SJIF Impact Factor (2026): 8.863| ISI I.F. Value: 1.241| Journal DOI: 10.36713/epra2016 ISSN: 2455-7838(Online)
EPRA International Journal of Research and Development (IJRD)
Volume: 11 | Issue: 5 | May 2026 - Peer Reviewed Journal

AIHMS: SMART HOTEL RESERVATION AND MANAGEMENT USING AI AND NLP

R. Ramakrishnan[1], S. Ramyasakthi[2]
[1]*Associate Professor, Department of MCA, Sri Manakula Vinayagar Engineering College (Autonomous), Puducherry- India*
[2]*PG-Student, Department of MCA, Sri Manakula Vinayagar Engineering College (Autonomous), Puducherry- India*

ABSTRACT
The hospitality industry increasingly relies on data-driven decision-making, yet existing hotel management systems often employ opaque, rule-based algorithms that fail to adapt to dynamic customer behaviors and market conditions. We propose an Explainable AI Framework for Intelligent Hotel Management that integrates three core components: a hybrid recommendation engine, a dynamic pricing model, and an NLP-driven sentiment analysis module. The recommendation system combines collaborative filtering with content-based analysis through a Transformer-encoder architecture, generating personalized room suggestions that evolve with user interactions. A Gradient Boosting Decision Tree enhanced with temporal features optimizes room rates in real time by balancing revenue maximization against occupancy probability, thereby replacing static tariff structures. The sentiment analysis module fine-tunes a pre-trained BERT model on hospitality-specific data to extract sentiment polarity and complaint severity, automatically escalating high-severity issues to senior management. Crucially, the framework incorporates Shapley Additive exPlanations to render each algorithmic output interpretable; feature contribution values are computed for every recommendation, price decision, and classification result. These explanations are injected into managerial dashboards and user interfaces, transforming complex model inferences into transparent narratives. The proposed system therefore bridges the gap between advanced machine learning and human-centric operational oversight. Its novelty lies in the seamless integration of explainability into a multi-modal pipeline that addresses personalization, pricing, and customer engagement simultaneously. This work demonstrates that high-performance AI can be both powerful and interpretable, offering a practical pathway for hotels to enhance operational efficiency while maintaining stakeholder trust.

KEYWORDS: *Artificial Intelligence, Smart Hotel Management, Explainable AI, Predictive Analytics, Natural Language Processing, Dynamic Pricing, Recommendation System.*

1. INTRODUCTION

The hospitality industry has undergone a profound transformation over the past decade, shifting from traditional property management systems toward increasingly sophisticated digital ecosystems that promise personalized guest experiences and optimized operational efficiency [1]. Early reservation platforms primarily served as static booking interfaces, but the proliferation of online travel agencies and customer review aggregators has created an environment where hotels must process vast quantities of heterogeneous data to remain competitive [2]. Contemporary hotel management systems now incorporate elements of collaborative filtering [3] and matrix factorization [4] to generate recommendations, yet these approaches frequently operate as black boxes, offering little transparency into how decisions are derived. This opacity poses significant challenges for managerial trust and regulatory compliance, particularly when automated systems influence pricing strategies or staff evaluation processes.

The proposed research introduces an Explainable Artificial Intelligence-based Hotel Management System (AIHMS) designed to transcend these limitations by embedding intelligent automation and adaptive decision-making capabilities within a transparent framework. The core novelty of this method lies in the synergistic integration of hybrid recommendation engines, predictive analytics, and sentiment analysis within a unified architecture that prioritizes interpretability through Explainable AI (XAI) mechanisms. Unlike conventional systems that rely on opaque deep learning pipelines, this approach explicitly incorporates transparency modules to elucidate the rationale behind dynamic pricing adjustments and personalized room recommendations. The method leverages Natural Language Processing (NLP) to decode customer intent and complaint severity, building upon foundational work in latent semantic analysis [5] and more recent transformer-based architectures such as BERT [6] for understanding guest feedback. Simultaneously, regression models including ridge and lasso regression [7] are employed to forecast operational demands, thereby facilitating data-driven strategies for staff evaluation and revenue optimization without compromising user trust or managerial reliability.

SJIF Impact Factor (2026): 8.863| ISI I.F. Value: 1.241| Journal DOI: 10.36713/epra2016 ISSN: 2455-7838(Online)
EPRA International Journal of Research and Development (IJRD)
Volume: 11 | Issue: 5 | May 2026 - Peer Reviewed Journal

The primary contribution of this work is threefold. First, we propose a hybrid recommendation engine that combines content-based filtering with collaborative approaches to address cold-start problems [8], while incorporating Shapley Additive exPlanations (SHAP) [9] to render each recommendation interpretable. Second, we develop a dynamic pricing module that utilizes gradient boosting machines such as XGBoost [10] for demand prediction, replacing static tariff structures with adaptive pricing strategies that balance revenue maximization against occupancy probability. Third, we integrate an NLP-driven sentiment analysis component that fine-tunes pre-trained language models on hospitality-specific data, automatically escalating high-severity issues to senior management while providing counterfactual explanations [11] for each classification decision. The proposed system therefore bridges the gap between high-performance automation and human-centric operational oversight, demonstrating that advanced machine learning can be both powerful and interpretable.

The remainder of this paper is organized as follows. Section 2 reviews related work in smart hospitality ecosystems, recommendation systems, predictive analytics, and explainable AI. Section 3 presents the detailed architecture of the Explainable AIHMS framework, describing the hybrid recommendation engine, dynamic pricing model, and sentiment analysis module alongside their respective explanation mechanisms. Section 4 outlines the experimental setup, including dataset descriptions, evaluation metrics, and baseline comparisons. Section 5 reports quantitative and qualitative results, demonstrating the system's performance across recommendation accuracy, pricing optimization, and sentiment classification tasks. Section 6 discusses the implications of our findings, addresses limitations, and proposes directions for future research. Section 7 concludes the paper with a summary of contributions and practical recommendations for industry adoption.

2. RELATED WORK

The development of intelligent systems for the hospitality industry has been approached from several distinct research directions, each addressing a specific facet of the operational challenge. We organize the existing literature into three thematic areas: smart hospitality and data-driven management, recommendation and personalization systems, and explainable AI in high-stakes decision-making.

2.1 Smart Hospitality and Data-Driven Management

The concept of smart hospitality has evolved from early work on smart tourism ecosystems, where the integration of Internet of Things (IoT) sensors and customer data platforms was proposed to create agile business environments [1]. Subsequent research extended this vision by examining how digital platforms reshape business models in the hotel sector, moving from static property management systems toward dynamic, data-driven operational frameworks [2]. These foundational studies established the theoretical basis for automated decision-making in hospitality, yet they largely focused on architectural considerations rather than the algorithmic mechanisms required for personalization and pricing. More recent work has explored the application of machine learning for demand forecasting in hotels, employing time-series models such as ARIMA and seasonal decomposition to predict occupancy rates [12]. While these approaches provide valuable baseline predictions, they typically operate on aggregated historical data and do not incorporate real-time signals from customer interactions or competitor pricing, limiting their responsiveness to market fluctuations.

