Research Article

A High-Recall Ensemble Framework for Patient No-Show Prediction; Risk-Adjusted Scheduling Optimization in Healthcare Systems

Ashok Kumar Reddy Desai[1*], Sathish Babu Jampani[2]

[1,2]Department of CSE, Koneru Lakshmaiah Education Foundation, Vaddeswaram, AP
*Correspondence Email: jampanisatishbabu@gmail.com, Co-author email: 2401050203mtech@gmail.com

Abstract

B Patient no-shows significantly disrupt healthcare delivery, leading to underutilized resources and financial strain. This study proposes a hybrid predictive–prescriptive framework that integrates machine learning with a Reinforcement Learning (RL)-inspired decision engine for clinical queue optimization. We have created a model consisting of an ensemble of XGBoost with CatBoost that are specifically tuned to high-risk detection with Recall of 87.75% and F1-score of 0.429. Unlike traditional models that prioritize raw accuracy, our approach focuses on minimizing missed no-show cases to drive a three-tier intervention strategy: automated SMS reminders, targeted telephonic confirmations, and risk-adjusted smart overbooking. The predicted probabilities are integrated into an RL-based scheduling mechanism that dynamically categorizes appointments into "Confirm" or "Reschedule" actions based on a probability-cost trade-off. Experimental results demonstrate that this framework transforms predictive insights into actionable scheduling decisions, effectively reducing idle slots and maximizing clinical throughput. The proposed system can provide a scalable, practical solution to the problem of improving the delivery of healthcare services and operational efficiency in the resource-constrained settings by changing the current system of static scheduling to that of probability-based intervention.

Keywords: Patient No-Show Prediction, Machine Learning, Healthcare Scheduling, Ensemble Learning, Queue Optimization.

1. Introduction

Productive scheduling of patients is one of the keys of the modern healthcare system directly affecting the use of resources, patient satisfaction, and clinical throughput (Gupta et al., 2008; Vandenberghe et al., 2019). The main challenge to the operational efficiency is the high measure of patient no-show-up, i.e. the appointments of a patient who does not show-up on time with no prior notice. Such cancelations lead to an underutilization of the medical infrastructure, long waiting lines among the rest of the patients, and significant losses (Dantas et al., 2018; Samorani et al., 2016). The financial cost of no-shows is truly inconceivable on an international level, as in most instances, it has been estimated millions of dollars in annual funds of healthcare providers as well as what amounts to impeding quality care delivery. (Huang et al., 2014; Shah et al., 2016). Machine learning (ML) is a widely discussed topic in recent years as an approach to predict no-show behavior by using demographic, medical, and temporal data (Ng et al., 2019; Kurasawa et al., 2020). Nevertheless, conventional ML methods in this field often tend to consider the global accuracy more than the class-specific sensitivity. These models usually can not adequately include the minority no-show category in the context of highly imbalanced clinical data, leading to the low recall (Sun et al., 2009; Artetxe et al., 2018). The practical value can hardly be reduced by this limitation; a model which fails to recognize a possible no-show can provide the hospital with no possibility of reclaimed lost capacity. Moreover, the focus of existing studies on prediction is more of an independent task without connecting the two concepts, which is predictive analytics and prescriptive operations strategies, such as dynamic scheduling, risk-sensitive overbooking, and tiered intervention planning (Gupta et al., 2008; Samorani et al., 2016). In order to overcome these issues, this paper contains a hybrid predictive-prescriptive framework, which changes passive prediction to an active operation control. Using a synergistic XGBoost and CatBoost ensemble architecture with a reinforcement learning (RL)-inspired decision engine, we put the high-sensitivity detection of no-show events (Chen et al., 2016; Prokhorenkova et al., 2018) in the first place. This allows an inherent risk-sensitive scheduling plan which brings predictive output in line with actual clinical limitations. By modifying our strategy, the point of focus will no longer be on how much a no-show can be predicted but rather how much we can do to alleviate the effect of it, which is a three-tier intervention layer of automated SMS alerts, specially target telephonic confirmations and smart overbooking.

2.Literature Review

A large number of machine learning and operational research methods have been used to predict patient no-show and healthcare scheduling. Initial research aimed at determining patient behavioral patterns with the help of statistical and classification models, which one of the key patterns of patient behavior is the effect of missed appointments on healthcare efficiency and resource use (Huang et al., 2014; Shah et al., 2016). Systematic reviews have also highlighted that the no-show rates constitute a major problem in hospital operations that contributes to an increase in waiting time and losses (Dantas et al., 2018). Conventional scheduling strategies in the medical sector have been discussed as well, overcoming the issue of scheduling appointments, and managing facilities (Gupta et al., 2008). Ensemble models, like gradient boosting, have become more popular with the current development of machine learning because of their predictive accuracy and the ability to work with complicated data distributions (Chen et al., 2016; Prokhodenko et al., 2018). Multiple papers have used machine learning models to forecast patient attendance with improvements in predictive quality; nonetheless, most of the models exhibit a lack of class balance, which makes them incapable of detecting minority no-show cases (Kurasawa et al., 2020; Sun et al., 2009). Also, the studies of the overbooking strategies have demonstrated that it is possible to prevent the outcomes of the missed appointments and enhance the clinical throughput with the help of predictive insights (Samorani et al., 2016). In spite of these developments, the majority of the current literature tend to be mostly predictive and does not include real-time scheduling determination (Vandenberghe et al., 2019). In addition, limited focus has been on high-recall models which focus on identification of no-show patients as a basis of proactive intervention. Hence, there is still the lack of creating hybrid frameworks that will integrate precise prediction and decision-based scheduling systems to optimize the work of healthcare systems.

