Machine Learning in Hospitality: Interpretable Forecasting of Booking Cancellations

Uday kumar pawar
PG
, Department of Computer science and engineering, Sree Dattha institute of Engineering & Science, Hyderabad, Telangana, India - 501510
Email: udaykumarpawar1@gmail.com

Mr Ahmed Pasha
Associate
Professor, Department of Computer science and engineering, Sree Dattha institute of Engineering & Science, Hyderabad, Telangana, India - 501510
Email: hasnainahsan050@gmail.com

Abstract*— Booking cancellations present a critical challenge in the hospitality sector, distorting demand forecasts and causing approximately 20% revenue loss. Accurate cancellation prediction is vital for optimizing hotel operations, pricing strategies, and resource management. The Hotel Booking Demand dataset, containing 32 attributes from resort and city hotels, was utilized, encompassing both categorical and numerical features related to booking behavior, customer type, and reservation details. Preprocessing involved handling missing values, eliminating duplicates, and applying RandomUnderSampler to balance class distribution. Multiple classification algorithms were implemented, including Logistic Regression, Decision Tree, Random Forest, XGBoost, Gradient Boosting, LGBMClassifier, Support Vector Machine (SVM), and Multi-Layer Perceptron (MLP), with GridSearchCV for hyperparameter tuning. A Stacking Classifier combining XGBoost, Random Forest, and Gradient Boosting with Logistic Regression as the meta-model enhanced predictive performance. To further improve accuracy, Label Encoding was applied to categorical features, and Recursive Feature Elimination (RFE) identified key predictors influencing cancellations. A Voting Classifier integrating Decision Tree and MLP models improved accuracy to 98.7%. Additionally, Explainable AI (XAI) methods such as LIME and SHAP were employed to interpret feature importance, while a Flask-based user interface enabled interactive, real-time prediction for effective decision support in hotel management systems.

"Keywords— Cancellation forecasting, hotel booking, artificial intelligence, machine learning, revenue management, explainable artificial intelligence".

I. Introduction

Tourism serves as a major engine of global economic growth, fostering employment, infrastructure development, and cultural exchange. According to the World Tourism Organization (UNWTO), international tourist arrivals reached approximately 1.3 billion in 2023—recovering to 88% of pre-pandemic levels—and are expected to surpass these levels by 2024 [1]. Beyond direct visitor spending, tourism stimulates multiple industries and strengthens long-term economic resilience through sustained demand for hospitality services [2]. Within this ecosystem, hotels play a pivotal role as the primary interface between service providers and consumers. However, the perishable nature of hotel inventory—where unsold rooms equate to lost revenue—makes accurate demand forecasting essential for maintaining operational efficiency and profitability [3].

Despite its strategic importance, forecasting customer behavior in hospitality remains complex due to volatile external influences such as geopolitical instability, environmental disruptions, and shifting traveler preferences [4]. The increasing digitalization of bookings through online travel platforms has further intensified this challenge by enabling flexible reservations and easy cancellations. Consequently, cancellation rates now range between 10% and 30%, leading to revenue losses of up to 20% and complicating operational planning [5], [6]. While extensive studies have examined demand forecasting, relatively few have focused on predicting booking cancellations at the individual level or explaining the behavioral factors that drive such actions [7]. This research gap limits hotel managers' ability to proactively address uncertainties and optimize revenue strategies.

To mitigate these issues, this study proposes an interpretable machine learning (ML) framework that balances predictive performance with explainability. Although Artificial Intelligence (AI) and ML have revolutionized the hospitality sector—enhancing personalization, automating services, and optimizing operations—most existing models lack interpretability, especially in predicting cancellation behavior [8]. Understanding the rationale behind a model's prediction is critical for building managerial trust and enabling data-driven decisions [9]. Therefore, this research emphasizes model transparency to provide actionable insights into cancellation drivers while maintaining computational efficiency and adaptability.

The study makes three key contributions. First, it advances hotel demand forecasting by developing an interpretable framework for customer-level cancellation prediction. Second, it introduces a computationally efficient model that achieves high accuracy with minimal features, ensuring scalability across varied operational contexts. Third, it enhances managerial decision-making through interpretable outputs that clarify causal factors influencing cancellations. Collectively, these contributions promote the adoption of transparent, AI-driven analytics in hospitality revenue management, enabling resilient and data-informed operations that sustain competitive advantage in a dynamic global tourism market [10].

