UNIVERSITY OF CATANIA

DEPARTMENT OF MATHEMATICS AND COMPUTER SCIENCE

INTERNATIONAL PH.D IN COMPUTER SCIENCE XXXVIII cycle

SUSANNA SAITTA

Intelligent Pricing Systems for Hotels: From Prediction to Optimization

PH.D THESIS

Supervisor: prof. Giovanni Maria Farinella

Academic Year 2024–2025

"Without data, you're just another person with an opinion."

William Edwards Deming

Abstract

Dynamic pricing, initially successful in the airline industry, has become a cornerstone of revenue optimization in the hospitality sector. Traditionally, this strategy relies on manual interventions, where Revenue Managers (RMs) continuously monitor market conditions, key performance indicators (KPIs), and external factors to define optimal pricing for future days of stay (DOS). However, this process is increasingly time-consuming and difficult to scale.
In this context, machine learning (ML) techniques offer promising opportunities to process the growing volume of data generated by accommodations and to enhance pricing decisions.
This thesis proposes a structured approach to dynamic pricing prediction, starting with the design of a domain-specific dataset that reflects the informational context in which RMs operate. Due to the lack of publicly available datasets in this field, the proposed dataset is built from real-world hotel data, incorporates static, dynamic and engineered features, and is further enriched through experts annotations. Using this dataset, five ML models are benchmarked to predict the price that a Revenue Manager would assign to an entry-level room over a 90-day horizon. Since the system is intended to support non-expert users, who may struggle to trust the output of AI models, the framework also integrates Explainable Artificial Intelligence (XAI) techniques. In particular, SHapley Additive exPlanations (SHAP) are used to provide clear and accessible explanations of model predictions, both at the global and local level. Moreover, the thesis presents the architecture of the entire system, showing how the proposed solution can be integrated into a Revenue Management System. This work has already moved beyond the research phase: the system is currently in production within the RMS of the partner company, actively informing pricing decisions and driving innovation in operational practice.

Building on this foundation, ongoing research explores a more ambitious direction: the development of an autonomous agent capable of suggesting optimal prices by interacting with the environment and learning from outcomes rather than past decisions. This transition from imitation to optimization aims to overcome the limitations of human heuristics and unlock more effective pricing strategies, particularly for small and medium-sized accommodations that lack access to dedicated revenue management expertise.

Contents

List of Figures

List of Tables

To myself.

To the countless hours, the quiet resilience, and the unwavering determination to grow beyond what I thought possible. To the version of me who kept going—who believed, doubted, fell, and stood up again. This work is a tribute to my journey, my discipline, and my constant desire to improve and push past my limits.

Chapter 1

Introduction

In recent years, dynamic pricing has emerged as a pervasive and strategically significant approach across a broad spectrum of industries, underscoring its adaptability and efficacy in revenue optimization. This methodology has been increasingly implemented in domains such as wireless telecommunications (Elreedy et al., 2018), ticketing for sporting events (Sahin and Erol, 2017; Sahin, 2019), real estate valuation (Ragapriya et al., 2023), and digital advertising, with particular emphasis on the allocation of billboard space (Lak et al., 2015). Despite the diversity of these sectors, the underlying objective remains consistent: to dynamically adjust prices in response to real-time fluctuations in demand, consumer behavior, and contextual variables, thereby enhancing profitability while preserving competitive market positioning.

Specifically, in the hospitality sector, the adoption of dynamic pricing has evolved through several key phases, largely inspired by developments in the airline sector. In the 1980s, airlines pioneered revenue management techniques to optimize seat pricing based on demand fluctuations, laying the groundwork for similar strategies in hotel operations. Hotels began experimenting with dynamic pricing in the 1990s, initially through rule-based systems and manual segmentation. These early approaches relied heavily on historical booking data and managerial intuition. The widespread use of online travel agencies (OTAs) and the emergence of real-time booking platforms in the early 2000s marked a turning point, dramatically increasing market transparency and competitive pressure. In response, hotels began to integrate automated pricing tools capable of dynamically adjusting rates based on occupancy levels, booking lead times, competitor pricing, and other market signals.

The first academic work to rigorously address dynamic pricing in hospitality is attributed to Sheryl E. Kimes, whose paper (Kimes, 1989) introduced the concept of revenue optimization tailored to hotel environments. This foundational study framed pricing as a strategic lever for managing perishable inventory — namely, room nights — and influenced decades of subsequent research. Today, dynamic pricing is a core component of hotel revenue management systems, supported by sophisticated algorithms that balance profitability, demand forecasting, and perceived fairness.

1.1 Motivation

Having established that dynamic pricing constitutes a core component of hotel revenue management systems, it is necessary to delve into the operational mechanisms through which Revenue Managers (RMs) implement this strategy. This clarification serves as a foundational step for the analytical framework of this thesis.

In the hospitality sector, dynamic pricing refers to a revenue optimization strategy whereby room rates are continuously adjusted over time to maximize the financial performance of a lodging facility. Rather than relying on fixed prices, hotels adopt flexible pricing schemes that respond to fluctuations in demand, competitive pressure, and broader market conditions. This adaptive approach enables hoteliers to capitalize on emerging revenue opportunities and mitigate potential losses due to underpricing or unsold inventory. The effectiveness of dynamic pricing in enhancing hotel revenue has been empirically supported by studies such as (Alshakhsheer et al., 2017) and (Abrate, Nicolau, and Viglia, 2019), which demonstrate its superiority over static pricing models.

In practice, the implementation of dynamic pricing requires RMs to engage in ongoing monitoring of a wide array of variables. These include internal performance metrics — commonly referred to as Key Performance Indicators (KPIs) — such as occupancy rates, average daily rate (ADR), and revenue per available room (RevPAR), as well as external factors like competitor pricing, market demand, and macroeconomic trends. Moreover, temporal elements such as seasonality, public holidays, and local events (e.g., cultural festivals, sports tournaments) play a significant role in shaping pricing decisions. Many works in the literature have focused on the factors that influence dynamic price changes. Some studies have examined these variables independently of the specific application domain (Deksnyte and Lydeka, 2012), while others have concentrated explicitly on the hospitality sector (El-Nemr, Canel-Depitre, and Taghipour, 2019; Zhang et al., 2017), highlighting the unique challenges and opportunities inherent to hotel pricing.

The complexity of this task is exemplified in Figure 1.1, which shows a screenshot of the daily calendar interface from the Revenue Management System (RMS) of the partner company that supported this research.

On the left-hand side, the interface displays real-time data on weather conditions, KPIs, demand forecasts, and scheduled events across multiple dates, while the right-hand side provides the corresponding room rates segmented by room type. This decision-making process is repeated daily — and often multiple times per day — as RMs evaluate whether to increase, decrease, or maintain current prices for future booking dates. A robust dynamic pricing strategy, therefore, must not only account for the complexity of influencing factors but also address the operational challenges associated with real-time price adjustments. While this manual approach allows for granular

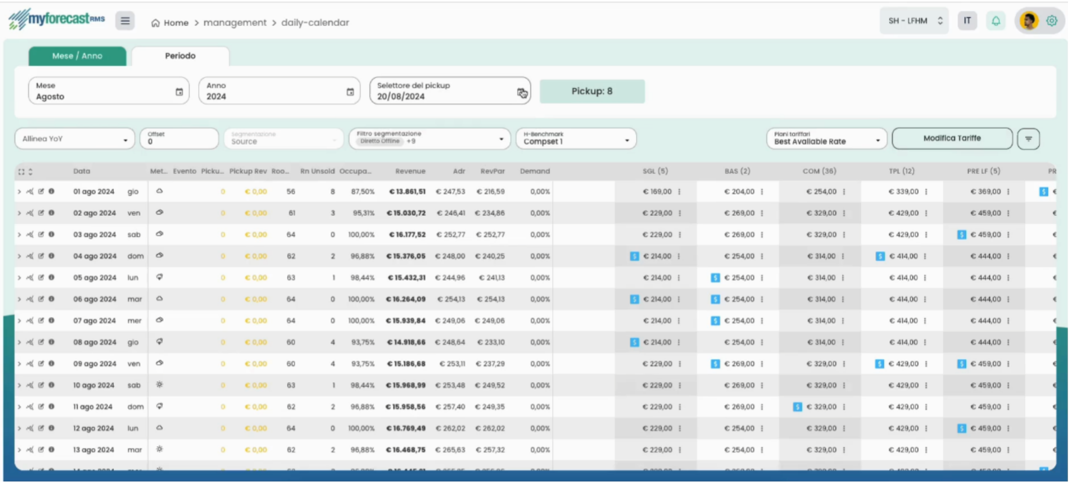

FIGURE 1.1: Daily calendar interface used by Revenue Managers to track key performance indicators and adjust pricing strategies for future dates of stay (DOS).

control, it is inherently labor-intensive, time-consuming, and prone to inefficiencies. Consequently, the need for automated systems capable of supporting or replacing human decision-making has become increasingly evident.

1.2 Objectives

In this context, Machine Learning (ML) techniques offer significant potential to process the growing volume of data generated by accommodation facilities, thereby enhancing the dynamic pricing process.

Although ML and traditional data analysis methods have long been available, their adoption within the hospitality sector has been relatively slow. As noted by Pande, 2020, the integration of ML into hospitality operations is a relatively recent development. Furthermore, Mariani and Wirtz, 2023 observed a marked increase in research works related to analytics in the hospitality and tourism domains, indicating a growing interest in data-driven decision-making. Despite this trend, Goli and Haghighinasab, 2022 identified a notable gap in the literature, particularly regarding dynamic pricing in the B2B segment, and their findings suggest that academic contributions on this topic within Italy are still relatively rare and fragmented. Building on these premises, the objective of this thesis is to investigate and propose a Machine Learning–based approach to address the challenge of dynamic pricing in the hospitality industry. Specifically, we adopt a supervised learning framework designed to predict the daily room rate that a Revenue Manager (RM) would assign for the upcoming 90 days of stay. This predictive task is grounded in real-world data collected from accommodation facilities, with pricing labels provided by experienced revenue management professionals. The availability of expert-annotated decisions enabled the construction of a robust supervised dataset, reflecting authentic pricing behavior. Figure 1.2 illustrates an example of the target variable, displaying the pricing decisions made for two consecutive dates of stay—June 7th and 8th, 2023. This example highlights a key characteristic of hotel pricing: each calendar date maintains its own distinct price history, shaped by a combination of internal performance metrics and external market dynamics. As shown in the figure, even adjacent dates can exhibit markedly different pricing trajectories. Indeed, when examining the data from bottom to top, it becomes evident that selling prices vary significantly across the two dates, underscoring the dynamic and often non-linear nature of pricing decisions. This observation further reinforces the need for a data-driven approach capable of capturing fine-grained temporal variations in pricing behavior. In response to this challenge, the primary objective of this thesis was to create a task-specific dataset and to propose a Machine Learning–based solution that supports Revenue Managers with accurate and adaptive price forecasting.

Beyond predictive accuracy, interpretability is a key consideration in our approach. To foster transparency and trust in the model's outputs, a second objective was to integrate an Explainable Artificial Intelligence (XAI) technique — SHapley Additive exPlanations (SHAP) — which provides insights

2023/06/07

Data immissione tariffa	Standard € 209,00
03/06/2023 19.42.30	-
22/05/2023 20.07.56	-
19/05/2023 09.18.09	€ 209,00
18/05/2023 18.04.15	-
12/05/2023 19.49.42	-
11/05/2023 08.45.50	€ 189,00
10/05/2023 19.50.00	-
24/04/2023 08.38.03	€ 179,00
31/01/2023 12.00.56	€ 159,00
31/10/2022 11.18.37	€ 149,00
10/10/2022 10.25.19	€ 139,00
03/10/2022 19.36.00	-
25/05/2022 11.54.08	**€ 119,00**

2023/06/08

Data immissione tariffa	Standard € 149,00
07/06/2023 16.27.32	-
07/06/2023 16.07.05	€ 149,00
06/06/2023 19.44.00	-
05/06/2023 14.05.33	€ 169,00
03/06/2023 19.42.30	-
01/06/2023 19.55.19	-
29/05/2023 19.55.05	-
22/05/2023 09.27.18	-
19/05/2023 09.18.10	€ 189,00
05/05/2023 20.04.39	-
04/05/2023 19.51.00	-
24/04/2023 08.38.04	€ 179,00
14/04/2023 19.54.09	-
27/03/2023 19.45.13	-
18/11/2022 10.05.54	€ 169,00
31/10/2022 11.18.37	€ 149,00
03/10/2022 19.36.00	-
25/05/2022 11.54.08	**€ 139,00**

FIGURE 1.2: Price history for two consecutive dates of stay (DOS).

into the relative contribution of each feature to the predicted price. This enables Revenue Managers not only to receive automated pricing suggestions, but also to understand the rationale behind each recommendation.
Once the predictive prototype demonstrated satisfactory performance, a third objective was pursued: the design of a complete system architecture to enable the real-world deployment. A subsequent goal was to integrate the proposed solution into the Revenue Management System (RMS) of the partner company that co-financed this research, thereby bridging the gap between academic innovation and industrial application.
As a natural progression of the work, the thesis introduces a Preliminary Methodological Design to support the transition from supervised learning to reinforcement learning. This framework lays the foundation for future development of autonomous pricing agents capable of learning optimal strategies through interaction with the environment, rather than relying solely on historical data. By shifting from prediction to goal-oriented optimization, the proposed design opens new avenues for dynamic pricing systems that are adaptive, scalable, and increasingly self-sufficient.
In line with the foundational purpose of the doctoral scholarship, the research responds directly to the innovation needs of enterprises. Through the achievement of these objectives, the company is now equipped with a tool that reflects the current state-of-the-art and strengthens its competitive position within the revenue management systems sector.

1.3 Contributions

The main contributions of this thesis can be summarized as follows:

- We present a method for constructing a domain-specific dataset designed to support the predictions of dynamic pricing in the hospitality sector.

- We benchmark multiple Machine Learning models to address the challenge of automated dynamic pricing, with the goal of integrating predictive capabilities into a Revenue Management System (RMS).
- We incorporate an Explainable Artificial Intelligence (XAI) technique—specifically SHapley Additive exPlanations (SHAP)—to enable non-expert users to interpret, trust and understand the outputs of the predictive model.
- We validate the proposed approach using real-world data from operational hotel environments, demonstrating its applicability and effectiveness in practice.
- We design the end-to-end system architecture required to transition the predictive prototype into a deployable software component, specifically tailored for integration within an existing RMS.
- We propose a Preliminary Methodological Design that outlines the transition from a supervised learning approach to a reinforcement learning framework. This design sets the foundation for future development of autonomous pricing agents capable of learning optimal strategies through interaction with the environment, rather than relying solely on historical labels.

1.4 Publications

The findings and contributions of this thesis have been disseminated through the following publications in international journals and conferences:

- S. Saitta, V. D'Amico and G.M. Farinella, "Predictions of hotel dynamic pricing for Revenue Management System". Accepted by Communications in Computer and Information Science book series.