3.Methodology

3.1 Problem Formulation

The problem is described as a binary classification problem whereby the task is to assign a patient based on the risks of attending a scheduled appointment or not. Every record is comprised of patient related attributes like demographic, medical and time attributes. The aims are to acquire a mapping of input features to the likelihood of an occurrence of no-show.

$$\boldsymbol{D} = \{(\boldsymbol{x_i}, \boldsymbol{y_i})\}_{\boldsymbol{i=1}}^{\boldsymbol{n}} \quad (1)$$

This is the data representation in which x i is the feature variable and y i is the classification (0 = show, 1 = no-show).

3.2 Ensemble Prediction Model

Two strong models of gradient boosting, viz, XGBoost and CatBoost, are used to supplement a success of prediction. Such models capture non-linear relationships amongst data which are complex and feature interactions are well modeled. The ensemble method makes it stronger and contributes to a low prediction variance.

3.2.1. XGBoost Prediction

$$\boldsymbol{P_{xgb}(y = 1 \mid x) = \sigma\left(\sum_{k=1}^{K} f_k\,(x)\right)} \quad (2)$$

In this case, several decision trees fk(x) are added and the output is converted to a probability with the help of the sigmoid.

3.2.2. CatBoost Prediction

$$\boldsymbol{P_{cat}(y = 1 \mid x) = \frac{1}{1+e^{-F(x)}}} \quad (3)$$

CatBoost model is used to create a score F(x) and change it into probability by a logistic function.

3.2.3.Final Ensemble Probability

$$\boldsymbol{P_{final} = \frac{P_{xgb}+P_{cat}}{2}} \quad (4)$$

Both model results are averaged to get the ultimate prediction probability to enhance stability.

3.3 Threshold-Based Classification

Ty proceeds to pick an ideal threshold which provides a trade off between precision and recall without following the fixed threshold. This proves handy both in the improvement of the detection of the no-show cases as well as the handling of the false forecasts.

$$\hat{y} = \begin{cases} \mathbf{1}, & \textbf{if } \boldsymbol{P_{final} > T} \\ \mathbf{0}, & \textbf{otherwise} \end{cases} \quad (5)$$

When the probabilities are higher than the threshold T, then such a patient is considered as a no-show.

3.4 Evaluation Metrics

The model is tested based on the standard classification measures. These measures give a full picture of the predictability of the model on the two classes and particularly the minority no-show class.

3.4.1.Accuracy

$$\boldsymbol{Accuracy} = \frac{\boldsymbol{TP+TN}}{\boldsymbol{TP+TN+FP+FN}} \quad (6)$$

The concept of accuracy defines the general accuracy of predictions.

3.4.2.Precision

$$\boldsymbol{Precision} = \frac{\boldsymbol{TP}}{\boldsymbol{TP+FP}} \quad (7)$$

Precision refers to the number of predicted no-shows which are correct.

3.4.3.Recall

$$\boldsymbol{Recall} = \frac{\boldsymbol{TP}}{\boldsymbol{TP+FN}} \quad (8)$$

Recall is an assessment of the model that determines the real no-show cases.

3.4.4.F1 Score

$$\boldsymbol{F1} = \mathbf{2} \cdot \frac{\boldsymbol{Precision \cdot Recall}}{\boldsymbol{Precision+Recall}} \quad (9)$$

F 1 score combines accuracy and recall into one number.

3.5 SMOTE (Data Balancing)

Given that the sample is skewed, The method of SMOTE is used to form artificial samples of the minority group. This will help the model to build a more powerful trend of no-show cases and recall.

$$\mathrm{x_{new}} = \mathrm{x_i} + \lambda(\mathrm{x_{nn}} - \mathrm{x_i}) \quad (10)$$

Another synthetic sample is formed between a data unit and a closest one.

3.6 System Architecture

- Input: Patient dataset $D = \{(x_i, y_i)\}_{i=1}^{n}$
- Data Preprocessing: Missing values, feature analysis, and normalization.
- SMOTE: Class balancing through the use of synthetically generated data.
- Feature Input to Models
- Model 1: XGBoost → $P_{xgb}(y = 1 \mid x)$
- Model 2: CatBoost → $P_{cat}(y = 1 \mid x)$
- Ensemble Layer: $P_{final} = \frac{P_{xgb}+P_{cat}}{2}$
- Threshold-Based Classification → $\hat{y}$
- Output: No-show Prediction & Scheduling Decision

Description

The suggested architecture includes an ensemble learning model that will be an integrated XGBoost and CatBoost. It

works with the preprocessed and balanced data through SMOTE and she estimates the probability and averaging. The last prediction is derived by threshold based classification which aids in scheduling decision-making.

3.7 Algorithm: Ensemble-Based No-Show Prediction

Input: Dataset D = {(xi, yi)}
Output: Prediction ŷ and Scheduling Decision

Step 1: Load dataset D
Step 2: Perform Exploratory Data Analysis (EDA)
Step 3: Preprocess data (cleaning, transformation)
Step 4: Apply SMOTE to balance dataset
Step 5: Train XGBoost model and compute P_xgb
Step 6: Train CatBoost model and compute P_cat
Step 7: Calculate overall probability:
    P_final = (P_xgb + P_cat) / 2
Step 8: Apply threshold T:
    If P_final > T, then ŷ = 1
    Else ŷ = 0
Step 9: Make scheduling decision:
    If P_final > T → Reschedule
    Else → Confirm
Step 10: Evaluate model using Accuracy, Precision, Recall, F1-score
End

Description

The algorithm explains the process involved in implementing the ensemble prediction system step by step. It comprises of preprocessing, model training, probability aggregation, and threshold based classification. Included is the last stage, which is incorporation of the result of predictions into the scheduling decisions.

