Prices Do Matter: Modeling Price Competitiveness for Online Hotel Industry

Ruitao Zhu*
Shanghai Jiao Tong University
Shanghai, China
sjtu_zrt@sjtu.edu.cn

Wendong Xiao*
Alibaba Group
Hangzhou, China
xunxiao.xwd@alibaba-inc.com

Yao Yu
Alibaba Group
Hangzhou, China
sichen.yy@alibaba-inc.com

Yangsu Liu
Shanghai Jiao Tong University
Shanghai, China
liu_yangsu@sjtu.edu.cn

Zhenzhe Zheng†
Shanghai Jiao Tong University
Shanghai, China
zhengzhenzhe@sjtu.edu.cn

Shuqi Zhang
Alibaba Group
Hangzhou, China
zhangshuqi.zsq@alibaba-inc.com

Dong Li
Alibaba Group
Hangzhou, China
shiping@taobao.com

Fan Wu
Shanghai Jiao Tong University
Shanghai, China
fwu@cs.sjtu.edu.cn

Abstract

Broad adoption of Online Travel Platforms (OTPs) has led to increasing interest in accurately predicting users' hotel purchase behavior, with price being a key influencer in user decision-making and receiving significant focus. In examining the hotel purchasing process, we identify a pervasive trend that users make extensive price comparisons before making decisions. Existing research primarily focuses on a hotel's own price, neglecting the complex dynamics of market-driven price competition. In this paper, we propose the concept of Marketplace-oriented Hotel Price Competitiveness (MHPC) to model a hotel's pricing competitiveness within the marketplace. Being independent of specific user preferences, MHPC can be applied to and improve various downstream operations in the online hotel industry, such as hotel ranking and pricing, ultimately benefiting hoteliers, users, and OTPs. Furthermore, a novel Hotel Price Competitiveness-aware Purchase Prediction Model (HP^3M) is constructed by incorporating MHPC and demand dynamics into a multi-task learning framework, featuring three distinct submodules to encompass the tri-dimensional facets of MHPC. Extensive offline and online experiments demonstrate HP^3M's effectiveness in predicting hotel purchase probability and enhancing the performance of hotel ranking and pricing compared to the state-of-the-art methods. HP^3M has been fully deployed on Fliggy, a leading OTP in China, serving thousands of hoteliers and tens of millions of users.

*Equal contribution.
†Corresponding author.

KDD '25, Toronto, ON, Canada

ACM ISBN 979-8-4007-1245-6/25/08
https://doi.org/10.1145/3690624.3709420

CCS Concepts

• **Information systems → Recommender systems**.

Keywords

Price Competitiveness, Purchase Probability Prediction, Online Travel Platforms

ACM Reference Format:
Ruitao Zhu, Wendong Xiao, Yao Yu, Yangsu Liu, Zhenzhe Zheng, Shuqi Zhang, Dong Li, and Fan Wu. 2025. Prices Do Matter: Modeling Price Competitiveness for Online Hotel Industry. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.1 (KDD '25), August 3–7, 2025, Toronto, ON, Canada.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3690624.3709420

1 Introduction

Online Travel Platforms (OTPs), such as Booking, Airbnb, and Fliggy[1], have become one of the most prevalent hotel booking channels connecting hoteliers and users. To improve the engagement of both sides, OTPs provide search ranking services for users by prioritizing hotels based on individual preferences, and pricing services for hoteliers by maximizing hotelier revenue through supply-demand analysis and purchase probability prediction. Accurately modeling user purchase behavior is crucial for these services. Numerous studies have established price as a key determinant of user purchasing behavior [13, 21], and prices directly affect the overall revenue of OTPs and hoteliers. However, existing research predominantly simplifies price as a static item feature within purchase probability prediction models, neglecting its dynamic nature and the competitive context within the platform.

In practice, faced with an abundance of choices, users commonly engage in comparing the prices of similar hotels as well as different travel dates to ultimately select the most cost-efficient options. To capture such price comparison patterns, studies like PCNet [18] for personalized ranking, and Airbnb's research [16] for search page revenue maximization, introduce deep feature crosses between user

[1]www.fliggy.com, one of the most popular online travel platforms in China.

profiles and item features. As depicted in Figure 1, these preference-oriented studies indirectly capture the competitive relation of hotels through users' varying degrees of preference for different hotels. However, for the online hotel industry, the price competition among hotels is determined by the overall marketplace demand, rather than individual user preferences. Existing models, deeply integrating user preference with the retrieved hotel features within a query, encounter challenges in deriving the hotel price competitiveness independent of specific user preferences. Consequently, these models are confined to specific downstream tasks.

Figure 1: Comparison between preference-oriented and our proposed marketplace-oriented hotel price competitiveness.

In this work, we propose the concept of *Marketplace-oriented Hotel Price Competitiveness* (MHPC) to represent a hotel's pricing competitiveness within the entire marketplace. MHPC captures the interplay between hotels and the marketplace, which remains independent of individual user preferences, enabling its extensive application to enhance various downstream tasks in the online hotel industry. Specifically, MHPC not only aids hoteliers in establishing competitive pricing to boost their revenue and attract more users but also enhances user experience by achieving more precise hotel purchase probability predictions. Such advancements in predictive modeling facilitate improved performance for downstream applications such as hotel ranking and pricing, thus contributing to increased platform revenue. Therefore, the wide application of MHPC serves to optimize the utilities of hoteliers, users, and OTPs, fortifying the social welfare of the online hotel industry.

We decompose MHPC into three dimensions, *i.e.*, *Intra-platform, Inter-platform and Temporal Competitiveness*, and propose the **H**otel **P**rice Competitiveness-aware **P**urchase **P**rediction **M**odel (**HP^{3}M**) to assimilate MHPC into the predictive framework. HP^3M consists of three submodules to represent the three dimensions of MHPC and incorporates demand dynamics, fluctuations in market demand for hotels, as an additional pivotal factor influencing hotel purchase probability. HP^3M is developed within a multi-task learning framework that concurrently addresses price competitiveness and demand dynamics.

Our main contributions in this work are summarized as follows

• We propose the concept of *Marketplace-oriented Hotel Price Competitiveness* (MHPC), independent of individual user preferences, to enhance various tasks in the online hotel industry, such as ranking and pricing, thereby benefiting hoteliers, users, and OTPs.

• We propose a novel Hotel Price Competitiveness-aware Purchase Prediction Model (HP^3M), developed in a multi-task learning framework incorporating MHPC and demand dynamics modeling. This model features three meticulously engineered submodules capturing MHPC's tripartite dimensions.

• Extensive experiments on large-scale industrial data and online A/B tests show that HP^3M significantly outperforms existing methods for predicting hotel purchase probability, hotel ranking, and hotel pricing. Specifically, our model gained 3.2% relative GMV improvement and 4.9% relative CVR improvement in the four-week online A/B test. HP^3M has been successfully deployed on Fliggy, serving thousands of hoteliers and tens of millions of users.

2 PRELIMINARIES

This section formalizes the concept of MHPC and the integrated hotel purchase prediction problem, aligning with the input features and outputs of our designed modules in Section 3. Our main notations are summarized in Table 1.

Table 1: Main notations.

