Sinopfbd, 2026; 11(1): 356-377 https://doi.org/10.33484/sinopfbd.1795688

Sinop Üniversitesi Fen Bilimleri Dergisi
Sinop Uni J Nat Sci
E-ISSN: 2564-7873

https://dergipark.org.tr/tr/pub/sinopfbd

Predicting Furniture Prices Using Machine Learning Methods

Kenan KILIÇ

How to cite: Kılıç, K. (2026). Predicting furniture prices using machine learning methods. *Sinop Üniversitesi Fen Bilimleri Dergisi*, *11*(1), 356-377. https://doi.org/10.33484/sinopfbd.1795688

Open Access

Research Article

Corresponding Author
Kenan KILIÇ
kenan.kilic@bozok.edu.tr

ORCID of the Author
K.K: 0000-0003-1607-9545

Received: 02.10.2025
Accepted: 05.05.2026

Abstract

This study aims to predict furniture prices using machine learning-based regression models and to compare their performance within a unified experimental framework. Accurate price prediction plays a critical role in decision-making processes such as pricing strategies, inventory management, and market analysis in the furniture sector. A dataset containing both numerical and categorical features related to furniture products was analyzed. Various machine learning algorithms, including Decision Tree, Random Forest, Gradient Boosting, XGBoost, K-Nearest Neighbors (KNN), Lasso, and Ridge, were implemented. Model performance was evaluated using multiple metrics, including Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), Mean Absolute Percentage Error (MAPE), and the coefficient of determination (R^2). The results demonstrate that ensemble-based boosting methods significantly outperform other models. Among all models, XGBoost achieved the best overall performance, with the lowest error values (MAE = 2.46, RMSE = 3.32, MAPE = 1.08%) and the highest R^2 value (0.9993). Gradient Boosting also produced comparable results but required a longer computation time. In contrast, Random Forest, linear models (Lasso and Ridge), and KNN exhibited relatively lower performance. Overall, the findings highlight the effectiveness of machine learning approaches for furniture price prediction and indicate that XGBoost provides the most suitable balance between accuracy and computational efficiency. The proposed framework offers a reliable and data-driven solution that can support decision-making processes in the furniture industry.

Keywords: Machine learning, furniture industry, price prediction, wood industry, regression analysis

Makine Öğrenme Yöntemleri Kullanarak Mobilya Fiyatlarını Tahmin Etme

Yozgat Bozok University, Yozgat Vocational School, Design Department, Yozgat, Türkiye

Özet

Bu çalışma, makine öğrenmesi tabanlı regresyon modelleri kullanarak mobilya fiyatlarını tahmin etmeyi ve bu modellerin performanslarını bütünleşik bir deneysel çerçeve içerisinde karşılaştırmayı amaçlamaktadır. Doğru fiyat tahmini, mobilya sektöründe fiyatlandırma stratejileri, stok yönetimi ve pazar analizi gibi karar alma süreçlerinde kritik bir rol oynamaktadır. Çalışmada, mobilya ürünlerine ait hem sayısal hem de kategorik özellikleri içeren bir veri seti analiz edilmiştir. Karar Ağacı, Rastgele Orman, Gradient Boosting, XGBoost, K-En Yakın Komşular (KNN), Lasso ve Ridge gibi çeşitli makine öğrenmesi algoritmaları uygulanmıştır. Model performansı; Ortalama Mutlak Hata (MAE), Kök Ortalama Kare Hata (RMSE), Ortalama Mutlak Yüzde Hata

(MAPE) ve Belirleme Katsayısı (R^2) gibi metrikler kullanılarak değerlendirilmiştir. Elde edilen sonuçlar, topluluk tabanlı boosting yöntemlerinin diğer modellere kıyasla belirgin şekilde daha üstün performans sergilediğini göstermektedir. Tüm modeller arasında XGBoost en iyi performansı sağlamış olup, en düşük hata değerleri (MAE = 2.46, RMSE = 3.32, MAPE = %1.08) ve en yüksek R^2 değeri (0.9993) elde edilmiştir. Gradient Boosting benzer doğruluk seviyelerine ulaşmasına rağmen daha uzun hesaplama süresi gerektirmiştir. Buna karşılık, Random Forest, doğrusal modeller (Lasso ve Ridge) ve KNN daha düşük performans göstermiştir. Sonuç olarak, makine öğrenmesi yöntemlerinin mobilya fiyat tahmininde oldukça etkili olduğu görülmüş ve XGBoost modelinin doğruluk ile hesaplama verimliliği açısından en uygun çözüm olduğu belirlenmiştir. Önerilen yaklaşım, mobilya sektöründe karar destek süreçlerini destekleyebilecek güvenilir ve veri odaklı bir çözüm sunmaktadır.

This work is licensed under a Creative Commons Attribution 4.0 International License

Anahtar Kelimeler: Makine öğrenmesi, mobilya endüstrisi, fiyat tahmini, ağaçişleri endüstrisi, regresyon analizi

Introduction

The furniture sector represents a strategically important domain for many national economies, contributing significantly to industrial production, employment, and value creation. Companies operating in this sector consider various, such as profitability, market dynamics, and competitive conditions when determining product prices. From the consumer perspective, the price of a product should be aligned with the benefits it provides, and this balance plays a critical role in purchasing decisions [1]. Previous studies indicate that price is one of the most influential factors affecting consumer preferences, followed by customization and delivery-related factors [2]. Furthermore, is directly related to its capacity to respond to consumer demands and changes in individuals' lifestyles. Predicting sales volumes in this area is of critical importance in terms of understanding market trends and predicting future demand. These predictions provide a strategic basis for analysing changes in consumer trends, managing the competitive environment in the sector, effectively controlling inventory, and planning production processes efficiently. To carry out this strategic planning, sales trends must be examined using time-dependent data [3]. Forecasting is a methodology used to determine the future performance of a business or organisation [4]. This process plays a critical role in short, medium and long-term decisions. Since companies face dynamic internal and external factors, they often resort to qualitative forecasting methods [5]. However, qualitative methods are generally less effective than quantitative forecasting models due to their structural deficiencies and limited extrapolation capabilities, which can often negatively impact production planning processes. In the literature, forecasting is applied to various business processes, especially demand forecasting and sales forecasting [6-9]. A sales forecast can be defined as a forecast based on the analysis of past performance and future market conditions of a business [4]. This forecast serves as the basis for business planning [10]. Most management decisions are directly affected by current market conditions [4]. Additionally, forecast accuracy directly impacts the business's ability to respond quickly to market demand, future costs, and potential benefits [10]. Artificial intelligence is a computer science discipline that mimics human mental abilities and utilises algorithms and mathematical models to generate solutions to complex problems [11]. In recent times, particularly, rapid technological advances in this field have underscored the potential of artificial intelligence to deliver innovative and effective solutions to real-world problems [12]. The furniture industry is highly sensitive to changes in consumer preferences, lifestyle trends, and economic conditions. Therefore, accurate price prediction is essential for understanding market dynamics, optimizing pricing strategies, and improving competitiveness. Traditional statistical approaches are often limited in capturing nonlinear and complex relationships between variables such as cost, demand,

discounts, and product characteristics [5]. In addition, model interpretability has become an important aspect in machine learning applications, particularly for understanding the contribution of input features to prediction outcomes. Machine learning has emerged as a powerful approach for solving complex prediction problems by learning patterns from historical data [13-15]. Among these methods, decision tree-based approaches and ensemble learning techniques have gained significant attention due to their strong predictive performance and flexibility [16-18]. In addition, instance-based methods, such as K-Nearest Neighbors (KNN) and regularization-based linear models including Lasso and Ridge regression have been widely applied in regression tasks [19-22]. Furthermore, boosting-based algorithms, such as Gradient Boosting and XGBoost have demonstrated superior performance in modeling nonlinear relationships and high-dimensional data [23, 24]. These models iteratively improve prediction accuracy by minimizing residual errors and have been widely adopted in various forecasting applications. Recent studies have demonstrated that machine learning-based approaches significantly outperform traditional models in terms of accuracy and robustness [25-30]. In particular, ensemble-based models such as Gradient Boosting and XGBoost have shown superior performance in modeling nonlinear relationships and high-dimensional datasets, although their effectiveness may vary depending on the dataset characteristics [31-35]. Recent research has also focused on sector-specific applications. For instance, Çardak et al. [36] analyzed consumer preferences for furniture products on e-commerce platforms using machine learning techniques and demonstrated that behavioral data significantly improves predictive performance. Similarly, Akyüz et al. [37] applied machine learning-based regression models in the forest products sector, highlighting the effectiveness of data-driven approaches in economic modeling. In addition, recent studies confirm that machine learning-based approaches provide higher predictive accuracy and better generalization performance compared to classical models [38-42]. Despite these advancements, existing studies generally focus on limited datasets or specific algorithms and fail to provide a comprehensive comparison of multiple regression models under a unified experimental framework, particularly in the context of price prediction in the furniture sector. Moreover, the combined impact of preprocessing techniques, feature types, and evaluation metrics on model performance remains insufficiently explored. In this context, the present study aims to predict furniture prices using various machine learning-based regression models and to provide a comprehensive comparative evaluation of their performance. The study considers several algorithms, including Decision Tree, Random Forest, Gradient Boosting, XGBoost, K-Nearest Neighbors (KNN), Lasso, and Ridge, and evaluates their effectiveness using multiple performance metrics. Unlike many existing studies that focus on a limited number of models or specific applications, this study provides a comprehensive comparison of multiple machine learning regression algorithms within a unified experimental framework. In addition, the study incorporates both numerical and categorical features in a structured dataset, enabling a more realistic representation of real-world furniture data. The use of multiple evaluation metrics enables a robust and multi-dimensional assessment of model performance. The main contributions of this study can be summarized as follows; high-accuracy prediction of furniture prices using machine learning techniques, a comprehensive comparative analysis of widely used regression algorithms applied to price prediction tasks, a structured data preprocessing and modeling framework adaptable to different application domain, a multi-dimensional evaluation using MAE, MSE, RMSE, MAPE, and R^2 metrics, transparent reporting of hyperparameter settings to ensure reproducibility. The remainder of this paper is organized as follows. Section 2 presents the materials and methods, Section 3 provides the results and discussion, and Section 4 concludes the study with key findings and future research directions.