II. Related Work

Recent advances in data-driven analytics have significantly improved hotel booking cancellation prediction; however, challenges persist regarding interpretability and generalizability. Herrera et al. [11] developed a machine learning-based framework using ensemble models for cancellation forecasting, achieving strong predictive accuracy but lacking transparency in feature-level explanations. Hikmawati et al. [12] employed data mining techniques to explore patterns in Average Daily Rate (ADR), cancellations, and booking behavior, highlighting the role of pricing trends but offering limited operational applicability due to the absence of a unified predictive framework.

Artificial Intelligence (AI) applications in hospitality have gained traction, with Limna [13] providing a comprehensive overview of AI's transformative impact on service automation and customer engagement. However, this study addressed AI broadly without focusing on predictive modeling challenges in cancellation management. Sekhon and Ahuja [14] reviewed various machine learning techniques for managing cancellations, identifying ensemble models such as Random Forest and Gradient Boosting as effective. Despite their success, these models often lacked interpretability, hindering managerial understanding of predictions. Cai et al. [15] explored customer perceptions of AI-enabled systems, emphasizing that trust and brand association influence technology acceptance—reinforcing the importance of explainable and reliable models in fostering user trust.

In related areas of hospitality analytics, Gómez-Talal et al. [16] utilized big data to minimize restaurant food waste through predictive management tools, demonstrating the potential of analytics for operational optimization. Within hotel cancellations, Chen et al. [17] integrated probabilistic feature interactions with machine learning to enhance transparency in predictions, yet their approach required extensive feature engineering, limiting adaptability. Nababan et al. [18] applied K-nearest neighbors (K-NN) with SMOTE to handle class imbalance, improving recall but sacrificing precision, indicating a trade-off between accuracy and computational efficiency.

Beyond hospitality, ensemble and meta-learning strategies have shown robust predictive capabilities. Barton and Lennox [19] demonstrated that model stacking enhances variable importance robustness in industrial systems, while Sahu et al. [20] proved that stacking classifiers outperform individual models in medical diagnostics, suggesting cross-domain applicability to hotel forecasting. Collectively, these studies underscore ensemble learning's strength in boosting accuracy and stability, though its integration with interpretability remains limited in hotel cancellation prediction.

Overall, prior research reveals two key gaps: (1) insufficient interpretability in existing models and (2) limited development of adaptive, lightweight frameworks balancing accuracy and transparency. Addressing these challenges, the current study proposes an interpretable ensemble-based model that combines predictive strength with explainability, offering a transparent, scalable, and operationally efficient solution for hotel booking cancellation forecasting.

III. Materials And Methods

The proposed system aims to predict hotel booking cancellations with high accuracy and interpretability using advanced machine learning and ensemble-based frameworks. The Hotel Booking Demand dataset, comprising 32 categorical and numerical attributes from city and resort hotels, serves as the foundation for modeling customer behavior and reservation dynamics. The workflow encompasses systematic preprocessing, including data cleaning, label encoding, and class balancing through RandomUnderSampler to ensure fair model training. Feature relevance is enhanced using Recursive Feature Elimination (RFE) to identify the most influential predictors of cancellations. A Stacking Classifier integrating XGBoost, Random Forest, and Gradient Boosting as base learners with Logistic Regression as the meta-model is implemented to improve generalization and robustness. Additionally, a Voting Classifier combining Decision Tree and Multi-Layer Perceptron (MLP) further strengthens predictive performance. Model transparency is achieved using SHAP and LIME for explainable AI (XAI)-based interpretation, while a Flask-based interface enables real-time deployment, facilitating proactive decision-making and improved revenue management in hotel operations.