Notation	Description
$\mathcal{H} = \{h_1, h_2, \cdots, h_{\lvert\mathcal{H}\rvert}\}$	Set of hotels in the marketplace
$\mathcal{U}$	Set of users
p_h^t	Price of hotel h at time t
$\boldsymbol{x}_h^b$	Embedding vector of hotel h's basic hotel features
$\boldsymbol{x}_h^d$	Embedding vector of hotel h's demand dynamics features
$\boldsymbol{x}_{\mathcal{U}}^c, \boldsymbol{x}_{\mathcal{U}}^p$	Historical click and purchase of users
$\boldsymbol{p}_h^o$	Set of hotel h's historical prices collected from other OTPs
$\boldsymbol{P}_h^{S_t}$	Sequence consisting of T nearest historical prices of hotel h to time t
$\boldsymbol{V_I}$	Intra-platform competitiveness of hotel h
$\boldsymbol{V_E}$	Inter-platform competitiveness of hotel h
$\boldsymbol{V_S}$	Temporal competitiveness of hotel h
S_h^p	Price competitiveness-based purchase probability of hotel h (MHPC)
S_h^d	Demand-based purchase probability of hotel h
S_h^w	Overall purchase probability of hotel h

2.1 Hotel Price Competitiveness Modeling

Let $\mathcal{H} = \{h_1, h_2, \cdots, h_{\lvert\mathcal{H}\rvert}\}$ denote the set of hotels in the marketplace and $\mathcal{U}$ denote the set of users. Each hotel $h \in \mathcal{H}$ has basic hotel feature embedding $\boldsymbol{x}_h^b$ (*e.g.*, location, brand, hotel star) and selling price p_h^t at time t.

MHPC reflects a hotel's relative pricing advantage over its competitors; accordingly, accurate modeling of MHPC demands a thorough identification of a hotel's potential competitors. Based on the analysis of user price comparison behavior data on Fliggy, we decompose MHPC into the following three dimensions.

Intra-platform comparison. In e-commerce platforms, user search behavior is guided by specific preferences and intents, particularly on OTPs where users often search for hotels by filtering for desired destinations. This pattern yields a set of candidate hotels with substantial similarity and shared market focus. Within this competitive set, pricing adjustments by any hotel can affect the booking probability of its peers. To capture this competitive relationship within the same platform, we present Definition 2.1.

Definition 2.1. **Intra-platform competitiveness**. Let $\boldsymbol{x}_{\mathcal{U}}^c$ and $\boldsymbol{x}_{\mathcal{U}}^p$ denote the features extracted from historical click and purchase of users, and $\boldsymbol{X}_{\mathcal{H}}^b = \{\boldsymbol{x}_{h_1}^b, \boldsymbol{x}_{h_2}^b, \cdots, \boldsymbol{x}_{h_{\lvert\mathcal{H}\rvert}}^b\}$ denote the basic features of all hotels within our platform. Then the Intra-platform competitiveness of the target hotel h is defined as:

$$\boldsymbol{V_I} = \mathcal{F}_I(\boldsymbol{x_h^b}, \boldsymbol{x_{\mathcal{U}}^c}, \boldsymbol{x_{\mathcal{U}}^p}, \boldsymbol{X_{\mathcal{H}}^b}),$$

where $\mathcal{F}_I$ denotes the submodule **ICRM**, detailed in Section 3.1.1.

Inter-platform comparison. A hotel may be listed on multiple OTPs, with price discrepancies potentially resulting from the distinct marketing strategies of each platform. Assuming equal visibility across OTPs, users are inclined to purchase from the platform offering the most favorable deals. Based on the above analysis, we present Definition 2.2.

Definition 2.2. **Inter-platform competitiveness**. Let $\boldsymbol{p_h^o}$ denote hotel h's historical prices collected from other OTPs. Then the Inter-platform competitiveness of hotel h is defined as:

$$\boldsymbol{V_E} = \mathcal{F}_E(p_h^t, \boldsymbol{p_h^o}),$$

where $\mathcal{F}_E$ denotes the submodule **ECRM**, detailed in Section 3.1.2.

Temporal comparison. Hotel prices undergo periodic variations driven by shifts in market demand and regular promotional campaigns by OTPs. Users typically consider hotel prices over a period before finalizing their travel dates. To characterize the impact of the temporal price dynamics on users' purchasing decisions, we present Definition 2.3.

Definition 2.3. **Temporal competitiveness**. Let $\boldsymbol{P_h^{S_t}} = \left\{(t_h^1, p_h^{t_h^1}), (t_h^1, p_h^{t_h^1}), \cdots, (t_h^T, p_h^{t_h^T})\right\}$ denote the sequence consisting of T nearest historical prices of hotel h to time t. Then the temporal competitiveness of hotel h is defined as:

$$\boldsymbol{V_S} = \mathcal{F}_S(p_h^t, t, \boldsymbol{P_h^{S_t}}),$$

where $\mathcal{F}_S$ denotes the submodule **SCRM**, detailed in Section 3.1.3.

2.2 Integrated Hotel Purchase Prediction

Since MHPC represents a hotel's price competitiveness within the marketplace and lacks a ground-truth label, it is incompatible with direct supervised learning. Hence, we embed MHPC in the task of predicting hotel purchase probability, and define MHPC as the the purchase probability of hotel h with price p_h^t, given price competition within the hotel set $\mathcal{H}$:

$$S_h^p = p(y_h^t = 1 \mid \boldsymbol{x_h^b}, p_h^t, t, \boldsymbol{V_I}, \boldsymbol{V_E}, \boldsymbol{V_S}), \tag{1}$$

where y_h^t denotes whether hotel h is purchased at time t.

Distinct from other e-commerce platforms, hotel offerings on OTPs are characterized by stringent, time-sensitive capacity constraints [23, 24]. Consequently, beyond prices, demand dynamics, *i.e.*, fluctuations in market demand, also affect purchase probability. The demand-based purchase probability is formalized as:

$$S_h^d = p(y_h^t = 1 \mid \boldsymbol{x_h^b}, \boldsymbol{x_h^d}, t), \tag{2}$$

where $\boldsymbol{x_h^d}$ denotes the embedding vector of demand dynamics features (*e.g.*, historical sales, historical exposure unique visitors).

By Incorporating both MHPC and demand dynamics, we formally define the integrated hotel purchase prediction problem. We aim to predict the probability of hotel h is purchased at time t:

$$S_h^w = p(y_h^t = 1 \mid \boldsymbol{x_h^b}, \boldsymbol{x_h^d}, p_h^t, t, \boldsymbol{V_I}, \boldsymbol{V_E}, \boldsymbol{V_S}). \tag{3}$$

3 Methodology

This section presents the details of the **H**otel **P**rice Competitiveness-aware **P**urchase **P**rediction **M**odel (**HP^{3}M**). The central idea of HP^3M is to establish an effective framework to incorporate MHPC and demand dynamics in predicting the hotel purchase probability. As shown in Figure 2, we construct HP^3M within a multi-task learning architecture, featuring three submodules to delineate the three dimensions of MHPC: the Internal, External, and Self Competitiveness Representation Modules (*ICRM, ECRM, and SCRM*), yielding representations $\boldsymbol{V_I}$, $\boldsymbol{V_E}$, and $\boldsymbol{V_S}$. An aggregation layer then computes their respective impact weights, culminating in the unified representation $\boldsymbol{V_C}$. Then, a monotonicity-based sigmoid-shaped scoring function is employed to predict MHPC. Finally, to correct the prediction bias linked to the endogeneity between demand and price features, we introduce an auxiliary feature correlation loss to refine the accuracy of MHPC scoring.

3.1 MHPC-based Hotel Purchase Prediction

As demonstrated in Section 2.1, we decompose MHPC into three dimensions, each represented by a distinct submodule. We elaborate on the three submodules in this section.

3.1.1 Internal Competitiveness Representation Module (ICRM). The goal of ICRM is to represent the intra-platform competitiveness defined in Definition 2.1. Therefore, we aim to solve two problems in this section: (1) how to identify potential competitive hotels; and (2) how to characterize the price competitiveness over competitors.