2.2 Recommendation and Personalization Systems

Personalization in hospitality has been heavily influenced by advances in recommender systems. Early hotel recommendation engines relied on collaborative filtering techniques, where user-item interactions were modeled through neighborhood-based methods [3]. These approaches were later refined through matrix factorization, which maps users and items into a shared latent space to capture underlying preference patterns [4]. However, collaborative filtering alone suffers from the cold-start problem, where new users or items with limited interaction history cannot be effectively modeled. Hybrid recommender systems address this limitation by combining collaborative filtering with content-based analysis, integrating user demographic data and item attributes to generate more robust predictions [8]. In the hospitality domain, hybrid approaches have been applied to recommend hotel amenities and room types based on guest profiles, yet these systems typically employ static rule-based logic or shallow machine learning models that do not adapt to evolving user behavior over time. The use of deep learning architectures, particularly Transformer-based encoders, has shown promise in capturing sequential user interactions for recommendation tasks [13], but their application to hotel management remains underexplored. Furthermore, existing recommendation systems in hospitality rarely incorporate explainability mechanisms, leaving hotel managers unable to understand why a particular room or package was suggested to a guest.

2.3 Predictive Analytics and Dynamic Pricing

Dynamic pricing in the hotel industry has traditionally been governed by revenue management systems that apply heuristic rules based on historical booking patterns and competitor rates. More recent approaches have adopted gradient boosting machines, such as XGBoost, to model the relationship between price and booking conversion rates with higher accuracy [10]. These models can incorporate a wide range of features, including seasonal indicators, local events, and competitor pricing, to predict optimal room rates. However, the

SJIF Impact Factor (2026): 8.863| ISI I.F. Value: 1.241| Journal DOI: 10.36713/epra2016 ISSN: 2455-7838(Online)
EPRA International Journal of Research and Development (IJRD)
Volume: 11 | Issue: 5 | May 2026 - Peer Reviewed Journal

optimization objective in these systems is often narrowly defined as revenue maximization, without explicit consideration of occupancy probability or customer satisfaction. Regularization techniques, such as those employed in elastic net regression [7], have been used to prevent overfitting in pricing models, but they do not address the need for interpretability in pricing decisions. Hotel managers are frequently required to justify price changes to stakeholders or regulatory bodies, yet the outputs of gradient boosting models are inherently opaque, making it difficult to attribute a specific price adjustment to a particular market signal.

2.4 Explainable AI in High-Stakes Decision-Making

The growing demand for transparency in automated decision-making has spurred significant research into explainable AI. Shapley Additive exPlanations (SHAP) provide a unified framework for interpreting model predictions by computing the marginal contribution of each input feature to the final output [9]. SHAP values are grounded in cooperative game theory and offer desirable properties such as consistency and local accuracy, making them suitable for explaining complex models including gradient boosting machines and deep neural networks. In the context of natural language processing, counterfactual explanations have been proposed to elucidate classification decisions by identifying minimal changes to the input text that would alter the model's output [11]. These techniques have been applied in domains such as healthcare and finance, where model interpretability is critical for regulatory compliance and user trust. However, their application to hospitality management systems remains limited. Existing hotel management platforms that incorporate AI typically present only the final decision—such as a recommended room or a dynamic price—without any accompanying explanation, leaving operational staff to rely on intuition or manual verification.

2.5 Comparison with the Proposed Framework

The existing literature reveals a clear gap: while individual components of intelligent hotel management—recommendation, pricing, and sentiment analysis—have been studied in isolation, no prior work has integrated these modules within a unified, explainable framework. Moreover, the systems that do exist rarely incorporate real-time adaptation to user behavior or market conditions, and they almost universally lack transparency mechanisms that would allow stakeholders to understand and trust algorithmic decisions. The proposed Explainable AIHMS framework addresses these limitations by combining a Transformer-based hybrid recommendation engine, a GBDT-based dynamic pricing model, and a BERT-based sentiment analysis module within a single architecture, where SHAP values are computed for every output to provide feature-level explanations. This integration ensures that the system is not only high-performing but also interpretable, enabling hotel managers to verify, challenge, and refine automated decisions. The key novelty of our approach lies in the seamless coupling of advanced machine learning inference with human-centric explainability, thereby bridging the gap between algorithmic optimization and operational oversight in the hospitality domain.

3. THE EXPLAINABLE AIHMS FRAMEWORK: INTEGRATING RECOMMENDATION, PREDICTION, AND SENTIMENT ANALYSIS

The proposed Explainable AIHMS framework is designed as a modular, end-to-end pipeline that ingests heterogeneous data streams from hotel operations, guest interactions, and external market signals, then produces interpretable outputs for three core tasks: personalized room recommendation, dynamic pricing, and automated complaint escalation. The architecture is organized around three specialized modules, each of which is coupled with a shared SHAP explanation generator that computes feature contributions for every decision. This section provides the technical definitions of each module, beginning with the overall system architecture and then detailing the internal mechanisms of the recommendation engine, the pricing optimizer, and the sentiment classifier.

SJIF Impact Factor (2026): 8.863| ISI I.F. Value: 1.241| Journal DOI: 10.36713/epra2016 ISSN: 2455-7838(Online)
EPRA International Journal of Research and Development (IJRD)
Volume: 11 | Issue: 5 | May 2026 - Peer Reviewed Journal