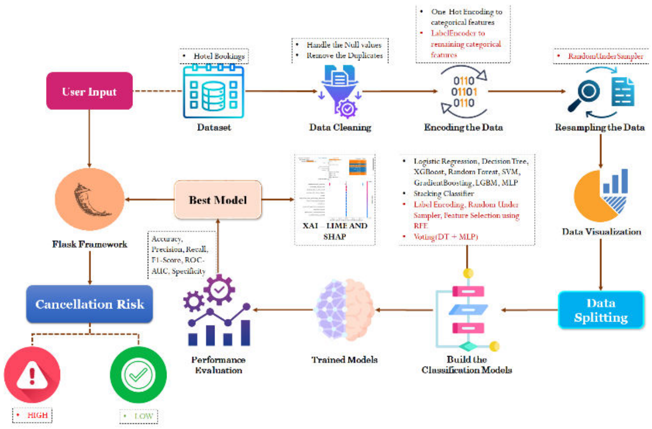

Fig. 1. System Architecture

The system architecture illustrates a machine learning framework for predicting hotel booking cancellations. It begins with booking data undergoing preprocessing, splitting, and stratified 5-fold cross-validation with random undersampling to balance classes. Multiple models, including Logistic Regression, SVM, Decision Tree, Random Forest, Gradient Boosting, XGBoost, LightGBM, and MLP, are trained and optimized through hyperparameter tuning. The best-performing models are integrated into a stacking ensemble with meta-model selection. Performance is evaluated using accuracy, recall, precision, F1-score, specificity, and AUC, with SHAP ensuring interpretability and transparency.

i) Data Set Collection

The Hotel Booking Demand dataset, obtained from the Kaggle public repository, comprises 119,390 samples with 32 categorical and numerical attributes representing customer demographics, reservation details, and booking behavior for both resort and city hotels. The target variable, is_canceled, indicates whether a reservation was canceled or honored. The dataset includes diverse features such as lead time, booking channels, room types, and customer types, reflecting real-world variability in hotel operations. Minor data imbalance and missing values in attributes like agent, company, and country enhance the dataset's complexity and analytical relevance. Its large scale, heterogeneity, and temporal richness make it highly suitable for developing robust, generalizable machine learning models for cancellation prediction and demand forecasting in the hospitality sector.

	hotel	is_canceled	lead_time	arrival_date_year	arrival_date_month	arrival_date_week_number	arrival_date_day_of_month	stays_in_weekend_nights	stays_in_week_nights	adults
0	Resort Hotel	0	342	2015	July	27	1	0	0	2
1	Resort Hotel	0	737	2015	July	27	1	0	0	2
2	Resort Hotel	0	7	2015	July	27	1	0	1	1
3	Resort Hotel	0	13	2015	July	27	1	0	1	1
4	Resort Hotel	0	14	2015	July	27	1	0	2	2

5 rows × 32 columns

Fig. 2. Hotel Booking Demand Dataset

ii) Pre-Processing

The preprocessing phase ensures data quality and model readiness through systematic cleaning, encoding, resampling, feature selection, visualization, and partitioning, thereby enhancing the reliability and performance of subsequent predictive modeling.

a) Data Cleaning: The dataset underwent a comprehensive data cleaning process to ensure accuracy and consistency. Missing values in key attributes such as country, agent, company, and children were imputed with appropriate substitutes, while duplicate records were identified and removed to eliminate redundancy. This step was critical to maintaining data integrity and preventing biases in model learning. Clean and complete data enhanced the reliability of subsequent feature extraction, resampling, and model training stages.

b) Data Encoding: To transform categorical variables into a numerical format suitable for machine learning algorithms, a hybrid encoding approach was applied. One-hot encoding was utilized for multi-class categorical attributes, while label encoding was applied to ordinal or binary categorical variables. This transformation ensured that all features were represented quantitatively without introducing ordinal bias. Encoding facilitated algorithmic interpretation of categorical information, thereby improving model convergence, interpretability, and predictive accuracy across diverse data attributes.

c) Data Resampling: To address the imbalance between canceled and non-canceled bookings, the dataset was balanced using a random under-sampling strategy. This technique reduced the size of the majority class, ensuring an equal representation of both classes. By achieving class balance, the model was prevented from developing a bias toward the dominant class, leading to improved fairness and generalization. The resampling process was essential for enhancing classification performance and maintaining robustness in cancellation prediction.

d) Feature Selection: Feature selection was implemented using a recursive elimination-based approach to identify the most influential predictors of booking cancellation. The process iteratively evaluated and retained twelve key attributes that contributed significantly to the model's predictive capability. By reducing redundant and irrelevant features, the dimensionality of the data was minimized, improving computational efficiency and model interpretability. This step ensured that the final learning model focused on the most impactful and discriminative factors.

e) Data Visualization: Data visualization techniques were employed to provide an intuitive understanding of booking cancellation patterns and dataset balance. Graphical representations, such as pie and bar charts, were used to illustrate the proportion of canceled versus non-canceled bookings both before and after resampling. These visual insights enabled the identification of class imbalance and the verification of sampling effectiveness. Visualization facilitated better comprehension of data distribution, supported informed preprocessing decisions, and enhanced interpretability for subsequent analytical modeling.

iii) Train & Test:

The processed dataset was divided into training and testing subsets using an 80:20 ratio while preserving class distribution through stratified sampling. This ensured that both subsets maintained proportional representation of canceled and non-canceled bookings. The training set facilitated model learning, whereas the testing set provided an unbiased evaluation of predictive performance. This structured data partitioning was vital for assessing model generalization, minimizing overfitting, and validating the system's reliability on unseen data.

iv) Integration of XAI & Flask Framework:

Explainable Artificial Intelligence (XAI) techniques were incorporated to enhance the interpretability and transparency of the hotel booking cancellation prediction model. LIME (Local Interpretable Model-Agnostic Explanations) was employed to generate local explanations by analyzing feature contributions for individual predictions, providing insight into how each attribute influences the model's output. Additionally, SHAP (SHapley Additive exPlanations) was used to compute global feature importance, enabling a deeper understanding of the model's behavior across multiple predictions. These methods ensure that decision-making remains interpretable and trustworthy for hotel management.

The predictive system was deployed through a Flask-based web interface, allowing users to interact with the trained model in real-time. This integration enables seamless data input, instant prediction of booking cancellation probabilities, and immediate visualization of explanatory insights generated by LIME and SHAP. Flask serves as a lightweight, efficient deployment framework that bridges the machine learning backend and end-users, providing a user-friendly environment for operational decision support.

v) Algorithms:

Logistic Regression is a statistical classification algorithm that models the probability of a binary outcome using a logistic function. It maps input features to probabilities between 0 and 1, effectively distinguishing between classes through a linear decision boundary. The algorithm's interpretability, computational efficiency, and resistance to overfitting make it suitable for baseline performance assessment. Its ability to estimate feature importance enhances model transparency, contributing to a clear understanding of decision-making processes in predictive analytics.

$$P(y = 1 \mid X) = \frac{1}{1 + e^{-(W^T x + b)}} \quad (1)$$

The Decision Tree algorithm classifies data by recursively partitioning the feature space into homogeneous subsets based on the most significant attributes. Each internal node represents a decision rule, while leaf nodes correspond to class outcomes. This hierarchical structure enables intuitive interpretation and handles both numerical and categorical data effectively. Its non-parametric nature captures nonlinear relationships, improving classification accuracy while maintaining transparency and ease of visualization in decision-making tasks.

Extreme Gradient Boosting (XGBoost) is an ensemble learning technique that constructs a series of decision trees in a sequential manner, where each subsequent tree corrects the errors of its predecessors. It employs gradient-based optimization to minimize loss functions efficiently, ensuring superior predictive performance. XGBoost integrates regularization mechanisms to prevent overfitting and leverages parallel processing for computational efficiency, making it highly robust and accurate for large-scale, high-dimensional datasets.

Random Forest is an ensemble classifier that combines multiple decision trees to improve prediction stability and accuracy. By introducing randomness in both feature selection and data sampling, it reduces overfitting and enhances generalization across unseen data. Each tree contributes a vote toward the final classification, ensuring robustness against noise and outliers. Its capability to handle diverse data types and assess feature importance makes it a widely adopted model in complex classification tasks.

$$Gini = 1 - \sum_{i=1}^{C} (Pi)^2 \quad (2)$$

Support Vector Machine (SVM) is a powerful supervised learning algorithm that constructs an optimal hyperplane to separate classes with the maximum margin. It transforms input data into higher-dimensional spaces using kernel functions, enabling effective classification of nonlinear patterns. SVM's margin-based optimization minimizes generalization error, ensuring strong performance in both balanced and imbalanced datasets. Its robustness and adaptability to complex feature distributions make it a reliable choice for high-dimensional classification problems.

Gradient Boosting is an ensemble method that builds models sequentially, where each new model minimizes the residual errors of the previous ones through gradient descent optimization. It combines weak learners, typically decision trees, to form a strong predictive model. By focusing on difficult-to-predict instances, Gradient Boosting achieves high accuracy and flexibility. Its iterative refinement process enhances model robustness and ensures superior performance on both structured and unstructured data.

$$F_m(x) = F_{m-1}(x) + \gamma_m h_m(x) \quad (3)$$

Light Gradient Boosting Machine (LGBM Classifier) is an advanced gradient boosting framework optimized for efficiency and scalability. It utilizes a leaf-wise tree growth strategy with depth constraints, resulting in faster convergence and improved accuracy. LGBM effectively handles large datasets with high-dimensional features and supports categorical variables natively. Its lightweight design, reduced memory consumption, and high prediction speed make it particularly suitable for real-time classification applications requiring both accuracy and computational efficiency.

