

Machine Learning-Driven Hotel User Behavior Analytics and Intelligent Decision Support: A State-of-the-Art Review

Zhengyi Zhang*

School of Computer Science and Artificial Intelligence, Guangdong University of Education, Guangzhou, China

**Corresponding author: Zhengyi Zhang.*

Abstract

Machine-learning studies in hotel analytics do not address one common unit of analysis. Some predict an event at booking level, some interpret language at review level, and others estimate value at customer level. This review examines how those separate tasks may contribute to hotel decisions. The papers considered cover order propensity, guest-experience mining, cancellation risk, customer lifetime value, mobile reservation behavior, and explainable modeling. Their samples, targets, and evaluation measures differ, so the review does not rank algorithms by reported accuracy. Instead, it asks what evidence each model produces, who could use that evidence, and what remains to be decided after a prediction is made. Read together, the studies point to a five-part process: customer records are assembled, relevant features are constructed, a task-specific estimate is produced, the estimate is interpreted, and a manager selects an action. This sequence also exposes gaps in the literature. Structured records and review text are not routinely connected; explanations are often technical rather than operational; changing market conditions receive less attention than model fit; and the effect of the eventual hotel intervention is seldom measured. A closed-loop framework is proposed to organize these gaps. The framework treats a later customer response as evidence for revising both the model and the decision rule. Machine learning, on this account, informs hotel judgment but does not replace it.

Keywords

machine learning, hotel user behavior analytics, intelligent decision support, hospitality analytics, explainable artificial intelligence

1. Introduction

The digital record surrounding a hotel stay begins well before check-in. A prospective guest may search, compare, abandon a page, return through a mobile platform, change a reservation, cancel it, or post a review afterward [1, 2]. These traces are often grouped under 'user behavior,' although they capture different stages and different levels of commitment. A search query indicates consideration; a paid booking records a choice; a review is a retrospective account. Combining them is analytically useful only if those differences are retained.

This distinction matters because conventional occupancy and revenue records mainly describe realized demand. They say less about an abandoned booking or the service experience behind a low rating. Platform information, interface design, reputation, social functions, and cancellation terms may all affect what is eventually observed [1]. The recent literature responds with separate models for order propensity, review content, cancellation, CLV, recommendation, and revenue decisions [3-8]. The resulting scores are not interchangeable, even when they are all described as customer analytics.

The fragmentation is therefore more than a matter of topic. The unit being predicted changes from a transaction to a text passage and then to a customer relationship. An order probability cannot be read in the same way as an aspect sentiment or a lifetime-value estimate [3-8]. More importantly, none of these quantities specifies a hotel action. Between prediction and action lie a threshold, a budget, a service rule, and a person accountable for the decision.

The sources also illuminate different sides of the same guest. Reservation and payment data show what occurred; review language may reveal what the guest valued or disliked [4, 8]. A joint model can exploit both, but joining data does not automatically produce understanding. A strong association in text, for example, is not proof that changing the associated service feature will change behavior. Questions of interpretation, privacy, temporal drift, staff use, and intervention cost therefore deserve separate evaluation rather than being treated as secondary details of model performance.

For that reason, this review uses the hotel decision as its point of comparison. Each study is read along a short chain: the observation collected, the feature represented, the outcome estimated, the explanation offered, and the managerial choice that might follow. This approach permits comparison across unlike methods without implying that their accuracy figures are directly comparable.

The analysis is organized around three questions. Which machine-learning methods have been used for hotel behavior data? At what point could their outputs affect marketing, service, revenue, or relationship management? What additional work is required before the resulting system is understandable, updateable, and usable in routine operations? After defining the central concepts, the paper reads the application studies comparatively, develops an organizing framework, and discusses the limits of current evidence.

2. Conceptual Background

Hotel user behavior refers to the observable and analyzable actions, preferences, intentions, evaluations, and interactions generated by customers before, during, and after hotel consumption. In digital hospitality environments, such behavior is reflected not only in final booking outcomes, but also in search activities, browsing sequences, price comparison, cancellation decisions, review writing, loyalty participation, and post-stay feedback. Mobile online travel agencies further expand this behavioral process by integrating hotel information, interface design, recommendation mechanisms, social interaction, and personalized service requests into the reservation journey [1]. Therefore, hotel user behavior analytics should consider the entire digital journey rather than only completed transactions.