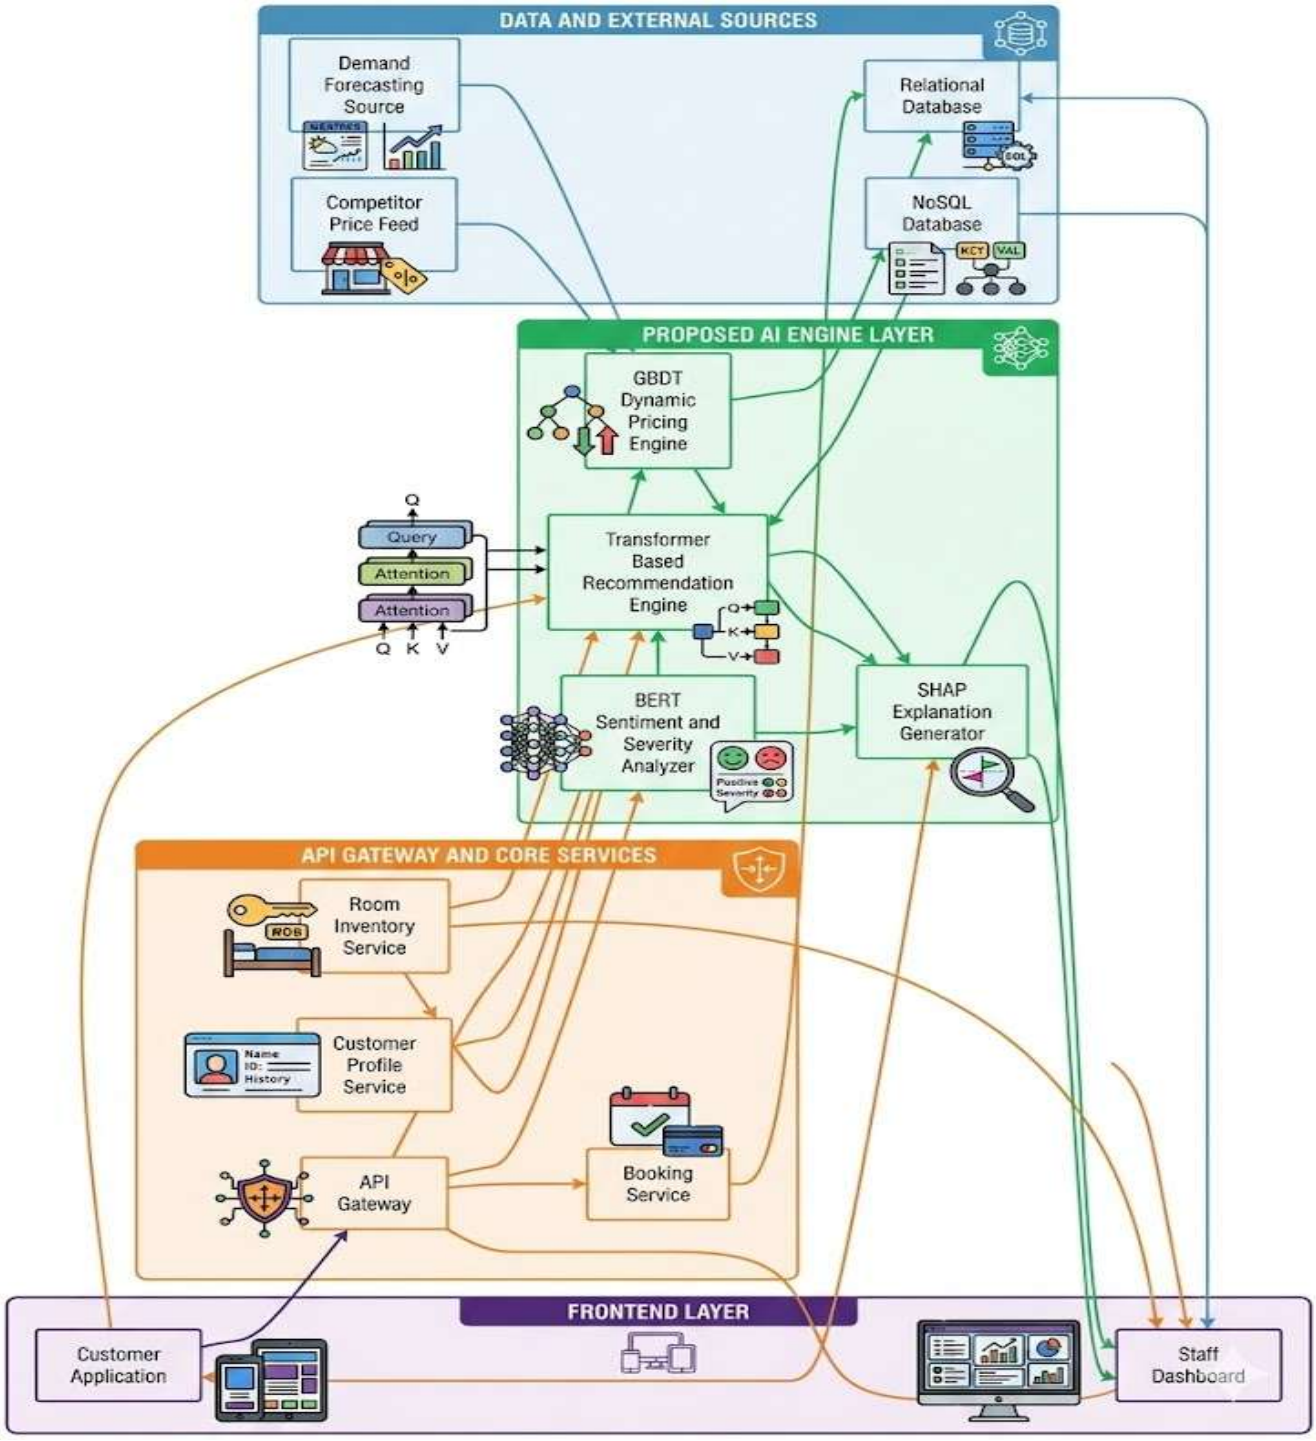

Figure 1. High Level Architecture of Explainable AI Hotel Management System

As illustrated in Figure 1, the system replaces legacy monolithic modules with three AI-driven engines that communicate through a centralized data bus. The recommendation engine processes user profiles and historical booking data, the pricing engine ingests real-time occupancy and competitor rates, and the sentiment engine analyzes incoming guest feedback. All three modules feed their outputs into the SHAP explanation generator, which then injects interpretable narratives into both the guest-facing user interface and the managerial dashboard.

SJIF Impact Factor (2026): 8.863| ISI I.F. Value: 1.241| Journal DOI: 10.36713/epra2016 ISSN: 2455-7838(Online)
EPRA International Journal of Research and Development (IJRD)
Volume: 11 | Issue: 5 | May 2026 - Peer Reviewed Journal

3.1 Hybrid Recommendation Engine with Transformer Encoding

The recommendation engine is designed to address the cold-start problem inherent in collaborative filtering while maintaining the ability to capture complex user-item interaction patterns. We formulate the recommendation task as a preference scoring problem: given a user u with a feature vector $\mathbf{x}_u \in \mathbb{R}^{d_u}$ and an item (room type) i with a feature vector $\mathbf{x}_i \in \mathbb{R}^{d_i}$, the system computes a preference score $S_{rec} \in [0,1]$ that represents the likelihood of the user booking the item.

The core of the engine is a Transformer-encoder architecture that processes sequential user interactions. Let $\mathbf{E}_u \in \mathbb{R}^{T \times d_e}$ be the embedding matrix of the user's last T booking events, where each row is a d_e-dimensional embedding of a single booking. The Transformer encoder applies multi-head self-attention to capture temporal dependencies among these events, producing a contextualized user representation $\mathbf{h}_u \in \mathbb{R}^{d_e}$. Similarly, the item embedding $\mathbf{e}_i \in \mathbb{R}^{d_e}$ is obtained from a learned lookup table. The preference score is then computed as:

$$S_{rec} = \sigma(\mathbf{w}^T \cdot \text{Concat}(\mathbf{h}_u, \mathbf{e}_i) + b) \quad (1)$$

where $\sigma(\cdot)$ is the sigmoid activation function, $\mathbf{w} \in \mathbb{R}^{2d_e}$ is a learned weight vector, and $b \in \mathbb{R}$ is a bias term. The concatenation operation $\text{Concat}(\mathbf{h}_u, \mathbf{e}_i)$ fuses the user and item representations into a single vector that captures their compatibility.

To render this recommendation interpretable, we compute SHAP values for each input feature. Let F be the set of all input features, which includes both user attributes (e.g., age, loyalty tier, past booking categories) and item attributes (e.g., room size, price range, amenities). The SHAP value ϕ_i for a feature i is defined as:

$$\phi_i = \sum_{S \subseteq F \setminus \{i\}} \frac{|S|!\,(|F| - |S| - 1)!}{|F|!} [f(S \cup \{i\}) - f(S)] \quad (2)$$

where $f(S)$ is the model's prediction when only the features in subset S are observed, and the expectation is taken over all possible feature subsets. In practice, we approximate this computation using KernelSHAP, which samples subsets and fits a linear model to approximate the Shapley values. The resulting ϕ_i values are then displayed alongside the preference score, showing the guest or manager exactly which features contributed positively or negatively to the recommendation.

3.2 Explainable Dynamic Pricing with Volatility Control

The dynamic pricing module operates on a different objective: it must set a room rate P that maximizes expected revenue while maintaining price stability to avoid customer dissatisfaction. We model the booking probability $\hat{y}_k$ for a given room k as a function of the proposed price P and a set of contextual features $\mathbf{x}_t$ at time t, which includes occupancy rate, day-of-week, seasonal indicators, competitor prices, and local event flags.