- S. Saitta, V. D'Amico and G.M. Farinella, "Dynamic Price Prediction for Revenue Management System in Hospitality Sector". In Proceedings of the 13° International Conference on Data Science, Technology and Applications (DATA24). (Saitta, D'Amico, and Farinella, 2024)

Additionally, the following works, not directly related to the topics of this thesis, were produced during the PhD:

- V. Minio et al., "Towards a monitoring system of the sea state based on microseism and machine learning". Environmental Modelling & Software, Elsevier – September 2023. (Minio et al., 2023)
- Flavio Cannavò et al., "Machine Learning and Microseism as a Tool for Sea wave Monitoring". Presented at EGU 2023. (Cannavò et al., 2023)

1.5 Thesis Outline

The remainder of this thesis is organized into 7 Chapters:

- Chapter 2 provides an in-depth review of the state-of-the-art methodologies related to dynamic pricing and its application in the hospitality sector;
- Chapter 3 Presents the supervised learning framework developed for dynamic price prediction, with a focus on the dataset design, the theoretical rationale behind the chosen Machine Learning models, and the criteria used for performance evaluation.
- Chapter 4 outlines the experimental setup and presents the results obtained on real-world data.
- Chapter 5 explores Explainable Artificial Intelligence (XAI) techniques, with particular emphasis on SHapley Additive exPlanations (SHAP), used to interpret model predictions and support decision-making.

- Chapter 6 describes the design of the end-to-end system architecture, illustrating how the predictive prototype was transformed into a deployable software component within a Revenue Management System.
- Chapter 7 Presents the ongoing research efforts aimed at determining the optimal dynamic pricing policy through Reinforcement Learning.
- Chapter 8 summarizes the key findings of the research, discusses its implications, and outlines potential directions for future work.

Chapter 2

Related Works

This chapter aims to provide an overview of the existing literature on dynamic pricing in the hospitality sector. This analysis aims to provide a thorough understanding of existing methodologies and key findings from relevant studies in this field, laying the foundation for subsequent discussions of our research contributions. First, three main lines of research can be distinguished for addressing the problem of dynamic pricing. The main differences among the works depend on the considered target variable, on the exploitation of the price-elasticity coefficient in the process of determining the dynamic price and on the specific market in which dynamic prices are applied. These research directions are detailed below.

2.1 ADR as target variable

Some works address the problem of estimating the Average Daily Rate (ADR) as the target variable. Studies in this context have tried to forecast ADR at the city level. In Al Shehhi and Karathanasopoulos, 2018 is presented a method which exploits conventional time-series and machine learning models to forecast ADR. Luxury and upscale hotels of eight cities of the Middle East and North Africa have been considered in the study. The results show the usefulness of machine learning models for forecasting prices. Al Shehhi and Karathanasopoulos, 2020 have employed data from five cities in the Persian Gulf for the ADR prediction. Also in this case luxury and upscale hotels have been considered to implement a price forecasting using both traditional statistical models and artificial intelligence models. Zhang et al., 2019 have focused on ADR prediction at the hotel level, proposing a dynamic pricing system which considers three main steps: a base price is set considering competitor prices, the future occupancy is predicted with a sequence learning model which combines DeepFM and the seq2seq model, and finally a DNN is employed to predict the ADR for each hotel and date.

2.2 Focus on Price-elasticity coefficient and demand forecast

A second group of works considers as fundamental the estimation of the coefficient of the price-elasticity of demand in the process that leads to the set of the dynamic price. Zhu et al., 2022 have proposed a model which predicts the price elasticity coefficient, both at the hotel and room type level, taking into account competitors, temporal and hotel-specific factors. Once the coefficient is obtained, it is used to estimate the occupancy for each eligible price via a specific demand function. The optimal price will be then the one that maximizes the expected revenue. Shintani and Umeno, 2022 have presented a method that enables simulating the magnitude of changes in demand as a result of a change in price. Specifically, they have used a time-rescaling regression to forecast the demand and have introduced a parametric learning model that allows the price elasticity of demand coefficient to be estimated from historical data. Once a new price rate has been chosen, this is used together with the previously obtained coefficient to update the booking curve and to compute the new demanded quantity. Bayoumi et al., 2013 proposed to study the dynamic pricing process based on four price multipliers. The optimal parameters for the multipliers have been computed via the Covariance Matrix Adaptation Evolutionary Strategy (CMA-ES) using simulated data as input together with the Monte Carlo method. To determine how the influence of their method on the reference price set by the RM will affect the demand, the authors have inserted a module that computes a demand index, which is also used to appropriately moderate the new simulations. The simulation and optimization loop is repeated until the multipliers parameters that maximize the revenue are found. Vives, Jacob, and Aguiló, 2019 have proposed a data transformation system and the application of a demand model based on log-linear regression to estimate the price-elasticity coefficient for each predefined season and booking period. This approach has also been considered in Vives and Jacob, 2020. The authors have adapted the online transient demand function to two mathematical models (one deterministic and one stochastic) to estimate prices and quantities that maximize the revenue along distinct booking horizons and seasons using Lagrange multipliers. In Vives and Jacob, 2021 the method using the deterministic model has been applied to several hotels of Spain. Bandalouski et al., 2021 have proposed to disaggregate the demand into categories and forecast it using time-series methods. The result of this step is used to estimate the two coefficients of the defined demand function. They then obtain the optimal price rates optimizing a concave quadratic objective function with linear constraints. Once the optimal prices are obtained, they can also be used to estimate the optimal quantity to sell via the demand function. Shadiqurrachman, Ridwan, and Kusuma, 2019 have proposed a pricing policy system in which multiple linear regression are used. Each linear regression aims to capture the relationship between the average price existing in one part of the planning horizon and the average price related to the other parts. The target variable of these multiple regressors is the quantity sold in the considered part of the planning horizon. The estimated coefficients are given as input to an integer nonlinear programming method to find the optimal pricing policy for the entire planning period. Once the optimal prices have been found, these are used as input of the demand model to compute the quantity of rooms that would be sold by adopting the estimated prices. The components employed for the dynamic pricing model of the latter work have been first presented by Shakya et al., 2012, who however used neural networks for the demand model and an evolutionary algorithm to find the pricing policy that maximizes the revenue along the planning horizon.

2.3 Pricing for peer-to-peer platforms

A third group of studies considers the problem of setting prices of Airbnb listings. Rather then estimating room prices, in this case the goal is the price estimation for the proposed accommodation, such as an entire house, a cabin, a boat and much more. Although this problem is not a market perfectly comparable to the one which consider the dynamic price estimation of hotel rooms, it is important to mention the studies in this context. Indeed, Ye et al., 2018 have introduced five evaluation metrics which have been widely adopted to evaluate the goodness of the dynamic prices predicted by machine learning models (Zhang et al., 2019; Zhu et al., 2022). In addition to the introduction of these metrics, the authors have proposed a pricing system composed of three steps. First, the booking probability for each listing is predicted per night performing a binary classification with Gradient Boosting. Secondly, this probability is used as input, together with other features, to predict the price for each listing-night through a regression model. As last step, a customised logic is applied to generate the final price to be used. Kalehbasti, Nikolenko, and Rezaei, 2019 used features of the rentals, owner characteristics and reviews with various machine learning models to predict the prices of Airbnb listings in Amsterdam. Peng, Qin, and Li, 2020 have made use of numerical, geospatial and textual data for Principal Component Analysis. The first six principal components have been employed as predictors together with XGBoost model. This last machine learning method has been also used by Liu, 2021.

From the above analysis, it appears there is not yet a standard procedure in the literature for dynamic pricing in the hospitality sector. Furthermore, there are no datasets that serve as benchmarks for measuring the progress of various approaches in this specific task.

Chapter 3

Dynamic Price Predictions: Proposed Approach

In this chapter, we present the background that motivated the development of an automated dynamic pricing mechanism, beginning with an analysis of how a Revenue Management System (RMS) operates within the hospitality sector. This operational overview highlights the complexity and frequency of pricing decisions, which in turn underscores the need for intelligent automation. Building on this foundation, we introduce the rationale behind the design of the dataset used in our analysis, including a detailed description of the features collected and their relevance to the pricing task. We then provide a theoretical overview of the Machine Learning models considered in this research, outlining their key characteristics and suitability for addressing the problem. Finally, we discuss the evaluation metrics adopted to assess model performance, which serve both as a basis for comparing predictive approaches and for constructing ad hoc error matrices in subsequent chapters.

3.1 Background

The primary objective of a Revenue Management System (RMS) is to transform raw hotel data into strategic insights that support data-driven pricing decisions while providing real-time visibility into the hotel's commercial performance. The operational flow of an RMS (see Figure 3.1) is structured around a network of interconnected systems that exchange data in real time. At the foundation of this ecosystem lies the Property Management System (PMS), a software platform used by hotels to manage core operational functions such as reservations, check-ins and check-outs, room inventory, and transactional records. The PMS continuously collects and updates data related to hotel sales, availability, and guest activity. At the center of the architecture is the Revenue Management System, which interfaces directly with the PMS and, in the case of the RMS examined in this study, with two additional external data sources. These sources typically include rate shopping tools or market intelligence platforms that provide real-time information on competitor pricing and market trends. The RMS aggregates and processes this data to generate a set of Key Performance Indicators (KPIs) — quantitative metrics that reflect the hotel's performance in areas such as occupancy, average daily rate (ADR), revenue per available room (RevPAR), and price positioning. Based on these KPIs, the Revenue Manager is able to make strategic pricing decisions. For each room type, the manager sets optimal

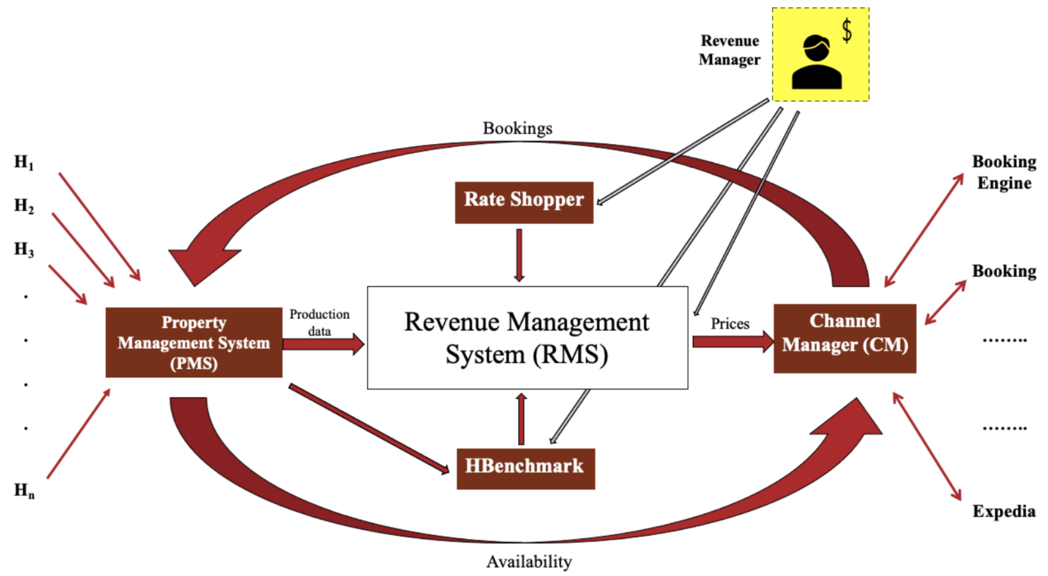

FIGURE 3.1: Revenue Management System flow.

prices that reflect both internal performance metrics and external market conditions. These prices are then transmitted to the Channel Manager (CM), a distribution platform responsible for synchronizing rates and availability across all online sales channels, including Online Travel Agencies (OTAs), the hotel's official website, and Global Distribution Systems (GDS). When a booking is made through one of these channels, the information flows back through the CM to the PMS, which updates the hotel's inventory in real time. The PMS then communicates the updated availability to the CM, which propagates the changes across all connected platforms, ensuring consistency and preventing overbooking. This cyclical data flow — from PMS to RMS, through CM, and back — forms the backbone of modern revenue management in the hospitality sector. It enables hotels to respond dynamically to market changes, optimize pricing strategies, and maintain operational efficiency across multiple distribution channels.

3.2 Dataset

Given the operational framework of a Revenue Management System (RMS), particular attention must be directed toward the figure of the Revenue Manager (RM), whose decisions are central to the dynamic pricing process. Revenue management, in its essence, is the discipline of optimizing revenue by strategically managing a limited supply—such as hotel room inventory — through dynamic price adjustments. This involves continuously modifying selling prices in response to internal performance metrics, fluctuations in demand, and temporal factors such as booking lead time or seasonality. The Revenue Manager performs this task by analyzing a variety of indicators and contextual signals, adjusting prices in real time to maximize revenue. This practice, known as dynamic pricing, is a cornerstone of modern revenue optimization strategies in the hospitality sector. The goal of this thesis is to replicate this decision-making mechanism computationally, by predicting the dynamic price that a Revenue Manager would set under specific conditions. To achieve this, the design of the dataset plays a critical role.

3.2.1 Construction logic

Our objective was to construct a data structure that faithfully reproduces the informational context in which a Revenue Manager operates. Notably, to the best of our knowledge, the existing literature does not provide a benchmark dataset specifically tailored for dynamic price prediction in the hospitality domain. This absence motivated the development of a novel dataset architecture, explicitly designed to capture the temporal and contextual dynamics of pricing decisions. The proposed dataset integrates the logic of time series

e_d			s_d		$x_{t-n,d}$					$y_{t,d}$
$l_{t,d}$	WD	…	SR	…						
0	0	…	SR_0	…	$x_{0-n,0}$	…	$x_{0-2,0}$	$x_{0-1,0}$	$x_{0-0,0}$	$y_{0,0}$
1	0	…	SR_1	…	$x_{0-n,1}$	…	$x_{0-2,1}$	$x_{0-1,1}$	$x_{0-0,1}$	$y_{0,1}$
2	0	…	SR_2	…	$x_{0-n,2}$	…	$x_{0-2,2}$	$x_{0-1,2}$	$x_{0-0,2}$	$y_{0,2}$
…	…	…	…	…	…	…	…	…	…	…
90	0	…	SR_{90}	…	$x_{0-n,90}$	…	$x_{0-2,90}$	$x_{0-1,90}$	$x_{0-0,90}$	$y_{0,90}$
0	0	…	SR_1	…	$x_{1-n,1}$	…	$x_{1-2,1}$	$x_{1-1,1}$	$x_{1-0,1}$	$y_{1,1}$
…	…	…	…	…	…	…	…	…	…	…
89	0	…	SR_{90}	…	$x_{1-n,90}$	…	$x_{1-2,90}$	$x_{1-1,90}$	$x_{1-0,90}$	$y_{1,90}$
90	0	…	SR_{91}	…	$x_{1-n,91}$	…	$x_{1-2,91}$	$x_{1-1,91}$	$x_{1-0,91}$	$y_{1,91}$
…	…	…	…	…	…	…	…	…	…	…
0	1	…	SR_7	…	$x_{7-n,7}$	…	$x_{7-2,7}$	$x_{7-1,7}$	$x_{7-0,7}$	$y_{7,7}$
…	…	…	…	…	…	…	…	…	…	…
90	1	…	SR_{97}	…	$x_{7-n,97}$	…	$x_{7-2,97}$	$x_{7-1,97}$	$x_{7-0,97}$	$y_{7,97}$
…	…	…	…	…	…	…	…	…	…	…

FIGURE 3.2: Dataset structure.

analysis within a tabular format, making it compatible with a wide range of machine learning models while preserving the temporal dependencies inherent in pricing behavior. The guiding principle behind its construction was to emulate the structure of the data that a Revenue Manager consults daily when deciding whether to increase, decrease, or maintain the selling price for a given Day of Stay (DOS). Each record in the dataset encapsulates both the current state of the accommodation facility on the target DOS and the historical evolution of relevant variables over the preceding n days. This dual perspective enables the model to learn from both static and dynamic signals, reflecting the way human decision-makers incorporate short-term trends and contextual shifts into their pricing strategies. Formally, each record r in the dataset is represented as a tuple:

$$r = (e_d, s_d, x_{t-n,d}, y_{t,d}) \tag{3.1}$$

where:

- d refers to the target Day of Stay (DOS) for which the price prediction is made;
- t denotes the date on which the prediction is requested;
- n indicates the number of days prior to t from which historical data is drawn;
- e_d and s_d represent the engineered and static features associated with the DOS;
- $x_{t-n,d}$ is the sequence of dynamic features observed over the n days preceding t for the same DOS;
- $y_{t,d}$ is the target variable, i.e., the price to be predicted for the DOS on date t.