Material and Method

This section presents the dataset, data preprocessing procedures, and machine learning methods employed in the study. The dataset is first introduced, followed by a detailed description of the data preparation process and the modeling stages.

Data Set

This dataset was created to analyse furniture sales data using multiple regression techniques. It contains 2.500 rows with 15 columns, 7 numerical columns and 7 categorical columns, and a target variable (price) representing the total revenue from furniture sales. The dataset captures various aspects of furniture sales, such as pricing, cost, sales volume, discount percentage, inventory levels, delivery time, and different categorical attributes such as furniture type, material, colour, and store location [31]. In this study, the price variable represents the selling price of furniture products as provided in the dataset. It is not derived from a specific formula; rather, it reflects market-based values influenced by multiple factors such as cost, demand, discount rate, and product characteristics. It should be emphasized that the price variable does not contain negative values. Negative values observed in the dataset are related to revenue or profit-based variables and represent financial losses under certain conditions. It should be noted that some variables in the dataset, such as revenue or profit-related features, may contain negative values. These values do not indicate data errors but rather represent financial losses under certain conditions, such as high costs, discounts, or low sales volume. Importantly, the target variable (price) does not contain negative values. Therefore, the presence of negative values in the dataset is consistent with real-world economic scenarios and does not affect data validity. The characteristics of the first 5 products in the dataset are given in Table 1. The statistical properties and distributions of the numerical variables in the dataset are given in Table 2.

Table 1. *Features of the first 5 products in the dataset*

Features	Product 1	Product 2	Product 3	Product 4	Product 5
Price	218.54	477.82	379.4	319.4	120.21
Cost	181.61	385.03	276.74	281.84	69.74
Sales	40	7	32	48	19
Profit Margin (%)	16.9	19.42	27.06	11.76	41.98
Inventory	105	192	59	45	35
Discount (%)	27.8	26.94	21.95	11.01	3.18
Delivery Days	9	6	2	2	9
Category	Bed	Chair	Table	Table	Chair
Material	Plastic	Glass	Metal	Glass	Glass
Color	Red	Blue	Black	Green	Brown
Location	Rural	Rural	Suburban	Rural	Rural
Season	Spring	Summer	Fall	Summer	Fall
Store Type	Online	Online	Online	Retail	Online
Brand	BrandA	BrandD	BrandD	BrandD	BrandD
Revenue	3949.17	-3521.0	14285.56	12261.07	-4588.26

Table 1 presents the characteristics of the first five products in the dataset across multiple numerical and categorical features. The features include price, cost, sales volume, profit margin, inventory status, discount rates, delivery times, categories, material types, colours, locations, seasons, store types, brands, and revenues. When products are categorised differently (e.g., bed, chair, table), significant differences are observed between prices and costs. While profit margin and revenues reflect the financial performance of the products, features such as discount rates and delivery times may also influence the customer experience.

Table 2. *Statistical properties and distributions of numerical variables in the data set*

Statistic	Price	Cost	Sales	Profit margin	Inventory	Discount percentage	Delivery days	Revenue
count	2500	2500	2500	2500	2500	2500	2500	2500
mean	274.50	191.93	24.92	30.21	97.72	14.95	4.89	5926.85
std	130.90	98.59	14.05	11.54	57.62	8.62	2.56	6902.74
min	50.70	26.51	1.00	10.02	0.00	0.01	1.00	-14214.57
25%	159.10	106.40	13.00	20.23	50.00	7.76	3.00	1216.72
50%	277.64	189.34	25.00	30.30	94.00	14.92	5.00	5523.23
75%	387.38	263.19	37.00	40.13	147.00	22.29	7.00	10233.54
max	499.87	447.02	49.00	50.00	199.00	29.99	9.00	32922.08

Upon examining the statistical properties of the numerical variables in the dataset, it is observed that there are 2.500 observations for each variable. The average value of the Price variable is approximately 274.50 USD, while the average value of the Cost variable is 191.93 USD. The average value of the Sales variable is 24.92, and the average Profit Margin is 30.21%. These averages provide information about the general price and cost structure of the products. The standard deviation (Std) values, which indicate the spread of the data, show that the prices have a spread of 130.90 USD, while the costs have a spread of 98.59 USD. The lowest price in the dataset was recorded as 50.70 USD, and the highest price was 499.87 USD. Similarly, the lowest value in the Cost variable was 26.51 USD, and the highest value was 447.02 USD. The minimum value of the Sales variable was 1, while the maximum value was 49. When the quartile values are examined, the 25th percentile quartile of the Price variable is calculated as 159.10 USD, the 50th percentile (median) is 277.64 USD, and the 75th percentile quartile is 387.38 USD. Similarly, quartile values are calculated for other variables, and these values help us understand the data distribution in more detail. For example, for Profit Margin, the 25th percentile value is 20.23%, the median value is 30.30%, and the 75th percentile value is 40.13%. The presented statistical analyses enable us to understand the distribution of price, cost, sales, profit margin, and other variables of the products in the dataset. This information is beneficial for evaluating the economic performance of the products and making strategic decisions. In the study, a correlation matrix was created to examine the linear relationships between the variables. This matrix provided preliminary information about the direction and strength of the relationships between the variables and was also used to assess the risk of multicollinearity. Thus, before proceeding with the modelling process, the relationships between the variables were identified and their suitability for analysis was evaluated. The correlation matrix is presented in Figure 1.

Data Loading, Preparation and Preprocessing

The study was conducted in a cloud-based Python 3 environment provided by the Kaggle platform for data analysis and machine learning applications. The computational environment was based on the Kaggle/Python Docker image, which includes a wide range of pre-installed libraries for data processing, visualization, and modeling. This setup provides a suitable infrastructure for achieving high computational efficiency and reproducibility throughout the analysis process. During the analysis, various Python libraries serving different purposes were utilized. NumPy was used for numerical computations and matrix operations, Pandas for data manipulation and DataFrame-based analysis, Matplotlib and Seaborn for data visualization and correlation analysis, and Scikit-learn for data preprocessing, model training, hyperparameter optimization, and performance evaluation. In addition, the SciPy library was employed for statistical analysis and distribution-based operations. Data preprocessing is one of the most critical stages in machine learning applications, as it directly affects model performance. The dataset used in this study consists of both numerical and categorical variables,

which were processed using appropriate techniques to prepare them for the modeling phase. The preprocessing stage plays a crucial role in improving data quality, enhancing the generalization capability of the model, and ensuring a more stable learning process. In this study, the target variable, price, was directly obtained from the dataset and was not derived from any independent variable. Although variables such as cost, profit margin, and discount are related to price, there is no deterministic (directly computable) relationship among them. Instead, the price variable represents a complex structure resulting from the interaction of multiple factors [14, 15]. This characteristic transforms the prediction task from a simple linear modeling problem into a more complex and multidimensional learning problem. Consequently, the machine learning methods employed in this study were selected to capture complex relationships, variable interactions, and latent patterns within the dataset, thereby enabling a more effective representation of the underlying data structure.