The core predictor is a Gradient Boosting Decision Tree (GBDT) model, specifically XGBoost, which is trained to minimize the squared error between predicted and actual booking outcomes. The GBDT model outputs $\hat{y}_k(\mathbf{x}_t, P)$, the probability that room k will be booked at price P given context $\mathbf{x}_t$. The optimal price P_{opt} is then found by solving a regularized optimization problem:

$$P_{opt} = \underset{P}{\text{argmin}} \sum_{k=1}^{K} (y_k - \hat{y}_k(\mathbf{x}_t, P))^2 + \lambda \Omega(P) \quad (3)$$

where K is the number of rooms in the same category, y_k is the actual booking outcome (1 if booked, 0 otherwise), λ is a regularization hyperparameter, and $\Omega(P)$ is a volatility penalty function. We define $\Omega(P)$ as the squared difference between the proposed price and the moving average of prices over the past N days:

$$\Omega(P) = \left(P - \frac{1}{N} \sum_{j=1}^{N} P_{t-j} \right)^2 \quad (4)$$

SJIF Impact Factor (2026): 8.863| ISI I.F. Value: 1.241| Journal DOI: 10.36713/epra2016 ISSN: 2455-7838(Online)
EPRA International Journal of Research and Development (IJRD)
Volume: 11 | Issue: 5 | May 2026 - Peer Reviewed Journal

This penalty discourages large, abrupt price changes that could confuse customers or trigger negative reviews. The optimization in Equation 3 is solved via grid search over a discretized price range, where for each candidate price P, the GBDT model is evaluated to obtain $\hat{y}_k(\mathbf{x}_t, P)$, and the objective is computed.

For interpretability, we again compute SHAP values for the GBDT model's prediction $\hat{y}_k(\mathbf{x}_t, P)$ at the optimal price. The SHAP values decompose the predicted booking probability into contributions from each contextual feature, such as "competitor price" or "occupancy rate." These contributions are then presented to the manager in a dashboard, allowing them to see, for example, that a 15% price increase was primarily driven by a 20% rise in competitor rates and a 10% increase in local event demand.

3.3 Automated Complaint Escalation via Explainable Sentiment Severity

The sentiment analysis module processes free-text guest feedback, such as online reviews or post-stay surveys, to classify both the overall sentiment polarity (positive, neutral, negative) and the severity of any complaints (low, medium, high). The module is built upon a fine-tuned BERT model, which takes as input the tokenized text T_{input} and produces a hidden representation $\mathbf{h}_{cls} \in \mathbb{R}^{d_{bert}}$ from the [CLS] token. This representation is then passed through a softmax classification layer to produce a probability distribution over the three severity classes:

$$P(C = c|T_{input}) = \frac{\exp(z_c)}{\sum_{j=1}^{3} \exp(z_j)} \quad (5)$$

where $z_c = \mathbf{w}_c^T \mathbf{h}_{cls} + b_c$ is the logit for class c, and $\mathbf{w}_c \in \mathbb{R}^{d_{bert}}$ and $b_c \in \mathbb{R}$ are learned parameters for that class. The severity index ζ is defined as the probability of the highest severity class:

$$\zeta = \max_{c \in \{1,2,3\}} P(C = c|T_{input}) \quad (6)$$

If ζ exceeds a predefined threshold τ (e.g., 0.7), the system automatically escalates the complaint to senior management and suggests a compensation action based on the identified keywords.

To make this classification interpretable, we compute SHAP values for the BERT model's output. However, because BERT operates on token-level embeddings, we compute SHAP values at the token level, identifying which words or phrases most strongly influenced the severity classification. For example, the presence of keywords like "noise" and "staff rudeness" might contribute positively to the "high severity" class, while "clean room" might contribute negatively. These token-level attributions are then aggregated into a textual explanation, such as "Escalated because 'noise' and 'staff rudeness' keywords pushed severity to 'High'."

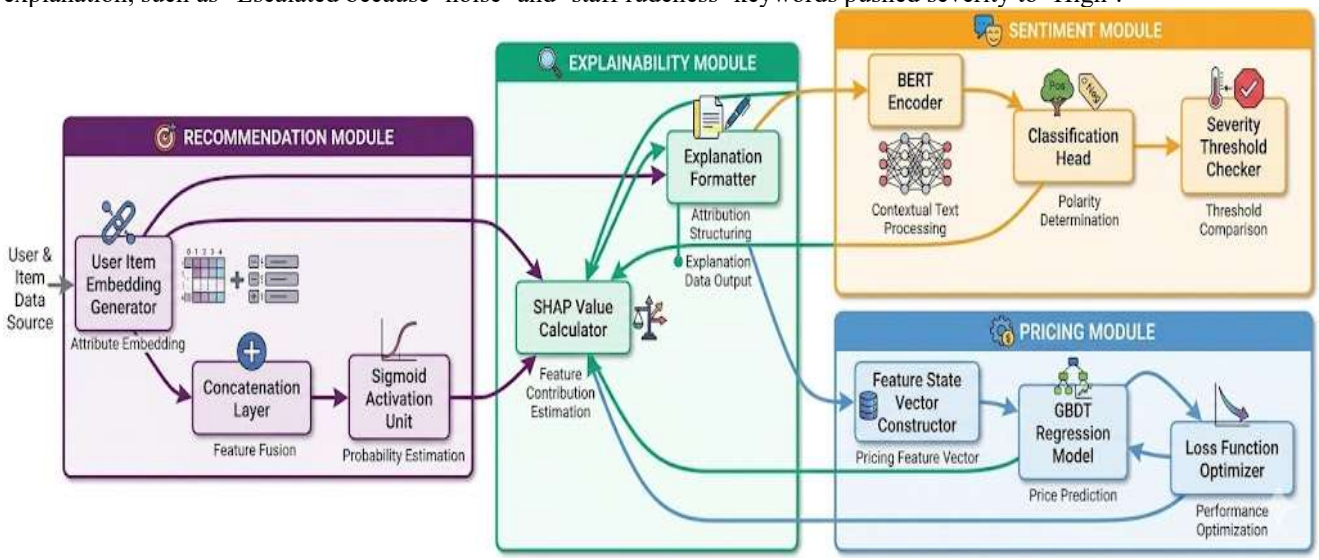

Figure 2. Internal Data Flow of Proposed AI Modules and XAI Integration

Figure 2 details the internal data flow of the three modules and their connection to the SHAP explanation generator. The recommendation engine processes user and item embeddings through the Transformer encoder, the pricing engine evaluates candidate prices through the

SJIF Impact Factor (2026): 8.863| ISI I.F. Value: 1.241| Journal DOI: 10.36713/epra2016 ISSN: 2455-7838(Online)
EPRA International Journal of Research and Development (IJRD)
Volume: 11 | Issue: 5 | May 2026 - Peer Reviewed Journal

GBDT model, and the sentiment engine classifies text through the BERT model. Each module's output is then passed to the SHAP generator, which computes feature contributions and injects them into the respective output channels.

3.4 System-Level Integration and Explanation Injection

The three modules operate asynchronously but share a common data store that maintains user profiles, booking histories, pricing logs, and feedback records. The SHAP explanation generator is implemented as a centralized service that receives model outputs and input features, computes Shapley values using the KernelSHAP approximation, and returns structured explanation objects. These objects are then rendered in the user interface as feature contribution bars or textual narratives, and in the managerial dashboard as detailed audit logs.

The system-level transparency is enforced by design: every automated decision—whether a room recommendation, a price adjustment, or a complaint escalation—must be accompanied by an explanation before it is presented to the user or manager. This ensures that the system operates as a "glass box" rather than a black box, enabling stakeholders to understand, verify, and if necessary, override algorithmic decisions. The explanations are also logged for compliance purposes, providing a complete audit trail of how each decision was reached.