Potential Competitors Identification. Identifying potential competitors relies on characterizing the hotel similarities. We conceptualize similarity across three dimensions: (1) attribute-based, clustering hotels by fundamental attributes like location, star rating, and price range; (2) click-based, clustering by user click behaviors; (3) purchase-based, clustering by user purchase behaviors. By leveraging the above similarities, we identify the top N hotels most similar to the target hotel h among candidate hotels within the same region, obtaining three distinct competitive groups: $\mathbb{G}_h^R, \mathbb{G}_h^C, \mathbb{G}_h^P$.

Firstly, hotels are clustered by basic hotel features using the K-means algorithm, with K set to the number of commercial zones. Hotels within the same cluster are deemed competitors of comparable quality, where the pricing of one influences the others. The cluster containing target hotel h is denoted as $\mathbb{G}_h^R$.

Note that competitive dynamics among hotels often stem from shared market demand, identifying potential competitors entails utilizing historical marketplace-item interaction data to precisely delineate hotels' demand situation. Accordingly, we employ graph-based representation and item-based collaborative filtering [11] to capture purchase-based and click-based interactions, respectively.

Utilizing the graph-based approach, we construct a weighted bipartite graph with hotel and user nodes from the past 90 days of purchase data. Edge weights reflect the frequency of purchases. Notably: (1) hotels commonly connected by the same user indicate

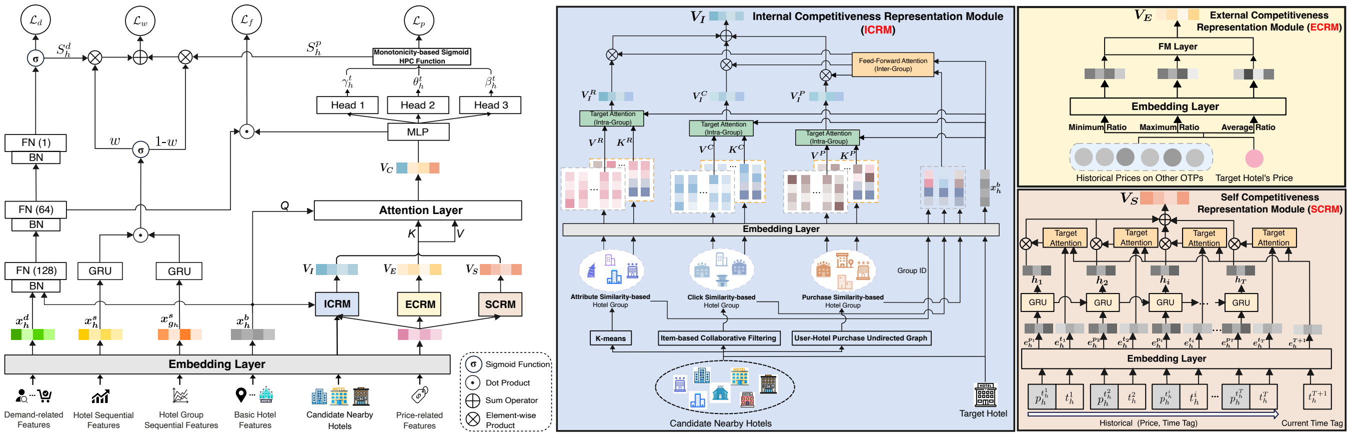

Figure 2: The architecture of Hotel Price Competitiveness-aware Purchase Prediction Model (HP^{3}M).

strong competitive relations due to shared demand, and (2) users linked to the same hotel suggest similar purchasing preferences. Employing *metapath2vec* [5] with metapaths hotel-user-hotel and user-hotel-user, we capture these patterns and derive the embedding of hotel. The similarity between hotels is computed via the inner product of embeddings, with the top N forming the competitive group $\mathbb{G}_h^P$ for hotel h. For item-based collaborative filtering, we leverage recent 30-day user click data, employing the vector inner product for similarity measurement. The top N hotels most similar to hotel h constitute the hotel group $\mathbb{G}_h^C$.

Modeling price competitiveness over competitors. To represent the hotel's price competitiveness over its potential competitors, we need to model two distinct types of competitive relationships: (1) *intra-group competition*; (2) *inter-group competition.*

The price comparison is expressed through the price ratio p_h^t/p_r^t between the target hotel h and each competitive hotel r. For each competitive group $k \in R, C, P$, we define $\boldsymbol{V^k} \in \mathbb{R}^{N\times d}$ and $\boldsymbol{K^k} \in \mathbb{R}^{N\times d}$ as matrices representing the price ratio and basic hotel features, respectively. The j-th row in these matrices corresponds to the embedding vectors for the price ratio and basic features of the j-th hotel in group k. Additionally, $\boldsymbol{G^k} \in \mathbb{R}^d$ denotes the embedding vector for the Group ID of the competitive group k.

As shown in Figure 2, ICRM constructs a target attention layer for each group, with the embedding vector of target hotel $\boldsymbol{x_h^b}$ as the query vector, $\boldsymbol{K^k}$ as the key vector, and $\boldsymbol{V^k}$ as the value vector, to learn the intra-group competition representation $\boldsymbol{V_I^k}$:

$$\boldsymbol{V_I^k} = softmax\left(\frac{\boldsymbol{x_h^b}(\boldsymbol{K^k})^T}{\sqrt{d}}\right)\boldsymbol{V^k}, k \in \{R, C, P\}.$$

Next, we adopt the feed-forward attention network to model the inter-group competition $\boldsymbol{V_I}$:

$$\boldsymbol{V_I} = softmax(w_R, w_C, w_P)[\boldsymbol{V_I^R}, \boldsymbol{V_I^C}, \boldsymbol{V_I^P}]^T,$$

where $w_k \in \mathbb{R}$ represents the competitiveness score computed through a multilayer perceptron (MLP) with $\boldsymbol{x_h^b}$ and $\boldsymbol{G^k}$ as input.

3.1.2 *External Competitiveness Representation Module (ECRM).*

A hotel often concurrently lists on various platforms, with prices differing due to distinct platform marketing strategies. Price-sensitive users tend to compare rates for the same hotel across multiple OTPs. Consequently, ECRM aims to represent the inter-platform competitiveness defined in Definition 2.2.

Hotels are more likely to receive bookings on the platform that offers a lower price. Accordingly, we employ the price ratio between the target hotel's price p_h^t and prices $\boldsymbol{p_h^o}$ from other OTPs to model this price discrepancy. Nevertheless, the availability of cross-OTP pricing data is often limited due to privacy concerns, resulting in data sparsity and challenges in establishing robust models. To mitigate these issues, we utilize DeepFM [7], adept at capturing feature interaction for sparse data. As depicted in Figure 2, the input of ECRM encompasses embeddings derived from the ratio of p_h^t to the minimum ($\boldsymbol{p_h^{min}}$), maximum ($\boldsymbol{p_h^{max}}$), and average ($\boldsymbol{p_h^{avg}}$) price within $\boldsymbol{p_h^o}$. To further minimize noises from other OTPs' price data, a selector block is implemented above the FM layer, yielding the representation of inter-platform competitiveness:

$$\boldsymbol{V_E} = \mathrm{FM}\left([\boldsymbol{p_h^{min}}, \boldsymbol{p_h^{max}}, \boldsymbol{p_h^{avg}}]\right) \odot \boldsymbol{\sigma}\left(\mathrm{FM}([\boldsymbol{p_h^{min}}, \boldsymbol{p_h^{max}}, \boldsymbol{p_h^{avg}}])\right).$$

where $\boldsymbol{\sigma}(\cdot)$ denotes the fully connected layer with sigmoidal activation and "$\odot$" denotes the element-wise product operator.