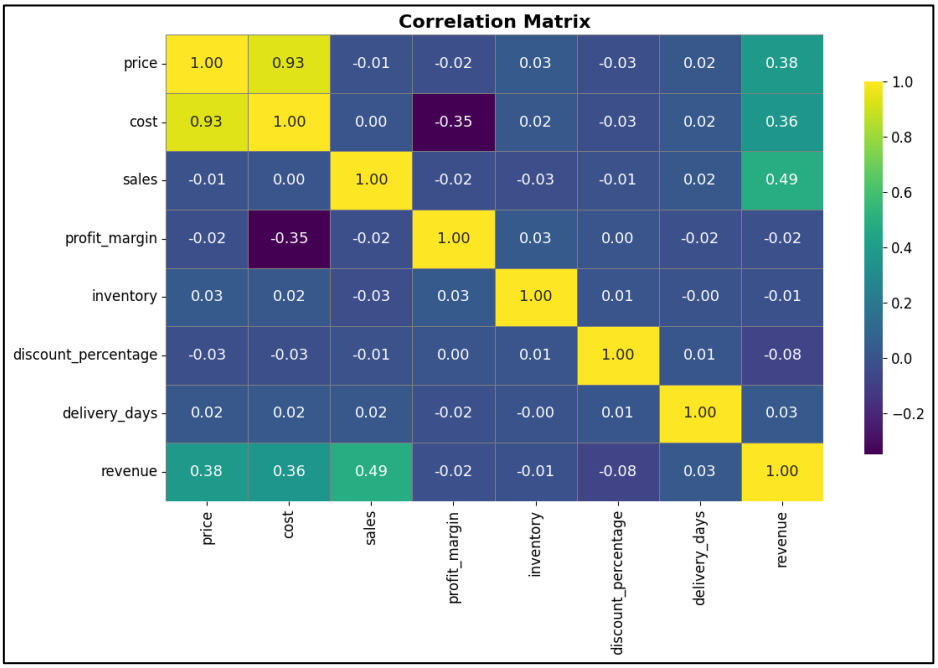


Figure 1. *Correlation Matrix*

Handling of Missing Values and Outliers

Prior to the modeling stage, the dataset was examined for missing values and outliers to ensure data quality and reliability. Missing value analysis was conducted using Pandas-based exploratory data analysis methods (e.g., isnull().sum()). The results indicated that there were no missing values in the dataset. Therefore, no imputation or data removal procedure was required. In addition, outlier detection was performed using the Z-score method for numerical variables. Z-scores were calculated for each feature, and observations with absolute Z-score values greater than 3 were evaluated as potential outliers, as commonly adopted in statistical analysis [14]. The analysis showed that extreme values in the dataset were limited and did not significantly affect the overall data distribution. Therefore, the dataset was used without applying an extensive data removal procedure, preserving its original structure.

Processing Numeric Variables

The dataset includes various numerical variables such as price, cost, sales volume, profit margin, inventory level, discount percentage, delivery days, and revenue. Among these variables, price was defined as the dependent (target) variable, whereas the remaining numerical variables were used as numerical input features for the model. Prior to the modeling stage, correlation analysis was performed

to better understand the relationships among numerical variables. In this context, a correlation matrix was constructed to reveal the direction and strength of linear relationships between variables, thereby contributing to the assessment of multicollinearity risk. To prevent variables with different scales from exerting disproportionate influence on the model, the numerical data were standardized using the StandardScaler method. This transformation ensures that each variable has a mean of zero and a standard deviation of one, allowing all features to contribute to the model on a comparable scale. The standardization process is defined as follows:

$$z = \frac{x-\mu}{\sigma} \quad (1)$$

Where x represents the original value, μ denotes the mean, and σcorresponds to the standard deviation. The standardization procedure was implemented within a Pipeline structure, where preprocessing steps are integrated into the modeling process. In this way, transformation parameters were learned exclusively from the training data, thereby preventing data leakage during cross-validation. This approach enhances the reliability and generalizability of the obtained performance results. The primary reason for selecting the StandardScaler method is its ability to produce more stable results across variables with different distributions [14] and its lower sensitivity to outliers compared to Min-Max scaling. This characteristic is particularly advantageous in datasets that include both economic and operational variables, as it contributes to improved model performance.

Coding of Categorical Variables

The categorical variables in the dataset include qualitative attributes such as product category, material type, color, location, season, store type, and brand. Since these variables are not inherently numerical, they must be transformed into a suitable format for machine learning algorithms. To this end, categorical variables were converted into numerical representations using the OneHotEncoder method. In this approach, each category is represented as a separate binary (0–1) variable, thereby preventing the introduction of any artificial ordinal relationship [14] among categories. This procedure is particularly important to avoid misleading effects of categorical variables on the model. During the encoding process, one reference category for each categorical variable was excluded from the dataset to reduce the risk of multicollinearity. This approach contributes to improved model stability and facilitates more reliable interpretation of the results. The transformation of categorical variables was performed alongside numerical preprocessing using ColumnTransformer and Pipeline structures. This integrated framework ensures that all preprocessing steps are embedded within the modeling process, thereby preventing data leakage during cross-validation. As a result, model performance is evaluated in a more reliable and generalizable manner.

Data Partitioning and Preparation

The dataset was divided into input variables (X) and the target variable (y) for the modeling process. While the target variable represents the product price, all other variables were treated as features used in the prediction process. Instead of creating a fixed train–test split, a cross-validation approach was adopted to ensure a more reliable and generalizable evaluation of model performance. In this context, 5-fold cross-validation was employed during the hyperparameter optimization stage. After determining the best-performing model, 10-fold cross-validation was applied to evaluate its overall performance. In this process, the dataset was partitioned into multiple subsets, and the model was tested on different portions of the data, with the final performance obtained by averaging the results across all folds. Data preprocessing steps were implemented within a Pipeline structure to ensure seamless integration with the modeling process. This approach ensures that all transformations are learned only from the corresponding training subsets, thereby preventing data leakage during cross-validation.

Prediction Algorithms

Decision Tree

It is a common machine learning algorithm used to classify or predict data. Decision trees recursively partition the data space based on feature values. Decision trees make the best splits using criteria such as Gini Index or Information Gain [16, 17]. However, it may need to be simplified using techniques such as pruning, as it carries the risk of overfitting.

Random Forest

It is an ensemble method consisting of multiple decision trees, each trained independently on the data. This method aims to obtain more accurate results by giving weighted votes to the prediction of each decision tree. Random Forest uses different data for each tree by randomly selecting subsets from the dataset, which increases the model's generalizability and reduces the risk of overfitting [18]. Additionally, this method offers high accuracy in both classification and regression problems.

K-Nearest Neighbours (KNN)

It is one of the supervised learning algorithms and is generally used in classification and regression problems. The KNN algorithm classifies or predicts new data based on the class or average value of the K nearest neighbours to make a prediction. This method evaluates each example in the dataset according to its distance to its neighbours; the Euclidean distance is generally used [19]. KNN makes predictions directly using the data without requiring any parameters and is easy to understand due to its simple structure. In KNN regression problems, the predicted value is calculated as the average of the nearest neighbours. In classification problems, the predicted class is selected from the most frequent neighbour class [20]. However, the performance of KNN depends on the selection of the correct distance measure and the number of neighbours, so it is important to choose the hyperparameters carefully.

Lasso (Least Absolute Shrinkage And Selection Operator)

It is a method used in linear regression models, which helps both to reduce the complexity of the model and to select important features. This technique adds a penalty to the sum of the absolute values of the regression coefficients using L1 regularisation. This added penalty causes some coefficients to be set to zero, ensuring that the model relies only on significant variables. This helps overfitting and increases prediction accuracy [21].

Ridge Regression

This is a regularisation technique used in linear regression models. This method adds a penalty to the sum of the squares of the regression coefficients to reduce the model's complexity. This added penalty helps prevent the model from overfitting. Ridge regression is particularly effective in dealing with situations such as multicollinearity (high correlation between independent variables) and provides more stable predictions in such cases [22]. Ridge regression focuses on making all coefficients small and does not set any coefficient to zero. This feature is especially preferred in cases where all variables are significant. However, it may be less effective than Lasso regression in terms of model interpretability. Ridge regression is an effective technique used to make more general and robust predictions by mitigating the complexity of the linear regression model [21].

Gradient Boosting

Gradient Boosting is an algorithm that aims to create a strong model by combining multiple weak learners. The advantages of this method include high accuracy and efficiency. When hyperparameters

are adjusted appropriately, very effective results can be obtained. However, it also has disadvantages such as long training time and the risk of overfitting [23].

XGBoost (Extreme Gradient Boosting)

It is a model that offers a faster and more effective classification method by optimising the Gradient Boosting algorithm. It provides high performance, particularly in large datasets and complex problems. Fast training time and low memory usage are important advantages of this model. However, the model also has disadvantages, such as the risk of overfitting and the complexity of parameter settings [24].

Hyperparameter Optimization

Hyperparameter optimization was performed to enhance model performance and to identify the most suitable parameter combinations for each algorithm. In this process, the GridSearchCV method was employed, which enables the systematic evaluation of different hyperparameter configurations. Within the GridSearchCV framework, a comprehensive set of predefined hyperparameter combinations was systematically evaluated. To ensure a reliable and robust performance estimation [14], a 5-fold cross-validation strategy was applied, where each hyperparameter configuration was tested on different subsets of the data. The R^2 score was used as the primary evaluation metric, and the parameter set yielding the highest performance was selected as the optimal solution. The hyperparameter search ranges defined for the machine learning algorithms used in this study are presented in Table 3. This search space includes the key parameters that directly influence the learning capacity and generalization performance of each model.

