Diabetes is a suitable test case because affected patients need long-term monitoring, often take several medicines, frequently have coexisting illnesses and are hospitalised repeatedly [3].

### A. Problem Statement

Hospitals need a reliable way to identify diabetic patients who are likely to return within 30 days after discharge so that additional follow-up can be planned. However, predicting such readmissions is difficult because electronic health records may contain missing or unknown values, categorical information and multiple encounters belonging to the same patient. In this study, only 7.35% of the prepared records represent readmissions within 30 days, creating a substantial class imbalance. As a result, a prediction model may show high overall accuracy while failing to identify many patients who are actually readmitted. The problem addressed in this research is therefore to develop and assess an explainable machine-learning model that can estimate 30-day readmission risk for diabetic patients, identify the factors associated with that risk, and provide performance measures that reflect the minority readmission class.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 7, October 2026
Impact Factor: 8.2

Predicting readmission is nevertheless hard. Raw records contain missing entries, categorical codes, repeated encounters of the same person and, above all, a strongly skewed outcome: in the data prepared here only about 7.35% of cases are early readmissions. A model can therefore look accurate while overlooking nearly every patient who actually returns. The aim of this work is to build an explainable Random Forest that flags patients at higher risk of returning within 30 days and reveals which factors drive that risk. The main contributions are:

a leakage-aware workflow that keeps one encounter per patient and covers cleaning, encoding, feature engineering and exploratory analysis;

an evaluation that reports precision, recall, F1-score, ROC-AUC and the confusion matrix next to accuracy;

a permutation-importance ranking of predictors, together with a test of whether 20 features can replace the full set; and

a threshold study showing how the alert rate can be tuned to the follow-up capacity of a hospital.

## II. RELATED WORK

Jencks et al. analysed rehospitalisation among Medicare beneficiaries and showed that returns within a month are frequent, which established 30-day readmission as a standard outcome for quality studies [1]. A later systematic review of risk-prediction models by Kansagara et al. found that most tools discriminate only moderately, partly because administrative data omit many patient-level influences [2]. That finding frames the present work: the goal is to extract useful patterns and rank predictors, not to claim near-perfect prediction.

The Diabetes 130-US Hospitals collection was examined by Strack et al., who related treatment and hospitalisation variables, including HbA1c measurement, to readmission among diabetic inpatients [3]. The same collection is used here. Random Forest [4] averages many decision trees grown on different samples and feature subsets. It copes with many inputs, models nonlinear effects and interactions, accepts encoded categorical variables, returns probabilities and supports importance estimates, which makes it a practical single algorithm for tabular clinical data.

Two gaps remain. Studies often stress accuracy although imbalance makes it uninformative, and precision–recall measures are more revealing in that setting [7]. Many studies also give a prediction without saying which variables support it. This paper tackles both gaps by combining imbalance-aware evaluation, threshold analysis and permutation importance [6] around one model.

## III. DATA AND METHODOLOGY

### A. Dataset

The study uses the publicly available, de-identified Diabetes 130-US Hospitals data, which holds 101,766 hospital encounters described by 50 variables. These cover demographics, admission and discharge details, diagnoses, laboratory procedures, medications, earlier hospital use, length of stay and diabetes-related information. The target is binary: 1 if the patient returned within 30 days and 0 otherwise. Table I summarises the data after preparation.

TABLE I: Summary of the Dataset

| Item | Value |
|---|---|
| Original encounters / variables | 101,766 / 50 |
| Records after first-encounter selection | 69,987 |
| Not readmitted within 30 days | 64,844 (92.65%) |
| Readmitted within 30 days | 5,143 (7.35%) |
| Training / testing records | 55,989 / 13,998 |
| Features after encoding | 138 |

### B. Cleaning and Patient-Level Selection

Missing and unknown categories were identified and treated, identifiers that carry no predictive meaning were dropped, categorical variables were converted to numeric form, and numeric values were validated. Because one person can

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 7, October 2026
Impact Factor: 8.2

appear in several encounters, using every row could place the same patient in both training and test sets and inflate the scores. The primary analysis therefore keeps only each patient's first encounter, giving the 69,987 records in Table I.

### C. Feature Engineering

Raw fields were turned into model-ready variables. Age bands were mapped to a numeric age, previous outpatient, emergency and inpatient contacts were combined into a total-visit count, and discharge destination was encoded in categories. Retained clinical counts include the number of earlier inpatient visits, diagnoses, medications and laboratory procedures, plus the length of stay. After encoding, 138 input features remained.

### D. Evaluation Metrics

Accuracy, precision, recall, F1-score and ROC-AUC were computed on the held-out test set. Precision is TP/(TP+FP), recall is TP/(TP+FN), and F1 is their harmonic mean, 2PR/(P+R). ROC-AUC measures how well the predicted probabilities separate the two classes over all thresholds; 0.50 means chance-level separation and 1.00 means perfect separation. Because positives are rare, recall and precision matter more here than accuracy.