3.1.3 *Self Competitiveness Representation Module (SCRM).*

As discussed in Section 2.1, users' hotel purchasing decisions are susceptible to price fluctuations driven by temporal factors like demand seasonality and platform promotions. SCRM aims to represent the competitiveness of the current price over historical prices, as defined in Definition 2.3. Specifically, SCRM focuses on two key temporal patterns: long-term global price consistency and short-term local price fluctuation. Long-term patterns indicate a hotel's general pricing trends and market standing, while short-term trends capture immediate price variations and market reactions.

As shown in Figure 2, we design a Gated Recurrent Units (GRU) [4] based attention mechanism to fuse the long-term global price consistency and short-term local price fluctuation. Given the historical price sequence $\boldsymbol{P_h^{S_t}} = \left\{(t_h^1, p_h^{t_h^1}),\ \ (t_h^1, p_h^{t_h^1}), \cdots, (t_h^T, p_h^{t_h^T})\right\}$, we transform the price ratio $p_h^t/p_h^{t_h^i}, i \in \{1, 2, \cdots, T\}$ and each corresponding time tag t_h^i to embedding vectors $\boldsymbol{e_h^{p_i}}$ and $\boldsymbol{e_h^{t_i}}$. The embedding vector of the current time tag is denoted as $\boldsymbol{e_h^{T+1}}$. Iteratively, the i-th GRU cell is defined as:

$$\boldsymbol{h_i} = GRU(\boldsymbol{h}_{i-1}, \boldsymbol{e_h^{p_i}}).$$

Furthermore, to capture the impact weight of each historical price on the current price, we adopt the target attention network with $\boldsymbol{e_h^{T+1}} \in \mathbb{R}^d$ as the query vector, $\boldsymbol{E_h^t} = [\boldsymbol{e}_h^{t_1}, \boldsymbol{e}_h^{t_2}, \cdots, \boldsymbol{e}_h^{t_T}]$ as the key vector, and $\boldsymbol{H_h} = [\boldsymbol{h}_1, \boldsymbol{h}_2, \cdots, \boldsymbol{h}_T]$ as the value vector. Thus, the representation of the temporal competitiveness $\boldsymbol{V_S}$ is formalized as:

$$\boldsymbol{V_S} = softmax\left(\frac{\boldsymbol{e_h^{T+1}}\boldsymbol{E_h^t}}{\sqrt{d}}\right)\boldsymbol{H_h}^T.$$

3.1.4 MHPC Representation Aggregation Layer. Note that the impact of $\boldsymbol{V_I}, \boldsymbol{V_E}, \boldsymbol{V_S}$ on users' purchasing decisions varies across different hotels. For example, for budget hotels, with an abundance of similar options readily available, the intra-platform competitiveness $\boldsymbol{V_I}$ often plays a more significant role. For resort hotels, which typically exhibit notable price changes over time, the temporal competitiveness $\boldsymbol{V_S}$ tends to be more influential.

The three dimensions of MHPC are aggregated into a unified representation by adopting the attention mechanism to account for the varying impact weights, formulated as:

$$\boldsymbol{V_C} = \sum_{i \in \{I,E,S\}} \alpha_i \boldsymbol{V_i}.$$

As previously indicated, the weight α_i varies contingent upon the hotel category, suggesting that the basic features of hotels provide essential information for user decision-making. Therefore, we employ the attention mechanism, with x_h^b as the query vector:

$$\alpha_i = \frac{\exp\left(\phi(\boldsymbol{x_h^b}, \boldsymbol{V_i})\right)}{\sum_{k \in \{I,E,S\}} \exp\left(\phi\left(\boldsymbol{x_h^b}, \boldsymbol{V}_k\right)\right)},$$

where the function ϕ encapsulates the similarity between the query vector x_h^b and the key vector $\boldsymbol{V}_i$, indexed by $i \in \{I, E, S\}$, and is parameterized by a MLP, which is formally defined as:

$$\phi(\boldsymbol{x_h^b}, \boldsymbol{V_i}) = MLP\left(\boldsymbol{x_h^b} + \boldsymbol{V_i} + \boldsymbol{x_h^b} \otimes \boldsymbol{V_i}\right),$$

where $\otimes$ denotes the element-wise product operator.

3.1.5 Monotonicity-based Sigmoid MHPC Scoring Function. Intuitively, a hotel's likelihood of purchase increases when it offers a lower price, assuming other factors remain constant. Previous research has characterized this monotonic relationship as a sigmoid curve [17, 18], and we observe a similar pattern from Fliggy's hotel purchase data. Additionally, as MHPC is employed for evaluating hotel quality and aiding merchants in pricing, the MHPC score must maintain a strict monotonic relationship with prices, thereby reinforcing hoteliers' confidence in the reasonableness and interpretability of OTP's predictive model. Furthermore, the sigmoid-shaped monotonic relationship provides robust a priori insights into empirical phenomena. This domain knowledge guides and restricts the modeling process, reducing the likelihood of overfitting and enhancing generalization capabilities, in contrast to arbitrary and uninformed learning approaches. Hence, we design the monotonicity-based sigmoid MHPC scoring function. The MHPC $S_h^p \in (0, 1]$ is defined as the purchase probability at price p_h^t:

$$S_h^p = min\left\{\frac{1}{1 + (1 + \gamma_h)^{-(\bar{p}_h - p_h^t - \beta_h)}} + \frac{1}{2} \times \theta_h, 1\right\}, \qquad (4)$$

where $\bar{p}_h$ denotes the historical average price of hotel h, and parameters $\gamma_h \in (0, 1), \beta_h \in (0, +\infty), \theta_h \in (0, 1)$ control the shape of the MHPC scoring function. These parameters are derived from a shared-bottom two-layer MLP with input $\boldsymbol{V_C}$ and three output heads. Specifically, the heads for γ_h and θ_h use a fully connected layer with sigmoid activation, while the head for β_h uses a fully connected layer with a ReLU activation.

3.2 Integrated Hotel Purchase Prediction

As discussed in Section 2.2, aside from pricing factors, demand dynamics significantly affect hotel purchase probability, with such effects arising from seasonal and cyclical demand patterns driven by variations in popularity. Illustratively, hot spring resorts typically see increased demand in winter relative to summer, whereas business hotels experience heightened occupancy during weekdays as compared to weekends. Therefore, we predict the demand-based hotel purchase probability by incorporating demand-related feature embedding $\boldsymbol{x_h^d}$ and basic hotel feature embedding $\boldsymbol{x_h^b}$:

$$S_h^d = \text{MLP}\left([\boldsymbol{x_h^d},\ \boldsymbol{x_h^b}]\right).$$

Note that $\boldsymbol{x_h^d}$ captures numerical demand-related features such as historical orders, which may fall short in capturing demand volatility. Events causing abrupt demand spikes, like concerts, can markedly diminish the role of price in purchasing decisions. To address this, we incorporate sequential features that track time-series variations in search and click Unique Visitors (UV), order volume, and occupancy rates. However, the constrained inventory of an individual hotel does not fully represent market demand. For instance, a hotel's room inventory may stand at 10, whereas a concert in the area could generate a potential demand for 100 rooms, making the increase from 0 to 10 insufficient to accurately reflect the true scope of demand.