Table 3. *Hyperparameter search space used in GridSearchCV*

Model	Hyperparameter	Values
Decision Tree	max_depth	[None, 5, 10, 20]
	min_samples_split	[2, 5, 10]
	min_samples_leaf	[1, 2, 4]
Random Forest	n_estimators	[100, 200]
	max_depth	[None, 10, 20]
	min_samples_split	[2, 5]
	min_samples_leaf	[1, 2]
	max_features	[sqrt, log2]
Gradient Boosting	n_estimators	[100, 200]
	learning_rate	[0.05, 0.1, 0.2]
	max_depth	[3, 5]
	min_samples_split	[2, 5]
	min_samples_leaf	[1, 2]
KNN	n_neighbors	[3, 5, 7, 9]
	weights	[uniform, distance]
	p	[1, 2]
Lasso	alpha	[0.001, 0.01, 0.1, 1, 10]
Ridge	alpha	[0.01, 0.1, 1, 10, 100]
XGBoost	n_estimators	[100, 200]
	learning_rate	[0.05, 0.1]
	max_depth	[3, 5]
	subsample	[0.8, 1.0]
	colsample_bytree	[0.8, 1.0]

The optimization process conducted based on the search space defined in Table 3 resulted in the identification of the optimal hyperparameter combinations for each model. The obtained optimal parameter values are presented in Table 4.

Table 4. *Optimal Hyperparameters obtained via GridSearchCV*

Model	Hyperparameter	Optimal Value
Decision Tree	max_depth	None
	min_samples_split	2
	min_samples_leaf	4
Random Forest	n_estimators	200
	max_depth	None
	min_samples_split	2
	min_samples_leaf	1
	max_features	sqrt
Gradient Boosting	n_estimators	200
	learning_rate	0.05
	max_depth	5
	min_samples_split	2
	min_samples_leaf	1
KNN	n_neighbors	9
	weights	distance
	p	2
Lasso	alpha	0.1
Ridge	alpha	1
XGBoost	n_estimators	200
	learning_rate	0.1
	max_depth	5
	subsample	0.8
	colsample_bytree	0.8

The best model (best estimator) obtained through the GridSearchCV process was refitted on the entire dataset. To more reliably evaluate its generalization performance, a 10-fold cross-validation strategy was applied. In this framework, the dataset was divided into ten equal subsets, each of which was used as a test set in turn, and the model predictions were generated using the cross_val_predict method. The resulting predictions were aggregated across all folds, and the model performance was calculated based on these combined predictions. This approach provides a more robust, reliable, and unbiased evaluation of model performance.

Evaluation Metrics

To evaluate model performance, various statistical metrics are commonly used in regression analysis. The coefficient of determination (R^2) is widely accepted as a key indicator of model explanatory power [25]. In addition, error-based metrics such as Mean Absolute Error (MAE), Root Mean Square Error (RMSE), and Mean Absolute Percentage Error (MAPE) are frequently used to assess prediction accuracy [26-28]. These metrics are particularly important in practical applications such as sales forecasting and demand estimation, where accurate predictions directly influence business decisions [29,30]. R-squared (R^2) is a statistical measure that indicates how well the model explains the dependent variable in relation to the independent variables. A value closer to 1 indicates that the model predictions

are more closely aligned with the actual values [25]. The R^2 formula, as shown in Equation 2, is used to measure the model's success.

$$R^2 = 1 - \frac{\sum_{i=1}^{m}(Y_i - \bar{Y})^2}{\sum_{i=1}^{m}(X_i - Y_i)^2} \quad (2)$$

Here, *m* represents the number of observations, X_i represents the value of the independent variable at the ith observation, Y_i represents the value of the true dependent variable at the ith observation, and $\overline{Y}$ represents the mean of the dependent variable. A high R^2 value indicates that the model's predictions are accurate. Mean Absolute Error (MAE) calculates the average absolute difference between the actual values and the values predicted by the model. MAE is used as a metric to measure the overall accuracy of the model, and ideally, its value should be close to zero [26]. As shown in the formula given in Equation 3:

$$\text{MAE} = \frac{1}{m}\sum_{i=1}^{m}|X_i - Y_i| \quad (3)$$

Here, Xi represents the actual values and Y_i represents the predicted values. MAE is beneficial for situations where the magnitude of the model error is critical. Root Mean Square Error (RMSE) takes the square root of the average of the squared differences between the true values and the predicted values. RMSE is often used in cases where large errors are important. Ideally, the closer the value of RMSE is to zero, the more accurate the model's predictions are [27]. The formula is:

$$\text{RMSE} = \sqrt{\frac{1}{m}\sum_{i=1}^{m}(X_i - Y_i)^2} \quad (4)$$

Here, X_i are the actual values and Y_i are the predicted values. RMSE is a critical evaluation tool for situations where large errors have a significant impact on the model's performance.

Mean Absolute Percentage Error (MAPE) calculates the average of the absolute percentage differences between the true values and the predicted values. This metric is handy when the percentage error is substantial. Although it is not possible for MAPE to be zero, a low MAPE value indicates good model performance [28]. According to the formula presented in Equation 5:

$$\text{MAPE} = \frac{1}{m}\sum_{i=1}^{m}\left|\frac{Y_i - X_i}{Y_i}\right| \quad (5)$$

MAPE is used to evaluate the model's performance for data at different scales and is often preferred in business applications, such as sales forecasting.

These metrics are used to evaluate the model's accuracy and its success in predicting future sales [30]. After the model is trained, predictions are made on the test data, and these predictions are compared with the actual values. A comprehensive analysis is conducted to evaluate the accuracy of the model, using performance metrics such as R-squared, MAE, RMSE, and MAPE. This analysis enables us to determine whether the model can accurately predict future sales.

Experimental Setup

The experimental process was conducted in a cloud-based Python 3 environment provided by the Kaggle platform for data analysis and machine learning applications. The computing environment is based on the Kaggle/Python Docker infrastructure, which includes a wide range of pre-installed libraries for data processing, modeling, and visualization. This setup ensures reproducibility and consistency of the computational processes. The experiments were carried out on a Linux-based system using Python 3.10. The computational infrastructure consisted of a processor with 2 physical cores (4 logical cores) and approximately 31 GB of RAM. This hardware configuration enabled the efficient execution of computationally intensive tasks such as data preprocessing, model training, and cross-validation. During

the modeling process, Pipeline and ColumnTransformer structures were employed to integrate data preprocessing and learning steps within a unified framework. Numerical variables were standardized using the StandardScaler method, while categorical variables were transformed into numerical representations using the OneHotEncoder method. This approach ensured that all data transformations were embedded within the modeling pipeline, thereby preventing data leakage during cross-validation. In this study, XGBoost was used as the primary model for interpretability analysis and solving the regression problem. In addition, several machine learning algorithms were evaluated under the same data preprocessing framework, and their performances were compared. All models were trained under identical preprocessing steps and experimental conditions to ensure a fair comparison. To improve model performance, hyperparameter optimization was performed using the GridSearchCV method. A 5-fold cross-validation strategy was applied during the hyperparameter search process, and the parameter combination yielding the highest performance was selected. The best model obtained from this optimization process was trained using the entire dataset. To evaluate the generalization performance of the model, a 10-fold cross-validation approach was employed. In this process, the model was repeatedly tested on different subsets of the dataset, and predictions were generated using the cross_val_predict method. The obtained predictions were aggregated, and performance metrics were calculated based on these results. This approach provides a more robust and unbiased evaluation of model performance. Model performance was evaluated using multiple metrics, including Mean Absolute Error (MAE), Mean Squared Error (MSE), Root Mean Squared Error (RMSE), Coefficient of Determination (R^2), and Mean Absolute Percentage Error (MAPE). These metrics enable a comprehensive assessment of both prediction accuracy and explanatory power. Furthermore, the SHAP (SHapley Additive exPlanations) [43] method was employed to interpret the model's decision-making process. This analysis examines the contribution of each feature to the model predictions and identifies the most influential variables. As a result, the interpretability of the model is enhanced, allowing for a more meaningful evaluation of the results. Within this framework, the overall workflow, including data preprocessing, modeling, hyperparameter optimization, and performance evaluation, is presented in Figure 2. Figure 2 illustrates the overall workflow of the proposed machine learning approach, including data preprocessing, modeling, hyperparameter optimization, and cross-validation processes. This framework ensures the systematic and reproducible development of the model and enables a reliable evaluation of model performance.