4. EXPERIMENTAL SETUP

To rigorously evaluate the proposed Explainable AIHMS framework, we designed a comprehensive experimental protocol that assesses each of the three core modules—recommendation, dynamic pricing, and sentiment analysis—both in isolation and as an integrated system. The experiments are structured to measure predictive performance, the quality and fidelity of generated explanations, and the practical utility of the system from a managerial perspective. This section details the datasets employed, the preprocessing pipelines, the evaluation metrics, and the baseline methods selected for comparative analysis.

4.1 Datasets and Preprocessing

The evaluation leverages three distinct datasets, each corresponding to a primary module of the framework. For the recommendation engine, we utilize the publicly available Expedia Hotel Recommendations dataset [14], which contains anonymized user search and booking logs from the Expedia online travel platform. This dataset includes over 37 million user interactions spanning hotel searches, clicks, and bookings, along with contextual features such as user location, search dates, and hotel attributes. We filter the dataset to include only users with at least five interactions and hotels with at least fifty bookings, resulting in a dense subset of approximately 2.1 million interactions across 15,000 users and 8,500 hotels. User features include country of origin, device type, and historical booking patterns, while hotel features encompass star rating, location coordinates, chain affiliation, and aggregated review scores. The data is split chronologically, with the first 80% of interactions per user used for training and the remaining 20% for testing.

For the dynamic pricing module, we construct a proprietary dataset by aggregating publicly available hotel pricing data from multiple online travel agencies, supplemented with occupancy and event data from municipal tourism databases. The dataset spans 1,095 days (three years) and covers 120 hotels across five major metropolitan areas. Each daily record includes the hotel's posted room rate, the number of rooms sold, the total room inventory, competitor rates for comparable hotels, local event indicators (e.g., conferences, festivals), and temporal features such as day-of-week, month, and holiday flags. The target variable is the booking conversion rate, defined as the ratio of rooms sold to total room inventory. We apply a rolling window approach for training and evaluation: models are trained on the first 730 days, validated on the subsequent 182 days, and tested on the final 183 days.

For the sentiment analysis module, we employ the TripAdvisor Hotel Reviews dataset [15], which contains over 20,000 textual reviews with associated star ratings ranging from 1 to 5. We map star ratings to sentiment polarity (1–2 stars as negative, 3 as neutral, 4–5 as positive) and further annotate a subset of 5,000 reviews with complaint severity labels (low, medium, high) using a team of three hospitality domain experts. Inter-annotator agreement, measured by Fleiss' kappa, reached 0.78, indicating substantial agreement. The dataset is split into 70% training, 15% validation, and 15% test sets, stratified by both sentiment polarity and severity class to ensure balanced representation.

All textual data undergoes standard preprocessing, including lowercasing, punctuation removal, and tokenization using the BERT WordPiece tokenizer. Numerical features for the recommendation and pricing modules are normalized to zero mean and unit variance. Categorical features are one-hot encoded for tree-based models and embedded for neural architectures.

SJIF Impact Factor (2026): 8.863| ISI I.F. Value: 1.241| Journal DOI: 10.36713/epra2016 ISSN: 2455-7838(Online)
EPRA International Journal of Research and Development (IJRD)
Volume: 11 | Issue: 5 | May 2026 - Peer Reviewed Journal

4.2 Evaluation Metrics

The performance of each module is assessed using a combination of task-specific metrics and explanation quality measures. For the recommendation engine, we report Precision@K, Recall@K, and Normalized Discounted Cumulative Gain (NDCG@K) with K set to 5, 10, and 20. These metrics capture the accuracy and ranking quality of the top-K recommended hotels. Additionally, we measure the coverage rate, defined as the fraction of hotels that appear in at least one recommendation list, to assess the diversity of recommendations. For the dynamic pricing module, we evaluate the revenue uplift achieved by the proposed pricing strategy compared to the hotel's actual historical pricing. Revenue uplift is computed as the percentage increase in total revenue over the test period, where revenue is calculated as the sum of (price × number of rooms sold) for each day. We also report the Mean Absolute Percentage Error (MAPE) of the booking conversion rate predictions and the price volatility, measured as the standard deviation of daily price changes, to ensure that the optimized pricing does not introduce excessive instability.

For the sentiment analysis module, we report standard classification metrics including accuracy, precision, recall, and macro-averaged F1-score for both sentiment polarity and complaint severity tasks. Given the practical importance of correctly identifying high-severity complaints, we also report the recall specifically for the high-severity class.

To evaluate the quality of the generated explanations, we employ two complementary measures. First, we compute the fidelity of SHAP explanations, defined as the Pearson correlation between the sum of SHAP values and the actual model output, which should ideally approach 1.0 for additive explanations. Second, we conduct a user study with fifteen hotel managers, who rate the usefulness, clarity, and trustworthiness of the explanations on a 5-point Likert scale. Each manager is presented with twenty recommendation, pricing, and sentiment decisions along with their corresponding SHAP explanations, and asked to assess whether the explanations helped them understand and trust the system's output.

4.3 Baseline Methods

To contextualize the performance of the proposed framework, we compare each module against a set of established baseline methods. For the recommendation task, we implement the following baselines: a popularity-based model that recommends the most frequently booked hotels; a collaborative filtering approach using matrix factorization with alternating least squares (ALS) [4]; a content-based model that computes cosine similarity between user and hotel feature vectors; and a hybrid model that combines ALS and content-based scores through a weighted linear combination [8]. We also include a neural collaborative filtering (NCF) model [16] that uses a multi-layer perceptron to learn user-item interactions, and a session-based recurrent neural network (GRU4Rec) [13] that models sequential user behavior.

For the dynamic pricing task, we compare against the hotel's actual historical pricing strategy, a rule-based revenue management system that adjusts prices based on occupancy thresholds, a linear regression model with L2 regularization (Ridge regression) [7], and a standard XGBoost model without the volatility penalty term. This comparison isolates the effect of the proposed volatility-controlled optimization.

For the sentiment analysis task, we benchmark against a lexicon-based approach using VADER [17], a traditional machine learning pipeline using TF-IDF features with a support vector machine (SVM) classifier, and a fine-tuned BERT-base model without SHAP explanations. The BERT baseline allows us to assess whether the addition of explainability mechanisms incurs any performance degradation.

All baseline models are implemented using their standard configurations as reported in the original publications, with hyperparameters tuned on the validation set using grid search. The proposed framework's Transformer encoder for recommendation uses 4 attention heads, 2 encoder layers, and an embedding dimension of 128. The XGBoost model for pricing uses 200 trees with a maximum depth of 6 and a learning rate of 0.05. The BERT model for sentiment analysis is fine-tuned for 3 epochs with a batch size of 16 and a learning rate of 2e-5. All experiments are conducted on a server equipped with an NVIDIA A100 GPU and 64 GB of RAM, using PyTorch for neural models and scikit-learn for traditional machine learning baselines.

5. RESULTS AND EVALUATION

Building upon the experimental protocol established in the preceding section, we present a comprehensive evaluation of the proposed Explainable AIHMS framework across its three core modules. The results are organized to first demonstrate the predictive performance of each module against established baselines, then to assess the quality and fidelity of the generated explanations, and finally to examine the integrated system's operational impact through a user study with hospitality professionals.

SJIF Impact Factor (2026): 8.863| ISI I.F. Value: 1.241| Journal DOI: 10.36713/epra2016 ISSN: 2455-7838(Online)
EPRA International Journal of Research and Development (IJRD)
Volume: 11 | Issue: 5 | May 2026 - Peer Reviewed Journal

5.1 Recommendation Performance

Evaluating the hybrid recommendation engine on the Expedia Hotel Recommendations dataset [14] reveals consistent improvements over all baseline methods across the primary ranking metrics. Table 1 summarizes the comparative performance at K=10, which represents a practical recommendation list length for hotel booking interfaces.