To effectively capture regional demand shifts, we utilize features of hotels in the target hotel's commercial area. Specifically, let $\boldsymbol{x_h^s} = (x_h^{s_1}, x_h^{s_2}, \cdots, x_h^{s_D})$ denote historical sequence of demand-related features for hotel h spanning the last D days, and $\boldsymbol{x_{g_h}^s} = (x_{g_h}^{s_1}, x_{g_h}^{s_2}, \cdots, x_{g_h}^{s_D})$ denote the corresponding sequence for hotels within the same commercial area as h. These two sequences are encoded via two GRU structures to ascertain the impact weights of demand-based hotel purchase probability, defined as:

$$w_h = \sigma(\text{GRU}(\boldsymbol{x_h^s}) \cdot \text{GRU}(\boldsymbol{x_{g_h}^s})),$$

where the GRU utilizes the same architecture described in Section 3.1.3. Then the integrated hotel purchase probability corresponding to Equation (3) is formulated as:

$$S_h^w = w_h \times S_h^d + (1 - w_h) \times S_h^p.$$

3.3 Loss Function and Model Training

Our model is trained in the multi-task framework to minimize the combined MHPC-based and demand-based purchase prediction losses. Given M samples, we adopt the negative log-likelihood loss function to train all parameters of HP^3M:

$$\mathcal{L}_j = -\frac{1}{M}\sum_{i=1}^{M}\left(y_i \log(S^j_{h_i}) + (1-y_i)\log(1-S^j_{h_i})\right),\ j \in \{p, d, w\},$$

where $y_i \in \{0, 1\}$ denotes the purchase label for the i-th sample.

The separate optimization of demand and price competitiveness-based purchase probabilities by two modules fails to capture the inherent correlations between price and demand features, introducing bias in the MHPC scoring task. To address this endogeneity [17], we introduce a feature correlation loss using the inner product of the intermediate representations from the two sub-networks as an auxiliary loss:

$$\mathcal{L}_f = -MLP(\boldsymbol{V_D}) \odot MLP(\boldsymbol{V_C}), \tag{5}$$

where $\mathcal{L}_f$ is minimized when $\boldsymbol{V_D}$ and $\boldsymbol{V_C}$ are totally uncorrelated. The effectiveness of the proposed auxiliary loss will be illustrated in Section 4.2.1 through an ablation study.

The total loss of the overall framework is defined as:

$$\mathcal{L} = \mathcal{L}_w + \lambda\mathcal{L}_d + \gamma\mathcal{L}_p + \beta\mathcal{L}_f,$$

where λ, β, γ are hyper-parameters [2] to balance these tasks.

4 Experiments

In this section, we conduct extensive offline experiments and online A/B tests to evaluate the effectiveness of the proposed HP^3M. We mainly focus on the following questions:

- **RQ1**: How accurate is the MHPC score obtained via HP^3M?
- **RQ2**: How do different modules of HP^3M contribute to the model performance?
- **RQ3**: How does HP^3M perform on the **hotel pricing task** compared to baseline methods?
- **RQ4**: How does HP^3M perform on the **hotel ranking task** compared to baseline methods?
- **RQ5**: How does our proposed HP^3M perform on the platform's live production environment?

4.1 Experiment Setup

4.1.1 Dataset. Given the absence of publicly accessible datasets in online hotel industry for the hotel pricing and ranking tasks, we construct a real-world dataset[3] based on hotel transaction data from *Fliggy*. The dataset is split into **Regular Season** and **Peak Season**, as shown in Table 2. Specifically, the Regular Season dataset contains data from June 1 to June 14, 2023, excluding Chinese holidays. Conversely, the Peak Season dataset, spanning from June 15 to June 28, 2023, coincides with a Chinese holiday travel peak. Peak Seasons exhibit higher hotel demand and price volatility, leading to distinct data distributions compared to Regular Seasons. We use the first 7-day data for model training and the remaining 7-day data for testing.

[2]In the offline experiments of this study, $\lambda = 0.1, \gamma = 0.1, \beta = 0.00017$.
[3]https://tianchi.aliyun.com/dataset/184364

Table 2: Statistics of datasets.

Dataset		#Users	#Hotels	#Bookings
Regular Season	training	3.43 M	240,596	900,176
	testing	2.99 M	242,470	780,152
Peak Season	training	4.02 M	237,532	954,425
	testing	3.12 M	239,213	943,538

4.1.2 Baseline Methods. To compare the proposed model's performance across two distinct tasks, we introduce task-specific baselines, including state-of-the-art (SOTA) methods for both.

Hotel Pricing. The hotel pricing task aims at finding the optimal price P^{t*}_h for hotel h at time t to maximize expected revenue, independent of user characteristics or query features:

$$P^{t*}_h = \arg\max_{P^t_h} P^t_h \cdot r^t_h(P^t_h), \tag{6}$$

where r^t_h denotes the predicted purchase probability.

- *Meituan pricing* [20] adopts a DNN model to regress the suggested price for a certain room at a specific date.
- *Airbnb pricing* [19] employs a customized regression model to obtain the suggested price and adjust the price by incorporating hotelier goals and special events.
- *PEM pricing* [23] is the SOTA hotel pricing model, which introduces the elasticity of demand function into hotel pricing to determine suggested prices dynamically.

Hotel Ranking. The hotel ranking task entails ordering retrieved hotels by purchase likelihood, accounting for user preferences, search intent, item attributes, and contextual features.

- *PCNet* [18] is the SOTA hotel purchase probability prediction model, incorporating price competitiveness assessment.
- *PCNet-PC* is a variant of PCNet which removes all price competitiveness modules.

4.1.3 Performance Metrics.

- **Rank Biased Overlap (RBO) [15]** quantifies the similarity between two ranked lists with the given weight of each position:

$$\text{RBO}(R, I) = (1-p)\sum_{i=1}^{\infty} p^{i-1}\frac{|R_{1:i} \cap I_{1:i}|}{i}, \tag{7}$$

where $R_{1:i}$ and $I_{1:i}$ denote the top i items in list R and list I, respectively. The parameter p determines how steeply the weights fall, *i.e.*, the smaller p, the higher weight for top positions. RBO is utilized to evaluate the effectiveness of the derived MHPC score in modeling price competition among hotels.

For the hotel pricing task, the absence of optimal price labels renders direct comparison of predictions to ground truth infeasible. Fortunately, the monotonic relationship between price and purchase probability allows us to ascertain the range within which the optimal price lies. For instance, given a historical hotel sample priced at P^t_h that failed to result in a transaction, the theoretical optimal price lies in $[\infty, P^t_h]$ (lower price, higher purchase probability). Based on the above analysis, we employ the following two evaluation metrics used in [19, 20, 25] to compare the performance of various hotel pricing strategies.

Denote the actual price of hotel h as P_h, the suggested price as P^{sug}_h, and the number of samples in each case in Table 3 as a, b, c, d.

Table 3: Number of samples in each case.

	Booking	Non-Booking
$P_h^{sug} \geq P_h$	a	b
$P_h^{sug} < P_h$	c	d

• **Price Decrease Recall (PDR)** measures the percentage of samples where the suggested prices are lower than actual prices among all non-booking samples.

$$PDR = \frac{d}{b+d}. \tag{8}$$

• **Price Increase Recall (PIR)** measures the percentage of samples where the suggested prices are no lower than actual prices among all booking samples.

$$PIR = \frac{a}{a+c}. \tag{9}$$

A higher value of PDR and PIR indicates that the model's suggested price is closer to the theoretically optimal price. However, it should be noted that the singular performance of PDR or PIR does not adequately reflect a model's pricing effectiveness. Consider a toy case where the model always suggests a price of zero, resulting in PDR=1 and PIR=0; focusing solely on PDR would yield a misleading conclusion. Therefore, to account for both metrics, we adapt the concept of F1-score from statistical learning to pricing tasks, introducing the following comprehensive evaluation metric.