This structure, illustrated in Figure 3.2, enables the model to capture the temporal evolution of pricing-relevant variables while maintaining compatibility with supervised learning frameworks. It reflects the operational logic of revenue management and provides a foundation for learning predictive models that approximate the pricing decisions of human experts.

3.2.2 Proposed variables

Building upon the dataset structure described in the previous section, it becomes essential to examine the nature of the features that populate the input space of the predictive model. Since the dataset was designed to emulate the decision-making context of a Revenue Manager, the features included must reflect the informational dimensions that influence pricing strategies on a daily basis. To this end, the input variables have been categorized into three main groups — static, dynamic, and engineered features —each contributing distinct types of information. This classification not only facilitates a clearer understanding of the data architecture but also supports the development of models capable of capturing both structural and temporal patterns relevant to dynamic pricing. These three categories are described in detail below:

- **Static features** (s_d): These features represent intrinsic attributes of the accommodation facility that are associated with a specific Day of Stay (DOS) but remain constant over time. They typically encode baseline conditions or initial configurations relevant to pricing decisions. A representative example is the starting selling rate, which is defined at the beginning of a new financial business season for each DOS. Static features serve as reference points and help the model anchor its predictions within the structural pricing framework of the hotel.

- **Dynamic features** ($x_{t-n,d}$): This category includes features that evolve over time and reflect the temporal dynamics of demand, booking behavior, and inventory status. Dynamic features include both daily variations and aggregated values. The latter reflect the cumulative behavior of relevant variables from the beginning of the financial business season up to the date on which the prediction is made, specifically for the Day of Stay (DOS) being evaluated. For instance, the number of Room Nights (RNs) booked on a specific day for a given DOS captures short-term booking activity, while the cumulative number of RNs booked from the start of the financial season up to the current date provides a broader view of demand accumulation. These features are essential for modeling trends, seasonality, and short-term fluctuations that influence pricing decisions.

- **Engineered features** (e_d): Engineered features are derived through data transformation processes or constructed explicitly to enrich the input space with domain-specific insights. They often combine raw variables or encode categorical information in a format suitable for machine learning algorithms. Examples include the Lead Time ($l_{t,d}$), which quantifies the number of days between the prediction date and the target DOS, and the WeekendDay indicator (WD), which flags whether the DOS falls on a weekend. Additional engineered features may include binary flags for holidays or proximity to local events. These features enhance the model's ability to capture latent patterns and contextual nuances that are not directly observable from raw data alone.

In addition to the input features described above, the dataset includes the target variable ($y_{t,d}$), which corresponds to the price dynamically set by the Revenue Manager for a given DOS on a specific prediction date t. This price represents the outcome of a complex decision-making process shaped by all the factors previously discussed.

A representative sample of the collected dataset is shown in Figures 3.3 and 3.4. Figure 3.3 highlights a subset of the input features, whereas Figure 3.4 focuses on the target variable used for prediction.

Index	Previsione_A	Demand	rnoDellaSettim	Giorno	Mese	WeekEndDay	Season	ttimanaDellAn	estivitaPasqu	Evento	Festivita	TariffaIniziale
0	0	0.8587	6	1	1	1	2	1	0	0	1	199
1	1	0.6598	7	2	1	0	2	1	0	0	0	144
2	2	0.713	1	3	1	0	2	2	0	0	0	144
3	3	0.6439	2	4	1	0	2	2	0	0	0	144
4	4	0.5421	3	5	1	0	2	2	0	0	0	144
5	5	0.5942	4	6	1	0	2	2	0	0	1	144

RN_t-5	RN_agg_t-5	RN_Unsold_t-5	Revenue_t-5	evenue_Agg_t-	ADR_t-5	ADR_Agg_t-5	REVPAR_t-5	EVPAR_Agg_t-	Occupancy_t-5	cupancy_Agg_	Delta_t-5
0	22	5	0	5572.77	0	253.31	0	206.44	0	0.81	0
0	17	10	0	2614.04	0	153.77	0	96.8	0	0.63	0
0	18	9	0	2346.68	0	130.37	0	86.9	0	0.67	0
0	18	9	0	2346.69	0	130.37	0	86.91	0	0.67	0
1	11	16	128	1569.31	128	142.66	4.74	58.11	0.04	0.41	0
2	19	8	358.18	2842.37	179.09	149.6	13.27	105.28	0.07	0.7	0

FIGURE 3.3: Sample of Input Features from the Collected Dataset.

Index	Price_t-0
0	259
1	199
2	169
3	169
4	144
5	159

FIGURE 3.4: Sample of Target Variable from the Collected Dataset

Its structure has been specifically designed to enable the predictive model to replicate the pricing behavior of the Revenue Manager, capturing the underlying logic and adaptive nature of human decision-making within a structured, data-driven framework.

3.3 Models

To address the task of dynamic price prediction, we selected five Machine Learning algorithms that represent a range of modeling paradigms, from linear to non-linear and from interpretable to more complex architectures. The choice of models was guided by their theoretical relevance, practical effectiveness in forecasting tasks, and suitability for structured data typical of the hospitality domain. Specifically, we included Multiple Linear Regression (MLR) as a baseline model due to its simplicity and interpretability. In addition, we considered three ensemble-based methods built upon decision trees: Random Forest, Gradient Boosting (GB) and Light Gradient Boosting Machine (LGB). These models are known for their ability to capture non-linear relationships, handle heterogeneous feature spaces, and deliver strong predictive performance in real-world applications. Finally, we incorporated a Multi-Layer Perceptron (MLP), a feedforward neural network capable of modeling complex patterns through multiple layers of abstraction.
In the following sections, we provide a theoretical overview of these algorithms, discussing their underlying mechanisms, assumptions, and suitability for the dynamic pricing problem addressed in this thesis.

3.3.1 Multiple Linear Regression

Multiple Linear Regression (MLR) is a foundational statistical model that assumes a linear relationship between the dependent variable and the independent variables. Formally, the MLR model is defined as:

$$y_i = \beta_0 + \sum_{j=1}^{p} \beta_j X_{ij} + \varepsilon_i, \tag{3.2}$$

where y_i is the predicted value for observation i, β_0 is the intercept term, β_j are the coefficients associated with each feature β_{ij}, and ε is the model's error term. The objective of MLR is to estimate the β coefficients by minimizing the residual sum of squares (RSS) between the observed target values and those predicted by the model. This is typically achieved through Ordinary Least Squares (OLS) optimization, which finds the best-fitting hyperplane in the multidimensional feature space.

Relevance to Dynamic Pricing

Although its capacity to model complex, non-linear interactions is limited, MLR remains a valuable benchmark due to its interpretability and ease of implementation. Moreover, the estimated coefficients offer direct insights into the marginal effect of each variable, which can be useful for Revenue Managers seeking explainable pricing strategies.

In the context of dynamic pricing, Multiple Linear Regression (MLR) offers a transparent and interpretable framework for understanding how individual features — such as booking lead time, day of the week, occupancy rate, and competitor pricing —contribute linearly to the final room rate. By assuming a linear relationship between input variables and the target outcome, MLR enables a straightforward estimation of the marginal effect of each feature, which can be particularly valuable for Revenue Managers seeking explainable and auditable pricing strategies. In fact, the estimated coefficients provide direct insights into feature relevance, facilitating decision-making in environments where interpretability is prioritized. Despite its simplicity, MLR remains a widely adopted benchmark model due to its ease of implementation and interpretability. However, the model's capacity to capture non-linear interactions and complex dependencies is inherently limited. MLR is also sensitive to multicollinearity among predictors and may struggle with outliers or missing data, which are common in real-world pricing scenarios. Furthermore, its performance tends to degrade in high-dimensional or rapidly evolving environments, where more flexible models are required to adapt to changing market dynamics. For all the above considerations, MLR is best suited for moderately dynamic pricing environments, where transparency and simplicity are prioritized over modeling complexity.

3.3.2 Ensemble Method with Homogeneous Learners

Ensemble learning is a powerful machine learning paradigm that constructs a strong predictive model by aggregating the outputs of multiple weak learners. The core intuition behind this approach is that a collection of diverse models can outperform any single constituent model, particularly in terms of generalization and robustness. In this work, we focus on homogeneous ensembles, where all base learners are decision trees — chosen for their interpretability, flexibility, and ability to capture non-linear relationships.

When employing decision trees within ensemble frameworks, it is essential to distinguish between two principal methodologies: bagging and boosting. These differ fundamentally in how the individual models are trained and combined (see Figure 3.5). In our experimental setup, we adopted one

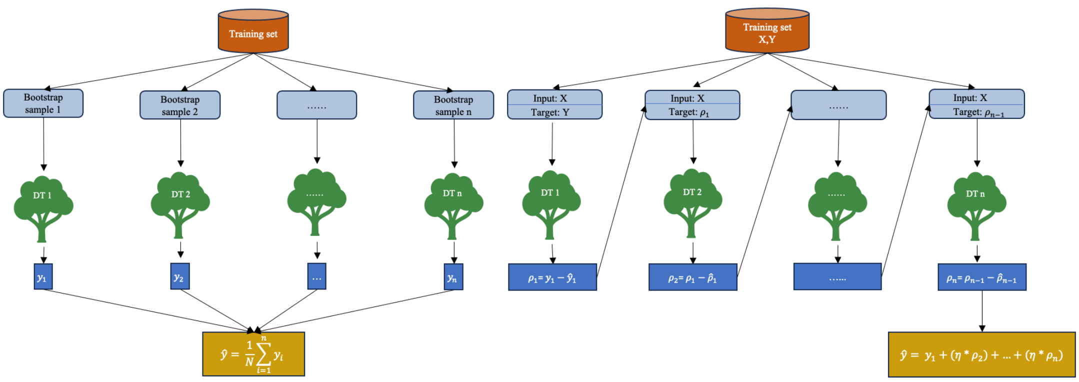

FIGURE 3.5: Bagging (left) vs Boosting (right) technique.

bagging-based model — Random Forest (RF) — and two boosting-based models: Gradient Boosting (GB) and Light Gradient Boosting Machine (LGB).

Random Forest: Bagging-Based Ensemble

Random Forest (RF) is a classic example of a bagging algorithm, where multiple decision trees are trained in parallel on different subsets of the training data. These subsets are generated via bootstrap replicas, created through sampling with replacement. Each tree is built independently, and for regression tasks, the final prediction for a new data point is computed as the average of the predictions from all trees in the ensemble. This parallel structure offers several advantages:

- Variance reduction: By averaging across multiple trees, RF mitigates overfitting and improves stability.
- Feature randomness: At each split, RF selects a random subset of features, further enhancing model diversity.
- Scalability: The parallel nature of RF allows efficient training on large datasets.

Gradient Boosting and LightGBM: Boosting-Based Ensembles

In contrast to bagging, boosting constructs models sequentially and additively, with each new tree trained to correct the residual errors of the ensemble built so far. These algorithms are named for their use of the gradient descent procedure to minimize the loss function as trees are added. Specifically:

- The first tree is fitted to the training data.
- Subsequent trees are added one at a time, each focusing on the errors made by the current ensemble.

- The final prediction is the sum of the outputs from all individual trees, weighted appropriately.

Although both Gradient Boosting (GB) and Light Gradient Boosting (LGB) follow this general boosting framework, they differ significantly in their tree growth strategies (see Figure 3.6). In fact, GB grows trees level-wise, prioritizing nodes closest to the root, while LGB uses a leaf-wise strategy, selecting leaves that most significantly reduce the loss function, often resulting in better accuracy. However, it may also produce deeper and more complex trees, which require careful regularization to avoid overfitting. Furthermore, LGB

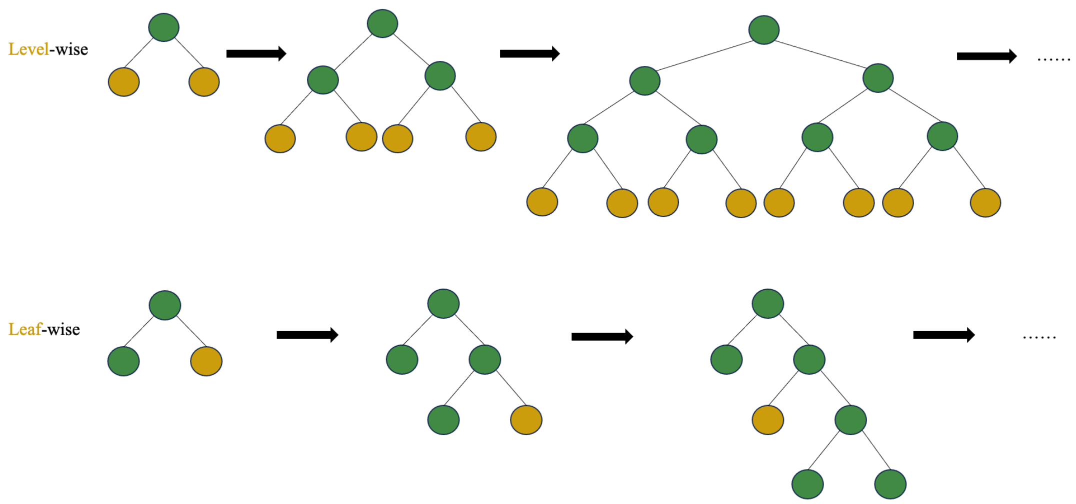

FIGURE 3.6: Leaf-wise vs level-wise tree growing method.

employs two techniques that enhance its speed compared to other boosting models:

- Gradient-based One-Side Sampling (GOSS). This technique filters data to find a split value by selecting all records with large gradients - those contributing most to the error - and randomly sampling instances with small gradients. This selective sampling preserves the integrity of gradient estimation while reducing the number of data points processed.

- Exclusive Feature Bundling (EFB). EFB identifies mutually exclusive features, i.e., features that never take non-zero values simultaneously, and bundles them into a single composite feature. This shrinks the dimensionality of the input space, thus reducing the model complexity, and speed up training phase without sacrificing predictive performance.

These optimizations make LGB particularly well-suited for high-dimensional and large-scale datasets, where training speed and memory efficiency are critical.