Customer relationship management provides a theoretical basis for linking user behavior analytics with long-term business value. From a CRM perspective, hotels should not treat customers as isolated transactions, but as dynamic relationships shaped by repeated stays, loyalty engagement, service experiences, feedback behavior, and changing travel needs. Customer lifetime value is closely related to this perspective because it helps hotels evaluate future customer value based on booking frequency, spending level, stay duration, cancellation behavior, review sentiment, and loyalty participation [7, 8]. In hospitality contexts, CLV analysis is particularly important because customer behavior is often non-contractual, irregular, and affected by seasonality, travel purpose, and destination preference.

Intelligent decision support refers to the use of data, analytical models, explanation mechanisms, and managerial rules to assist decision-making under uncertainty. In hotel management, decision support is needed for pricing, marketing, room inventory control, customer retention, service recovery, and resource allocation. Machine learning can enhance this process by transforming heterogeneous hotel data into predictive and interpretable insights. However, prediction alone is not sufficient. To become managerially useful, model

outputs must be connected with business indicators, operational rules, explainable interpretation, and human judgment so that analytical results can be translated into practical hotel management decisions.

3. Literature Review

3.1 Customer Ordering Prediction and Precision Marketing

Order prediction addresses a decision made under incomplete information: whether scarce marketing attention should be directed to a particular customer. Xiong et al. combine random forest and support vector machine models in a two-stage procedure [3]. The design is suited to interactions that a simple linear segmentation may miss. Its output can separate customers by estimated booking propensity, but it does not establish that a discount causes an additional booking. That second question would require evidence from an intervention or campaign comparison.

This reading changes the appropriate performance test. Accuracy describes prediction, whereas a marketing team needs incremental bookings net of campaign cost. A false positive may waste an incentive; a false negative may exclude a responsive customer; an unnecessary discount may reduce revenue without changing behavior. Segment lift, contact cost, offer redemption, and later retention are consequently more informative for deployment than one aggregate accuracy figure.

3.2 Hotel Experience Mining and Review-Based Sentiment Intelligence

A completed reservation contains little detail about the meaning of the stay to the guest. Reviews add descriptions of cleanliness, room comfort, location, staff conduct, food, facilities, and emotion. Yet the review is not a single attitude. Praise for location can coexist with criticism of check-in, making an overall positive-negative label too coarse for many service decisions.

Jia et al. use a CNN-BiLSTM-attention model to recover hotel-experience characteristics from review text [4]. Convolution detects local expressions, the bidirectional sequence component retains context, and attention highlights influential parts of the comment. The combination helps separate experience attributes that an average rating would merge. Attention, however, should be read cautiously: it identifies text used by the model, not necessarily the cause of the guest's experience.

The useful managerial object is therefore a recurring aspect in a defined place and period. Repeated complaints about bathroom cleanliness suggest an inspection problem; criticism of room photographs points instead to information accuracy; slow check-in may implicate staffing or process design. Linking aspect, sentiment, property, and time can direct service recovery or training. Without that last link, review mining remains a description of customer voice rather than evidence for a specific operational response.

3.3 Booking Cancellation Forecasting

Cancellation prediction has an unusually visible cost of error. If risk is understated, rooms may remain unsold; if it is overstated, aggressive overbooking can displace a confirmed guest. Staff planning, confirmation messages, deposits, and room release are also affected before arrival. The relevant model output is thus a calibrated probability that can be paired with interventions of different cost, not simply a cancel-or-arrive label.

Jishan et al. use Bayesian models to express uncertainty around cancellation [5]. Lead time, stay details, room type, special requests, guest information, and commitment signals contribute evidence without determining the outcome individually. The probabilistic form is valuable because a hotel can apply one threshold for a low-cost reminder and a higher threshold for a deposit or inventory action. In this setting, uncertainty is part of the decision rather than noise to be hidden.

Xiao et al. examine a different weakness: patterns that change over time. Their temporal reinforcement-learning and policy-enhanced LSTM framework accounts for shifts associated with holidays, events, seasonality, and demand conditions [6]. The two cancellation studies should therefore not be treated as direct competitors. Reference [5] concentrates on uncertainty in an estimate, whereas reference [6] concentrates on adaptation of predictions and policy. A deployable system would need both properties.