Multilayer Perceptron (MLP) is a type of feedforward neural network that maps input data to output predictions through interconnected layers of neurons. It employs nonlinear activation functions and backpropagation-based weight optimization to learn complex feature representations. MLP effectively captures intricate relationships within data, offering strong generalization and adaptability. Its deep learning architecture enables high predictive accuracy, making it suitable for modeling nonlinear and multidimensional classification problems.

Stacking Classifier is an ensemble learning strategy that combines predictions from multiple base classifiers using a meta-model to enhance overall predictive accuracy. In this approach, diverse base estimators such as XGBoost, Random Forest, and Gradient Boosting generate primary predictions, which are then integrated by a Logistic Regression meta-model. This hierarchical structure exploits the strengths of individual models while compensating for their weaknesses, leading to improved generalization and robustness in classification tasks.

$$\hat{y} = g(Y_{base}) = g\big(f_1(x), f_2(x), \ldots, f_m(x)\big) \quad (5)$$

The Voting Classifier serves as an ensemble learning method that consolidates predictions from multiple base models—specifically the Decision Tree and Multilayer Perceptron—to achieve enhanced classification performance. For this classifier, several preprocessing extensions were incorporated to optimize accuracy and generalization. Label Encoding was applied to convert categorical features into numerical form, Random Under-Sampling was used to balance class distributions, and Recursive Feature Elimination (RFE) selected the most relevant attributes. These enhancements collectively improved the Voting Classifier's efficiency, robustness, and interpretability across diverse data conditions.

$$\hat{y} = argmax_c \left(\sum_{i=1}^{n} II\ (\hat{y}_i = c) \right) \quad (6)$$

IV. Experimental Results

Accuracy: The accuracy of a test is its ability to differentiate the patient and healthy cases correctly. To estimate the accuracy of a test, we should calculate the proportion of true positive and true negative in all evaluated cases. Mathematically, this can be stated as:

$$Accuracy = \frac{\text{TP} + \text{TN}}{\text{TP} + \text{FP} + \text{TN} + \text{FN}} \quad (7)$$

Precision: Precision evaluates the fraction of correctly classified instances or samples among the ones classified as positives. Thus, the formula to calculate the precision is given by:

$$Precision = \frac{\text{True Positive}}{\text{True Positive} + \text{False Positive}} \quad (8)$$

Recall: Recall is a metric in machine learning that measures the ability of a model to identify all relevant instances of a particular class. It is the ratio of correctly predicted positive observations to the total actual positives, providing insights into a model's completeness in capturing instances of a given class.

$$Recall = \frac{\mathrm{TP}}{\mathrm{TP} + \mathrm{FN}} \quad (9)$$

F1-Score: F1 score is a machine learning evaluation metric that measures a model's accuracy. It combines the precision and recall scores of a model. The accuracy metric computes how many times a model made a correct prediction across the entire dataset.

$$F1\ Score = 2 * \frac{Recall\ \mathrm{X}\ Precision}{Recall + Precision} * 100 \quad (10)$$

AUC-ROC Curve: The AUC-ROC Curve is a performance measurement for classification problems at various threshold settings. ROC plots the True Positive Rate against the False Positive Rate. AUC quantifies the overall ability of the model to distinguish between classes, where a higher AUC indicates better model performance.

$$AUC = \sum_{i=1}^{n-1} (FPR_{i+1} - FPR_i) \cdot \frac{TPR_{i+1} + TPR_i}{2} \quad (11)$$

Specificity: It is calculated by identifying the number of people who test negative for a condition and dividing it by the total number of people without the condition, which includes those who tested negative and the false positives—those who tested positive but did not have the disease.

$$Specificity = \frac{TN}{(TN + FP)} \quad (12)$$

Table. 1. Performance Evaluation

ML Model	Accuracy	Precision	Recall	F1-Score	ROC-AUC	Specificity
LR	0.978	0.956	0.999	0.977	0.956	0.999
DT	0.909	0.920	0.901	0.910	0.953	0.899
XGB	0.978	0.957	1.000	0.978	0.978	1.000
RF	0.960	0.929	0.991	0.959	0.983	0.992
SVM	0.934	0.868	1.000	0.929	0.956	1.000
GB	0.931	0.935	0.928	0.932	0.931	0.928
LGBM	0.855	0.843	0.865	0.853	0.929	0.868
MLP	0.946	0.937	0.955	0.946	0.961	0.956
Stacking	0.952	0.905	1.000	0.950	1.000	1.000
Voting	0.987	0.973	1.000	0.987	1.000	1.000

The Table.1 presents performance metrics of various classification models, showing that the extended Voting Classifier achieved the highest overall predictive accuracy and robustness.