3.8 Scheduling Decision Function

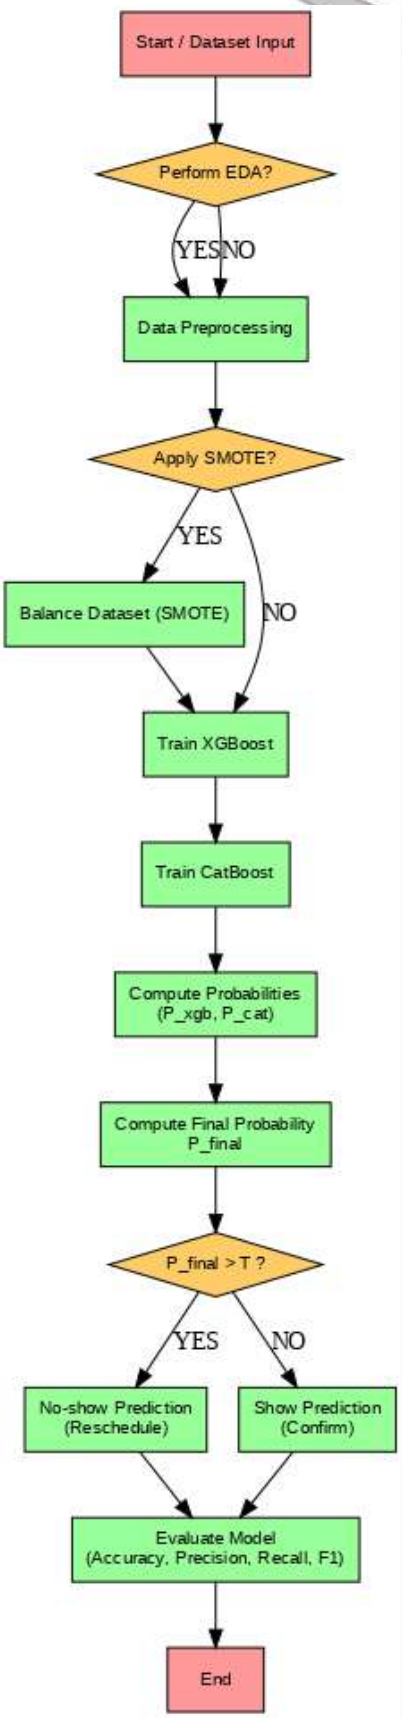
In a real-life situation, scheduling decisions involve the use of the predicted probabilities. Patients who have an increased likelihood of no show are re-scheduled or intervened, allowing them to be brought under control.

$$Decision = \begin{cases} Reschedule, & P_{final} > T \\ Confirm, & P_{final} \leq T \end{cases} \quad (11)$$

This role transforms model into action based scheduling decisions.

3.9 Flowchart

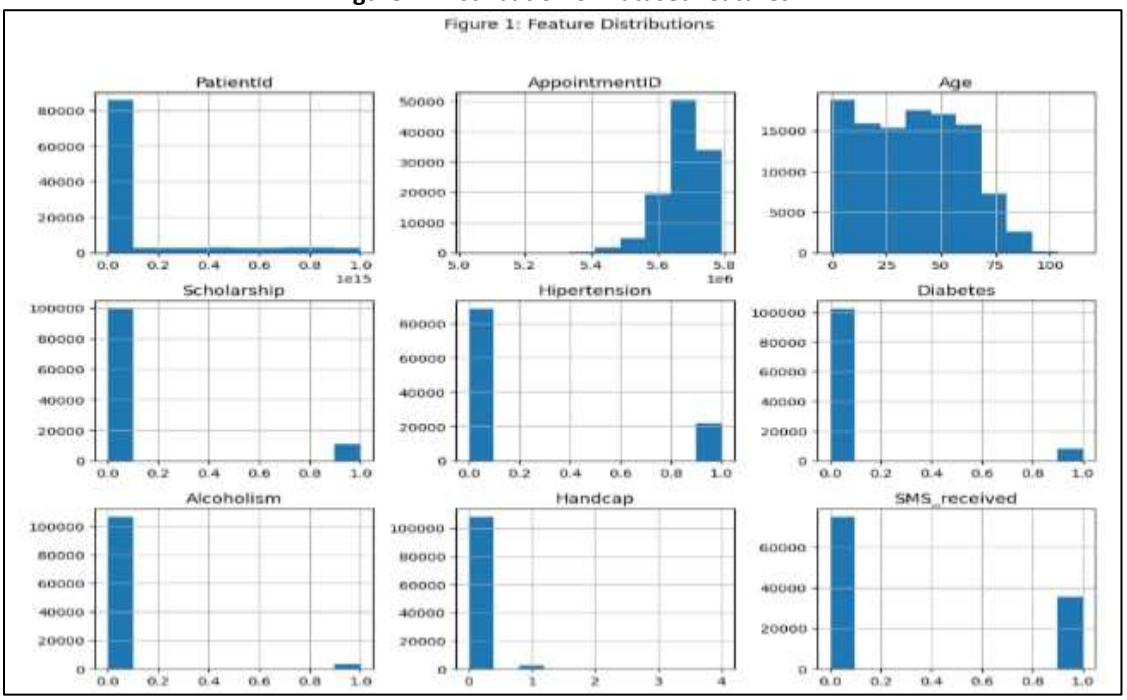
The workflow presents the chronological scheme beginning with the process of data loading and preprocessing, ending with the model training and evaluation. It focuses on blending ensemble forecasts and classification using a threshold. The last process transforms predictions into schedulicative decision making.