Relevance to Dynamic Pricing

Unlike linear models, ensemble methods—such as Random Forest (RF), Gradient Boosting (GB), and Light Gradient Boosting Machine (LGB) — are capable of capturing non-linear relationships and complex feature interactions. This is especially valuable in dynamic pricing scenarios, where factors such as demand, booking patterns, seasonality, and competitor pricing interact in intricate ways. The key benefits of tree-based ensemble models in dynamic pricing include:

- Non-linear modeling capability: Decision tree ensembles can capture complex interactions among features such as demand elasticity, time-based patterns, and competitive pricing signals.
- Feature importance metrics: Provide actionable insights into the drivers of price sensitivity, supporting strategic decision-making.
- Robustness to noise and missing data: Tree-based methods are inherently resilient to irregularities in real-world datasets, including outliers and incomplete records.

- Real-time adaptability: Boosting models like LGB can be incrementally updated, enabling near real-time responsiveness to market fluctuations.

However, these advantages come with certain trade-offs. Tree-based ensembles typically offer lower interpretability compared to linear models, as their internal decision processes are more complex and less transparent. In addition, they can be computationally intensive, especially when dealing with large ensembles or high-dimensional datasets. Careful hyperparameter tuning is often required to prevent overfitting and ensure optimal performance. Despite these limitations, ensemble methods are a compelling choice for dynamic pricing systems, where precision, scalability, and interpretability are essential for optimizing revenue.

3.3.3 Multi-layer Perceptron

The Multi-Layer Perceptron (MLP) is a class of feed-forward artificial neural networks characterized by a unidirectional flow of information, where each layer transmits its output exclusively to the subsequent layer. It is composed by multiple layers of interconnected neurons. In particular, an MLP has (see Figure 3.7):

- An input layer, which receives the raw features of the dataset.
- One or more hidden layers, which perform intermediate transformations. Specifically, each neuron in the hidden layer processes the values from the previous layer by performing a weighted linear summation, adding a bias term, and applying a non-linear activation function.
- An output layer which produces the final prediction, that is, a class label for classification tasks and a continuous value for regression tasks.

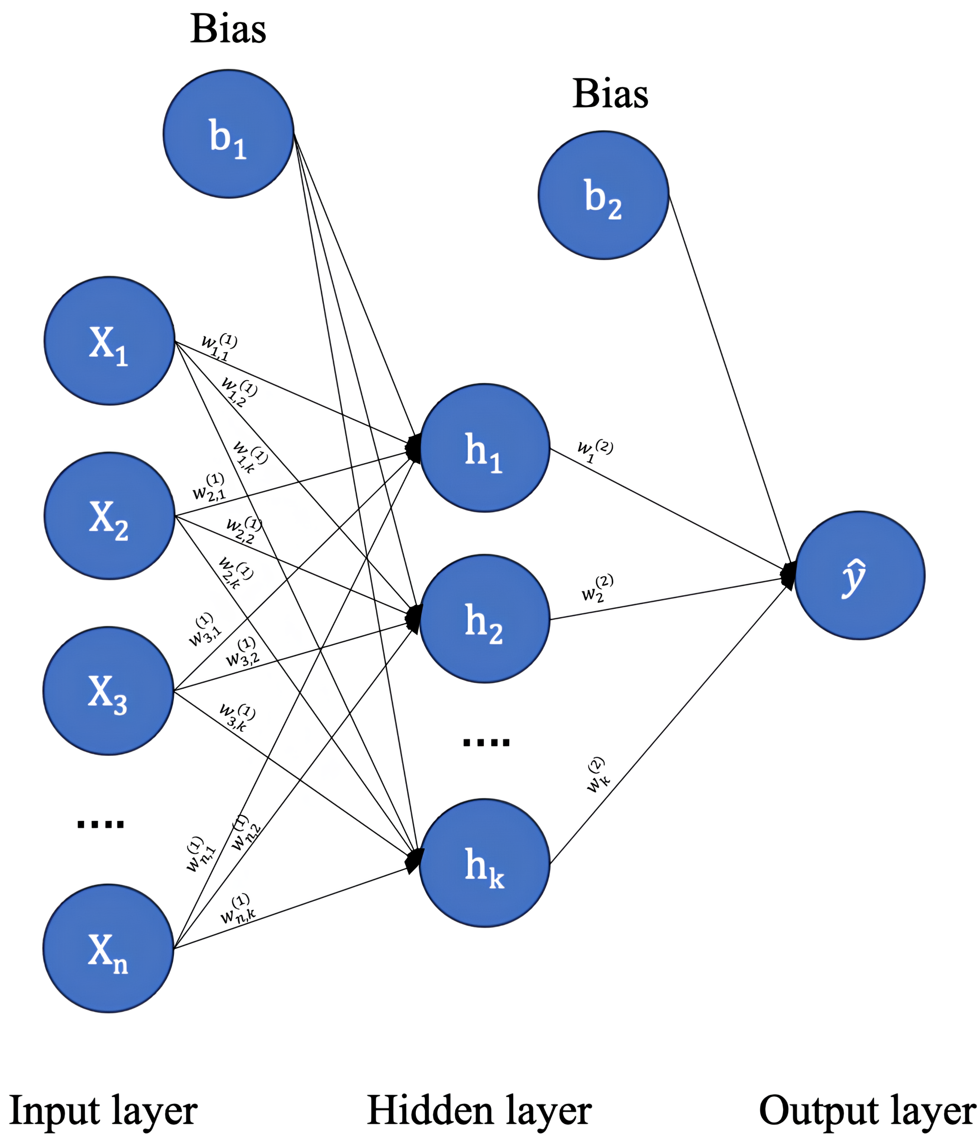

FIGURE 3.7: Multi-layer Perceptron structure for regression tasks.

The model is trained using the backpropagation algorithm, in conjunction with an optimization method such as stochastic gradient descent (SGD) or one of its variants (e.g. Adam, RMSprop). Backpropagation operates by computing the gradient of the loss function with respect to each weight in the network, propagating the error backward from the output layer to the input layer. These gradients are then used to iteratively update the weights, with the objective of minimizing the overall loss function — commonly the mean squared error for regression or cross-entropy for classification.

Relevance to Dynamic Pricing

MLPs are capable of learning highly non-linear relationships between input features and target variables. Unlike Multiple Linear Regression (MLR), which assumes a linear contribution of each feature, MLPs can capture intricate dependencies and interactions that are essential for accurate price forecasting. Compared to tree-based ensemble methods, MLPs offer greater architectural flexibility and can be more easily adapted to heterogeneous and high-dimensional datasets. While ensemble models rely on decision rules and feature splits, MLPs learn continuous representations, which can be advantageous when dealing with subtle behavioral patterns or temporal dynamics. The key benefits of MLPs in dynamic pricing include:

- Universal function approximation. MLPs are capable of modeling complex, high-dimensional relationships that go beyond linear or rule-based structures.
- Architectural flexibility. Unlike models with rigid structures — such as linear regression or decision trees —MLPs can be configured with varying numbers of layers and neurons, allowing them to scale with the complexity of the problem and the richness of the data. This flexibility

makes them particularly suitable for dynamic pricing, where the input data can be highly heterogeneous.

- Continuous refinement. MLPs supports iterative updates through back-propagation and gradient-based optimization. In dynamic pricing contexts, this training mechanism supports continuous refinement. As new data become available, the model can be retrained or fine-tuned to reflect the latest market conditions. This enables real-time responsiveness, allowing pricing strategies to evolve in sync with customer behavior and external influences. Compared to static models like MLR, which require retraining from scratch to incorporate new data, and ensemble methods, which may be computationally heavier to update, MLPs offer a more fluid and scalable approach to learning. This makes them particularly valuable in fast-paced environments where pricing decisions must be both accurate and timely.

Despite these strengths, MLPs also present some limitations. Their internal mechanisms are less transparent than those of MLR or ensemble models, which may hinder interpretability and managerial trust. While MLR provides direct insights into the marginal effect of each variable, and ensemble methods offer feature importance metrics, MLPs operate as black-box systems, making it more difficult to explain individual predictions. Moreover, MLPs are sensitive to overfitting and require large volumes of high-quality data to generalize effectively. Their performance depends heavily on careful hyperparameter tuning and the use of regularization techniques. In contrast, ensemble methods tend to be more robust to noise and missing data, and MLR remains computationally efficient and easy to implement. In summary, MLPs are particularly well-suited for highly dynamic pricing systems that prioritize predictive accuracy, adaptability, and scalability, especially in data-rich environments. While they may not offer the same level of interpretability as MLR or the robustness of ensemble models, their ability to learn complex mappings and respond to evolving inputs makes them a powerful asset in modern pricing architectures.

3.4 Evaluation Metrics

To assess and compare the predictive performance of the models employed in this thesis, we adopted three widely used evaluation metrics: Mean Absolute Error (MAE), Mean Absolute Percentage Error (MAPE), and the Coefficient of Determination (R^2). Each of these metrics captures different aspects of model accuracy and generalization, and their combined use provides a more comprehensive understanding of how well the models perform in different scenarios and data sets.

Mean Absolute Error (MAE)

The Mean Absolute Error (MAE) measures the average deviation between predictions and actual values. Consequently, the lower it is, the better is the model. It is defined as

$$MAE = \frac{1}{N} \sum_{i=1}^{N} |y_i - \hat{y}_i| \tag{3.3}$$

where $\hat{y}_i$ is the value predicted by the model for observation i, whereas y_i is the corresponding target value. By relying on absolute values, the metric disregards the direction of individual errors, thereby preventing positive and negative deviations from offsetting each other. This property contributes to its robustness, particularly in the presence of outliers, when compared to squared-error metrics such as Mean Squared Error (MSE). MAE is expressed in the same unit as the target variable, making it intuitively interpretable. For this reason, it is particularly useful when evaluating model performance on a single accommodation facility, as it reflects the average deviation in monetary terms (e.g. euros). However, it is less suitable for cross-hotel comparisons, since the scale of the target variable (e.g., average room rate) may vary significantly between properties. A low MAE in one hotel may not be equivalent to a low MAE in another, unless the underlying price distributions are comparable. Therefore, the interpretation of MAE must be contextualized with respect to the range and variability of the target variable.

Mean Absolute Percentage Error (MAPE)

The Mean Absolute Percentage Error (MAPE) measures the average relative error between predicted and actual values, expressed as a percentage. It is defined as

$$MAPE = \frac{100}{N} \sum_{i=1}^{N} \left| \frac{y_i - \hat{y}_i}{y_i} \right| \tag{3.4}$$

Like the Mean Absolute Error (MAE), the Mean Absolute Percentage Error (MAPE) is not sensitive to the direction of the error—positive and negative deviations contribute equally to the final score. It is also relatively robust to outliers, as it relies on absolute differences rather than squared errors. As with MAE, lower values indicate better model performance. MAPE addresses one of the key limitations of MAE: its lack of scale independence. Since MAPE expresses error as a percentage of the actual value, it enables meaningful comparisons across datasets with different price levels. This property makes it particularly suitable for evaluating model performance across hotels or accommodation facilities with heterogeneous pricing structures. Despite these advantages, MAPE presents a notable limitation when actual values y_i are close to zero. In such cases, the denominator in the error term becomes very small, which causes even minor deviations in prediction to result in disproportionately large percentage errors. This can lead to instability in the metric and potentially distort the overall evaluation of model performance. For example, predicting 3€ instead of 1€ yields a 200% error, while the same 2€ deviation on a 100€ room rate results in only a 2% error. However, this limitation is generally less relevant in the context of dynamic pricing, where predicted prices tend to remain well above zero. In most real-world applications—such as hotel room pricing—the target variable rarely approaches zero, thereby mitigating the risk of distortion and preserving the reliability of MAPE as a comparative metric.

Coefficient of Determination

The coefficient of determination, denoted as R^2, quantifies the proportion of variance in the dependent variable that is explained by the independent variables included in the model. As such, it serves as a measure of the goodness of fit in regression tasks. It is formally defined as:

$$R^2 = 1 - \frac{\sum_{i=1}^{N}(y_i - \hat{y}_i)^2}{\sum_{i=1}^{N}(y_i - \bar{y}_i)^2} \tag{3.5}$$

where $\bar{y}_i$ represents the mean of the observed values. The R^2 metric ranges from 0 to 1, with higher values indicating a better fit. An R^2 of 1 implies perfect prediction, while a value of 0 indicates that the model does not explain any of the variability in the target variable. Beyond its role as a performance metric, R^2 also serves as a diagnostic tool for assessing the validity of the underlying data structure. A high R^2 suggests that the input features capture meaningful patterns in the target variable, whereas a low R^2 may indicate that the model lacks explanatory power or that the dataset does not contain sufficient signal.

Chapter 4

Experiments and Results

This section presents the experimental framework and the results obtained from the implementation of the methodology described Chapter 3. It encompasses all phases of the empirical analysis, from data preparation to model evaluation, and aims to validate the proposed approach for dynamic pricing prediction in the hospitality domain. We begin by describing the data collection process and the characteristics of the dataset used in the experiments. Since the structure and design rationale of the dataset have already been discussed in the previous section, the focus here is placed on how the data was gathered, prepared, and integrated into the modeling pipeline. This is followed by an ablation study, conducted to identify the optimal configuration of input variables by varying the parameter n in Equation (3.1), which determines the temporal depth of the historical data considered for each prediction. Once the most informative feature set has been established, we proceed to the hyperparameter optimization phase, where each machine learning model is fine-tuned to achieve its best predictive performance. The selected models are then evaluated using appropriate metrics, and their results are compared to assess their effectiveness in replicating the pricing behavior of the Revenue Manager. To gain deeper insight into the nature of the prediction errors, we introduce a custom error decomposition matrix, which breaks down the model's performance across different time horizons and price ranges. This allows for a more granular analysis of where and when the models tend to underperform or excel. Finally, we present a set of reflections and practical considerations based on the results obtained. These include observations on model behavior, implications for real-world deployment, and potential avenues for future research and improvement.

4.1 Data collection

The experimental analysis conducted in this study is based on real-world data of three distinct hotel properties extracted from a Revenue Management System (RMS). These hotels were selected to ensure diversity in terms of geographical location, operational scale, and pricing strategy, thereby allowing for a more robust evaluation of the proposed predictive framework (refer to Figure 4.1).

Feature/ Hotel	Hotel 1	Hotel 2	Hotel 3
Place	Italy	Italy	Switzerland
N° Operative Rooms	27	78	107
N° Entry Level Rooms	17	14	10
Revenue Manager	RM 1	RM 2	RM 3

FIGURE 4.1: Characteristics of the Three Hotels Selected for the Experiments.

The focus of the experiments was placed on the entry-level room category, defined as the most basic and affordable double room offered by each accommodation facility. This choice was motivated by two main considerations. First, in standard revenue management practice, the selling price of the entry-level room is typically set directly by the Revenue Manager, while prices for higher-tier rooms are derived from this base rate through predefined pricing rules or differential markups. Second, limiting the analysis to a homogeneous room category facilitates a more consistent and meaningful comparison across hotels, avoiding confounding effects due to differences in room features, amenities, or pricing structures.