Fig. 3. Comparison Graph

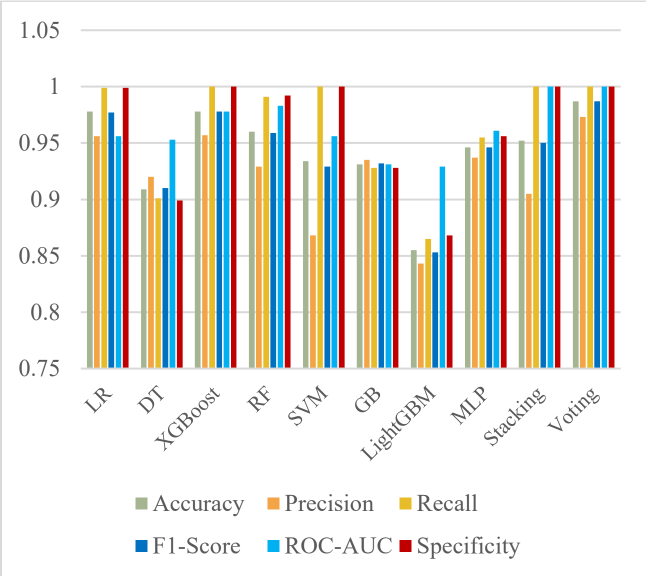

The comparative performance graph.1 illustrates that the Extension Voting model achieves superior accuracy, recall, and F1-score, outperforming all other machine learning classifiers consistently.

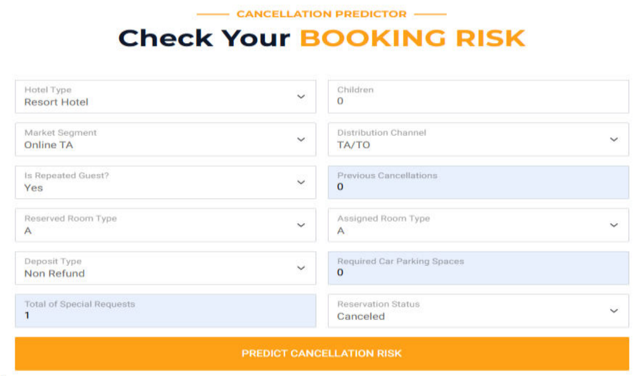

Fig.4 Enter Input Data – 1

In Fig.4, the user enters booking-related input data through the web interface to predict potential hotel booking cancellations.

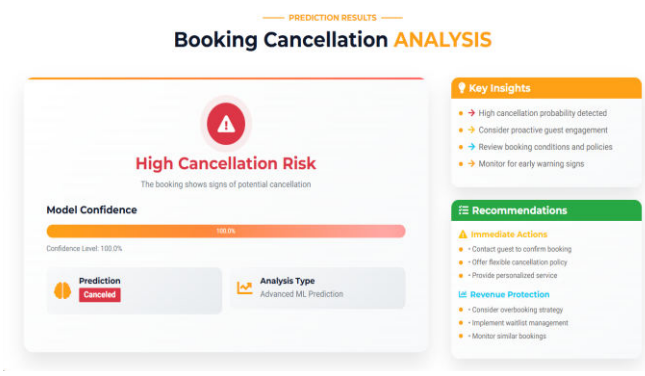

Fig.5 Booking Cancellation Analysis – 1

In Fig.5, the system displays the prediction result for the provided input values, indicating a "High Cancellation Risk."

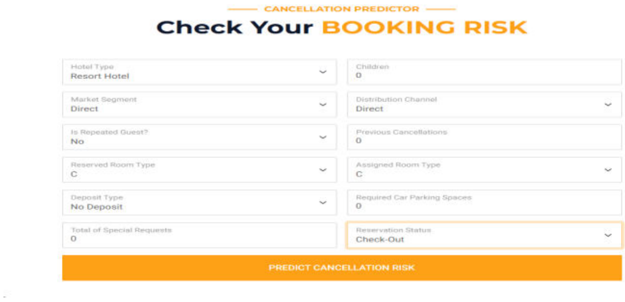

Fig.6 Enter Input Data – 2

In Fig.6, the user provides booking details through the interactive interface to analyze and predict possible hotel booking cancellations.