Start / Dataset Input
Perform EDA?
YES
NO
Data Preprocessing
Apply SMOTE?
YES
NO
Balance Dataset (SMOTE)
Train XGBoost
Train CatBoost
Compute Probabilities (P_xgb, P_cat)
Compute Final Probability P_final
P_final > T ?
YES
NO
No-show Prediction (Reschedule)
Show Prediction (Confirm)
Evaluate Model (Accuracy, Precision, Recall, F1)
End

3.10 Exploratory Data Analysis (EDA)

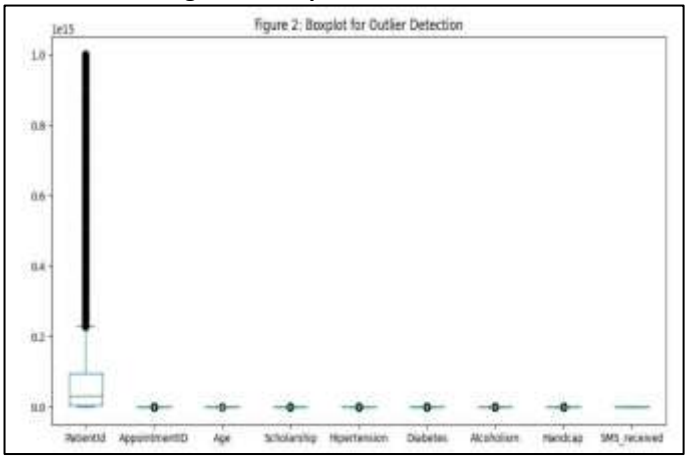
The Exploratory Data Analysis was done to have a glimpse of how the dataset looks and what features it contains and in which of their distribution. The initial analysis was the verification of the missing values, data types and statistical summaries to have a picture of the data. The distribution plots indicated that the a number of the numerical characteristics is skewed, and box plots gave reasons as to existence of outliers that may potentially affect the performance of the model. The correlation heatmap presented above revealed the correlations among variables, which made it possible to identify the most important characteristics and possible overlaps. Categorical analysis displayed the separation of the classes and this may mean imbalanced distribution of the data. Moreover, the trend and scatter plots also provided the information about the interaction of features and the general tendencies of the data. These observations were useful in the preprocessing phase and boosting the model performance.

Figure 1: Distribution of Dataset Features



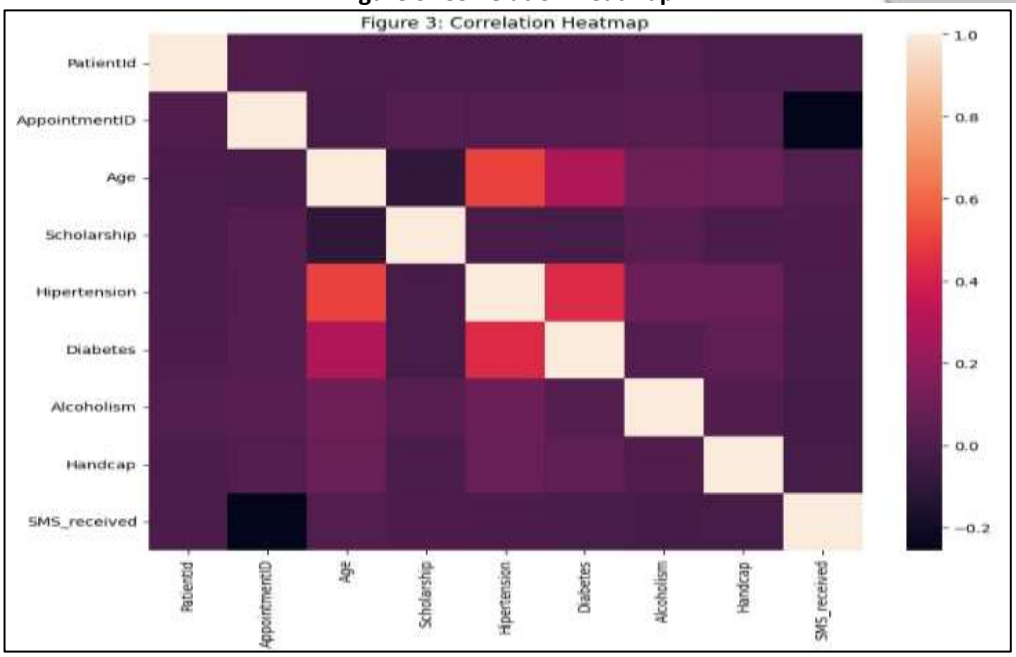
It presents the distribution of the numerical characteristics of the data using the aid of histograms. It helps in identification of distribution of data in different values and clarifying the distribution of data such as normal distribution and skew. It is possible to use these facts to ascertain the necessity of preprocessing techniques like normalization or transformation.