As previously mentioned, the three selected hotels exhibit substantial differences in terms of size, geographical location, and pricing strategy. More precisely:

- Hotel_1, located in Italy, operates with 27 rooms, of which 17 belong to the entry-level category.
- Hotel_2, also based in Italy, has a total of 78 operative rooms (OPRs), including 14 entry-level units.
- Hotel_3, situated in Switzerland, is the largest of the three, with 107 OPRs, but only 10 designated as entry-level.

Each hotel is managed independently and relies on its own Revenue Manager, whose pricing decisions reflect distinct strategic approaches and market sensitivities. This heterogeneity adds depth to the experimental evaluation, as the models must adapt to different pricing behaviors and operational contexts. Following the data extraction phase, the datasets were individually pre-processed to ensure consistency and quality. This included handling missing values, normalizing variables, and aligning temporal records. Subsequently, the data for each hotel was partitioned into training, validation, and test sets to support model development and performance assessment. Each dataset comprises 59,423 samples, covering the period from January 1, 2022, to October 15, 2023, thereby capturing a wide range of seasonal patterns, demand fluctuations, and pricing adjustments. Details regarding the data split

TABLE 4.1: Distribution of Training, Validation, and Test Sets per Hotel.

SET	PERIOD	N° OF SAMPLE
Training	from 2022/01/01 to 2022/12/31	33215
Validation	from 2023/01/01 to 2023/05/31	13741
Test	from 2023/06/01 to 2023/10/15	12467
TOT. Samples		59423

for each hotel are provided in Table 4.1, which outlines the distribution of samples across the different subsets used in the experiments.

4.2 Hyperparameters optimization

As detailed in Chapter 3, each record in the dataset is structured as a tuple of the form described in Equation (3.1). A key design choice in this structure concerns the parameter n, which appears in the dynamic feature component $(x_{t-n,d})$. This parameter defines the temporal depth of the input data—specifically, the number of days preceding the prediction date t that are considered when observing changes related to the target Day of Stay (DOS), denoted by d. To determine the optimal value of n, we conducted a preliminary evaluation by varying its value within the range $n \in \{0, 1, 2, 3, 4, 5\}$. The case $n = 0$ corresponds to a configuration in which only the KPIs recorded on the current day (t) are used to predict the price for future DOS, without incorporating any historical data. Conversely, values of $n = 1$ to $n = 5$ progressively extend the temporal window, allowing the model to incorporate KPIs recorded for the same DOS over the previous one to five days. This approach enables the model to capture short-term trends and fluctuations that may influence pricing decisions. For this analysis, we trained a Random Forest model using default hyperparameter settings and evaluated its performance on the validation set of each hotel, employing the metrics described in Chapter 3.4. The feature set corresponding to the value of n that yielded the best validation performance was then selected as the input configuration for the subsequent hyperparameter optimization phase. Given that each hotel operates under its own dynamic pricing logic—shaped by distinct market conditions, managerial strategies, and operational constraints—the hyperparameters of each machine learning model were optimized individually for each hotel. For this purpose, we identified the relevant hyperparameters for each model (as listed in Table 4.2) and defined a search space of candidate values.

TABLE 4.2: Hyperparameters Considered in Grid Search

MODEL	HYPERPARAMETERS OPTIMIZED
MLR	positive
RF	n_trees max_features max_depth
GB	n_trees max_features max_depth learning_rate
LGB	n_trees colsample_bytree max_depth num_leaves learning_rate
MLP	hidden_layer_size solver learning_rate

A grid search procedure was then employed to explore these configurations systematically. The optimal hyperparameters were selected based on their ability to maximize predictive accuracy and minimize error, as measured by the evaluation metrics: Mean Absolute Error (MAE), Mean Absolute Percentage Error (MAPE), and the coefficient of determination (R^2). All evaluations were performed on the validation sets to ensure that the selected configurations generalize well to unseen data.

4.3 Results

The experimental evaluation begins with the selection of the optimal value for the parameter n, which determines the temporal depth of the dynamic features $(x_{t-n,d})$ used in the predictive model. As shown in Table 4.3, the performance of the models on the validation sets was assessed for each hotel using three key metrics: Mean Absolute Error (MAE), Mean Absolute Percentage Error (MAPE), and the coefficient of determination (R^2).

TABLE 4.3: Results on the validation set per each hotel changing the value for n.

	Hotel_1			Hotel_2			Hotel_3		
n/Measures	**MAE**	**MAPE**	**R^2**	**MAE**	**MAPE**	**R^2**	**MAE**	**MAPE**	**R^2**
0	16.0821	9.6780	0.0350	22.7524	14.5762	-0.1676	16.5275	5.2303	0.2887
1	2.8754	1.7757	0.9410	3.9295	2.5144	0.9419	3.2977	0.9301	0.8885
2	2.9065	1.7976	0.9396	3.8572	2.4675	0.9454	2.9252	0.8428	0.9180
3	**2.8089**	**1.7383**	**0.9411**	3.6714	2.3433	0.9499	2.5034	0.7479	0.9432
4	2.8494	1.7608	0.9385	3.8116	2.4315	0.9477	2.561	0.7543	0.9370
5	2.8748	1.7776	0.9389	**3.6363**	**2.3132**	**0.9513**	**2.3755**	**0.7065**	**0.9449**

For Hotel_1, the best results were obtained when the model considered data from the three days preceding the prediction date ($n = 3$). However, the performance differences across values of n were relatively small, except for the case $n = 0$, which consistently underperformed due to the lack of historical context. In contrast, for Hotel_2 and Hotel_3, the highest predictive accuracy was achieved with $n = 5$, indicating that incorporating a broader temporal window of historical KPIs significantly improves model performance. Based on this observation, and to maintain consistency across experiments, $n = 5$ was selected as the optimal value for all subsequent analyses. With the feature set corresponding to $n = 5$ established, the next phase involved training and evaluating the machine learning models, each optimized via grid search on the validation sets. The final performance results on the test sets are reported in Table 4.4.

TABLE 4.4: Results on test set for Hotel_1, Hotel_2 and Hotel_3

	Hotel_1			Hotel_2			Hotel_3		
Model/Measures	**MAE**	**MAPE**	**R^2**	**MAE**	**MAPE**	**R^2**	**MAE**	**MAPE**	**R^2**
MLR	3.1820	1.9087	0.9189	**3.4445**	**2.0197**	**0.9037**	**0.9380**	**0.3771**	**0.9662**
RF	2.5146	1.4629	0.9232	6.7323	3.4384	0.8423	1.8742	0.7391	0.9112
GB	1.9479	1.1307	0.9240	5.9982	3.2247	0.8429	1.7649	0.7051	0.9124
LGB	**1.8845**	**1.1000**	**0.9340**	5.9894	3.1863	0.8343	1.7659	0.6711	0.8676
MLP	2.7272	1.6201	0.9182	3.5753	2.1117	0.9080	1.0482	0.4136	0.9650

To ensure comparability across hotels, the predicted and actual prices for Hotel_3, which operates in Swiss francs, were converted to euros using the exchange rate in effect at the time of the experiments (1.04 CHF per EUR). Analysis of the test results reveals that no single model consistently outperforms the others across all hotels, highlighting the influence of hotel-specific pricing dynamics and data characteristics. For Hotel_1, the Light Gradient Boosting (LGB) model achieved the best performance, demonstrating its ability to capture complex nonlinear relationships in smaller-scale settings. Conversely, for Hotel_2 and Hotel_3, the Multiple Linear Regression (MLR) model yielded superior results, suggesting that in certain contexts, simpler models may generalize better and offer more stable predictions. Among the three hotels, Hotel_3 recorded the most accurate predictions overall, with a MAPE of 0.3771 and an R^2 of 0.9662, indicating a high degree of alignment between predicted and actual prices. This result may be attributed to the hotel's larger dataset, more consistent pricing behavior, or clearer seasonal patterns, which enhance the model's ability to learn and generalize. These findings underscore the importance of tailoring model selection and feature design to the specific characteristics of each accommodation facility. They also highlight the potential of data-driven approaches to emulate human revenue management decisions with a high degree of precision.

4.3.1 Error Analysis by Time Horizon and Price Band

Given the nature of dynamic pricing strategies, it is essential not only to assess the overall performance of predictive models, but also to investigate how errors vary across different time horizons and price ranges. This granular analysis enables a deeper understanding of the model's behavior under varying forecasting horizons and pricing contexts, and may reveal systematic biases or limitations that are not evident from aggregate metrics alone. To this end, the test datasets were segmented along two dimensions:

- Forecasting horizon (H_i): representing the number of days between the prediction date and the target Day of Stay (DOS).
- Price band (B_i): representing the relative position of the predicted price within the distribution of prices for each hotel.

A closer look at each segmentation strategy is provided below, outlining the operational motivations and statistical foundations behind the chosen intervals.

Time Horizon Segmentation

The division of the prediction horizons was carried out in consultation with domain experts in revenue management, ensuring that the intervals reflect meaningful operational planning windows. The following four segments were defined:

H_1: short-term forecasts (0–7 days ahead)

H_2: near-term forecasts (8–15 days ahead)

H_3: mid-term forecasts (16–30 days ahead)

H_4: long-term forecasts (31–90 days ahead)

This segmentation allows for the evaluation of model performance across different planning horizons, which may be subject to varying levels of uncertainty and managerial intervention.

Price Band Segmentation

Due to the substantial variability in price distributions across the three hotels (see boxplots in Figure 4.2), a fixed or arbitrary definition of price ranges would have been inappropriate.

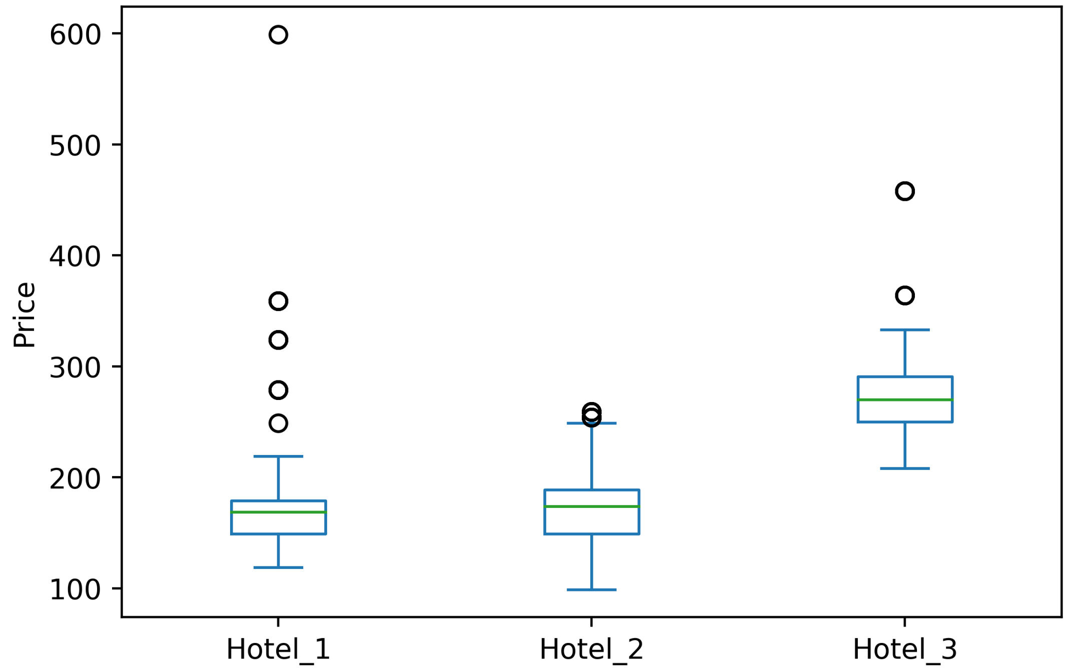

FIGURE 4.2: Boxplots of Price Distributions in the Test Set per Hotel

Instead, a data-driven approach was adopted, using key statistical indices to define price bands that are both consistent and adaptable to each hotel's pricing structure. Based on the minimum, 25th percentile, median (50th percentile), 75th percentile, and maximum values of each hotel's price distribution (see Table 4.5), the following four price bands were established:

B_1: prices between the minimum and the 25th percentile.

B_2: prices between the 25th and 50th percentile.

TABLE 4.5: Statistical indices of the target variable computed on the test set for Hotel_1 (a), Hotel_2 (b), and Hotel_3 (c)

STATISTICAL INDEX	VALUE
Max	599
75 p.	179
50 p.	169
25 p.	149
Min	119

(A)

STATISTICAL INDEX	VALUE
Max	259
75 p.	189
50 p.	174
25 p.	149
Min	99

(B)

STATISTICAL INDEX	VALUE
Max	458
75 p.	291
50 p.	270
25 p.	250
Min	208

(C)

B_3: prices between the 50th and 75th percentile.

B_4: prices between the 75th percentile and the maximum.

This method ensures that the segmentation reflects the actual distribution of prices, allowing for a fair comparison of model performance across low, medium, and high pricing tiers.

Combined Error Matrix

By jointly considering the four time horizons and four price bands, we constructed error matrices that report the Mean Absolute Percentage Error (MAPE) for each combination of horizon and price range. These matrices, illustrated in Figure 4.3, provide a detailed view of how prediction accuracy varies depending on both temporal and pricing dimensions.

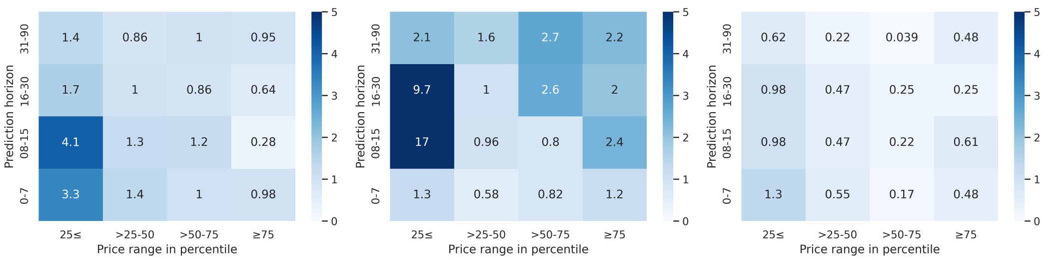

(A) (B) (C)

FIGURE 4.3: MAPE Matrices Across Time Horizons and Price Band for Hotel_1 (a), Hotel_2 (b) and Hotel_3 (c)

Such a breakdown is particularly valuable for identifying patterns of underperformance — e.g., whether the model struggles more with long-term forecasts for high-priced rooms, or short-term predictions in low-price segments. These insights can inform future improvements in feature engineering, model selection, and operational deployment.

4.4 Discussion and implications

The experimental analysis conducted to determine the optimal value for the parameter n — which governs the temporal depth of the dynamic input features — has provided strong empirical support for the effectiveness of the dataset construction methodology adopted in this study. As illustrated in Figure 4.4, the performance gap between the configuration with no historical input ($n = 0$) and the one incorporating up to five days of prior data ($n = 5$) is substantial across all three accommodation facilities. This confirms the critical role of temporal context in enhancing predictive accuracy within dynamic pricing scenarios.