3.4 Customer Lifetime Value Estimation

CLV changes the question from the next reservation to the expected course of the customer relationship. Hotel stays are episodic, and a period of inactivity may reflect destination or travel purpose rather than permanent attrition. Webb, Cho, and Legg frame CLV for this non-contractual hospitality setting [7]. Their contribution is primarily a business and measurement perspective: the time horizon and meaning of inactivity must be defined before a complex model is selected.

Nikolić et al. add an explanatory source that conventional recency-frequency-monetary variables do not contain [8]. Booking records quantify frequency, spending, and stay patterns, while reviews may disclose satisfaction, frustration, or attachment. Multimodal estimation is most persuasive where the two sources disagree. A guest with strong historical value and deteriorating review language, for instance, represents a different retention case from a guest who scores highly on both dimensions.

Chamberlain et al. use embeddings to summarize large customer histories in e-commerce [9]. The method is relevant to platforms that observe long sequences of searches and transactions, but its transfer to hotels is an inference, not a result established by that study. Destination dependence, sparse repeat stays, and seasonal travel can change the meaning of a behavioral sequence. Hotel validation against simpler CLV baselines is therefore necessary before embeddings determine loyalty status or retention expenditure.

3.5 Mobile OTA Behavior, Multimodal Learning, and Explainable Artificial Intelligence

Mobile OTA data deserve separate treatment because the platform both records and shapes behavior. Sun, Law, and Hyun identify information provision, social functions, layout, and customer-request features in mobile reservation behavior [1]. An abandoned booking may thus reflect interface friction rather than weak demand for the hotel. Reading every click as a stable customer preference would confound the guest with the design of the channel through which the guest was observed.

This problem helps explain the appeal of multimodal models. Dates, prices, room types, and transactions are naturally structured, whereas reviews preserve evaluation in language. Nikolić et al. combine those forms for CLV [8]. The analytical gain is clearest when text changes the interpretation of a behavioral score, not merely when it raises an accuracy measure by a small amount. That distinction should be tested explicitly in future hotel studies.

Explanation becomes consequential when a score changes price, promotional access, or the handling of an individual guest. SHAP, LIME, and attention displays can reveal influential inputs [8], but an attribution is not yet a managerial explanation. It must be translated into the decision at hand, accompanied by uncertainty, and prevented from being mistaken for a causal effect. A useful explanation should also give the manager grounds to reject the recommendation.

4. Synthesis of Key Findings

4.1 An Integrated Architecture for ML-Driven Hotel Decision Support

No single reviewed paper implements the entire path from raw hotel data to an evaluated managerial intervention. The five-layer architecture developed here is therefore an organizing synthesis, not a validated system copied from one source. It separates data assembly, feature representation, task-specific prediction, explanation, and decision. The separation makes dependencies visible: better modeling cannot repair an incorrectly linked customer history, and a clear explanation cannot make an unsuitable intervention fair or economical.

The first two layers determine what the model is allowed to know. Reservation, transaction, CRM, loyalty, mobile, and review records must be matched and converted into behavioral, temporal, textual, contextual, or value features. The prediction layer answers a deliberately narrow question, such as cancellation risk or CLV. Explanation then exposes the evidence behind the estimate. Only after those stages should operating constraints, service standards, capacity, and cost enter the choice of action. Conflating the layers makes it difficult to identify where an error originated.

Table 1 should consequently be read as an accountability map. Data teams can be asked about coverage and identity matching; model developers about calibration and subgroup performance; interface designers about the intelligibility of explanations; and hotel managers about the proportionality and outcome of an intervention. Reported predictive accuracy addresses only one part of that chain.

Table 1: Synthesis of Key Layers in ML-Driven Hotel Decision Support

Layer	Primary material	Main analytical work	Decision-stage contribution
Data	Reservations; transactions; CRM; loyalty; mobile events; reviews	Identity matching, quality checks, integration, access control	Creates a traceable evidence base
Feature	Behavioral, temporal, textual, contextual, and value signals	Feature construction, embeddings, aspect and sentiment extraction	Expresses customer histories in model-ready form
Prediction	Features linked to observed outcomes	Classification, regression, probabilistic, sequence, and multimodal models	Estimates ordering, cancellation, experience, preference, or value
Explanation	Predictions, feature effects, uncertainty, and relevant text	Attribution, local explanation, attention display, interpretable rules	Makes the evidence and limits of a result reviewable
Decision	Explained result plus cost, capacity, policy, and service constraints	Threshold rules, optimization, KPI review, and staff judgment	Guides a bounded action and records its outcome

4.2 From Predictive Analytics to Managerial Action

The studies enter hotel work at different moments. An order score may narrow a campaign list; aspect-level review evidence may direct an inspection; a cancellation probability may trigger confirmation; and CLV may affect a retention offer. These are bounded choices, not general claims that machine learning improves management. Naming the decision, the responsible team, and the alternative action makes the claimed value testable.