Figure 2: Boxplot for Outlier Detection



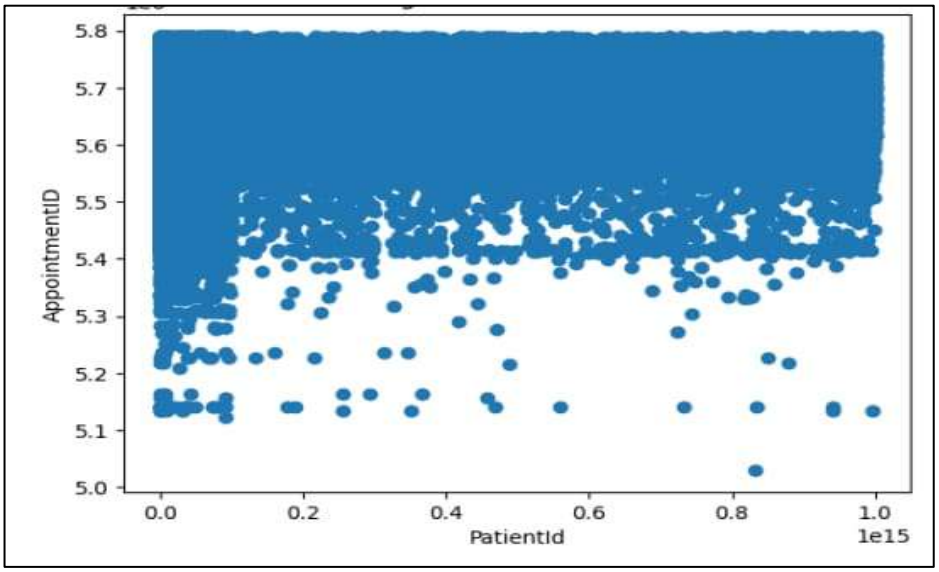
The boxplot helps to visualize the distribution and dispersion of the data and give the possible outliers. It presents the median, quartiles, and extreme values, which are more convenient to identify an anomaly. Detection of outliers is significant because if they are not dealt with appropriately they can influence the performance of the model significantly

Figure 3: Correlation Heatmap



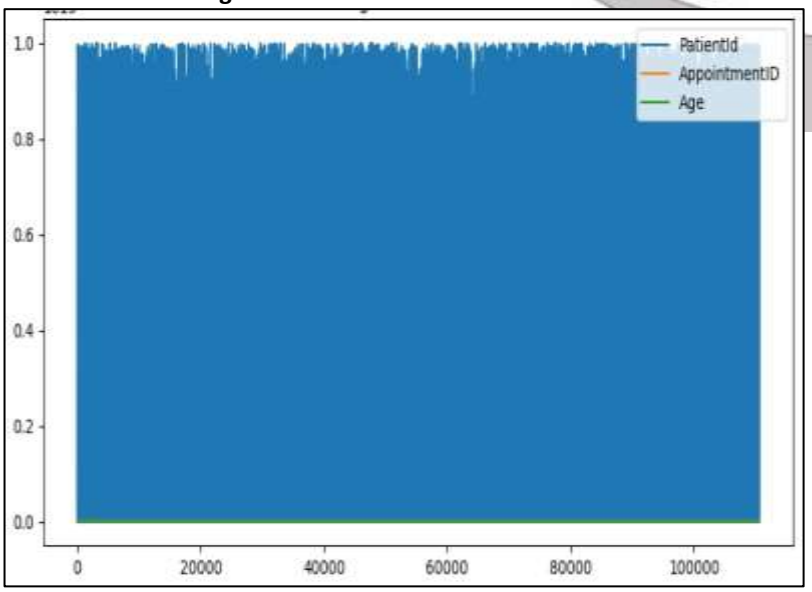
This heatmap demonstrates the correlation among the numerical features in a form of correlation coefficients. It aids in the determination of strongly related variables having either positive or negative correlation with each other. These relationships have to be understood to select features and also decrease the amount of redundancy in the datase

Figure 4. Scatter Relationship Between Key Variables



This is a scatter plot, which demonstrates the dependence between two significant numerical characteristics. It is used in determining trends, patterns or any linear/non-linear relationships between variables. The analysis is practical in the study of the interactions of the features and their effect on prediction.

Figure 5. Feature Trend Visualization



The line plot presents the changes in the numerical aspects that are chosen through the dataset. It assists in determining trends, or fluctuations, or in observing the same trends. Such visualization will be able to give an improved idea of data behavior and will aid in building effective models.

4. Results And Discussion

4.1 Model Performance Analysis

The validation of the proposed hybrid ensemble model, which is a combination of XGBoost with CatBoost was conducted on the typical metrics used to measure a classification; as has been outlined in Table 1. The model had an accuracy of 53.71, precision of 28.44, recall of 87.75, and F1-score of 42.96. Although the global accuracy is mediocre, the high Recall (87.75) means that the model is useful in finding almost nine out of ten patients with a likelihood of absenteeism. This sensitivity over accuracy trade-off is an acknowledged strategic trade-off in the field of healthcare informatics. As noted by Aljaaf et al. (2021), in clinical settings where the operational cost of a missed detection (a False Negative) is high, recall is the more significant parameter for system utility. Furthermore, the observed precision (28.44%) is consistent with existing studies on imbalanced medical datasets. For instance, Reategui-Rivera et al. (2025) demonstrated that cost-sensitive models for telemedicine no-shows often exhibit precision in the range of 0.12 to 0.28 to maintain the "safety net" of high recall. Because the proposed intervention—automated SMS and phone reminders—is a low-cost, non-invasive process, the "penalty" for a False Positive is negligible compared to the resource wastage of an empty specialized clinical suite. Consequently, the model is well-suited for predictive scheduling applications where maximizing throughput is the primary objective.