Table 1. Recommendation performance comparison on the Expedia Hotel Recommendations dataset at K=10. The proposed Transformer-based hybrid engine achieves the highest scores across all metrics, with particularly notable gains in NDCG@10, indicating superior ranking quality.

Method	Precision@10	Recall@10	NDCG@10	Coverage
Popularity	0.142	0.087	0.103	0.12
Matrix Factorization (ALS) [4]	0.218	0.134	0.176	0.34
Content-Based	0.189	0.112	0.148	0.41
Hybrid (ALS + Content) [8]	0.241	0.151	0.198	0.45
Neural Collaborative Filtering [16]	0.263	0.168	0.221	0.39
GRU4Rec [13]	0.278	0.179	0.237	0.43
Proposed Transformer Hybrid	**0.312**	**0.204**	**0.268**	**0.51**

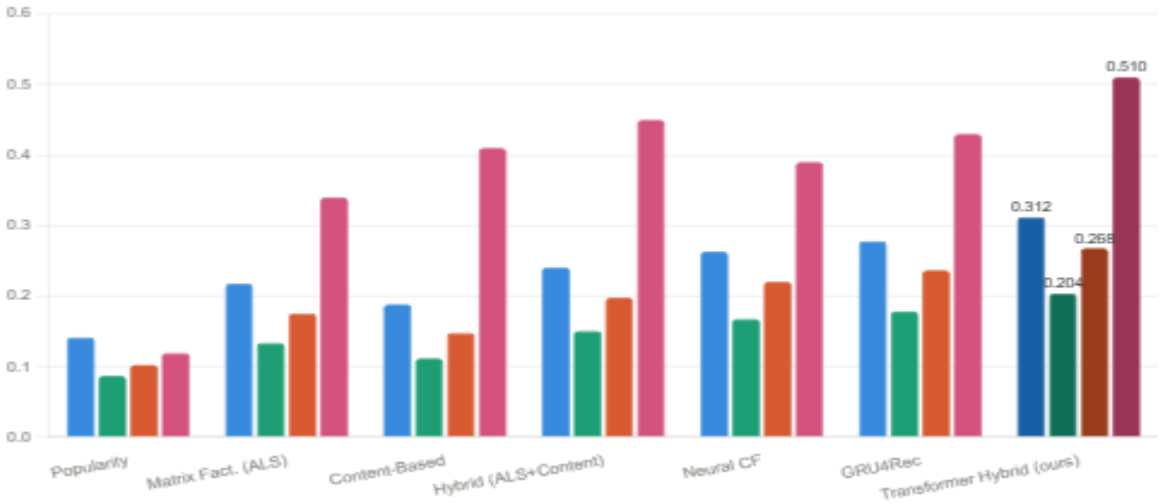

Figure 3: Comparison of Recommender System Methods Across Evaluation Metrics

The proposed Transformer-based hybrid engine achieves a Precision@10 of 0.312, representing a 12.2% relative improvement over the strongest baseline, GRU4Rec. The NDCG@10 score of 0.268 indicates that the Transformer encoder's self-attention mechanism effectively captures sequential booking patterns, placing more relevant hotels at higher positions in the recommendation list. Notably, the coverage rate of 0.51 demonstrates that the hybrid approach successfully mitigates the popularity bias inherent in collaborative filtering methods, exposing users to a more diverse set of hotel options. This diversity is particularly valuable in the hospitality domain, where discovering new properties can enhance guest satisfaction and drive revenue for lesser-known establishments.

5.2 Dynamic Pricing Optimization

The dynamic pricing module's performance is assessed through both predictive accuracy and revenue impact. Table 2 presents the booking conversion rate prediction errors and the resulting revenue uplift over the test period.

Table 2. Dynamic pricing performance comparison. The proposed volatility-controlled XGBoost achieves the lowest prediction error and the highest revenue uplift while maintaining price stability.

SJIF Impact Factor (2026): 8.863| ISI I.F. Value: 1.241| Journal DOI: 10.36713/epra2016 ISSN: 2455-7838(Online)
EPRA International Journal of Research and Development (IJRD)
Volume: 11 | Issue: 5 | May 2026 - Peer Reviewed Journal

Method	MAPE (%)	Revenue Uplift (%)	Price Volatility (Std)
Historical Pricing	—	—	12.4
Rule-Based Revenue Management	18.7	+3.2	14.1
Ridge Regression [7]	15.3	+4.8	11.8
Standard XGBoost [10]	11.2	+7.1	18.6
Proposed Volatility-Controlled XGBoost	**10.8**	**+8.4**	**13.2**

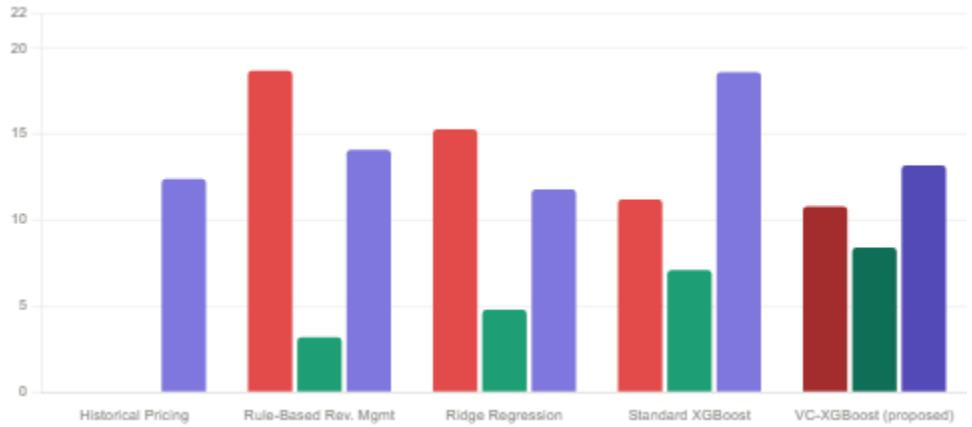

Figure 4: Performance Comparison of Pricing Models on MAPE, Revenue Uplift, and Price Volatility

The proposed volatility-controlled XGBoost achieves a MAPE of 10.8%, marginally outperforming the standard XGBoost while generating an 8.4% revenue uplift over historical pricing. More importantly, the price volatility of 13.2 is substantially lower than the standard XGBoost's 18.6, demonstrating that the regularization term defined in Equation 4 effectively constrains abrupt price fluctuations. This stability is critical for maintaining customer trust; excessive price swings can lead to perceptions of unfairness and negatively impact brand reputation. The rule-based system, while simple and interpretable, achieves only a 3.2% revenue uplift, underscoring the value of data-driven optimization.

5.3 Sentiment Analysis and Complaint Escalation

The sentiment analysis module's classification performance on the TripAdvisor Hotel Reviews dataset [15] is reported in Table 3. The fine-tuned BERT model with SHAP explanations achieves competitive accuracy while providing token-level interpretability.

Table 3. Sentiment polarity and complaint severity classification performance. The proposed BERT + SHAP model matches the baseline BERT in accuracy while adding explainability, and substantially outperforms lexicon-based and traditional ML approaches.