Consider a cancellation probability of 0.70. During a high-demand event it may justify an inventory adjustment, whereas on an ordinary low-demand date it may justify only a reminder. The number is unchanged; the cost of acting is not. The same problem arises with CLV: a high estimate sets no obvious benefit, budget, or contact strategy. Thresholds must therefore be chosen with reference to scarcity, intervention cost, service promises, fairness, and the damage caused by a wrong decision.

Several forms of knowledge are needed at this point. Data specialists can judge model validity, revenue and marketing teams understand commercial trade-offs, and frontline employees observe service conditions absent from the dataset. An explanation can support discussion across these roles, but it cannot transfer accountability to the model. Overrides should be possible and recorded, because disagreement between staff and the system may itself reveal missing context.

4.3 A Closed-Loop Framework for Hotel User Behavior Analytics

The proposed loop begins with observation and ends with another observation. Hotel records are represented as features; a task-specific model produces an estimate; and the estimate is reviewed in light of the operating situation. The hotel then chooses an intervention in marketing, service, retention, pricing, or inventory. This is the point at which most predictive studies stop, even though it is also the point at which managerial value begins.

Closing the loop requires the intervention and later response to be retained as linked evidence. A confirmation message, service-recovery action, or retention offer may be followed by a different booking or review outcome. That sequence can evaluate the policy as well as the model. Retraining on later data is not sufficient if the hotel cannot tell whether behavior changed because the market shifted or because its own intervention worked.

5. Challenges and Future Research Directions

5.1 Data, Privacy, and Governance Challenges

The practical data problem is continuity. Reservation, CRM, loyalty, payment, service, and review systems may assign different identifiers and retain records for different periods. A complete customer history can therefore be difficult to reconstruct even within one hotel group. The Hotel Booking Demand Dataset supports reproducible cancellation and demand studies [2], but it lacks the persistent identity, post-stay evaluation, loyalty activity, and complete revenue history required for longitudinal CLV and intervention analysis. A public benchmark is useful, but it cannot answer every managerial question.

Record quality and legitimate use must be considered together. Duplicate profiles, missing values, incompatible codes, and selective review behavior can bias an estimate before model choice becomes relevant. Profiling and dynamic offers then raise separate questions about consent and unequal treatment. A credible study should state where data came from, who could access them, how long they were retained, and whether performance differs across customer groups. Federated learning may keep raw data within participating hotels, but it does not correct poor local records or decide which uses are acceptable.

5.2 Model Interpretability and Adaptive Decision-Making

Aggregate accuracy can coexist with a poor individual decision. This is especially problematic for deep or multimodal models whose internal representations are difficult to relate to a service attribute. If a guest is denied an offer or marked for special treatment, the explanation must address that particular action. Feature importance alone is insufficient. The user needs the relevant evidence, the uncertainty around the result, and a warning when the relationship is associational rather than causal.

A model also has a date. Seasons, events, economic conditions, platform competition, and crises can change both demand composition and cancellation behavior. Validation on an earlier sample says little about current calibration unless temporal tests are reported. Future studies should specify drift indicators, update schedules, and conditions for reverting to a simpler rule. Incremental revenue, empty-room reduction, complaint rates, and intervention cost should be monitored beside model metrics.

5.3 Future Directions for Human-Centered Hotel Intelligence

The next research step need not always be a larger model. A more informative study may ask whether adding review text changes a retention decision, or whether mobile sequences improve cancellation management enough to justify their collection. Preference-aware systems face a similar test because guests value price, location, room attributes, amenities, service, and flexibility differently. Neurosymbolic methods and preference-centric optimization can represent rules and competing objectives [10, 11], but hotel-specific evaluation is still required. Complexity should be earned by a measurable decision benefit.