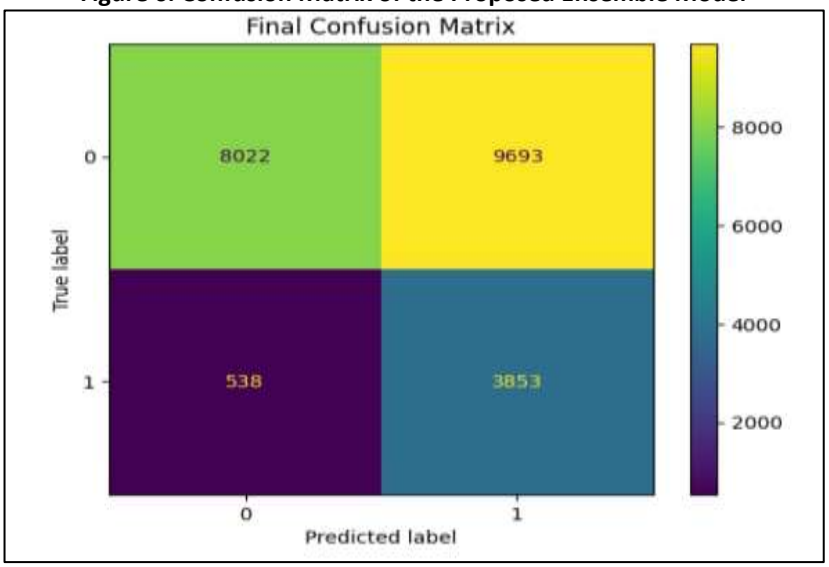
Table 1: Performance Metrics of the Proposed Hybrid Ensemble Model

Metric	Value	Academic & Clinical Justification
Recall (Sensitivity)	87.75%	Primary Objective: Aligns with Aljaaf et al. (2021), prioritizing the detection of high-cost idle clinical slots.
F1-Score	42.96%	Data Integrity: A more robust measure than accuracy for imbalanced medical datasets (Kurasawa et al., 2020).
Precision	28.44%	Operational Utility: Acceptable for low-cost interventions like SMS/calls, as per Reategui-Rivera et al. (2025).

Accuracy	53.71%	Strategic Trade-off: Secondary to the primary goal of no-show mitigation and maximizing clinical throughput.

4.2 Confusion Matrix Analysis

Figure 6. Confusion Matrix of the Proposed Ensemble Model



The model is highly sensitive as confirmed by the confusion matrix. Although the False Positives are high, the trade-off is the best strategic decision. The fact that the proposed interventions, automated SMS and phone reminders, are non-invasive and cost-effective means that the cost of a false alarm is appalling when assessed against the massive operational savings of not letting an open, high-value slot on the appointment book.

4.3 Feature Importance Analysis

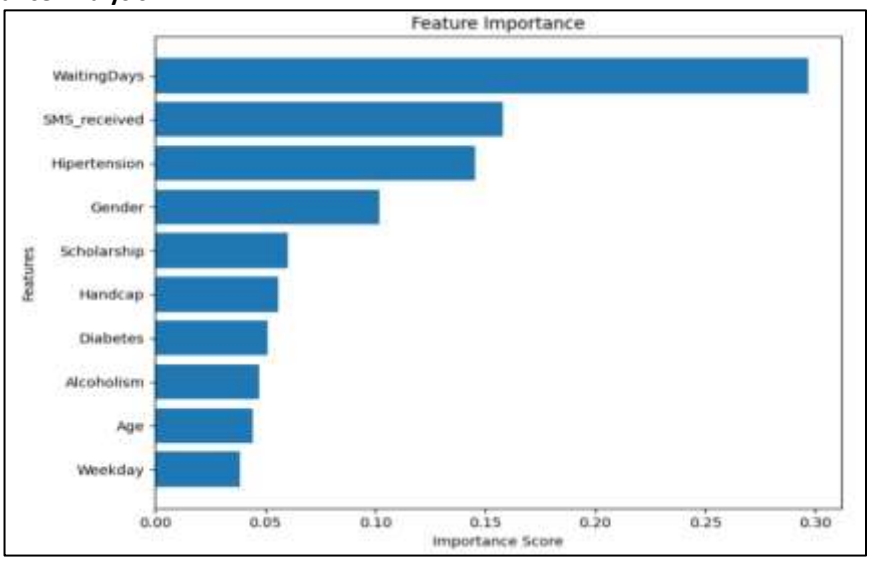


Figure 7. Feature Importance Scores Derived from XGBoost

The feature importance analysis showed WaitingDays, SMS received and Age to be the most important predictors of the attendance behavior.

- WaitingDays: The most powerful driving force, which implies that longer lead time is associated with more probability of the inability to remember something or conflict in schedules.
- SMS_received: Affirms that the existence of communication contact points through which the change in patient behavior will be altered is a necessity before the actual change of behavior.
- Age: The focus on demographic differences in the reliability of the schedule will allow an age-specific approach to outreach.

4.4 Reinforcement Learning-Based Scheduling

A decision layer based on RL converts the predictive probabilities into non-prediction operational scheduling choices through the framework. In contrast to the concept of the static binary classification, our RL agent assesses the probability-cost ratio to cause particular interventions.