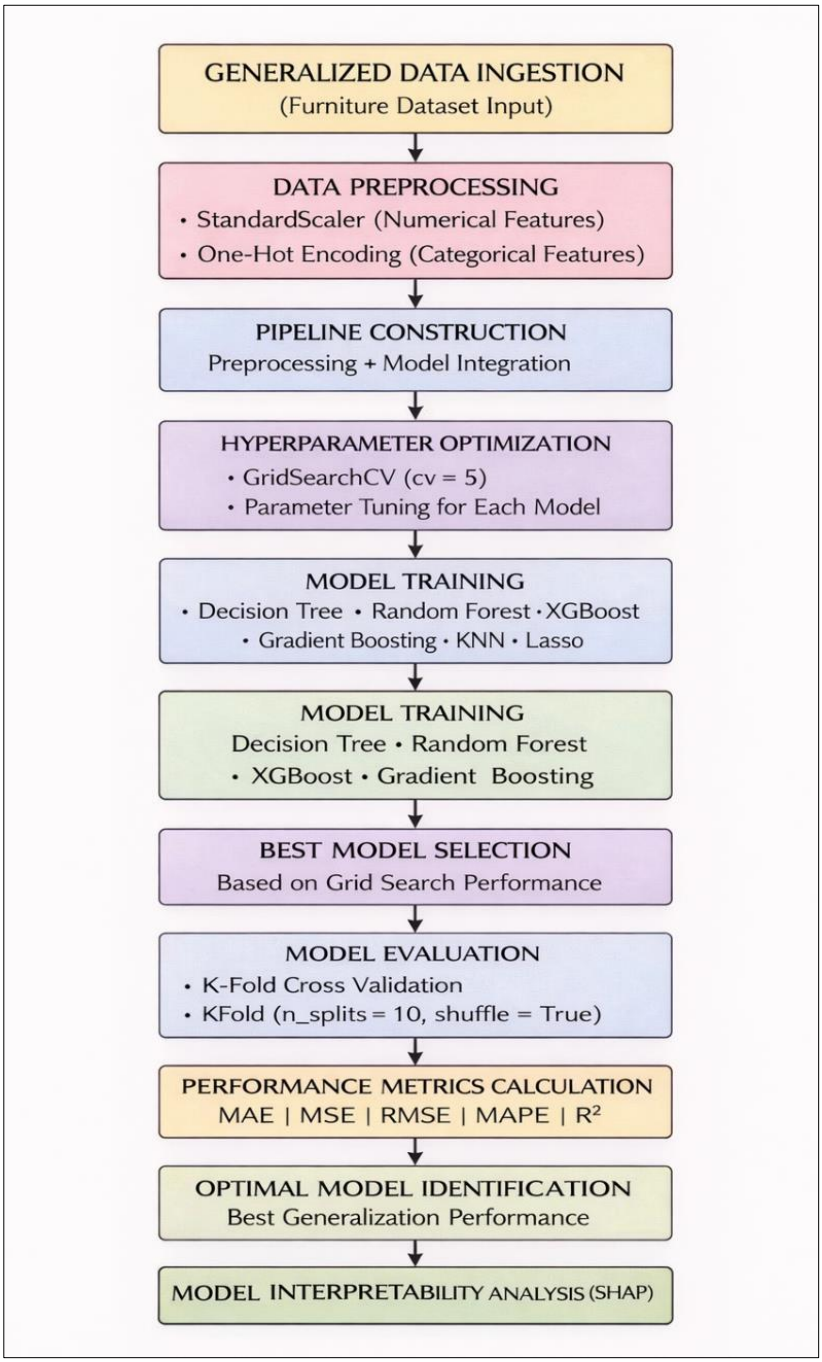


Figure 2. *Workflow of the proposed machine learning framework for furniture price prediction*

Result and Discussion

Recent studies have increasingly focused on the use of machine learning and hybrid approaches for forecasting and prediction tasks. Public datasets and real-world applications have been widely utilized to evaluate model performance [31]. Ensemble and hybrid models have demonstrated improved accuracy compared to traditional approaches in forecasting competitions and real-world scenarios [32-35]. In addition, deep learning and hybrid machine learning approaches have shown promising results in capturing complex patterns within data [38-41]. Recent research also highlights the importance of feature engineering and data-driven modeling strategies in improving prediction performance [42]. In the context of the furniture and forest products industry, machine learning-based approaches have been successfully applied to predict consumer preferences and economic indicators [36,37]. Furthermore,

model interpretability techniques such as SHAP have been introduced to better understand the contribution of input features to model predictions [43]. The analyses for furniture price prediction were performed using the Python programming language within an appropriate computational environment. A wide range of variables, including product price, cost, sales volume, profit margin, inventory level, discount rate, delivery time, category, material type, color, location, season, store type, and brand, were taken into account during the prediction process. Accurate price prediction plays a critical role in the furniture industry by supporting pricing strategies, optimizing inventory management, and enhancing competitive advantage. Accordingly, the performance of the developed models was comprehensively evaluated using multiple performance metrics. The comparative results in terms of regression performance and computational time are presented in Table 5.

Table 5. *Comparative regression performance of machine learning models*

Model	MAE	MSE	RMSE	MAPE	R^2 Score	Time
Decision Tree	5.96	60.81	7.79	2.53	0.9964	5.25 s.
Random Forest	27.19	1157.51	34.02	14.83	0.9323	117.50 s.
KNN	51.27	3921.29	62.62	29.49	0.7709	1595.79 s.
Lasso	16.35	472.80	21.74	9.92	0.9723	3.51 s.
Ridge	16.38	473.48	21.75	9.95	0.9723	2.00 s.
Gradient Boosting	2.56	11.69	3.42	1.09	0.9993	115.65 s
XGBoost	2.46	11.02	3.32	1.08	0.9993	15.80 s

An examination of Table 5 reveals that ensemble-based boosting methods significantly outperform the other models. The extremely high R^2 values (≈ 0.999) obtained in this study should be interpreted with caution. While these values indicate excellent predictive performance, they may also be influenced by strong relationships between certain input variables and the target variable. In particular, the SHAP analysis revealed that the cost variable had a dominant effect on the model predictions. This suggests that the model captures strong underlying relationships within the dataset, which may contribute to the very high accuracy. In addition, the use of 10-fold cross-validation ensures that the model performance is evaluated across different subsets of the dataset. The consistent results obtained across folds indicate that the model is stable and not dependent on a specific data partition. However, despite the robustness of the cross-validation results, the presence of strong relationships between variables may limit the generalization capability of the model. Therefore, future studies should validate the model using independent and more diverse datasets. Among all models, XGBoost achieved the highest performance, exhibiting the lowest error values (MAE = 2.46, RMSE = 3.32, MAPE = 1.08%) and the highest coefficient of determination ($R^2 = 0.9993$). Similarly, the Gradient Boosting model produced highly comparable results (MAE = 2.56, RMSE = 3.42, MAPE = 1.09%, $R^2 = 0.9993$). These findings indicate that boosting-based methods are highly effective in capturing complex and nonlinear relationships within the dataset. An important observation is that the XGBoost model also provides a significant advantage in terms of computational efficiency. While the Gradient Boosting model required 115.65 seconds, XGBoost achieved better results in only 15.80 seconds. This makes XGBoost the most suitable model in terms of both accuracy and efficiency. The Decision Tree model also demonstrated strong performance, with relatively low error values (MAE = 5.96, RMSE = 7.79) and a high R^2 value (0.9964). These results indicate that even a single tree structure can explain a large portion of the data variance. However, its performance remains lower than that of boosting methods, particularly in modeling complex relationships. In contrast, the Random Forest model exhibited lower-than-expected performance (MAE = 27.19, RMSE = 34.02, $R^2 = 0.9323$). Despite being an ensemble method, this result suggests that boosting-based approaches are better suited to the structure of the dataset.

Additionally, its relatively high computational time (117.50 seconds) reduces its efficiency. The Lasso and Ridge regression models produced very similar results. Both models showed moderate error levels (RMSE ≈ 21.75, MAPE ≈ 10%) and an R^2 value of 0.9723. This indicates that linear models can capture general trends but are insufficient for modeling complex and nonlinear relationships in the dataset. However, their main advantage lies in their low computational time (2–3.5 seconds). The KNN algorithm demonstrated the weakest performance among all models. The high error values (MAE = 51.27, RMSE = 62.62, MAPE = 29.49%) and relatively low R^2 value (0.7709) clearly indicate that the model is inadequate for this problem. Furthermore, the extremely high computational time (1595.79 seconds) makes KNN unsuitable in terms of both accuracy and efficiency.

Overall, the results can be summarized as follows:

Boosting-based models (XGBoost and Gradient Boosting) achieved the highest accuracy

The Decision Tree model demonstrated strong performance despite its simplicity

Lasso and Ridge models were limited due to their linear nature

KNN was not suitable for this problem.

These findings are consistent with the existing literature. In particular, boosting algorithms are known to achieve higher performance in complex data structures due to their ability to iteratively minimize residual errors. From a practical perspective, XGBoost emerges as the most suitable solution. Due to its high accuracy and low computational cost, it can be effectively applied in the furniture industry for price prediction, inventory management, and strategic decision-support systems. The relationship between the predicted values and the actual values is presented in Figure 3. To better understand the decision-making mechanism of the best-performing model, feature importance and SHAP (SHapley Additive exPlanations) analyses were conducted using the XGBoost model. This approach enhances the interpretability of the results by revealing the contribution of each feature to the model predictions. The results of the feature importance analysis are presented in Figure 4. As shown in Figure 4, the variable with the greatest influence on the model predictions is cost, followed by profit margin. The contributions of the other variables remain relatively low, indicating that the model predominantly exhibits a cost-driven structure. To further investigate the contribution of variables to the model in more detail, SHAP analysis was conducted, and the results are presented in Figure 5. As illustrated in Figure 5, higher cost values strongly increase the predicted price, whereas lower cost values have a negative effect. Similarly, the profit margin variable also exerts a positive influence on the model output. In contrast, variables such as sales volume, inventory, revenue, discount percentage, and categorical features (category, location, season, material, brand, color, and store type) have a limited impact on the model predictions. The near-zero SHAP values associated with these variables indicate that they are not significant contributors to the model. Overall, the results indicate that the XGBoost model primarily bases its predictions on cost and profit margin variables, while the remaining features play a secondary role. However, the strong relationship between the cost variable and the target variable (price) may have contributed to the high model performance. Therefore, it is recommended that future studies evaluate the generalizability of the model using different datasets. The findings of this study clearly demonstrate that machine learning-based regression models can be used with high accuracy for furniture price prediction. In particular, the superior performance of boosting-based models (XGBoost and Gradient Boosting) compared to other models indicates their strong capability to capture complex and nonlinear relationships. This finding is consistent with previous studies reporting that boosting algorithms provide strong generalization performance on high-dimensional and heterogeneous datasets [24, 34, 35].