• **Price Recall F1-score (PRF1)** measures the overall performance of a pricing algorithm.

$$PRF1 = \frac{2 \times PDR \times PIR}{PDR + PIR}. \tag{10}$$

For evaluating revenue and ranking performance, we employ the following commonly used metrics:

• **Gross Merchandise Volume (GMV)** measures the total revenue of hotels sold on the platform.

• **AUC**, **Logloss**, **NDCG@K** measures the ranking performance.

4.1.4 Implementation Details. In HP^3M, all feature embeddings are sized at 16, and the number of units of GRU is set to 64. The feedforward networks in ICRM, ECRM, SCRM, and the aggregation layer have two hidden layers with 64 and 32 neurons. Each of $\gamma_h^t, \beta_h^t, \theta_h^t$ defined in Equation (4) is calculated by a three-layer MLP with sizes of 128, 64, and 1. HP^3M is trained through the Adam optimizer with a learning rate of 0.001 and a batch size of 512. To avoid overfitting, we add dropout with a probability of 0.2.

For the baseline methods, the regression module of Meituan pricing and Airbnb pricing method is set to a three-layer MLP with sizes of 256, 64, and 1. The multi-sequence fusion model of PEM pricing contains three convolution kernels (sizes: 64, 128, 64) and a two-layer Bi-GRU module (number of units: 64, 32). Other parameters follow the original papers' recommendations.

4.2 Offline Experiments

4.2.1 Accuracy of MHPC Score (RQ1) & Ablation Study (RQ2). Given that the MHPC score is designed to indicate a hotel's price competitiveness and lacks a ground-truth label, we assess the predicted MHPC score's accuracy through a comparison of the relative hotel rankings within the same commercial district, as manually annotated by platform business operatives[4]. Furthermore, to illustrate the contributions of our proposed submodules and the auxiliary loss $\mathcal{L}_f$, we construct the following variants for comparison.

• **HP^{3}M-ICRM**: a variant of HP^3M which deletes ICRM;

• **HP^{3}M-ECRM**: a variant of HP^3M which deletes ECRM;

• **HP^{3}M-SCRM**: a variant of HP^3M which deletes SCRM;

• **HP^{3}M $-$ $\mathcal{L}_f$**: a variant of HP^3M which is trained without the auxiliary loss $\mathcal{L}_f$ defined in Equation (5).

• **HP^{3}M-simple**: a simple baseline to model price competitiveness, which replaces SCRM, ECRM, and ICRM with a 3-layer MLP using price ratios of target hotels across inter-platform, intra-platform, and temporal prices as inputs.

Each offline experiment is repeated 5 times with different random seeds and each result is presented in the form of mean *pm* standard. The experimental results are shown in Table 4, from which we observe: (1) The proposed ICRM, ECRM, and SCRM notably boost model performance across the intra-platform, inter-platform, and temporal dimensions, confirming the effectiveness of the tri-dimensional MHPC modeling strategy; (2) Omitting $\mathcal{L}_f$ diminishes RBO by 1.3% and 1.5% in peak and regular seasons, respectively, confirming the importance of addressing the feature correlation issue in hotel purchase prediction. (3) HP^3M outperforms HP^3M-simple with an 11.9% and 10.6% increase in RBO during peak and regular seasons, respectively, demonstrating the necessity and effectiveness of finely modeling price comparisons.

Table 4: Performance test for MHPC score & Ablation study (mean±std).

	Peak Season		Regular Season	
	RBO (p = 0.5) ↑	RBO (p = 0.7)↑	RBO (p = 0.5)↑	RBO (p = 0.7)↑
HP^{3}M	**0.7112** ±0.0086	**0.8027** ±0.0070	**0.7120** ±0.0083	**0.8148** ±0.0062
HP^3M-ICRM	0.6293 ±0.0104	0.7194 ±0.0077	0.6321 ±0.0052	0.7379 ±0.0045
HP^3M-ECRM	0.6901 ±0.0066	0.7822 ±0.0068	0.6910 ±0.0072	0.7893 ±0.0053
HP^3M-SCRM	0.6755 ±0.0061	0.7592 ±0.0059	0.6721 ±0.0068	0.7687 ±0.0047
HP^3M-$\mathcal{L}_f$	0.7020 ±0.0070	0.7914 ±0.0060	0.7004 ±0.0067	0.8022 ±0.0053
HP^3M-simple	0.6323 ±0.0063	0.7210 ±0.0048	0.6388 ±0.0075	0.7422 ±0.0052

4.2.2 Performance on Hotel Pricing Task (RQ3) . For the hotel pricing task defined in Equation (6), the comparison results are detailed in Table 5, yielding the following insights:

(1) HP^3M attains optimal balance between PDR and PIR across both datasets, as evidenced by superior PRF1 performance, indicating that our model approximates the theoretical optimal price more closely than baseline methods. Specifically, HP^3M shows a significant improvement compared to all baselines on PDR, while approaching the performance of Meituan Pricing, which directly regresses the booking prices and thus naturally has higher PIR. Note that these metrics are merely tentative indicators of potential revenue growth. In Section 4.3, we further assess the effectiveness

[4]Business operatives within each commercial district identify comparable hotels based on expertise and rank them by price.

of our proposed model in optimizing platform revenue through online A/B tests.

(2) HP^3M exhibits demand-responsive adaptability, with higher PIR in peak season indicating price increases for revenue maximization during high demand, and higher PDR in regular season reflecting price reductions to attract users during lower demand phases. This adaptive pricing behavior underscores the model's capacity to balance profitability and market competitiveness across varying demand conditions.

4.2.3 Performance on Hotel Ranking Task (RQ4). HP^3M is designed to predict the integrated hotel purchase probability during the hotel selection stage on OTPs before user request arrivals, thus excluding user profiles from its input features. In contrast, PCNet, the comparative baseline, is applied during the subsequent hotel ranking stage, utilizing user profiles. Therefore, to isolate the influence of user profiles and assess HP^3M's contribution to refining price competitiveness in the hotel ranking task, we introduce PCNet-PC+HP^3M. This variant substitutes PCNet's price competitiveness-aware modules with the aggregated MHPC representation $\boldsymbol{V_C}$ from HP^3M, thus normalizing for user profile disparities.

As illustrated in Table 6, we obtain the following observations. (1) PCNet-PC+HP^3M and PCNet show marked improvements over PCNet-PC across all metrics, highlighting the contribution of price competitiveness modeling in predicting hotel purchase probabilities. (2) PCNet-PC+HP^3M significantly outperforms PCNet on all metrics, evidencing the benefits of finely-grained, tri-dimensional MHPC modeling with dynamic impact weight.

4.3 Online A/B Test (RQ5)

To evaluate the performance of HP^3M in the real world, we conduct four-week online A/B tests in the Fliggy hotel pricing and search ranking system from July 1 to July 28, 2023, covering both peak days (summer holidays and weekends) and regular days. The online deployment workflow for HP^3M is provided in Appendix A.1.

For the hotel pricing task, we utilize PEM, which outperforms other baselines in offline experiments, as the online baseline. Curve-1 (red line) in Figure 3 shows that compared to PEM, HP^3M achieves an average 3.2% relative improvement in GMV, demonstrating the superiority of our method in improving revenue. Note that Curve-1 exhibits cyclical fluctuations. This is expected, as GMV is directly correlated with market demand, which is notably variable and cyclical. Furthermore, the peaks in online results coincide with weekend demand surges, which is consistent with our offline experimental findings comparing peak and regular seasons. This trend underscores the demand-responsive adaptability of our proposed HP^3M.