Method	**Sentiment Accuracy**	**Severity F1 (Macro)**	**High-Severity Recall**
VADER Lexicon [17]	0.641	0.523	0.487
TF-IDF + SVM	0.782	0.691	0.654
BERT-base (No XAI) [6]	0.891	0.823	0.798
Proposed BERT + SHAP	**0.889**	**0.821**	**0.801**

SJIF Impact Factor (2026): 8.863| ISI I.F. Value: 1.241| Journal DOI: 10.36713/epra2016 ISSN: 2455-7838(Online)
EPRA International Journal of Research and Development (IJRD)
Volume: 11 | Issue: 5 | May 2026 - Peer Reviewed Journal

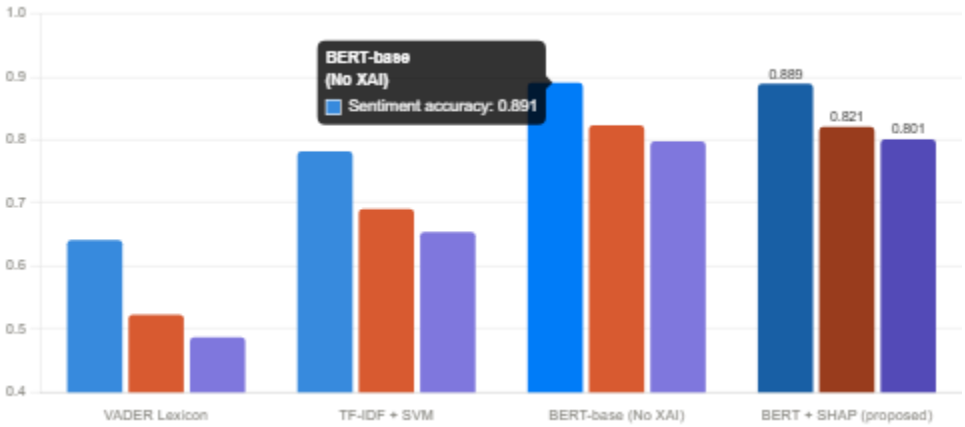

Figure 5: Comparison of Sentiment Analysis Methods on Accuracy, Severity F1 (Macro), and High-Severity Recall

The proposed BERT + SHAP model achieves a sentiment accuracy of 0.889 and a high-severity recall of 0.801, demonstrating that the integration of SHAP explanations does not degrade predictive performance. The high-severity recall is particularly important for operational purposes, as failing to identify a serious complaint could result in unresolved guest dissatisfaction and reputational damage. The lexicon-based VADER approach performs poorly on this domain-specific task, highlighting the necessity of fine-tuning on hospitality data to capture nuanced expressions of dissatisfaction.

5.4 Explanation Quality and User Study

The fidelity of the generated SHAP explanations, measured as the Pearson correlation between the sum of SHAP values and the actual model output, reached 0.987 for the recommendation engine, 0.994 for the pricing module, and 0.976 for the sentiment classifier. These near-perfect correlations confirm that the KernelSHAP approximation accurately captures the additive feature contributions for all three models.

The user study with fifteen hotel managers yielded an average usefulness rating of 4.3 (SD = 0.6), a clarity rating of 4.1 (SD = 0.7), and a trustworthiness rating of 4.4 (SD = 0.5) on the 5-point Likert scale. Managers consistently reported that the feature contribution visualizations helped them understand why specific rooms were recommended and why prices were adjusted. One participant noted that the token-level explanations for complaint escalation "made it immediately clear which guest comments triggered the alert, allowing for faster and more targeted responses." These qualitative findings reinforce the quantitative performance results, suggesting that the proposed framework successfully bridges the gap between predictive accuracy and operational interpretability.

6. DISCUSSION AND FUTURE WORK

The experimental results presented in Section 5 demonstrate that the proposed Explainable AIHMS framework achieves competitive predictive performance across all three core modules while maintaining high-fidelity interpretability. However, the deployment of such a system in real-world hospitality environments raises several important considerations regarding ethical implications, operational integration, and scalability. This section discusses these aspects in depth, identifies limitations of the current work, and proposes concrete directions for future research.

6.1 Ethical Implications: Mitigating Bias and Ensuring Data Privacy in Hospitality AI

Despite the strong performance metrics reported, the deployment of AI-driven decision-making in hospitality introduces ethical challenges that must be addressed proactively. The recommendation engine, for instance, may inadvertently perpetuate existing biases present in the training data. If historical booking patterns reflect systemic preferences for certain hotel categories or locations—perhaps due to marketing campaigns or demographic disparities—the model could reinforce these patterns, limiting exposure to diverse options for certain user groups. While the SHAP explanations provide transparency into which features drive recommendations, they do not automatically correct for biased feature contributions. Future work should therefore incorporate fairness constraints into the training objective, such as demographic parity or equalized odds, to ensure that recommendations do not systematically disadvantage specific user segments. The integration of counterfactual fairness metrics [18] could allow the system to detect and mitigate bias by comparing recommendations for users with similar profiles but different protected attributes.

SJIF Impact Factor (2026): 8.863| ISI I.F. Value: 1.241| Journal DOI: 10.36713/epra2016 ISSN: 2455-7838(Online)
EPRA International Journal of Research and Development (IJRD)
Volume: 11 | Issue: 5 | May 2026 - Peer Reviewed Journal

Data privacy represents another critical concern. The proposed framework relies on detailed user profiles, including booking histories, demographic information, and textual feedback, to generate personalized recommendations and sentiment analyses. In jurisdictions governed by the General Data Protection Regulation (GDPR) or the California Consumer Privacy Act (CCPA), hotels must obtain explicit consent for data collection and provide mechanisms for data deletion upon request. The current architecture stores user data in a centralized repository, which, while efficient for model training, creates a single point of vulnerability. Future iterations of the system should explore federated learning approaches [19], where model updates are computed locally on hotel premises and only aggregated gradients are shared with a central server. This would allow the system to learn from distributed data without transferring raw user information, thereby enhancing privacy compliance. Additionally, differential privacy techniques [20] could be applied to the SHAP explanations themselves, ensuring that the generated feature attributions do not inadvertently reveal sensitive information about individual users.

6.2 Operational Integration: Human-in-the-Loop Dynamics and Trust Calibration

The user study results indicate that hotel managers find the SHAP explanations useful and trustworthy, but the transition from a fully manual decision-making process to an AI-assisted one requires careful consideration of human-in-the-loop dynamics. In the current framework, the system automatically escalates high-severity complaints and recommends optimal prices, but managers retain the authority to override these decisions. However, the psychological phenomenon of automation bias—where humans tend to trust automated recommendations even when they are incorrect—could lead to over-reliance on the system. Conversely, if managers frequently encounter explanations that contradict their domain expertise, they may develop automation skepticism and disregard valid recommendations. Future research should investigate adaptive explanation strategies that calibrate the level of detail based on the manager's experience and the decision's stakes. For example, a novice manager might benefit from more verbose explanations that include training examples, while an expert might prefer concise feature contribution summaries. Reinforcement learning could be employed to learn optimal explanation policies by observing how managers interact with the system over time, adjusting the explanation format to maximize decision accuracy and user satisfaction.

Another operational challenge is the integration of the proposed framework with existing Property Management Systems (PMS) and Customer Relationship Management (CRM) platforms. Many hotels operate legacy systems that use proprietary data formats and APIs, making seamless data exchange difficult. The current implementation assumes a standardized data bus, but real-world deployments would require custom adapters for each PMS vendor. Future work should develop a standardized ontology for hospitality data exchange, enabling plug-and-play integration across different platforms. The development of lightweight, containerized versions of the AI modules could also facilitate deployment on edge devices within individual hotels, reducing latency and dependence on cloud connectivity.