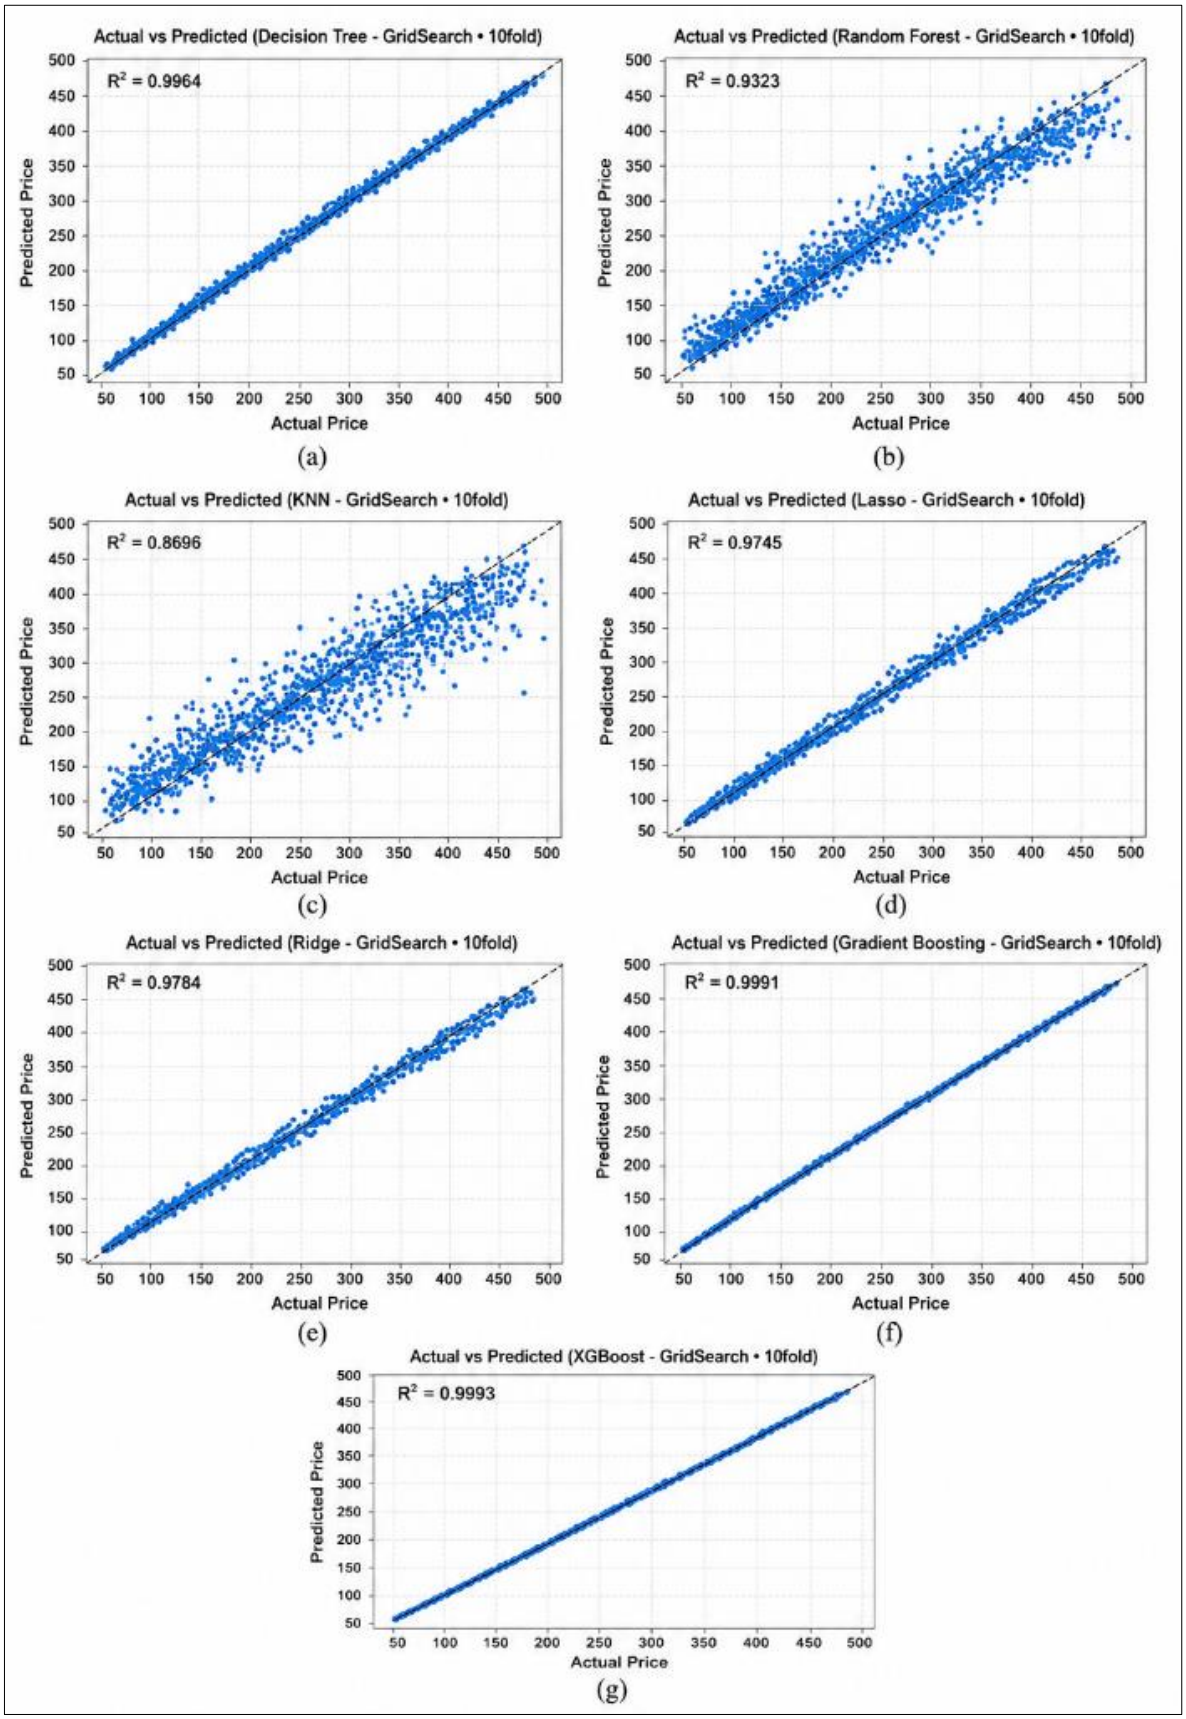


Figure 3. *Comparison of actual and predicted values of prediction models: (a) Decision Tree, (b) Random Forest, (c) KNN, (d) Lasso Regression, (e) Ridge Regression, (f) Gradient Boosting, (g) XGBoost*