For the hotel search ranking task, we utilize PCNet as the online baseline and evaluate HP^3M from the following two perspectives.

• We rerank hotels within the same commercial area on the PCNet-generated list according to their MHPC scores calculated by HP^3M. As depicted by Curve-2 in Figure 3, compared to ranking solely by PCNet, we achieved an average 1.4% increase in conversion rate (CVR) through reranking based on HP^3M.

• We deploy HP^3M within a specific Fliggy hotel search scenario, targeting price-sensitive users who utilize price range filters. Hotels of each commercial area are ranked based on the MHPC score from HP^3M, the purchase probability from PCNet, and the greedy method (ascending price order), respectively. As shown in Curve-3 and Curve-4 in Figure 3, over four weeks, our model outperformed PCNet by an average of 4.9% in CVR and surpassed the Greedy method by 13.2%, demonstrating that our proposed MHPC score accounts for both hotel quality and price, and markedly enhances the performance of hotel ranking.

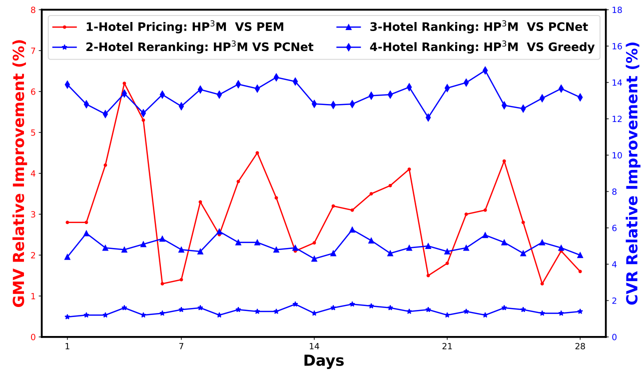

Figure 3: Four-week online A/B tests on two tasks at Fliggy.

5 Related Work

Price, an important factor influencing user decision-making [3, 8, 10], has been extensively studied within e-commerce platforms. Prevailing studies primarily incorporate price as a static item feature [1, 9], omitting granular price modeling. The importance of explicit price modeling are highlighted in many researches [2, 6, 12, 14, 21, 22] in order to refine predictive precision. Schafer *et al.* [12] first propose a framework elucidating price's impact on user behavior in recommender systems. Chen *et al.* [2] employ collective matrix factorization to derive user-item, item-price, and user-price matrices, mitigating the challenge of recommending to history-lacking users. Zheng *et al.* [22] introduce a price-aware Graph Convolution Network encompassing user, item, category, and price nodes. Greenstein-Messica [6] leverage Gaussian distribution to model the Willingness To Pay (WTP) of users by introducing user's price discount sensitivity, and combine the WTP distribution with seller reputation to calculate user-price matching score to improve prediction accuracy. Zhang *et al.* [21] design a customized heterogeneous hypergraph network to interlink prices and user preferences in session-based recommendations. Wan *et al.* [14] incorporate dynamic pricing of item and price elasticity into recommender systems via modeling users' decision making into three stages: category purchase, product choice and purchase quantity.

The aforementioned studies focus solely on modeling individual item pricing, disregarding the impact of inter-item price competition on users' purchasing behavior, leading to suboptimal purchase probability prediction performance. Wu *et al.* [18] propose PCNet, which accounts for item comparison prices to model users' price comparison behaviors, while Airbnb's study [16] views hotels on the same search page as competitors, setting prices based on estimated user value distributions for revenue optimization. Although existing efforts have demonstrated remarkable performance, these methods integrate user preferences with the retrieved hotels within

Table 5: Performance comparison on hotel pricing (mean±std).

	Peak Season			Regular Season		
Methods	PDR↑	PIR↑	PRF1↑	PDR↑	PIR↑	PRF1↑
Meituan Pricing	0.3286 ±0.0054	**0.6803** ±0.0193	0.4432 ±0.0022	0.4014 ±0.0036	**0.6598** ±0.0103	0.4991 ±0.0039
Airbnb Pricing	0.4488 ±0.0132	0.5648 ±0.0229	0.4996 ±0.0028	0.5145 ±0.0025	0.5449 ±0.0112	0.5294 ±0.0027
PEM Pricing	0.5581 ±0.0108	0.5002 ±0.0133	0.5274 ±0.0049	0.5895 ±0.0032	0.5418 ±0.016	0.5647 ±0.0124
$\boldsymbol{HP^3M}$	**0.6148** ±0.0146	0.6264 ±0.0141	**0.6202** ±0.0043	**0.6334** ±0.0134	0.5707 ±0.0166	**0.6001** ±0.0059

Table 6: Performance comparison on hotel ranking (mean±std).

	Peak Season				Regular Season			
Methods	AUC	Logloss	NDCG@5	NDCG@10	AUC	Logloss	NDCG@5	NDCG@10
PCNet	0.8082±0.0023	0.2334±0.0026	0.6925±0.0042	0.7410±0.0038	0.7874±0.0018	0.2982±0.0024	0.6956±0.0038	0.7105±0.0034
PCNet-PC	0.7964±0.0030	0.2811±0.0051	0.6466±0.0045	0.7184±0.0041	0.7762±0.0021	0.3353±0.0041	0.6520±0.0038	0.6876±0.0033
PCNet-PC+$\boldsymbol{HP^3M}$	**0.8120**±0.0028	**0.2295**±0.0044	**0.7102**±0.0042	**0.7523**±0.0036	**0.7899**±0.0020	**0.2901**±0.0039	**0.6988**±0.0040	**0.7190**±0.0030

a query, which overlooks the broader competitive market dynamics, thus limiting their applicability to particular downstream tasks.

6 CONCLUSION

Previous research underscores the significance of price in users' purchasing decisions, yet overlooks the competitive pricing among hotels. To address this gap, we introduce the concept of marketplace-oriented hotel price competitiveness (MHPC), encompassing Intra-platform, Inter-platform, and Temporal Competitiveness, to provide a comprehensive assessment of a hotel's competitive pricing within the market. This foundation enables our design of the Hotel Price Competitiveness-aware Purchase Prediction Model (HP^3M), a novel framework that integrates MHPC and demand dynamics into a multi-task learning approach, with dedicated sub-modules for each MHPC dimension. Furthermore, our approach tackles the challenge of feature correlation through an auxiliary loss and incorporates a monotonicity-based sigmoid function for robust MHPC scoring. Comprehensive offline and online evaluations establish HP^3M's superiority in pricing and ranking tasks over SOTA baselines. Specifically, a four-week online A/B test indicates a 3.2% relative GMV enhancement and a 4.9% relative CVR improvement. HP^3M has been fully deployed on Fliggy. Future works aim to extend MHPC to item-centric price competitiveness and its application to broader e-commerce contexts.

Acknowledgments

This work was supported in part by China NSF grant No. U2268204, 62322206, 62132018, 62025204, 62272307, 62372296, in part by Alibaba Group through Alibaba Innovative Research Program. The opinions, findings, conclusions, and recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the funding agencies or the government.

References

[1] Grigor Aslanyan, Aritra Mandal, Prathyusha Senthil Kumar, Amit Jaiswal, and Manojkumar Rangasamy Kannadasan. 2020. Personalized ranking in ecommerce search. In *Companion Proceedings of the Web Conference 2020*. 96–97.