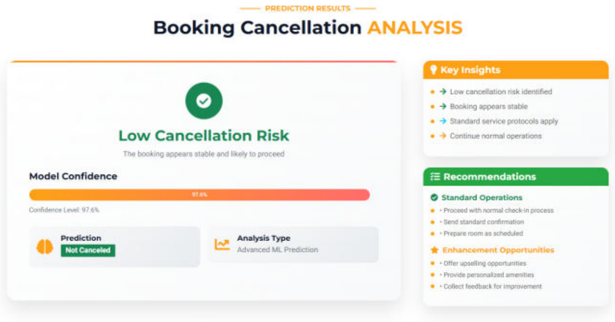

Fig.7 Booking Cancellation Analysis – 2

In Fig.7, the system generates the prediction output for the entered details, indicating a "Low Cancellation Risk."

V. CONCLUSION

In conclusion, the developed system effectively addresses the challenge of hotel booking cancellations using the Hotel Booking Demand dataset, which comprises 32 numerical and categorical features from both resort and city hotels. Comprehensive preprocessing, including handling missing values, removing duplicates, and applying RandomUnderSampler, ensured data quality and balance. Multiple machine learning algorithms—Logistic Regression, Decision Tree, Random Forest, XGBoost, Gradient Boosting, LGBMClassifier, Support Vector Machine (SVM), and Multi-Layer Perceptron (MLP)—were implemented and fine-tuned through GridSearchCV for optimal hyperparameter selection. The Stacking Classifier, integrating XGBoost, Random Forest, and Gradient Boosting with Logistic Regression as the meta-model, demonstrated strong predictive capability. To further enhance performance, Label Encoding was used for categorical features, and Recursive Feature Elimination (RFE) identified the most influential predictors affecting cancellations. The Voting Classifier, combining Decision Tree and MLP, achieved a superior accuracy of 98.7%, highlighting the effectiveness of ensemble learning. Additionally, Explainable AI (XAI) tools, including LIME and SHAP, provided interpretability by illustrating feature importance, while the Flask-based interface enabled real-time prediction and user interaction, ensuring practical applicability for efficient and data-driven hotel management decision-making.

The future scope of this work lies in expanding the predictive framework through integration with real-time booking systems and dynamic pricing platforms to enhance decision automation in hotel management. Incorporating additional behavioral, economic, and seasonal variables can further improve model adaptability to market fluctuations. Future developments may explore hybrid deep learning architectures for temporal pattern recognition and reinforcement learning for strategic cancellation mitigation. Moreover, deploying the model within cloud-based environments can enable scalable, continuous learning from live data streams. Enhancing interpretability and user interaction through advanced visualization tools will further strengthen its applicability in data-driven hospitality operations.

REFERENCES

[1] Kumar, A., Prasad, U., & Kumar Tiwari, R. (2025). Analysing the Impact of Exogenous Determinants on Booking Cancellations With the GNN. In AI and Green Technology Applications in Society (pp. 313-352). IGI Global Scientific Publishing.

[2] Luo, Z. (2025, August). Hotel Cancellation Rate Prediction: A Machine Learning Based Prediction Model. In 2025 3rd International Conference on Image, Algorithms, and Artificial Intelligence (ICIAAI 2025) (pp. 318-327). Atlantis Press.

[3] Pang, Z. (2025, August). Enhancing Demand Forecasting with Machine Learning: Insights from Hotel Booking Data. In 2025 13th International Conference on Traffic and Logistic Engineering (ICTLE) (pp. 163-167). IEEE.

[4] Alkan, T. A. H. A. (2025). A Comparative Study of Machine Learning and Deep Learning Approaches For Hotel Booking Cancellation Prediction. International Journal Of Scientific Research In Engineering & Technology, 5(03).

[5] Bhardwaj, A., Yadav, T., & Chaudhary, R. (2024, June). Predicting Hotel Booking Cancellations using Machine Learning Techniques. In 2024 15th International Conference on Computing Communication and Networking Technologies (ICCCNT) (pp. 1-6). IEEE.

[6] D. Gursoy and R. Cai, "Artificial intelligence: An overview of research trends and future directions," Int. J. Contemp. Hospitality Manage., vol. 37, no. 1, pp. 1–17, Jan. 2025.