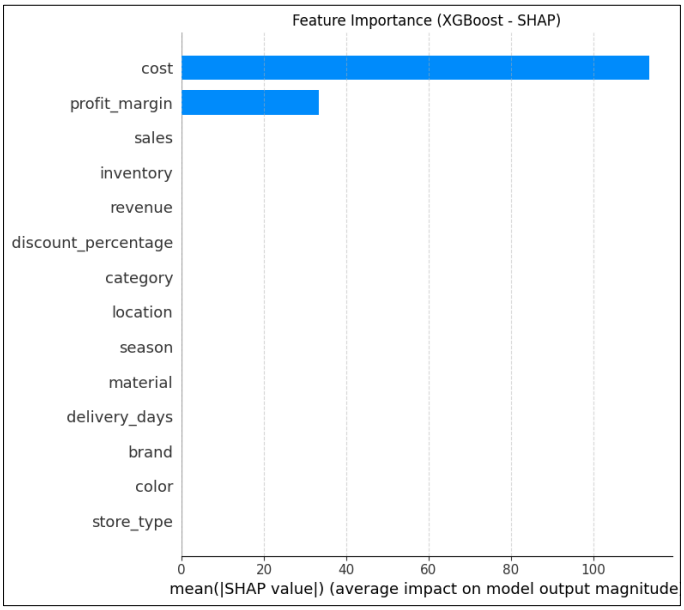


Figure 4. *SHAP-based feature importance for the XGBoost model*

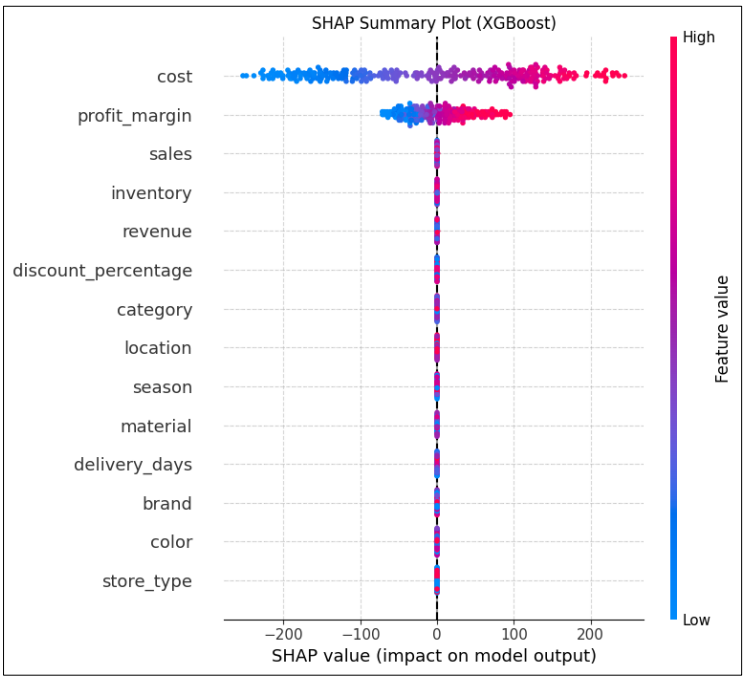


Figure 5. *SHAP summary plot for the XGBoost model*

The fact that the XGBoost model achieved the best performance with both low error values and short computational time indicates that it is advantageous not only in terms of accuracy but also in terms of computational efficiency. This supports its suitability for large-scale datasets and real-time applications. The literature also emphasizes that XGBoost achieves high accuracy and effectively controls overfitting due to its scalable structure and regularization mechanisms [24]. Although the Gradient Boosting model achieved accuracy levels similar to XGBoost, its longer training time indicates that model selection should consider not only accuracy but also computational cost. This perspective is also highlighted in the literature, where a multidimensional evaluation of model performance is recommended [32, 33]. The relatively lower performance of the Random Forest model is a noteworthy finding. Although Random Forest is generally considered a strong and stable model [18], it may exhibit lower performance compared to boosting-based methods in certain datasets. This suggests that boosting algorithms can

more effectively model complex relationships among variables. The comparatively lower performance of linear models such as Lasso and Ridge indicates that nonlinear relationships dominate the dataset. This finding is consistent with the machine learning literature, which states that linear models may be insufficient for capturing complex data structures [15, 14]. The KNN model exhibited the weakest performance, with high error values and low R^2, indicating that it is not suitable for this problem. Furthermore, its extremely high computational cost suggests that distance-based methods may suffer from the "curse of dimensionality" in high-dimensional datasets [20]. SHAP analysis further revealed that the model's decision-making process is largely driven by cost and profit margin variables. This finding aligns with studies emphasizing the importance of economic variables in price prediction [37, 36]. However, the strong dependence of the model on certain variables may also indicate that it has learned dataset-specific patterns. The extremely high R^2 values (≈ 0.999) indicate that the model explains the dataset very well. However, this may also suggest the presence of strong correlations among variables. As reported in the literature, very high accuracy values may sometimes arise from strong inter-variable relationships or dataset dependency [25]. Therefore, although the results are strong, they should be interpreted with caution. Future studies should evaluate the model on larger and more diverse datasets to better assess its generalization capability. Additionally, the use of deep learning and hybrid modeling approaches may further improve performance [40, 41]. The high R^2 values obtained in this study confirm the strong predictive capability of the models. However, such high accuracy may also result from strong relationships between input variables and the target variable. Although the target variable is not directly derived from these features, their interactions may have contributed to the model performance. Therefore, future studies should assess model generalizability using different datasets and feature selection techniques. From an application perspective, the findings of this study provide important implications for the furniture industry. Accurate price prediction can support decision-making processes such as pricing strategies, inventory management, and competitive positioning. In this context, the proposed approach offers an effective data-driven solution that can enhance operational efficiency and competitiveness.

Conclusion

The performances of various regression models for predicting furniture prices were comparatively evaluated in this study. The results clearly demonstrate that there are significant differences among the models in terms of prediction accuracy, error rates, and computational efficiency. Among all models, XGBoost achieved the best overall performance, exhibiting the lowest error values (MAE = 2.46, RMSE = 3.32, MAPE = 1.08%) and the highest coefficient of determination ($R^2 = 0.9993$). In addition, XGBoost provided a substantial advantage in terms of computational efficiency compared to other ensemble models, making it the most suitable approach for practical applications. The Gradient Boosting model also demonstrated very strong performance, achieving results comparable to XGBoost in terms of accuracy. However, its significantly longer computation time makes it less efficient for real-time or large-scale applications. The Decision Tree model showed relatively strong performance, with a high R^2 value (0.9964) and low error metrics. Despite its simplicity, it was able to capture a large portion of the variance in the data. However, its performance remains lower than that of boosting-based models when dealing with complex relationships. In contrast, the Random Forest model exhibited lower performance compared to boosting methods, with higher error values and reduced explanatory power ($R^2 = 0.9323$). This suggests that, for the given dataset, boosting-based approaches are more effective in modeling complex nonlinear relationships. The Lasso and Ridge regression models produced moderate results, indicating that linear models are limited in capturing nonlinear patterns in the dataset. Although these models offer computational efficiency, their predictive performance is lower compared to ensemble-based methods. The KNN algorithm showed the weakest performance, with the highest

error rates and the lowest R^2 value (0.7709), along with a very high computational cost. This confirms that distance-based methods are not suitable for this type of regression problem. Overall, the findings indicate that boosting-based models, particularly XGBoost, are the most reliable and effective methods for furniture price prediction. These models provide high accuracy, low error rates, and efficient computation, making them highly suitable for real-world applications involving complex and high-dimensional data.

Limitations and Future Work

Despite the promising results obtained in this study, several limitations were acknowledged. First, the dataset used in this research is relatively limited in size, which may restrict the generalizability of the findings. In addition, the dataset was obtained from a publicly available source and may not fully reflect real-world market dynamics. Another important limitation is the potential presence of strong correlations between input features and the target variable. As indicated by the SHAP analysis, certain variables such as cost have a dominant influence on the model predictions. While this contributes to high accuracy, it may also lead to model dependency on specific features, potentially affecting generalization performance. Furthermore, although cross-validation and hyperparameter optimization techniques were applied, the high R^2 values obtained in this study may still be influenced by dataset-specific patterns. Therefore, additional validation on independent datasets is necessary to confirm the robustness of the models. For future research, it is recommended to use larger and more diverse datasets that better represent real-world scenarios. In addition, incorporating advanced approaches such as deep learning models and hybrid machine learning frameworks may further improve prediction performance. Exploring feature selection techniques and reducing feature dominance may also enhance model interpretability and generalization.

Recommendations

Based on the findings of this study, several practical implications can be drawn. First, boosting-based algorithms, particularly XGBoost and Gradient Boosting, are strongly recommended for complex regression tasks such as furniture price prediction, as they demonstrate high predictive accuracy and robustness. Although Gradient Boosting achieves competitive performance, XGBoost is more suitable for real-world applications due to its significantly lower computational cost and faster training time. In contrast, distance-based methods such as K-Nearest Neighbors (KNN) are not appropriate for high-dimensional regression problems, as they tend to suffer from both reduced accuracy and increased computational burden. Similarly, linear models such as Lasso and Ridge regression should be used with caution; while they offer interpretability and may perform adequately on simpler datasets, they are generally insufficient for capturing complex nonlinear relationships. Furthermore, the results emphasize the critical importance of dataset characteristics in determining model performance. Comprehensive exploratory data analysis and effective feature engineering are therefore essential steps in selecting and optimizing the most suitable predictive model. Overall, the findings indicate that XGBoost provides the best balance between accuracy and computational efficiency, making it the most practical choice among the evaluated models. Future research should focus on the development of hybrid and deep learning-based approaches to further improve predictive performance.

Declarations

Acknowledgements: The author would like to thank the Kaggle platform for providing the dataset and computational environment used in this study.

Funding/Financial Disclosure: This study did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Ethics Committee Approval and Permissions: This study does not require ethics committee approval, as it does not involve human participants or animals.

Conflict of Interests: The author declares that there is no conflict of interest.

Authors Contribution: The author contributed to all stages of the study at a rate of 100%.

Artificial Intelligence (AI): Author declare that AI tools (such as ChatGPT and Gemini) were used for language correction, translation, and readability enhancement, as well as for creating tables, designing figures, and generating charts/graphics during the preparation of this article. The author assume full responsibility for the data accuracy, visual integrity, and originality of the entire content, including all AI-assisted outputs.