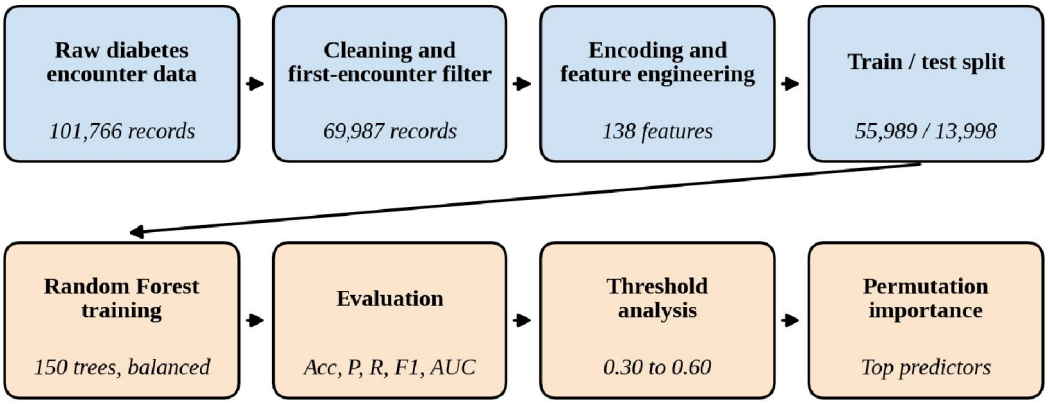



Fig. 1 Overall workflow of the study, from raw encounter data to explainability.

## IV. PROPOSED RANDOM FOREST MODEL

### A. System Overview

The proposed system accepts patient information, applies the same preprocessing and feature construction used in training, passes the resulting vector to the trained forest, and converts the predicted probability into a risk class through a chosen threshold. Conceptually it has four layers: a data layer holding the historical records, a processing layer for cleaning, encoding and engineering, a prediction layer containing the trained forest, and a presentation layer that displays the estimated risk. The model was implemented in Python with Pandas, NumPy and scikit-learn [5]; Matplotlib and Seaborn produced the plots, and Flask was used for the prediction interface.

### B. Random Forest Configuration

Each tree is grown on a bootstrap sample of the training data and considers only a random subset of features at every split, so the trees differ from one another. The forest's probability for a patient is the share of trees that vote for readmission. Table II lists the settings.

TABLE II: Random Forest Settings

| Parameter | Value |
|---|---|
| Number of trees | 150 |

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 7, October 2026
Impact Factor: 8.2

| | |
|---|---|
| Minimum samples per leaf | 10 |
| Class weighting | Balanced |
| Random state | 42 |
| Task / target | Binary / 30-day readmission |

### C. Imbalance Handling and Explainability

Balanced class weights make errors on the rare readmission class cost more during training. This does not remove imbalance, so results are judged with precision, recall, F1-score and ROC-AUC rather than accuracy alone. To explain the model, permutation importance shuffles one feature at a time on the test data and records how much ROC-AUC falls; a larger fall means the model leans more heavily on that feature [6].

## V. RESULTS AND DISCUSSION

### A. Exploratory Patterns

Early readmission is rare (7.35%), confirming strong imbalance. Readmission tends to rise with age, although age is better read as a proxy for disease burden and care needs than as a cause. Prior inpatient use shows the clearest gradient: the early-readmission rate climbs from about 5.8% for patients with no earlier inpatient stay to 27.3% for those with four or more. Primary diagnosis separates patients only modestly, from 6.24% for respiratory conditions to 8.47% for injury, with circulatory cases, the largest group (21,095 patients), at 7.78%. Gender and race differ little, and medication change and diabetes-medication use show slightly higher rates.

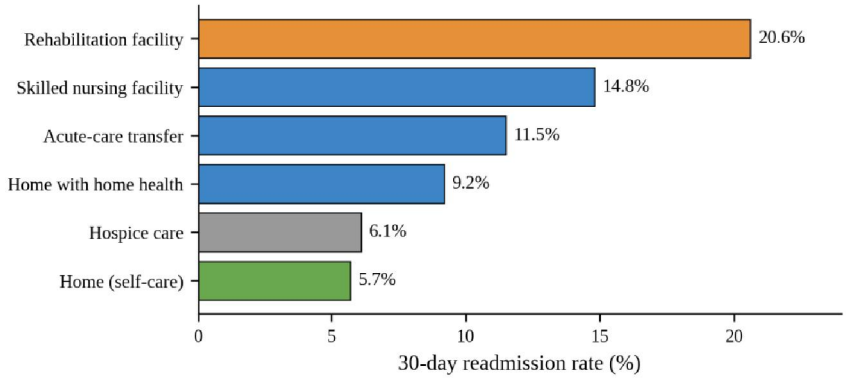



Fig. 2 Early readmission rate by discharge destination.