[2] Jia Chen, Qin Jin, Shiwan Zhao, Shenghua Bao, Li Zhang, Zhong Su, and Yong Yu. 2014. Does product recommendation meet its Waterloo in unexplored categories? No, price comes to help. In *Proceedings of the 37th international ACM SIGIR conference on Research & development in information retrieval*. 667–676.

[3] Shih-Fen S Chen, Kent B Monroe, and Yung-Chien Lou. 1998. The Effects of Framing Price Promotion Messages on Consumers' Perceptions and Purchase Intentions. *Journal of retailing* 74, 3 (1998), 353–372.

[4] Kyunghyun Cho, Bart van Merrienboer, Çaglar Gülçehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL*, Alessandro Moschitti, Bo Pang, and Walter Daelemans (Eds.). ACL, 1724–1734.

[5] Yuxiao Dong, Nitesh V Chawla, and Ananthram Swami. 2017. metapath2vec: Scalable representation learning for heterogeneous networks. In *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining*. 135–144.

[6] Asnat Greenstein-Messica and Lior Rokach. 2018. Personal Price Aware Multi-Seller Recommender System: Evidence from eBay. *Knowledge-Based Systems* 150, C (2018), 14–26.

[7] Huifeng Guo, Ruiming Tang, Yunming Ye, Zhenguo Li, and Xiuqiang He. 2017. DeepFM: A Factorization-Machine based Neural Network for CTR Prediction. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, Carles Sierra (Ed.). ijcai.org, 1725–1731.

[8] Sangman Han, Sunil Gupta, and Donald R Lehmann. 2001. Consumer price sensitivity and price thresholds. *Journal of retailing* 77, 4 (2001), 435–456.

[9] Jinhong Huang, Yang Li, Shan Sun, Bufeng Zhang, and Jin Huang. 2020. Personalized flight itinerary ranking at fliggy. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*. 2541–2548.

[10] Lakshman Krishnamurthi, Tridib Mazumdar, and SP Raj. 1992. Asymmetric Response to Price in Consumer Brand Choice and Purchase Quantity Decisions. *Journal of consumer research* 19, 3 (1992), 387–400.

[11] Badrul Sarwar, George Karypis, Joseph Konstan, and John Riedl. 2001. Item-based collaborative filtering recommendation algorithms. In *Proceedings of the 10th international conference on World Wide Web*. 285–295.

[12] J Ben Schafer, Joseph Konstan, and John Riedl. 1999. Recommender systems in e-commerce. In *Proceedings of the 1st ACM conference on Electronic commerce*. 158–166.

[13] Panniello Umberto. 2015. Developing a price-sensitive recommender system to improve accuracy and business performance of ecommerce applications. *International Journal of Electronic Commerce Studies* 6, 1 (2015), 1–18.

[14] Mengting Wan, Di Wang, Matt Goldman, Matt Taddy, Justin Rao, Jie Liu, Dimitrios Lymberopoulos, and Julian McAuley. 2017. Modeling Consumer Preferences and Price Sensitivities from Large-Scale Grocery Shopping Transaction Logs. In *Proceedings of the 26th International Conference on World Wide Web* (Perth,

Australia) *(WWW '17)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 1103–1112.
[15] William Webber, Alistair Moffat, and Justin Zobel. 2010. A similarity measure for indefinite rankings. *ACM Transactions on Information Systems (TOIS)* 28, 4 (2010), 1–38.
[16] Jiawei Wen, Hossein Vahabi, and Mihajlo Grbovic. 2019. Revenue Maximization of Airbnb Marketplace using Search Results. *arXiv preprint arXiv:1911.05887* (2019).
[17] Jeffrey M Wooldridge. 2015. *Introductory econometrics: A modern approach*. Cengage learning.
[18] Han Wu, Hongzhe Zhang, Liangyue Li, Zulong Chen, Fanwei Zhu, and Xiao Fang. 2022. Cheaper Is Better: Exploring Price Competitiveness for Online Purchase Prediction. In *2022 IEEE 38th International Conference on Data Engineering (ICDE)*. 3399–3412.
[19] Peng Ye, Julian Qian, Jieying Chen, Chen-hung Wu, Yitong Zhou, Spencer De Mars, Frank Yang, and Li Zhang. 2018. Customized regression model for airbnb dynamic pricing. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*. 932–940.
[20] Qing Zhang, Liyuan Qiu, Huaiwen Wu, Jinshan Wang, and Hengliang Luo. 2019. Deep learning based dynamic pricing model for hotel revenue management. In *2019 International Conference on Data Mining Workshops (ICDMW)*. 370–375.
[21] Xiaokun Zhang, Bo Xu, Liang Yang, Chenliang Li, Fenglong Ma, Haifeng Liu, and Hongfei Lin. 2022. Price does matter! modeling price and interest preferences in session-based recommendation. In *Proceedings of the 45th international ACM SIGIR conference on research and development in information retrieval*. 1684–1693.
[22] Yu Zheng, Chen Gao, Xiangnan He, Yong Li, and Depeng Jin. 2020. Price-Aware Recommendation with Graph Convolutional Networks. In *36th IEEE International Conference on Data Engineering, ICDE 2020, Dallas, TX, USA, April 20-24, 2020*. IEEE, 133–144.
[23] Fanwei Zhu, Wendong Xiao, Yao Yu, Ziyi Wang, Zulong Chen, Quan Lu, Zemin Liu, Minghui Wu, and Shenghua Ni. 2022. Modeling Price Elasticity for Occupancy Prediction in Hotel Dynamic Pricing. In *Proceedings of the 31th ACM International Conference on Information & Knowledge Management (CIKM '20)*.
[24] Ruitao Zhu, Detao Lv, Yao Yu, Ruihao Zhu, Zhenzhe Zheng, Ke Bu, Quan Lu, and Fan Wu. 2023. LINet: A Location and Intention-Aware Neural Network for Hotel Group Recommendation. In *Proceedings of the ACM Web Conference 2023*. 779–789.
[25] Ruitao Zhu, Wendong Xiao, Yao Yu, Yizhi Yu, Zhenzhe Zheng, Ke Bu, Dong Li, and Fan Wu. 2023. CANDY: A Causality-Driven Model for Hotel Dynamic Pricing. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*. 3616–3625.

A APPENDIX

A.1 Online Deployment of HP^{3}M on Fliggy

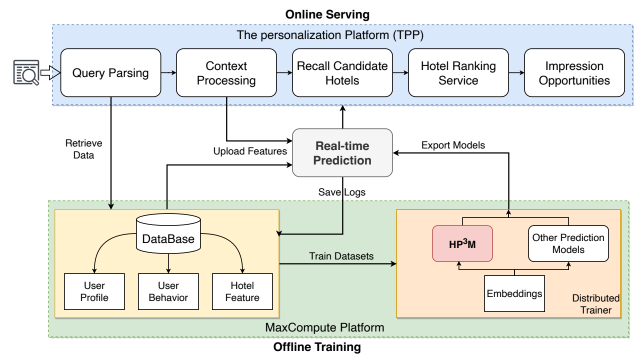

Figure 4: The workflow of online deployment of HP^3M in Fliggy's hotel search engine.

The proposed HP^3M has been fully deployed on Fliggy's hotel search engine and serves tens of millions of users daily. The deployment workflow of HP^3M is shown in Figure 4. During online serving, when a user request arrives, The Personalization Platform (TPP) first parses the query and processes context information into standard-format features. Next, the corresponding user profiles, user behaviors, and hotel features are retrieved from the database and are combined with the parsed query features to obtain the input features to the real-time predictor. Such features are used to recall candidate hotels, and then conduct hotel ranking services. The final ranking list is displayed to the user in the front end.