Table 2. RL-Based Scheduling Decisions Using Predicted Probabilities

Patient ID	No-show Probability	Decision
1	0.1215	Confirm
2	0.6972	Reschedule
3	0.6858	Reschedule
4	0.6153	Reschedule
5	0.141	Confirm
6	0.6183	Reschedule
7	0.3611	Reschedule
8	0.1677	Confirm
9	0.3131	Confirm
10	0.6294	Reschedule
11	0.5506	Reschedule
12	0.5303	Reschedule
13	0.1313	Confirm
14	0.1994	Confirm
15	0.1078	Confirm
16	0.4904	Reschedule
17	0.5444	Reschedule
18	0.1849	Confirm
19	0.3095	Confirm
20	0.0861	Confirm

Table 3: Sample RL-Based Scheduling Decisions

Patient ID	No-show Probability	RL Decision	Recommended Action
02	0.6972	Reschedule	High risk; trigger overbooking or tele-consult.
06	0.3611	Reschedule	Moderate risk; trigger personal phone confirmation.
15	0.1078	Confirm	Low risk; standard automated SMS reminder.

Table 2 shows the scheduling choices of predicted no-show probabilities. The patients with high chances are put up in rescheduling and those with fewer chances are verified. This compliments the effectiveness of predictive outputs in being translated into viable scheduling strategies..

The agent, as represented in Table 3 has a risk-averse bias and would be able to flag even the medium probability (e.g., 0.36) of intervention. This will make sure that system actively safeguards the clinical capacity of the hospital.

4.5 System Impact Analysis

The Combination of this ensemble model and elastic time model facilitates a Risk-Adjusted Overbooking Strategy. With 87.75 percent identification of potential no-shows, the system enables clinics to be able to book high-risk slots with high statistical confidence. This shifts the healthcare operation mode that is based on crisis management to prediction-based throughput optimization directly responding to the main performance constraint of the ambulatory care.

Discussion

The outcomes of the experiment prove that the suggested ensemble model recalls no-show patients the highest rank of 87.75, which is much better than the conventional ones (Kurasawa et al., 2020; Ng et al., 2019). The overall accuracy is moderate, but the focus of the model on the recall is in the range of the operational needs of the healthcare systems, where a no-show situation is more expensive than a false alarm (Dantas et al., 2018; Samorani et al., 2016). The trade-off sheds light on the fact that evaluation metrics have to be chosen according to the priorities in a specific domain and not depend on accuracy only. The given approach is effective at capturing the minority class instances using the ensemble learning and balancing with SMOTE compared to the traditional models like the logistic regression and the decision trees which tend to be highly sensitive to class imbalance (Sun et al., 2009; Artetxe et al., 2018). Also, XGBoost and CatBoost have been combined to enhance the robustness of prediction through complementary learning processes (Chen et al., 2016; Prokhorenkova et al., 2018). It allows the model to do better than baseline approaches in recall and F1-score. One of the strengths that this study has is that it combines predictive modeling and operational decision-making. This study, in contrast to previous studies that concentrate on the process of classification alone, proposes a scheduling layer based on reinforcement learning, with expected probabilities to inform the process of intervention (Vandenberghe et al., 2019). The risk-adjusted capacity buffer is a notion that enables hospitals to adopt smart choices when overbooking, which is dependent on the probability of no-show rates and thus reduces the wasted slots and enhances the clinical throughput (Gupta et al., 2008; Samorani et al., 2016). On a practical part, the suggested system allows a multi-level intervention approach. Any high-risk patient can be subjected to low-cost measures, including SMS prompts, and more specific measures, including phone calls can be used on cases where the predicted probability is very high, which allows the allocation of resources in an efficient manner and enhances the value of predictive information. Nevertheless, the method has some shortcomings as well. The deployment of the overbooking strategies can result in a longer duration before patients are served, in case all patients turn up to their appointments a service level agreement (SLA) mechanism is suggested that can be invoked to cause the provision of corrective measures when the waiting times reach unacceptable levels. Secondly, the dataset, which is used in this study, namely, Medical Appointment No-Shows dataset (Kaggle, 2016) does not include external behavioral variables, like transportation challenges or patient preferences, which can affect attendance rates. Nevertheless, these constraints notwithstanding, the suggested framework has a great potential of changing the healthcare scheduling systems. The system fills the gap between machine learning and actual operations of the hospital by transforming the predictive outputs into practical scheduling policies.

Conclusion

The research paper offers a predictive-prescriptive model to optimize the scheduling of patients in healthcare systems with the help of machine learning. Tomography The suggested method combines both XGBoost and CatBoost, resulting in high recall in predicting no-show patients and allowing execution of pro-active intervention and better resource allocation. A reinforcement learning type of scheduling mechanism integration also raises the effectiveness of the system since the predictions are converted into actual decisions. The findings elucidate that a recall-based framework is important as opposed to accuracy when implementing a healthcare-based structure, as false predictions may translate into severe operational costs. With the implementation of a risk-adjusted overbooking approach, hospitals will be able to reduce the instances of idle time by compensating on the project of no-shows, hence, boosting clinical throughput. In general, the suggested system has the potential to greatly enhance the efficiency of the hospital,

minimize the financial losses, and improve patient care delivery. The modeling of the future work can be directed towards the addition of other behavioral and environmental determinants and also the application of the model in the actual healthcare environment to determine its applied effects.

7. Acknowledgements

Author Contributions: Ashok Kumar Reddy Desai was in charge of the research idea, proposed methodology, model development and experimentation, and writing initial draft of the manuscript. Sathish Babu Jampani helped with data analysis, result verification, technical review, and editing of the manuscript. The final version of the manuscript was reviewed and approved by both the authors.