Notably, Figure 4.4c reveals that when past data is included, the models achieve comparable values of the coefficient of determination (R^2) across all hotels. This convergence suggests that the proposed feature engineering strategy enables the models to generalize effectively across heterogeneous pricing environments. In other words, the inclusion of dynamic variables - capturing the evolution of KPIs over time — allows the tabular dataset to emulate the temporal structure typically required by classical time-series models, while remaining compatible with machine learning algorithms designed for structured data. This hybrid approach bridges the gap between time series forecasting and tabular predictive modeling, offering a flexible framework that can be adapted to various operational contexts without sacrificing

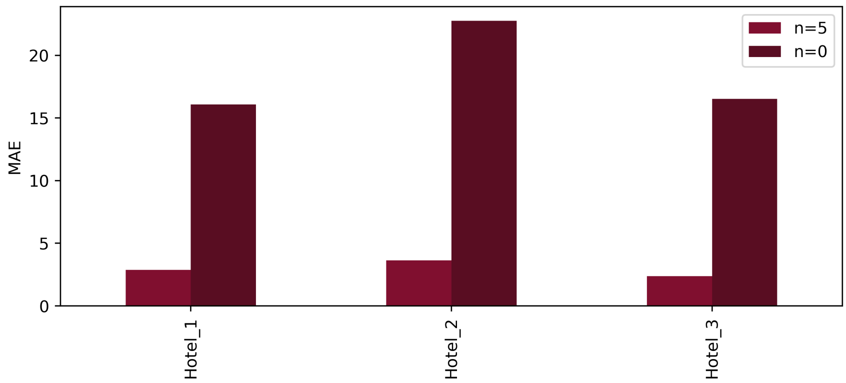

(A)

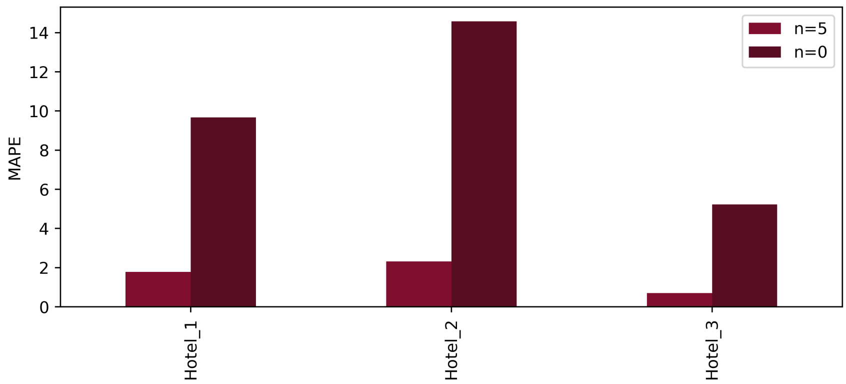

(B)

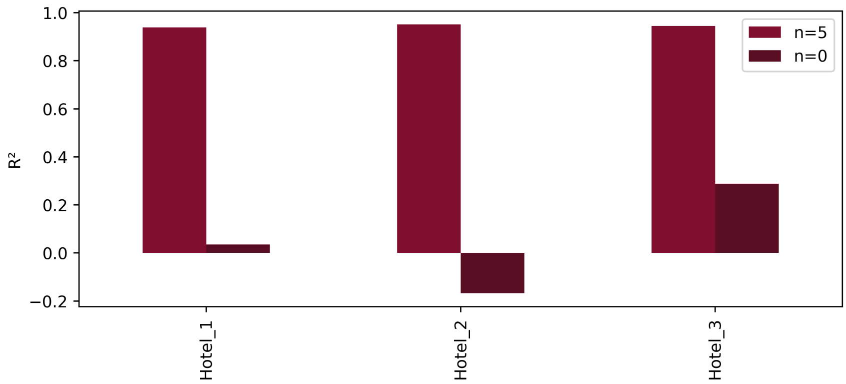

(C)

FIGURE 4.4: Variation of MAE (a), MAPE (b), and R^2 (c) with Respect to n

temporal richness. It also demonstrates that meaningful temporal dependencies can be encoded within a tabular format, enabling the use of powerful ML techniques without the need for specialized time-series architectures.

From a modeling perspective, the comparative analysis of algorithmic performance (see Table 4.4) highlights that Light Gradient Boosting (LGB) and Multiple Linear Regression (MLR) are the top-performing models for specific hotels, reflecting their ability to capture non-linear patterns or stable linear relationships depending on the data characteristics. However, when considering the overall average error across all facilities, the Multi-Layer Perceptron (MLP) emerges as the most robust model, as shown in Figure 4.5.

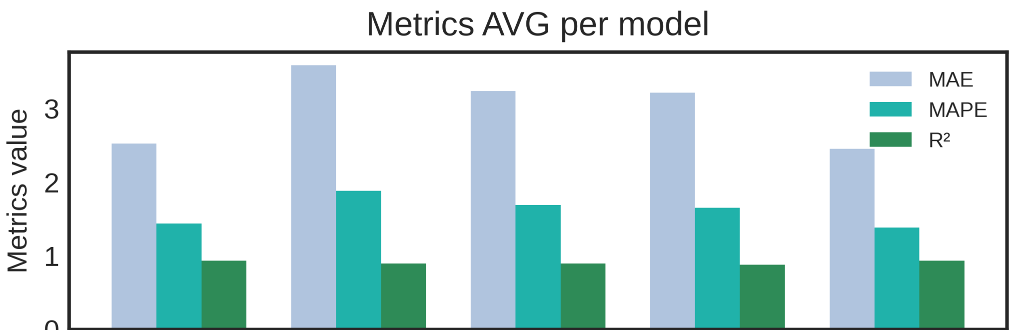

	MLR	RF	GB	LGB	MLP
MAE	2.5215	3.587	3.237	3.2133	2.4502
MAPE	1.4352	1.8801	1.6868	1.6525	1.3818
R^2	0.9296	0.8922	0.8931	0.8786	0.9304

FIGURE 4.5: Average Performance Metrics per Model Across Hotels.

This finding carries important practical implications. In scenarios where a Revenue Management System (RMS) seeks to deploy a unified predictive model across multiple properties — potentially with varying data distributions and pricing strategies — the MLP offers a promising solution. Its ability to generalize across different hotel contexts while maintaining competitive accuracy makes it a suitable candidate for scalable, centralized deployment.

An additional point of interest emerging from the experimental results concerns Hotel_2, which consistently exhibits the highest prediction errors in terms of Mean Absolute Error (MAE) and Mean Absolute Percentage Error (MAPE), as well as the lowest coefficient of determination (R^2) among the three properties (see Table 4.4 and Figure 4.3b). This discrepancy prompted a closer investigation into potential contextual factors that may have influenced model performance. A plausible explanation lies in a managerial transition that occurred at Hotel_2: starting from May 1, 2023, the property Hotel_2 underwent a managerial transition, adopting a new Revenue Manager whose pricing strategy likely diverges from that of their predecessor. The predictive model, however, was trained and validated almost entirely on data reflecting the previous manager's approach—only the final 30 days of the validation set correspond to the new pricing logic. As a result, the test set, which predominantly captures the decisions of the new manager, presents a distributional shift relative to the training and validation phases. This mismatch introduces a form of concept drift, whereby the underlying decision patterns evolve over time, reducing the model's ability to generalize effectively to the new context. Despite this shift, the model's performance remains within acceptable bounds, suggesting a degree of resilience and adaptability. In such cases, it is reasonable to expect that the model will gradually recalibrate as more data reflecting the new strategy becomes available. This highlights the importance of continuous learning and retraining in dynamic environments, especially when human decision-makers change or adjust their strategies. Further supporting this hypothesis is the observation that, while Hotel_1 and Hotel_3 show slight improvements in performance on the test set compared to the validation set, Hotel_2 experiences a deterioration of 1.0323 in MAE and 0.4218 in MAPE. This contrast reinforces the idea that consistency in managerial behavior contributes positively to model stability and accuracy. It is also worth noting that the modest improvements observed in the test sets of Hotel_1 and Hotel_3 are not unexpected, despite the validation sets being used for hyperparameter tuning. These improvements may be attributed to the seasonal composition of the test data, which includes the summer period — a time characterized by higher booking volumes and more frequent price adjustments. The model may have benefited from a richer and more representative training set, enabling it to better capture the pricing dynamics typical of high-demand periods. Finally, a deeper analysis of the error distribution matrices presented in Figure 4.3 reveals that the highest prediction errors are concentrated in the B_1 price band (i.e., the lowest quartile of prices) and within the H_1 and H_2 time horizons (i.e., predictions made 0–15 days before the Day of Stay). This pattern suggests that the model struggles most when forecasting short-term prices for lower-cost rooms. One possible explanation for this behavior is the absence of certain contextual variables that are particularly relevant in short-term booking scenarios. For instance, meteorological conditions — which can significantly influence last-minute travel decisions — were not included in the feature set. In fact, when prices are lower than average, demand may be more sensitive to external factors. The omission of such variables may limit the model's ability to capture these nuances, especially in the short-term forecasting window.

Chapter 5

Explainable Artificial Intelligence

Despite the remarkable advancements in artificial intelligence and its widespread adoption across industries, many end-users —particularly those without technical backgrounds — remain skeptical or even intimidated by the opaque nature of machine learning models. This concern becomes especially pronounced when AI systems are used to support decision-making processes that directly affect business outcomes, such as revenue optimization in the hospitality sector. In this context, Explainable Artificial Intelligence (XAI) emerges as a crucial paradigm aimed at bridging the gap between model complexity and user interpretability. The primary goal of XAI is to make the behavior of machine learning models transparent and understandable, even to non-expert users. For Revenue Managers (RMs), who are expected to rely on AI-generated predictions to guide pricing strategies, the ability to comprehend and validate model outputs is essential for building trust and ensuring informed decision-making. To address this need, we integrated an XAI technique into our predictive framework, enabling RMs to explore and interpret the rationale behind each prediction. This not only enhances the system's usability but also aligns with the broader objective of responsible and human-centered AI deployment.

5.1 Shapley Additive exPlanations

Among the various XAI methodologies available, we selected SHAP (SHapley Additive exPlanations), introduced by Lundberg and Lee, 2017, for its versatility and theoretical robustness. SHAP is a model-agnostic technique, meaning it can be applied to any machine learning model regardless of its internal architecture. Moreover, SHAP provides both global and local explanations, allowing users to understand feature importance at both the aggregate and individual prediction levels. The foundation of SHAP lies in cooperative game theory, where the goal is to fairly distribute the total payoff of a game among its players. In the context of machine learning, the "players" are the input features, and the "payoff" is the model's prediction. The SHAP value for a feature quantifies its contribution to the prediction by computing the average marginal impact across all possible feature combinations. Formally, given a set of features F and a value function v , the Shapley value for a feature i is defined as:

$$\phi(F, v) = \sum_{S \subseteq F \setminus \{i\}} \frac{|S|!(|F| - |S| - 1)!}{|F|!} [v(S \cup \{i\} - v(S)] \tag{5.1}$$

This equation computes a weighted average of the marginal contributions of feature i across all possible subsets S of the feature set F. This formulation ensures a fair and consistent attribution of influence, making SHAP a powerful tool for interpretability.

To better illustrate the interpretability offered by SHAP, below we delve into its two main applications: global explanations, which provide an overview of feature relevance across all predictions, and local explanations, which reveal the reasoning behind specific outputs.

5.1.1 Global Explanation

Global explanations aim to provide an overview of how each feature contributes to the model's predictions across the entire dataset. In our study, we employed the bee swarm plot to visualize global SHAP values. This plot displays the distribution of SHAP values for each feature, ranked by importance on the y-axis, with the magnitude of impact represented on the x-axis. Each dot corresponds to an individual observation and is color-coded based on the feature's value—highlighting how high or low values influence the prediction.

An example of a bee swarm plot is shown in Figure 5.1, referring to Hotel_1.

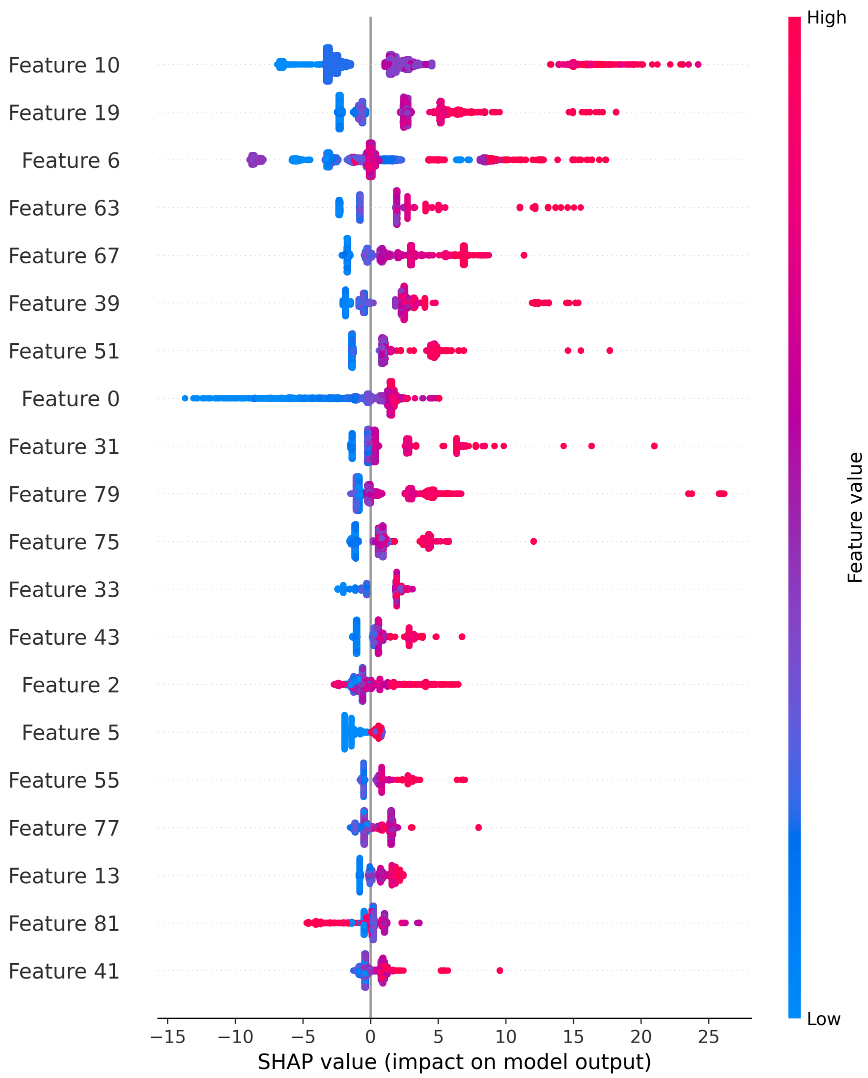

FIGURE 5.1: Bee Swarm Plot of SHAP Values for Hotel_1

It reveals that Feature 10 (starting selling rate) is the most influential variable. High values of this feature tend to increase the predicted price, while lower values exert a negative influence. Similarly, Feature 0 (LeadTime) demonstrates a nuanced behavior: shorter lead times generally reduce the predicted price, whereas longer lead times contribute positively.
This visualization enables Revenue Managers to assess whether the model's interpretation of feature importance aligns with their domain expertise and pricing logic. It serves as a diagnostic tool to validate the model's reasoning and identify potential inconsistencies.