References

[1] Raposo, A., Frade, S., Alves, M., & Marques, J. F. (2018). The neural bases of price estimation: Effects of size and precision of the estimate. *Brain and Cognition*, *125*, 157–164. https://doi.org/10.1016/j.bandc.2018.07.005

[2] Lihra, T., Buehlmann, U., & Graf, R. (2012). Customer preferences for customized household furniture. *Journal of Forest Economics*, *18*(2), 94–112. https://doi.org/10.1016/j.jfe.2011.11.001

[3] Şahin, O., & Çubukçu, B. (2024). Mobilya sektöründe satış tahmini: Yinelemeli sinir ağı modellerinin karşılaştırmalı analizi. *Mühendislik Bilimleri ve Tasarım Dergisi*, *12*(4), 686–706. https://doi.org/10.21923/jesd.1433624

[4] Gahirwal, M. (2013). *Inter time series sales forecasting*. *arXiv*. https://doi.org/10.48550/arXiv.1303.0117

[5] Luxhøj, J. T., Riis, J. O., & Stensballe, B. (1996). A hybrid econometric–neural network modeling approach for sales forecasting. *International Journal of Production Economics*, *43*(2–3), 175–192. https://doi.org/10.1016/0925-5273(96)00039-4

[6] Yip, D. H., Hines, E. L., & Yu, W. W. (1997). Application of artificial neural networks in sales forecasting. In *Proceedings of International Conference on Neural Networks (ICNN'97)* (Vol. 4, pp. 2121–2124). IEEE. https://doi.org/10.1109/ICNN.1997.614233

[7] Alon, I., Qi, M., & Sadowski, R. J. (2001). Forecasting aggregate retail sales: A comparison of artificial neural networks and traditional methods. *Journal of Retailing and Consumer Services*, *8*(3), 147–156. https://doi.org/10.1016/S0969-6989(00)00011-4

[8] Kuo, R. J., Wu, P., & Wang, C. P. (2002). An intelligent sales forecasting system through integration of artificial neural networks and fuzzy neural networks with fuzzy weight elimination. *Neural Networks*, *15*(7), 909–925. https://doi.org/10.1016/S0893-6080(02)00064-3

[9] Mahbub, N., Paul, S. K., & Azeem, A. (2013). A neural approach to product demand forecasting. *International Journal of Industrial and Systems Engineering*, *15*(1), 1–18. https://doi.org/10.1504/IJISE.2013.055508

[10] Fabianová, J., Kačmáry, P., Molnar, V., & Michalik, P. (2016). Using a software tool in forecasting: A case study of sales forecasting taking into account data uncertainty. *Open Engineering*, *6*(1), 73–86. https://doi.org/10.1515/eng-2016-0011

[11] Fetzer, J. H. (1990). *What is artificial intelligence? In artificial intelligence: Its scope and limits* (pp. 3–27). Springer. https://doi.org/10.1007/978-94-009-1900-6_1

[12] Öztürk, K., & Şahin, M. E. (2018). Yapay sinir ağları ve yapay zekâ'ya genel bir bakış. *Takvim-i Vekayi*, *6*(2), 25–36.

[13] Alpaydin, E. (2020). *Introduction to machine learning*. MIT Press. https://doi.org/10.7551/mitpress/10994.001.0001

[14] James, G., Witten, D., Hastie, T., & Tibshirani, R. (2013). *An introduction to statistical learning: with applications in R*. Springer. https://doi.org/10.1007/978-1-4614-7138-7

[15] Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The elements of statistical learning: Data mining, inference, and prediction* (2nd ed.). Springer. https://doi.org/10.1007/978-0-387-84858-7

[16] Breiman, L., Friedman, J., Olshen, R. A., & Stone, C. J. (2017). *Classification and regression trees*. CRC Press. https://doi.org/10.1201/9781315139470

[17] Quinlan, J. R. (1986). Induction of decision trees. *Machine Learning*, *1*(1), 81–106. https://doi.org/10.1007/BF00116251

[18] Breiman, L. (2001). Random forests. *Machine Learning*, *45*, 5–32. https://doi.org/10.1023/A:1010933404324

[19] Cover, T., & Hart, P. (1967). Nearest neighbour pattern classification. *IEEE Transactions on Information Theory*, *13*(1), 21–27. https://doi.org/10.1109/TIT.1967.1053964

[20] Altman, N. S. (1992). An introduction to kernel and nearest-neighbor nonparametric regression. *The American Statistician*, *46*(3), 175–185. https://doi.org/10.1080/00031305.1992.10475879

[21] Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B*, *58*(1), 267–288. https://doi.org/10.1111/j.2517-6161.1996.tb02080.x

[22] Hoerl, A. E., & Kennard, R. W. (1970). Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, *12*(1), 55–67. https://doi.org/10.1080/00401706.1970.10488634

[23] Friedman, J. H. (2001). Greedy function approximation: A gradient boosting machine. *Annals of Statistics*, *29*(5), 1189–1232. https://doi.org/10.1214/aos/1013203451

[24] Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 785–794). ACM. https://doi.org/10.1145/2939672.2939785

[25] Chicco, D., Warrens, M. J., & Jurman, G. (2021). The coefficient of determination R-squared is more informative than SMAPE, MAE, MAPE, MSE and RMSE. *PeerJ Computer Science*, *7*, e623. https://doi.org/10.7717/peerj-cs.623

[26] Chai, T., & Draxler, R. R. (2014). Root mean square error (RMSE) or mean absolute error (MAE)? *Geoscientific Model Development*, *7*(3), 1247–1250. https://doi.org/10.5194/gmd-7-1247-2014

[27] Ćalasan, M., Aleem, S. H. A., & Zobaa, A. F. (2020). RMSE calculation for photovoltaic models. *Energy Conversion and Management*, *210*, Article number 112716. https://doi.org/10.1016/j.enconman.2020.112716

[28] De Myttenaere, A., Golden, B., Le Grand, B., & Rossi, F. (2016). Mean absolute percentage error for regression models. *Neurocomputing*, *192*, 38–48. https://doi.org/10.1016/j.neucom.2015.12.114

[29] Hazır, E., Koç, K. H., & Esnaf, Ş. (2016). Türkiye mobilya satış değerlerinin örnek bir yapay zekâ uygulaması ile tahmini. *Selçuk-Teknik Dergisi*, 1172–1182.

[30] Nacar, E. N., & Erdebilli, B. (2021). Makine öğrenmesi algoritmaları ile satış tahmini. *Endüstri Mühendisliği, 32*(2), 307–320. https://doi.org/10.46465/endustrimuhendisligi.811183

[31] Agrawal, S. (2021). *Furniture sales data* [Veri seti]. Kaggle. https://www.kaggle.com/datasets/itsmesunil/furniture-sales-data

[32] Makridakis, S., Spiliotis, E., & Assimakopoulos, V. (2022). M5 accuracy competition: Results, findings, and conclusions. *International Journal of Forecasting, 38*(4), 1346–1364. https://doi.org/10.1016/j.ijforecast.2021.11.013

[33] Kourentzes, N., Barrow, D. K., & Crone, S. F. (2014). Neural network ensemble operators for time series forecasting. *Expert Systems with Applications, 41*(9), 4235–4244. https://doi.org/10.1016/j.eswa.2013.12.011

[34] Bandara, K., Bergmeir, C., & Smyl, S. (2020). Forecasting across time series databases using recurrent neural networks on groups of similar series: A clustering approach. *Expert Systems with Applications, 140*, 112896. https://doi.org/10.1016/j.eswa.2019.112896

[35] Natekin, A., & Knoll, A. (2013). Gradient boosting machines, a tutorial. *Frontiers in Neurorobotics, 7*, 21. https://doi.org/10.3389/fnbot.2013.00021

[36] Çardak, H., Bardak, S., Bardak, T., Çapraz, O., Özçetin, S., & Kızılırmak, S. (2025). Predicting consumer preferences for furniture products on e-commerce platforms: An analysis using machine learning and favorite listing data. *BioResources, 20*(4), 9768–9784. https://doi.org/10.15376/biores.20.4.9768-9784

[37] Akyüz, I., Polat, K., Bardak, S., & Ersen, N. (2024). Prediction of values of Borsa Istanbul forest, paper, and printing index using machine learning methods. *BioResources, 19*(3), 5141–5157. https://doi.org/10.15376/biores.19.3.5141-5157

[38] Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation, 9*(8), 1735–1780. https://doi.org/10.1162/neco.1997.9.8.1735

[39] Smyl, S. (2020). A hybrid method of exponential smoothing and recurrent neural networks for time series forecasting. *International Journal of Forecasting, 36*(1), 75–85. https://doi.org/10.1016/j.ijforecast.2019.03.017

[40] Makridakis, S., Spiliotis, E., & Assimakopoulos, V. (2018). The M4 competition: Results, findings, conclusion and way forward. *International Journal of Forecasting, 34*(4), 802–808. https://doi.org/10.1016/j.ijforecast.2018.06.001

[41] Bontempi, G., Ben Taieb, S., & Le Borgne, Y. A. (2012). Machine learning strategies for time series forecasting. In *European Big Data Management and Analytics Summer School* (pp. 62–77). Springer. https://doi.org/10.1007/978-3-642-36318-4_3

[42] Guyon, I., & Elisseeff, A. (2003). An introduction to variable and feature selection. *Journal of Machine Learning Research, 3*, 1157–1182.

[43] Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems, 30*. https://doi.org/10.48550/arXiv.1705.07874