Discharge destination matters strongly (Fig. 2). Patients sent to a rehabilitation facility were readmitted at 20.6% and those leaving for a skilled nursing facility at 14.8%, against 5.7% for patients who went home. These gaps should not be read as the facilities causing readmission; such patients usually have more serious conditions and heavier care needs.

### B. Model Performance

TABLE III: Test-Set Performance of the Random Forest

| Metric | Score |
|---|---|
| Accuracy | 0.851 |
| Precision | 0.174 |
| Recall | 0.276 |
| F1-score | 0.214 |
| ROC-AUC | 0.665 |

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 7, October 2026
Impact Factor: 8.2

Table III shows 85.1% accuracy, but this figure mostly reflects the large non-readmitted group. Only 17.4% of flagged patients were truly readmitted, and just 27.6% of actual readmissions were found. The ROC-AUC of 0.665 is clearly above chance yet only moderate, in line with earlier reviews of readmission models [2].

### C. Confusion Matrix

TABLE IV: Confusion Matrix on 13,998 Test Records

| Actual \ Predicted | Not readmitted | Readmitted |
|---|---|---|
| Not readmitted | 11,623 (TN) | 1,346 (FP) |
| Readmitted | 745 (FN) | 284 (TP) |

The model caught 284 of the 1,029 true readmissions and missed 745. These false negatives are the costliest errors because a high-risk patient may miss extra follow-up. The 1,346 false positives are patients flagged without returning, and they add unnecessary work for staff. Reducing one error type generally increases the other.

### D. Threshold Analysis

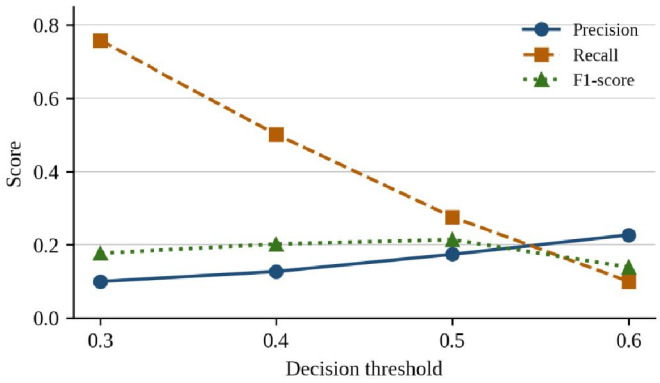



Fig. 3 Precision, recall and F1-score at different decision thresholds.

Fig. 3 shows the effect of moving the cut-off applied to the predicted probability. At 0.30 the model recovers 75.8% of early readmissions, but precision falls to 0.100. At 0.60 precision rises to 0.227 while recall drops to 0.100. The F1-score peaks at 0.50 (0.214). A hospital with ample follow-up capacity may prefer a low threshold, whereas a resource-limited one may prefer a high threshold; in either case the value should be chosen on separate validation data and operational needs.

### E. Key Predictors

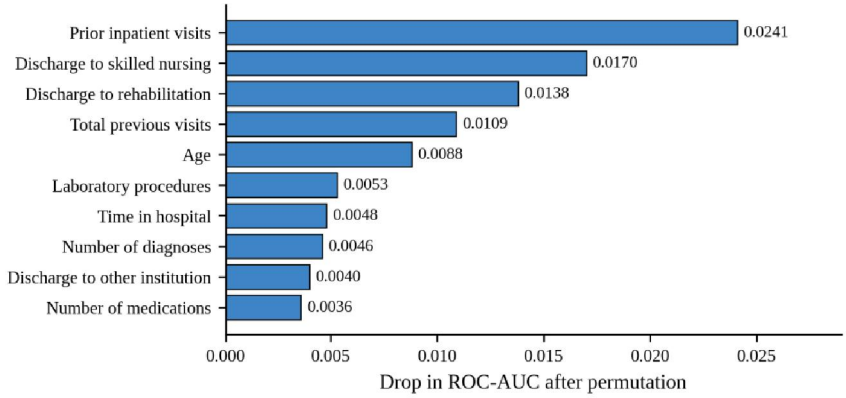



Fig. 4 Ten most influential predictors by permutation importance.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 7, October 2026
Impact Factor: 8.2

Fig. 4 ranks the predictors. The number of earlier inpatient visits is clearly first (AUC drop 0.0241), so a patient's hospital history carries the most information about future return. Discharge to a skilled nursing facility (0.0170) and to a rehabilitation facility (0.0138) follow, then total previous visits and age. Laboratory procedures, length of stay, number of diagnoses and number of medications form a weaker group. All of these are predictive indicators, not proven causes.