5.1.2 Local Explanation

While global explanations offer a high-level understanding, local explanations focus on individual predictions, revealing how specific features influenced the outcome for a given Day of Stay (DOS). For this purpose, we utilized the waterfall plot, which decomposes a single prediction into its constituent feature contributions.
At the base of the waterfall plot lies the expected value, representing the average prediction across the training set. This serves as a reference point when no input information is available. Each feature then adds (red bars) or subtracts (blue bars) from this base value, culminating in the final prediction shown at the top of the plot. Mathematically, the sum of all SHAP values equals the difference between the actual prediction $f(x)$ and the expected value $E[f(x)]$.
Figure 5.2 presents waterfall plots for three different DOS instances from Hotel_1. Notably, the base value remains constant across plots, while the magnitude and direction of feature contributions vary.

For example, Feature 10 (starting selling rate) increases the price in cases b) and c), but decreases it in case a), where it also ranks lower in importance.

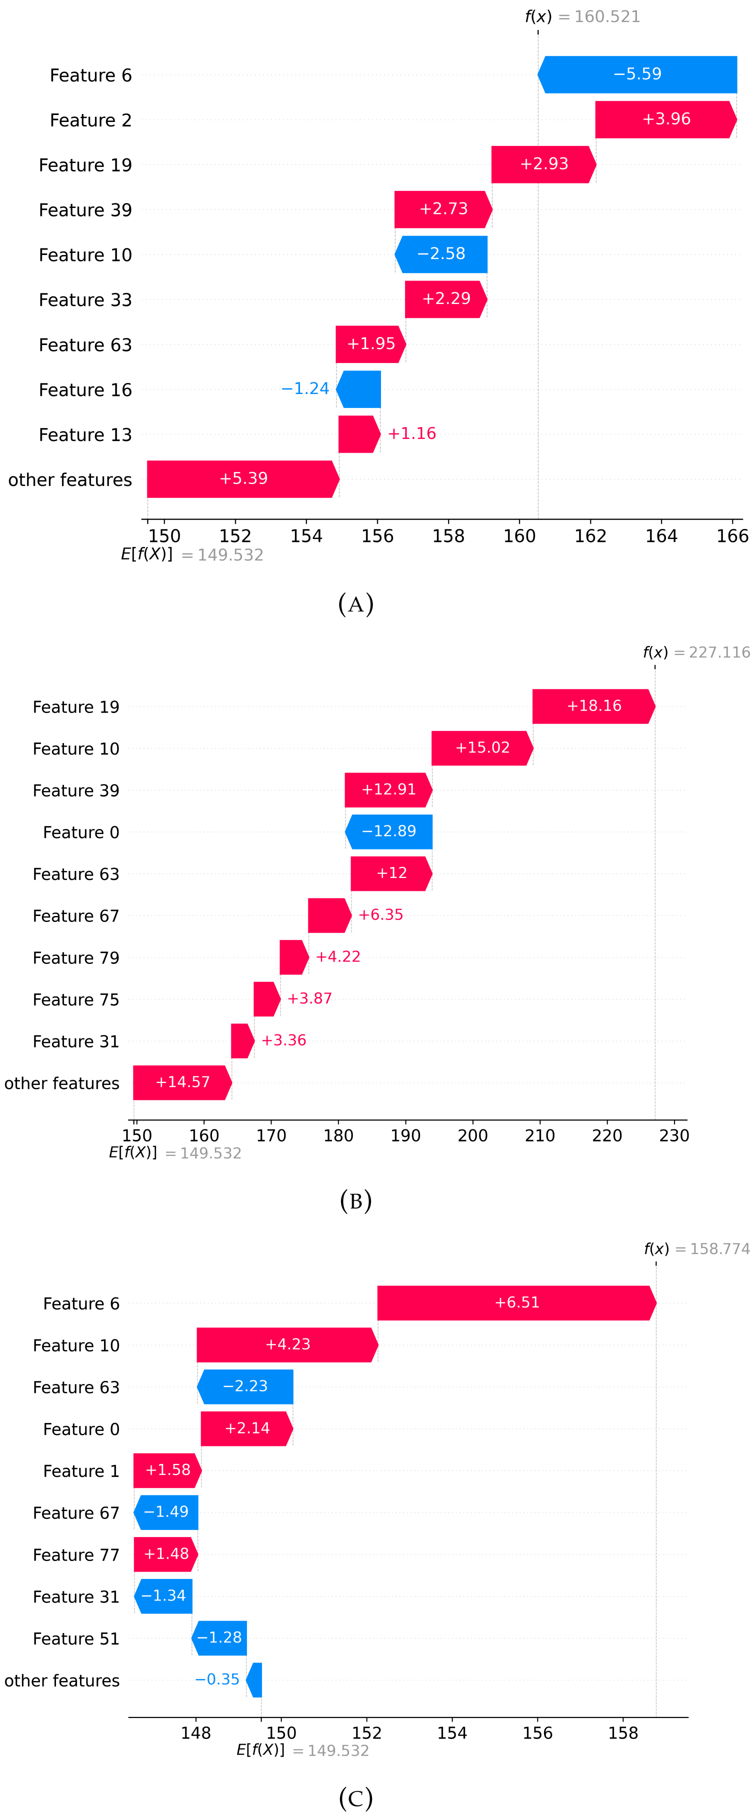

FIGURE 5.2: Waterfall Plots of SHAP Values for Three Days of Stay at Hotel_1

Feature 0, that is the LeadTime, is in the same position in b) and c) but in the first case it decreases the selling price of 12.89 euros while in the second case it adds 2.14 euros to the final prediction. At the same time, Leadtime is not even present among the most impactful features in a). Similarly, Feature 6 (week of the year) is the most impactful in a) and c), but with opposite effects—decreasing the price in a) and increasing it in c)—while it is irrelevant in b).

These insights underscore the granularity and flexibility of SHAP in capturing context-specific dynamics within predictive models. For Revenue Managers, such visualizations provide a transparent and intuitive view of the model's decision-making process, allowing them to identify not only which features influenced a given prediction, but also the magnitude of their impact for each specific Day of Stay. This level of interpretability fosters trust in the system and supports more informed, data-driven interventions.

Chapter 6

System Architecture

This section presents the architecture of the proposed system, designed to support predictive dynamic pricing in the hospitality domain. To enhance clarity and facilitate comprehension, the process is illustrated using standard flowchart symbols, where each shape corresponds to a specific task or decision point. The complete architecture is depicted in Figure 6.1, and is structured around two main operational flows: training time and inference time.

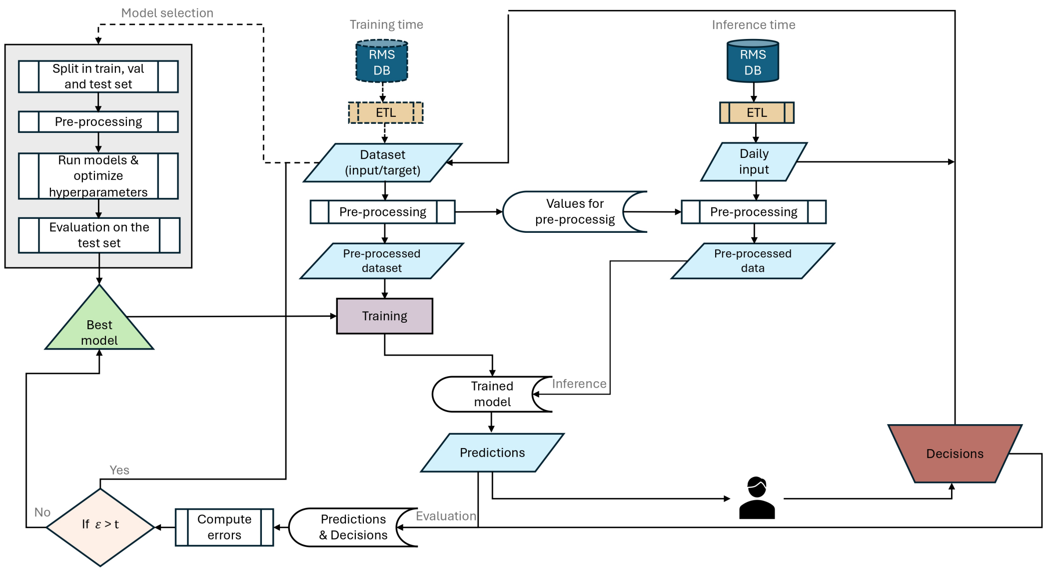

FIGURE 6.1: Process Architecture Overview.

In the following subsections, we provide a detailed description of both workflows: the Training Time Workflow, which governs model development and optimization, and the Inference Time Workflow, which manages daily prediction generation and integration into decision-making processes.

6.1 Training Time Workflow

The training phase begins with the acquisition of data from the source system, which in our study is a Revenue Management System (RMS) database. These raw data are processed through an Extract, Transform, and Load (ETL) pipeline, which performs cleaning, formatting, and structuring operations to produce the initial dataset, composed of input features and the target variable. It is important to note that the massive data extraction — covering two years of historical records — and its organization through the ETL process are required only during the initial execution of the system. Once the system is initialized, the dataset becomes available and is incrementally updated on a daily basis. The daily extraction and ETL application are then handled by the inference workflow. For this reason, the corresponding blocks in the flowchart are represented with dotted outlines, indicating that they are executed only once. The initial dataset is then passed to the model selection phase, which encompasses all the procedures described in Chapters 4.1 and 4.2.

This phase includes:

- Splitting the dataset into training, validation, and test sets.
- Preprocessing the data.
- Training ML models.
- Optimizing hyperparameters.
- Evaluating performance on the validation set.

The arrow connecting the dataset to the model selection block is also dotted, reflecting that this process is executed only during initialization or when re-training is triggered by performance degradation.
At the end of the model selection phase, the system identifies the best-performing model—i.e., the one that achieved the highest accuracy on the test set — and the corresponding optimal hyperparameter configuration. These outputs are passed to the next stage, where the entire initial dataset is preprocessed again and used to retrain the selected model using all available data. The final trained model is then saved and deployed for inference.

6.2 Inference Time Workflow

During inference, the system operates on a daily basis. Data related to the next 90 Days of Stay (DOS) are extracted from the RMS and processed through the ETL pipeline to generate the input features. These inputs are then pre-processed using the same parameters and transformations applied during training, ensuring consistency and compatibility with the deployed model.
The preprocessed data are fed into the trained model to generate price predictions, which are then presented to the Revenue Manager (RM). At this stage, the RM can either accept the predictions or modify them manually based on contextual knowledge or strategic considerations. The RM's decisions, along with the daily input data, are appended to the dataset and stored for future training cycles. In addition, both predictions and decisions are archived to enable periodic performance evaluation. During this evaluation, the system computes the prediction error ϵ. If the error exceeds a predefined threshold, the model selection process is triggered again in the next training cycle to identify a potentially better model. Conversely, if the error remains below the threshold, the system continues to use the previously selected model and its optimal hyperparameters.

6.3 Automation and Human Oversight

It is worth emphasizing that the manual intervention by the Revenue Manager — who reviews and potentially adjusts the predictions — can be eliminated once sufficient trust in the system has been established. In such cases, the entire workflow can be fully automated, allowing the model to operate independently and continuously update pricing recommendations without human oversight.

This architectural design ensures a balance between automation and interpretability, enabling the system to evolve over time while remaining transparent and adaptable to managerial input.

Chapter 7

Ongoing Work and Future Outlook: Optimal Dynamic Pricing

Having addressed the problem of dynamic pricing prediction in previous chapters — where the objective was to replicate the pricing behavior of the Revenue Manager (RM) through machine learning models — a natural progression is to move beyond imitation and toward optimization. The predictive approach, already implemented within the partner company's RMS, enables automation of the RM's decisions based on historical data. However, this strategy assumes that those decisions are inherently optimal, which may not always be the case. In contrast, the focus of this chapter is on optimal dynamic pricing, where the model is not trained to reproduce past choices, but to learn a pricing strategy that actively interacts with the environment and adapts to feedback, with the explicit goal of maximizing revenue and occupancy. This shift allows the system to overcome the limitations of human heuristics and potentially discover more effective pricing policies. The output is no longer a replication of previous decisions, but a price suggestion—an informed recommendation designed to improve business performance.

This transition is motivated by two key considerations:

- First, RM strategies do not always pursue strict revenue or occupancy maximization, as they may be influenced by subjective preferences, operational constraints or limited data availability.

- Second, many small and medium-sized accommodations lack the resources to employ a dedicated RM, and could benefit significantly from intelligent pricing suggestions that help them remain competitive.

By designing a model that learns from outcomes rather than decision, thus moving beyond imitation and toward optimization, this work aims to build a system that is both autonomous and goal-driven, capable of supporting a wide range of hospitality businesses.

7.1 Related Works

In the current research landscape, Reinforcement Learning (RL) has gained increasing attention as a viable and promising approach for dynamic pricing in the hospitality industry, offering adaptive, data-driven strategies that go beyond traditional rule-based or forecast-driven methods. Unlike supervised learning, which relies on historical labels, RL models learn optimal pricing policies through interaction with the environment, aiming to maximize cumulative rewards such as revenue or occupancy.
Several studies have explored this paradigm. For instance, *Tuncay et al.* formulate hotel room pricing as a Markov Decision Process (MDP) and propose a Q-learning–based model that introduces a custom reward function balancing profit and demand, while also addressing the cold-start problem through hotel-specific feature encoding that enables generalization to new properties (Tuncay et al., 2023). Similarly, *Singh et al.* employ a Deep Q-Network (DQN) agent trained on real hotel booking data to recommend daily pricing strategies, achieving a 15–20% increase in profit margins and a significant reduction in unoccupied rooms compared to fixed and rule-based pricing

(Singh, 2022). More recent contributions, such as the work by *Dietz* (Dietz, 2023), introduce hybrid frameworks that combine RL with monotonic constraints using Bernstein polynomials, allowing for controlled policy updates and ensuring pricing consistency over time. This integration enhances both the interpretability and stability of RL-based pricing systems, addressing key challenges in their deployment within real-world revenue management environments.

Overall, RL-based approaches represent a significant advancement in dynamic pricing, enabling continuous learning and adaptation in response to fluctuating market conditions, and paving the way for more autonomous and intelligent pricing strategies in hospitality. To provide a clearer overview of the methodological diversity and evaluation strategies adopted in recent research, Table 7.1 summarizes the main RL-based approaches applied to dynamic pricing in the hospitality sector.

TABLE 7.1: Comparative overview of RL-based dynamic pricing models in hospitality

Study	**RL Approach**	**Reward Function**	**Key Variables**	**Evaluation Metrics**	**Highlights**
Tuncay et al. (2023) Tuncay et al., 2023	Q-Learning	Profit + Demand balance	Hotel features, booking horizon	Cumulative reward	Addresses cold-start problem; generalizes to new hotels via feature encoding
Singh et al. (2022) Singh, 2022	Deep Q-Network (DQN)	Revenue maximization	Room type, length of stay, demand forecast	Profit increase, occupancy rate	Tested on real hotel data; reduces unsold rooms significantly
Dietz (2023) Dietz, 2023	RL + Bernstein constraints	Revenue with monotonicity	Price history, demand simulation	Policy stability, interpretability	Ensures consistent pricing behavior; improves trust and control

The comparison highlights the type of reinforcement learning algorithm used, the structure of the reward function, the key variables considered, and the metrics employed to assess performance. It also outlines the distinctive contributions of each study, offering a concise reference framework for understanding how these models address the complexity of pricing decisions in real-world hotel environments.