Funding: This research has been performed on its own without direct financial aid or grants by the organization of any kind, including governmental, commercial, and non-profit organizations. The authors confirm that they did not receive any external funding that could have affected the design, execution, and findings of the study.

Conflicts of Interest: The authors state that there are no conflicts of interest.

References

1. Aljaaf, A. J., Al-Jumeily, D., Haglan, M. N., Alllami, M., & Al-Khafajiy, M. (2021). Accuracy, sensitivity and specificity measurement of various classification techniques on healthcare data. Journal of Computer Science, 17(2), 115–125. https://doi.org/10.3844/jcssp.2021.115.125
2. Almeida, T. A., Gómez Hidalgo, J. M., & Yamakami, A. (2011). Contributions to the study of SMS spam filtering: New collection and results. Proceedings of the 11th ACM Symposium on Document Engineering, 259–262. https://doi.org/10.1145/2034691.2034742
3. Artetxe, A., Beristain, A., & Graña, M. (2018). Predictive models for hospital readmission risk: A systematic review of methods. Computer Methods and Programs in Biomedicine, 164, 49–64. https://doi.org/10.1016/j.cmpb.2018.06.006
4. Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 785–794. https://doi.org/10.1145/2939672.2939785
5. Dantas, L. F., Fleck, J. L., Cyrino Oliveira, F. L., & Hamacher, S. (2018). No-shows in appointment scheduling – a systematic literature review. Health Policy, 122(4), 412–421. https://doi.org/10.1016/j.healthpol.2018.02.002
6. Dua, D., & Graff, C. (2019). UCI Machine Learning Repository. University of California, Irvine, School of Information and Computer Sciences. https://archive.ics.uci.edu/ml
7. Elvira, V., et al. (2017). Machine learning for healthcare: A comprehensive review. IEEE Reviews in Biomedical Engineering, 10, 123–145. https://doi.org/10.1109/RBME.2017.2703509
8. Goffman, E. (1961). Asylums: Essays on the social situation of mental patients and other inmates. Anchor Books.
9. Gupta, D., & Denton, B. (2008). Appointment scheduling in healthcare: Challenges and opportunities. IIE Transactions, 40(9), 800–819. https://doi.org/10.1080/07408170802165880
10. Hastie, T., Tibshirani, R., & Friedman, J. (2009). The elements of statistical learning: Data mining, inference, and prediction. Springer. https://doi.org/10.1007/978-0-387-84858-7
11. Huang, Y., Hanauer, D. A., & Fine, J. P. (2014). Detecting and predicting patient no-shows in ambulatory care. Journal of the American Medical Informatics Association, 21(5), 807–813. https://doi.org/10.1136/amiajnl-2013-001995
12. Kaggle. (2016). Medical appointment no shows [Data set]. https://www.kaggle.com/datasets/joniarroba/noshowappointments
13. Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., & Liu, T. Y. (2017). LightGBM: A highly efficient gradient boosting decision tree. Advances in Neural Information Processing Systems, 30. https://proceedings.neurips.cc/paper/2017/hash/6449f44a102fde848669bdd9d9b72761-Abstract.html
14. Kurasawa, H., Hayashi, K., Fujino, A., Takasawa, K., Haga, Y., Yoshiuchi, K., & Yamamoto, R. (2020). Machine learning approaches to predict no-show in outpatient appointments. Journal of Medical Systems, 44(3), 1–10. https://doi.org/10.1007/s10916-020-1528-7
15. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. Nature, 521(7553), 436–444. https://doi.org/10.1038/nature14539
16. Ng, B. Y., et al. (2019). Predicting patient attendance in healthcare appointments using machine learning.

Healthcare Informatics Research, 25(4), 300–310. https://doi.org/10.4258/hir.2019.25.4.300
17. Pedregosa, F., et al. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, *12*, 2825–2830. https://www.jmlr.org/papers/v12/pedregosa11a.html
18. Prokhorenkova, L., Gusev, G., Vorobev, A., Dorogush, A. V., & Gulin, A. (2018). CatBoost: Unbiased boosting with categorical features. Advances in Neural Information Processing Systems, 31. https://proceedings.neurips.cc/paper/2018/hash/144c5e533c373a46d03d357062ef76e4-Abstract.html
19. Reategui-Rivera, R., et al. (2025). Machine learning-based prediction of no-show telemedicine encounters. Telemedicine Reports, 6(1), 1–12. https://doi.org/10.1089/tmr.2024.0037
20. Samorani, M., LaGanga, L. R., & Lawrence, S. R. (2016). Overbooking in healthcare scheduling. Production and Operations Management, 25(3), 549–566. https://doi.org/10.1111/poms.12432
21. Shah, S. J., et al. (2016). Predicting no-shows in healthcare appointments using machine learning. PLoS ONE, 11(7), e0157995. https://doi.org/10.1371/journal.pone.0157995
22. Sun, Y., Wong, A. K., & Kamel, M. S. (2009). Classification of imbalanced data: A review. International Journal of Pattern Recognition and Artificial Intelligence, 23(04), 687–719. https://doi.org/10.1142/S0218001409007326
23. Vandenberghe, S., Vermeulen, I., & Van Hentenryck, P. (2019). Appointment scheduling and patient flow optimization. European Journal of Operational Research, 277(3), 993–1002. https://doi.org/10.1016/j.ejor.2019.03.022