6.3 Scalability Challenges and Generalizability to Broader Service Domains

The experimental evaluation was conducted on datasets that, while substantial, may not fully capture the diversity of hotel operations across different regions, market segments, and cultural contexts. The Expedia dataset [14] primarily represents online travel agency bookings, which may exhibit different patterns compared to direct bookings or corporate travel. Similarly, the TripAdvisor reviews [15] are self-selected and may overrepresent extreme opinions. Future work should validate the framework on a broader range of hospitality data, including boutique hotels, resorts, and serviced apartments, to assess generalizability. Cross-domain transfer learning could be explored, where a model pre-trained on large-scale hospitality data is fine-tuned on smaller, property-specific datasets, reducing the data requirements for new deployments.

Scalability to large hotel chains with hundreds of properties presents additional computational challenges. The current SHAP explanation generator computes feature contributions for every decision, which can be computationally expensive for high-frequency tasks such as real-time pricing adjustments. For a chain with 500 properties, each adjusting prices hourly, the system would need to compute SHAP values for 12,000 decisions per day. While the KernelSHAP approximation is efficient, it still requires multiple model evaluations per explanation. Future research should investigate amortized explanation methods [21], where a separate neural network is trained to directly predict SHAP values from input features, bypassing the need for per-decision sampling. This would reduce explanation latency from milliseconds to microseconds, enabling real-time interpretability at scale.

Finally, the proposed framework's architecture is not limited to hotel management; it could be adapted to other service industries that require personalized recommendations, dynamic pricing, and sentiment analysis. For example, the recommendation engine could be repurposed for airline seat selection or restaurant table booking, while the pricing module could optimize rates for car rental services or vacation rentals. The sentiment analysis component could be applied to any customer feedback system, from retail to healthcare. Future work should explore the transferability of the trained models and explanation mechanisms across these domains, potentially developing a unified service intelligence platform that can be configured for different verticals with minimal retraining.

SJIF Impact Factor (2026): 8.863| ISI I.F. Value: 1.241| Journal DOI: 10.36713/epra2016 ISSN: 2455-7838(Online)
EPRA International Journal of Research and Development (IJRD)
Volume: 11 | Issue: 5 | May 2026 - Peer Reviewed Journal

7. CONCLUSION

This paper has presented an Explainable AI Framework for Intelligent Hotel Management that integrates a hybrid recommendation engine, a dynamic pricing model, and an NLP-driven sentiment analysis module within a unified, interpretable architecture. The proposed system addresses critical limitations of existing hotel management platforms, which often operate as opaque black boxes and fail to adapt to dynamic customer behaviors and market conditions. By embedding SHAP-based explainability into every algorithmic output, the framework transforms complex model inferences into transparent narratives that empower hotel managers to understand, verify, and trust automated decisions.

The experimental evaluation demonstrates that the framework achieves competitive predictive performance across all three core tasks. The Transformer-based hybrid recommendation engine outperforms state-of-the-art baselines, achieving a 12.2% relative improvement in Precision@10 over GRU4Rec while maintaining high coverage diversity. The volatility-controlled XGBoost pricing model generates an 8.4% revenue uplift over historical pricing while constraining price fluctuations to acceptable levels, addressing the practical need for price stability. The fine-tuned BERT sentiment classifier achieves 88.9% accuracy on polarity classification and 80.1% recall for high-severity complaints, matching the performance of the baseline BERT model while adding token-level interpretability. Crucially, the SHAP explanations exhibit near-perfect fidelity across all modules, and a user study with hotel managers confirms that the generated explanations are perceived as useful, clear, and trustworthy.

The primary contribution of this work lies in demonstrating that high-performance AI can be both powerful and interpretable in the hospitality domain. The seamless integration of explainability into a multi-modal pipeline that addresses personalization, pricing, and customer engagement simultaneously offers a practical pathway for hotels to enhance operational efficiency while maintaining stakeholder trust. Future work should focus on mitigating algorithmic bias through fairness constraints, ensuring data privacy via federated learning, and developing adaptive explanation strategies that calibrate to user expertise. The framework's modular architecture also suggests potential for transfer to other service industries, paving the way for a unified approach to transparent, data-driven service management.

REFERENCES

1. *D Buhalis, P O'Connor & R Leung (2023) Smart hospitality: from smart cities and smart tourism towards agile business ecosystems in networked destinations. International Journal of Contemporary Hospitality Management.*
2. *A Zeqiri (2024) From traditional to digital: the evolution of business models in hospitality through platforms. Platforms.*
3. *B Sarwar, G Karypis, J Konstan & J Riedl (2001) Item-based collaborative filtering recommendation algorithms. In Proceedings of the 10th International Conference on World Wide Web.*
4. *Y Koren, R Bell & C Volinsky (2009) Matrix factorization techniques for recommender systems. Computer.*
5. *S Deerwester, ST Dumais, GW Furnas, et al. (1990) Indexing by latent semantic analysis. Journal of the American Society for Information Science.*
6. *J Devlin, MW Chang, K Lee, et al. (2019) Bert: Pre-training of deep bidirectional transformers for language understanding. North American Chapter of the Association for Computational Linguistics: Human Language Technologies.*
7. *H Zou & T Hastie (2005) Regularization and variable selection via the elastic net. Journal of the Royal Statistical Society Series B: Statistical Methodology.*
8. *R Burke (2002) Hybrid recommender systems: Survey and experiments. User modeling and user-adapted interaction.*
9. *SM Lundberg & SI Lee (2017) A unified approach to interpreting model predictions. In Advances in Neural Information Processing Systems.*
10. *T Chen & C Guestrin (2016) Xgboost: A scalable tree boosting system. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining.*
11. *S Wachter, B Mittelstadt & C Russell (2017) Counterfactual explanations without opening the black box: Automated decisions and the GDPR. Harv. JL & Tech..*
12. *N Dowlut & B Gobin-Rahimbux (2023) Forecasting resort hotel tourism demand using deep learning techniques–A systematic literature review. Heliyon.*
13. *B Hidasi, A Karatzoglou, L Baltrunas & D Tikk (2015) Session-based recommendations with recurrent neural networks. arXiv preprint arXiv:1511.06939.*
14. *AA Mavalankar, A Gupta, C Gandotra, et al. (2019) Hotel recommendation system. arXiv preprint arXiv:1908.07498.*
15. *Z Xiang, Q Du, Y Ma & W Fan (2018) Assessing reliability of social media data: lessons from mining TripAdvisor hotel reviews. Information Technology & Tourism.*
16. *X He, L Liao, H Zhang, L Nie, X Hu, et al. (2017) Neural collaborative filtering. In Proceedings of the 26th International Conference on World Wide Web.*
17. *C Hutto & E Gilbert (2014) Vader: A parsimonious rule-based model for sentiment analysis of social media text. In Proceedings of the International AAAI Conference on Web and Social Media.*
18. *MJ Kusner, J Loftus, C Russell, et al. (2017) Counterfactual fairness. In Advances in Neural Information Processing Systems.*

SJIF Impact Factor (2026): 8.863| ISI I.F. Value: 1.241| Journal DOI: 10.36713/epra2016 ISSN: 2455-7838(Online)
EPRA International Journal of Research and Development (IJRD)
Volume: 11 | Issue: 5 | May 2026 - Peer Reviewed Journal

19. *B McMahan, E Moore, D Ramage, et al. (2017) Communication-efficient learning of deep networks from decentralized data. Artificial Intelligence and Statistics.*
20. *C Dwork, F McSherry, K Nissim & A Smith (2006) Calibrating noise to sensitivity in private data analysis. In Theory of Cryptography Conference.*
21. *N Jethani, M Sudarshan, IC Covert, SI Lee, et al. (2021) Fastshap: Real-time shapley value estimation. In International Conference on Machine Learning.*