### F. Reduced Feature Set

TABLE V: Effect of Feature Selection

| Feature set | Features | Accuracy | Precision | Recall | F1 | ROC-AUC |
|---|---|---|---|---|---|---|
| All features | 138 | 0.851 | 0.174 | 0.276 | 0.214 | 0.665 |
| Top 20 | 20 | 0.849 | 0.170 | 0.273 | 0.210 | 0.656 |

Cutting the inputs from 138 to 20 lowered ROC-AUC only from 0.665 to 0.656 (Table V). A compact feature set therefore keeps nearly all of the predictive value and is far easier to inspect and explain.

## VI. FINDINGS AND PRACTICAL IMPLICATIONS

The study led to four main findings. First, early readmission is a heavily imbalanced outcome, so accuracy alone is inadequate. Second, earlier inpatient use and discharge destination are the most important signals. Third, the Random Forest separates patients better than chance (ROC-AUC 0.665) but remains only moderately strong. Fourth, a 20-feature subset performs almost as well as the full set. Regarding the four hypotheses, those on prior inpatient use, discharge destination and ROC-AUC above 0.50 were supported, while the reduced-feature hypothesis was only partly supported because the smaller set matched, but did not beat, the full set.

For hospitals, patients with several earlier inpatient stays, or those discharged to rehabilitation or skilled nursing facilities, are sensible candidates for stronger transition-of-care coordination, medication review and prioritised follow-up calls. Risk thresholds should reflect available staff, and predictions must be combined with clinical judgement. For data scientists, reporting recall, precision, F1-score and ROC-AUC next to accuracy, splitting at patient level to avoid leakage, interpreting importance scores cautiously, tuning thresholds on validation data and running external validation before reuse in another hospital are all advisable.

## VII. LIMITATIONS AND FUTURE SCOPE

The data describe historical encounters from one set of hospitals, so current practice and other populations may differ. Social and behavioural factors such as medication adherence, family support, living conditions, transport and financial difficulty are absent, although they influence readmission. Moderate discrimination, together with both false negatives and false positives, limits operational use, and external validation is required before deployment. The model supports, and does not replace, clinicians.

Future work can use newer and more varied datasets, add socioeconomic and follow-up information, tune hyperparameters with cross-validation, and provide patient-level explanations that say why an individual was flagged. The model could also be linked to a hospital information system to give a risk score at discharge, with regular monitoring for drift in patient mix, treatment practice and data collection. Comparisons with other algorithms are left for a separate study.

## VIII. CONCLUSION

This paper developed an explainable Random Forest approach for predicting 30-day readmission among diabetic patients from 69,987 prepared records. The model attained 85.1% accuracy, 0.174 precision, 0.276 recall, 0.214 F1-score and 0.665 ROC-AUC, confirming that imbalance-aware metrics are essential. Prior inpatient visits emerged as the strongest predictor, followed by discharge destination, total earlier visits and age, and a 20-feature subset preserved most of the performance. The threshold study showed that the precision–recall balance can be matched to hospital

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 7, October 2026
Impact Factor: 8.2

resources. Used carefully, and validated externally, such a model can support follow-up planning for patients most likely to return.

## REFERENCES

1. S. F. Jencks, M. V. Williams, and E. A. Coleman, "Rehospitalizations among patients in the Medicare fee-for-service program," *N. Engl. J. Med.*, vol. 360, no. 14, pp. 1418–1428, 2009.
2. D. Kansagara *et al.*, "Risk prediction models for hospital readmission: A systematic review," *JAMA*, vol. 306, no. 15, pp. 1688–1698, 2011.
3. B. Strack *et al.*, "Impact of HbA1c measurement on hospital readmission rates: Analysis of 70,000 clinical database patient records," *BioMed Res. Int.*, vol. 2014, 2014.
4. L. Breiman, "Random forests," *Mach. Learn.*, vol. 45, no. 1, pp. 5–32, 2001.
5. F. Pedregosa *et al.*, "Scikit-learn: Machine learning in Python," *J. Mach. Learn. Res.*, vol. 12, pp. 2825–2830, 2011.

A. Altmann, L. Toloşi, O. Sander, and T. Lengauer, "Permutation importance: A corrected feature importance measure," *Bioinformatics*, vol. 26, no. 10, pp. 1340–1347, 2010.

6. T. Saito and M. Rehmsmeier, "The precision-recall plot is more informative than the ROC plot when evaluating binary classifiers on imbalanced datasets," *PLoS ONE*, vol. 10, no. 3, e0118432, 2015