[7] Y. Y. Febrian, D. R. Wijaya, and E. Ervina, "Hotel reservation cancellation prediction using boosting model," Proc. 2nd Int. Conf. Softw. Eng. Inf. Technol. (ICoSEIT), Feb. 2024, pp. 138–143.

[8] M. Yoo, A. K. Singh, and N. Loewy, "Predicting hotel booking cancelation with machine learning techniques," J. Hospitality Tourism Technol., vol. 15, no. 1, pp. 54–69, Jan. 2024.

[9] H. Lv, S. Shi, and D. Gursoy, "A look back and a leap forward: A review and synthesis of big data and AI literature in hospitality and tourism," J. Hospitality Marketing Manage., vol. 31, no. 2, pp. 145–175, Feb. 2022.

[10] Z. Liu, L. Zhang, W. Wang, and S. Chen, "Hospitality order cancellation prediction from a profit-driven perspective," Int. J. Contemp. Hospitality Manage., vol. 35, no. 6, pp. 2084–2112, May 2023.

[11] A. Herrera, P. Morales, and L. Torres, "Forecasting hotel cancellations through machine learning," Expert Syst., vol. 41, no. 9, p. 13608, Sep. 2024.

[12] N. K. Hikmawati, Y. Ramdhani, and W. Wartika, "Exploring ADR trends: A data mining approach to hotel room pricing, cancellations, and EDA," J. Appl. Data Sci., vol. 5, no. 1, pp. 189–202, 2024.

[13] P. Limna, "Artificial intelligence (AI) in the hospitality industry: A review article," Int. J. Comput. Sci. Res., vol. 7, pp. 1306–1317, Jan. 2023.

[14] G. Sekhon and S. Ahuja, "Review of machine learning models for managing hotel cancellations in the tourism industry," Proc. 3rd Int. Conf. Intell. Technol. (CONIT), Jun. 2023, pp. 1–6.

[15] R. Cai, L. N. Cain, and H. Jeon, "Customers' perceptions of hotel AI-enabled voice assistants: Does brand matter?," Int. J. Contemp. Hospitality Manage., vol. 34, no. 8, pp. 2807–2831, Jul. 2022.

[16] J. Gómez-Talal, M. Alvarez, and C. Lopez, "Avoiding food waste from restaurant tickets: A big data management tool," J. Hospitality Tourism Technol., vol. 15, no. 2, pp. 232–253, Mar. 2024.

[17] S. Chen, X. Liu, J. Ma, and H. Zhang, "Prediction of hotel booking cancellations: Integration of machine learning and probability model based on interpretable feature interaction," Decis. Support Syst., vol. 170, Jul. 2023, Art. no. 113959.

[18] A. A. Nababan, M. Jannah, and A. H. Nababan, "Prediction of hotel booking cancellation using K-nearest neighbors (K-NN) and SMOTE," INFOKUM, vol. 10, no. 3, pp. 50–56, 2022.

[19] M. Barton and B. Lennox, "Model stacking to improve prediction and variable importance robustness for soft sensor development," Digit. Chem. Eng., vol. 3, Jun. 2022, Art. no. 100034.

[20] B. Sahu, R. Pradhan, and A. Singh, "Performance analysis of state-of-the-art classifiers and stack ensemble model for liver disease diagnosis," Springer, 2022.

[21] A. Chatzimparmpas, R. Martins, and A. Kerren, "StackGenVis: Alignment of data, algorithms, and models for stacking ensemble learning using performance metrics," IEEE Trans. Vis. Comput. Graphics, vol. 27, no. 2, pp. 1547–1557, 2020.

[22] H. Duan, "Computational methods for estimating global sensitivity indices Shapley values," Ph.D. dissertation, Florida State Univ., 2023.

[23] S. M. Lundberg, G. Erion, and S. Lee, "From local explanations to global understanding with explainable AI for trees," Nature Mach. Intell., vol. 2, no. 1, pp. 56–67, Jan. 2020.

[24] A. Rai, "Explainable AI: From black box to glass box," J. Acad. Marketing Sci., vol. 48, no. 1, pp. 137–141, Jan. 2020.

[25] Z. A. Andriawan, T. W. Purboyo, and M. Z. Arifin, "Prediction of hotel booking cancellation using CRISP-DM," Proc. 4th Int. Conf. Informat. Comput. Sci. (ICICoS), Nov. 2020, pp. 1–6.