Human-AI collaboration should likewise be studied as observable work rather than as a general principle. Revenue managers and frontline staff need a way to question a recommendation, state why they overrode it, and see what followed. Their judgment is particularly important for exceptions, empathy, and service recovery. For a small or medium-sized hotel, a reliable dashboard with clear data checks may be more valuable than an opaque end-to-end platform.

6. Conclusion

The literature reviewed here addresses four recurring decisions: whom to contact, which service issue to prioritize, which booking deserves attention, and which customer relationship warrants investment. Order prediction, review mining, cancellation forecasting, and CLV contribute different evidence to those decisions [3-8]. Because the studies use different targets, samples, and metrics, their reported performance cannot support a simple ranking of methods. What can be compared is the distance between a model output and a defensible hotel action.

That comparison leads to two conclusions. First, a hotel decision-support system needs linked data, a narrowly defined prediction, an explanation suited to the decision, and a recorded intervention outcome.

Privacy and temporal monitoring apply throughout. Second, the evidence base remains limited: hotel-specific studies are heterogeneous, some methodological ideas are transferred from other industries, and most papers evaluate prediction more fully than implementation. The five-layer framework offered in this review is therefore a research agenda rather than proof of operational effectiveness. Field studies that report staff overrides, intervention cost, business outcome, and model performance together would provide the stronger evidence now needed.

References

- [1] Sun, S., Law, R., & Hyun, S. S. (2024). Exploration of hotel reservation through mobile online travel agencies. *International Journal of Tourism Research*, 26, e2734. <https://doi.org/10.1002/jtr.2734>
- [2] Antonio, N., Almeida, A., & Nunes, L. (2019). Hotel booking demand datasets. *Data in Brief*, 22, 41-49.
- [3] Xiong, H., Zhao, L. S., Yan, H. L., Zhao, X. L., & Jia, C. R. (2025). Research on precision marketing strategy of hotel customers based on machine learning. *Journal of Central China Normal University (Natural Sciences)*, 59(6), 843-854. (in Chinese)
- [4] Jia, T. Y., Zhang, C., Lu, F. Y., & Huang, Z. F. (2024). Analysis of two-factor experience characteristics and difference mechanism of city hotel based on deep learning. *Journal of Geo-information Science*, 26(2), 499-513. (in Chinese)
- [5] Jishan, M. A., Singh, V., Ghosh, A. K., Alam, M. S., Mahmud, K. R., & Paul, B. (2024). Hotel booking cancellation prediction using applied Bayesian models. *arXiv:2410.16406v2*.
- [6] Xiao, J., Abidin, S. Z., Vermol, V. V., & Gong, B. (2024). Dynamic temporal reinforcement learning and policy-enhanced LSTM for hotel booking cancellation prediction. *PeerJ Computer Science*, 10, e2442.
- [7] Webb, T., Cho, S. R., & Legg, M. P. (2022). Customer lifetime value: A data science approach for hospitality applications. *International Journal of Gaming Hospitality and Tourism*, 2(1).
- [8] Nikolić, M., Marjanović, M., & Radenović, Ž. (2026). Bridging behavioral and emotional intelligence: An interpretable multimodal deep learning framework for customer lifetime value estimation in the hospitality industry. *Mathematical and Computational Applications*, 31(2), 39.
- [9] Chamberlain, B. P., Cardoso, A., Liu, C. H. B., Pagliari, R., & Deisenroth, M. P. (2017). Customer lifetime value prediction using embeddings. *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1753-1762. <https://doi.org/10.1145/3097983.3098123>
- [10] Möller, M., Norlander, A., Zuidberg Dos Martires, P., & De Raedt, L. (2025). Neurosymbolic decision trees. *arXiv:2503.08762v1*.
- [11] Cao, L., Shi, M., & Shroff, N. B. (2025). Provably efficient multi-objective bandit algorithms under preference-centric customization. *arXiv:2502.13457v2*.

Funding

This research received no external funding.

Conflicts of Interest

The authors declare no conflict of interest.

Acknowledgment

This paper is an output of the science project.

Open Access

This chapter is licensed under the terms of the Creative Commons Attribution-NonCommercial 4.0 International License (<http://creativecommons.org/licenses/by-nc/4.0/>), which permits any noncommercial use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