7.2 Preliminary Methodological Design

As we already mentioned, the project aims to develop a Reinforcement Learning (RL) agent capable of suggesting optimal room prices in order to maximize both revenue and occupancy. At this stage, the methodological design is still preliminary, and the structure of the algorithm is defined at a high level. The core idea is to train an agent that, given a specific state (e.g., date, availability, seasonality, competitor prices), outputs a price suggestion. This price is then evaluated through a demand estimation model, which predicts the expected occupancy for that price. Based on the predicted occupancy, the system computes the reward, which reflects the revenue generated and the degree of occupancy achieved. The following pseudocode outlines the conceptual structure of the RL algorithm:

Algorithm 1 Dynamic Pricing Agent Training Procedure

Initialize agent with policy π
Initialize environment with demand estimation model
for each episode **do**
    for each time step t **do**
        Observe state s_t
        Suggest price $p_t = \pi(s_t)$
        Estimate occupancy $q_t = \texttt{DemandModel.predict}(s_t, p_t)$
        Compute revenue $r_t = p_t \cdot q_t$
        Define reward $R_t = f(r_t, q_t)$
        Update policy: $\pi \leftarrow \texttt{UpdatePolicy}(\pi, R_t, s_t, p_t)$
    end for
end for

The reward function, which plays a central role in guiding the agent's learning process, has not yet been formally defined. At this stage, its structure remains open and will be shaped through iterative experimentation. The challenge lies in designing a reward that appropriately balances revenue and occupancy, while remaining sensitive to business constraints and operational priorities. Several formulations are being considered, including linear combinations of estimated revenue and occupancy, or more complex weighted objectives that reflect strategic trade-offs.

The demand estimation model may be trained separately using historical booking data and can be updated periodically to reflect market changes.

This preliminary structure provides a foundation for future development, where the agent will be trained in a simulated environment and later validated on real-world data. The modularity of the design allows for flexibility in refining the state representation, the reward function, and the learning algorithm itself.

As a part of this preliminary design, a comparative evaluation is planned to assess the effectiveness of the proposed RL-based pricing agent. Although dynamic pricing prediction and optimal dynamic pricing pursue different objectives—the former replicating the Revenue Manager's decisions, and the latter aiming to learn a strategy that maximizes revenue and occupancy—it is nonetheless valuable to compare them. This comparison will help determine whether an autonomous agent, trained through interaction with the environment, can outperform the pricing decisions historically made by the RM. Such an evaluation will provide empirical insight into the potential of reinforcement learning to enhance pricing strategies and support more outcome-driven revenue management.

Chapter 8

Conclusions

This thesis presented a comprehensive approach to address the challenge of dynamic pricing in the hospitality sector, with the specific goal of assisting Revenue Managers (RMs) in their daily decision-making processes. In contexts where prices must be continuously monitored and adjusted, manual interventions become increasingly costly and time-consuming. To mitigate this issue, we designed a tailored dataset structure—reproducing the informational context in which a Revenue Manager operates—composed of static, dynamic, and engineered features. We then benchmarked five machine learning models (MLR, RF, GB, LGB, and MLP) to predict the price that a Revenue Manager would set for the entry-level room over a 90-day horizon.

The proposed approach was tested across three different hotels, demonstrating its generalizability and robustness. The results confirmed the effectiveness of the methodology, although some limitations emerged in specific segments — particularly in the 0–15 day time horizon and the lowest price band — where prediction errors were more pronounced. These findings suggest that future developments should consider the integration of additional contextual variables, such as meteorological data, which may influence short-term booking behavior and pricing decisions.

To ensure that the system is not only accurate but also interpretable and trustworthy, we incorporated Explainable Artificial Intelligence (XAI) techniques into the framework. Specifically, we adopted Shapley Additive exPlanations (SHAP), a model-agnostic method capable of providing both global and local insights into feature importance. This choice proved essential in enabling Revenue Managers— often non-expert users — to understand, validate, and confidently interact with the model's predictions.

In addition to the modeling and interpretability components, we designed a complete system architecture that transforms the prototype into a deployable product. This architecture supports both training and inference workflows, ensuring scalability, maintainability, and operational efficiency.

Importantly, the outcomes of this research have already been translated into practice: the predictive system developed during these years of research is currently in production within the RMS of the partner company, actively supporting pricing decisions. This successful deployment confirms the practical relevance and impact of the work, and demonstrates how academic research can directly contribute to business innovation.

In summary, the integration of dynamic features, the validation of the dataset design, the comparative modeling analysis, and the adoption of XAI techniques collectively demonstrate the feasibility and effectiveness of applying machine learning to dynamic pricing in hospitality. The methodological framework developed in this study offers a replicable blueprint for future research and real-world implementation, while also highlighting the importance of transparency, adaptability, and contextual awareness in AI-driven decision systems.

Looking ahead, ongoing research is exploring a more ambitious direction: the development of a reinforcement learning–based agent capable of suggesting optimal prices by interacting with the environment and learning from outcomes rather than past decisions. This transition from imitation to optimization aims to overcome the limitations of human heuristics and unlock more effective pricing strategies, particularly for small and medium-sized accommodations that lack access to dedicated revenue management expertise. The preliminary design of this agent lays the groundwork for future experimentation and comparative evaluation, opening new possibilities for autonomous and goal-driven pricing systems in hospitality.

Bibliography

Abrate, G., J. L. Nicolau, and G. Viglia (2019). "The impact of dynamic price variability on revenue maximization". In: *Tourism Management, vol. 74, pp. 224-233*. URL: https://doi.org/10.1016/j.tourman.2019.03.013.

Al Shehhi, M. and A. Karathanasopoulos (2018). "Forecasting Hotel Prices in Selected Middle East and North Africa Region (MENA) Cities with New Forecasting Tools". In: *Theoretical Economics Letters, vol. 8, pp. 1623-1638*. DOI: 10.4236/tel.2018.89104.

— (2020). "Forecasting hotel room prices in selected GCC cities using deep learning". In: *Journal of Hospitality and Tourism Management, vol. 42, pp. 40-50*. DOI: 10.1016/j.jhtm.2019.11.003.

Alshakhsheer, F. et al. (2017). "Improving Hotel Revenue through the Implementation of a Comprehensive Dynamic Pricing Strategy: A Conceptual Framework and Empirical Investigation of Jordanian Hotels". In: *Business Management Dynamics* 7, pp. 19–33. URL: https://www.researchgate.net/publication/323258769_Improving_Hotel_Revenue_through_the_Implementation_of_a_Comprehensive_Dynamic_Pricing_Strategy_A_Conceptual_Framework_and_Empirical_Investigation_of_Jordanian_Hotels_Business_Management_Dynamics_Vol7.

Bandalouski, A. M. et al. (2021). "Dynamic with demand disaggregation for hotel revenue management". In: *Journal of Heuristics*. URL: https://doi.org/10.1007/s10732-021-09480-2.

Bayoumi, A. et al. (2013). "Dynamic pricing for hotel revenue management using price multipliers". In: *Journal of Revenue and Pricing Management*. DOI: 10.1057/rpm.2012.44.

Cannavò, F. et al. (Apr. 2023). "Machine Learning and Microseism as a Tool for Sea wave Monitoring". In: *EGU General Assembly*. URL: https://meetingorganizer.copernicus.org/EGU23/EGU23-7492.html.

Deksnyte, I. and Z. Lydeka (2012). "Dynamic Pricing and its Forming Factors". In: *Journal of Business and Social Science*. URL: https://www.researchgate.net/publication/264974538_Dynamic_Pricing_and_Its_Forming_Factors.

Dietz, C.V. (2023). "Enhancing Dynamic Pricing in the Hotel Industry with Monotonic Constraints: A Reinforcement Learning Approach with Bernstein Polynomials". MA thesis. Politecnico di Milano. URL: https://www.politesi.polimi.it/handle/10589/218422.

El-Nemr, N., B. Canel-Depitre, and A. Taghipour (2019). "Determinants of hotel room rates determinants of hotel room rates". In: *Luxury Industries Conference*. URL: https://www.researchgate.net/publication/344942316_DETERMINANTS_OF_HOTEL_ROOM_RATES_DETERMINANTS_OF_HOTEL_ROOM_RATES_Luxury_Industries_Conference_London_2017.

Elreedy, D. et al. (Jan. 2018). "A framework for an agent-based dynamic pricing for broadband wireless price rate plans". In: *Journal of Simulation* 13, pp. 96–110. URL: https://www.tandfonline.com/doi/full/10.1080/17477778.2017.1418642.

Goli, F. and M. Haghighinasab (2022). "Dynamic Pricing: A Bibliometric Approach". In: *Iranian Journal of Management Studies, vol. 15, pp. 111-132*. DOI: 10.22059/IJMS.2021.315212.674336.

Kalehbasti, P., L. Nikolenko, and H. Rezaei (2019). “Airbnb Price Prediction Using Machine Learning and Sentiment Analysis”. In: *CD-MAKE, International Cross-Domain Conference for Machine Learning and Knowledge Extraction*. URL: https://doi.org/10.1007/978-3-030-84060-0_11.

Kimes, S.E. (Nov. 1989). “The basics of yield management”. In: *The Cornell Hotel and Restaurant Administration Quarterly, Elsevier* 30.3, pp. 14–19. URL: https://www.sciencedirect.com/science/article/pii/S0010880489800412.

Lak, P. et al. (2015). “Towards Dynamic Pricing for Digital Billboard Advertising Network in Smart Cities”. In: *ISC2, IEEE First International Smart Cities Conference*. DOI: 10.1109/ISC2.2015.7366211.

Liu, Y. (2021). “Airbnb Pricing Based Statistical Machine Learning Models”. In: *CONF-SPML, International Conference on Signal Processing and Machine Learning*. DOI: 10.1109/CONF-SPML54095.2021.00042.

Lundberg, S.M. and S. Lee (2017). “A Unified Approach to Interpreting Model Predictions”. In: 31*st Conference on Neural Information Processing Systems (NIPS)*. URL: https://arxiv.org/abs/1705.07874.

Mariani, M. and J. Wirtz (2023). “A Critical Reflection on Analytics and Artificial Intelligence-based in Analytics in Hospitality and Tourism Management Research”. In: *International Journal of Contemporary Hospitality Management, vol. 35, pp. 2929-2943*. DOI: 10.1108/IJCHM-08-2022-1006.

Minio, V. et al. (Sept. 2023). “A framework for an agent-based dynamic pricing for broadband wireless price rate plans”. In: *Environmental Modelling & Software, Elsevier* 167. URL: https://www.sciencedirect.com/science/article/pii/S1364815223001676?via%3Dihub.

Pande, R. (2020). “Investigation of the implementation of machine learning within the hospitality industry”. In: *CAUTHE Conference*. URL: https://www.researchgate.net/publication/342976829_Investigation_of_the_implementation_of_machine_learning_within_the_hospitality_industry.

Peng, N., Y. Qin, and K. Li (2020). “Leveraging Multi-Modality Data to Airbnb Price Prediction”. In: *ICEMME, 2nd International Conference on Economic Management and Model Engineering*. DOI: 10.1109/ICEMME51517.2020.00215.

Ragapriya, N. et al. (2023). “Machine Learning Based House Price Prediction Using Modified Extreme Boosting”. In: *AHAST, Asian Journal of Applied Science and Technology*. URL: https://doi.org/10.38177/ajast.2023.7105.

Sahin, M. (2019). “Optimization of dynamic ticket pricing parameters”. In: *Journal of Revenue Pricing Management, vol. 18, pp. 306-316*. URL: https://link.springer.com/article/10.1057/s41272-018-00183-1.

Sahin, M. and M. Erol (2017). “A Dynamic Ticket Pricing Approach for Soccer Games”. In: *Axioms, vol. 6*. URL: http://dx.doi.org/10.3390/axioms6040031.

Saitta, S., V. D’Amico, and G.M. Farinella (2024). “Dynamic Price Prediction for Revenue Management System in Hospitality Sector”. In: *Proceedings of the 13th International Conference on Data Science, Technology and Applications (DATA24), pp. 218-228*. DOI: 10.5220/0012707700003756.

Shadiqurrachman, S., A. Y. Ridwan, and A. Kusuma (2019). “Online Travel Agency Channel Pricing Policy based on Dynamic Pricing Model to Maximize Sales Profit Using Nonlinear Integer Programming Approach”. In: *ICOEMIS, Proceedings of the 1st International Conference on Engineering and Management in Industrial System*. DOI: 10.2991/icoemis-19.2019.48.

Shakya, S. et al. (2012). “Neural network demand models and evolutionary optimisers for dynamic pricing”. In: *13th SGAI International Conference on Innovative Techniques and Applications of Artificial Intelligence*. URL: http://dx.doi.org/10.1007/978-0-85729-130-1_16.

Shintani, M. and K. Umeno (2022). "General dynamic pricing algorithms based on universal exponentialn booking curves". In: *JSIAM Letters, vol. 14, pp. 49-52*. DOI: 10.14495/jsiaml.14.49.

Singh, I. (2022). "Dynamic Pricing using Reinforcement Learning in Hospitality Industry". In: *2022 IEEE Bombay Section Signature Conference (IBSSC)*, pp. 1–6. DOI: 10.1109/IBSSC56953.2022.10037523.

Tuncay, G. et al. (2023). "A Reinforcement Learning Based Dynamic Room Pricing Model for Hotel Industry". In: *INFOR: Information System and Opearational Research*, pp. 211–231. URL: https://doi.org/10.1080/03155986.2023.2235223.

Vives, A. and M. Jacob (2020). "Dynamic pricing for online hotel demand: the case of resort hotels in Majorca". In: *Journal of Vacation Marketing, vol. 26, pp. 1-16*. URL: https://doi.org/10.1177/13567667198673.

— (2021). "Dynamic pricing in different Spanish resort hotels". In: *Journal of Tourism Economics, vol. 27, pp. 1-14*. URL: https://doi.org/10.1177/1354816619870652.

Vives, A., M. Jacob, and E. Aguiló (2019). "Online Hotel Demand Model and Own-Price Elasticities: An Empirical Application to two Resort Hotels in a Mature Destination". In: *Tourism Economics, vol. 24, pp. 720-752*. DOI: 10.1177/1354816618800643.

Ye, P. et al. (2018). "Customized Regression Model for Airbnb Dynamic Pricing". In: 24^{th} *ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. URL: https://doi.org/10.1007/s10732-021-09480-2.

Zhang, Q. et al. (2019). "Deep Learning based Dynamic Pricing Model for Hotel Revenue Management". In: *ICDMW, International Conference on Data Mining Workshops*. DOI: 10.1109/ICDMW.2019.00061.

Zhang, Z. et al. (2017). "Key factors Affecting the Price of Airbnb Listings: A Geographically Weighted Approach". In: *Sustainability, vol. 9*. URL: https://www.mdpi.com/2071-1050/9/9/1635.

Zhu, F. et al. (2022). "Modeling Price Elasticity for Occupancy Prediction in Hotel Dynamic Pricing". In: *Proceedings of the* 31*st ACM International Conference on Information & Knowledge Management, pp. 4742-4746*. URL: https://dl.acm.org/doi/10.1145/3583780.3614800.