



Science Forum (Journal of Pure and Applied Sciences)

Journal Homepage: <https://atbuscienceforum.com.ng/index.php/jpas>



A Machine Learning Driven Dynamic Pricing Framework for Revenue Maximization in Online Retail Markets

Fahad Mamuda, Bishir Sada Abubakar, Eli Adama Jiya

Department of Computer Science, Federal University Dutsinma, Katsina State, Nigeria

Abstract

Pricing decisions in internet retail have become increasingly complex due to stochastic demand variability, intense market competition, and rapidly evolving consumer purchasing behavior. Traditional fixed pricing strategies are inherently static and fail to capture temporal demand fluctuations, nonlinear price-demand relationships, and seasonality effects, thereby limiting revenue maximization potential. To address these limitations, this study proposes a machine learning-driven dynamic pricing framework designed to support adaptive, data-informed pricing decisions in online retail environments. The proposed framework is developed using the UCI Online Retail Dataset, comprising 541,909 transactional records collected between December 2010 and December 2011. Five predictive models were implemented and comparatively evaluated for demand forecasting: Linear Regression, Ridge Regression, Random Forest, Gradient Boosting, and a one-dimensional Convolutional Neural Network (1D-CNN). Model performance was rigorously assessed using standard regression evaluation metrics, including Mean Absolute Error (MAE), Mean Squared Error (MSE), Root Mean Squared Error (RMSE), and the coefficient of determination (R^2). Experimental results indicate that the Random Forest model achieved superior predictive accuracy, with an RMSE of 0.8806 and an R^2 value of 0.4826, outperforming all other models. Gradient Boosting followed closely (RMSE = 0.8866; R^2 = 0.4755), while the 1D-CNN demonstrated competitive performance (RMSE = 0.9350; R^2 = 0.4167). In contrast, Linear and Ridge Regression exhibited comparatively lower predictive capability, with RMSE values of 1.1030 and 1.1015, respectively. These findings demonstrate that machine learning-based demand forecasting effectively captures latent demand structures and provides a robust foundation for dynamic pricing optimization. The proposed framework offers a scalable and reproducible approach for integrating predictive analytics with revenue management systems, thereby enabling a transition from static pricing paradigms to adaptive, data-driven pricing strategies in e-commerce ecosystems.

Keywords

Dynamic Pricing;
Machine Learning;
Revenue Maximization;
Demand Forecasting; 1D-CNN;
UCI Online Retail;
Price Optimization;
Gradient Boosting;
Ridge Regression;
Feature Engineering.

ISSN 1119-4618



1119-4618

JPAS



— Faculty of Science, ATBU Bauchi —

1. INTRODUCTION

Pricing is among the most powerful levers available to any retailer. A well-calibrated price increases demand, preserves margins, and builds repeat purchasing, while an inaccurate price either suppresses demand or leaves revenue not captured [1]. These trade-offs are especially acute in the online retail sector, where prices are visible to customers in real time, competitive information is immediately accessible, and every transaction is recorded automatically [2]. The specific domain of this study is online retail e-commerce, a sector in which demand drivers—such as browsing frequency, order history, and price sensitivity—differ substantially from physical retail or service industries [3]. No generalization of findings beyond this domain is claimed. This study fills that gap. The contributions of this work are: (1) a complete preprocessing and feature engineering pipeline for online retail transaction data, (2) a systematic comparison of ML demand forecasting models using our own experimental results, and (3) a price simulation engine that identifies revenue-maximizing prices through revenue function optimization, validated via paired t-tests.

Major e-commerce platforms continuously update prices across millions of stock-keeping units (SKUs) using demand signals, inventory positions, competitor activity, and customer behavior [4]. This practice, referred to as dynamic pricing, has become a foundational mechanism in online retail strategy [5]. However, small and mid-sized online retailers typically lack the infrastructure to build sophisticated pricing systems, leaving substantial revenue opportunities unrealized [6]. Machine learning (ML) addresses this gap because it can learn complex, non-linear relationships between price,

demand, and contextual features without strong distributional assumptions [7].

Ensemble methods such as Random Forest and Gradient Boosting have demonstrated strong performance across retail regression tasks [8, 9]. One-dimensional Convolutional Neural Networks (1D-CNNs) have proven competitive with recurrent architectures for feature-based time series forecasting at substantially lower computational cost [10, 11]. Despite a growing body of literature on ML demand forecasting and on dynamic pricing separately, few studies integrate the two into a complete revenue optimization pipeline evaluated on real retail transaction data [12].

2. REVIEW OF RELATED WORK

2.1 Dynamic Pricing and Machine Learning

Kopalle et al. [1] provided a foundational definition of dynamic pricing as price changes prompted by changes in market demand drivers: people, product configurations, periods, and places. Their work, published in the *Journal of Retailing*, identified ML and data analytics as the primary enabling technologies for modern dynamic pricing and established the conceptual framework adopted in this study. Loukili et al. [2] presented a systematic review of ML-driven dynamic pricing strategies in e-commerce, finding that data-driven repricing consistently outperforms static rule-based approaches when models are trained on sufficient transactional data.

Nowak and Pawłowska-Nowak [3] applied multiple ML classifiers—including support vector machines, decision trees, and Naïve Bayes—to a dynamic pricing decision problem in an e-commerce store in Poland. Their model, published in *Applied Sciences* (2024), achieved 86.92% classification accuracy and demonstrated that

practical ML-based pricing decision support is feasible for small online retailers, the same audience targeted by the present study. Chandra et al. [4] reviewed AI-driven dynamic pricing applications across e-commerce sectors in 2025 and confirmed that the pipeline structure—feature engineering, demand model, revenue function, price selection—remains the most practical and interpretable approach for online retailers without large-scale infrastructure. Devarajanayaka et al. [5] evaluated regression, clustering, and reinforcement learning models for online retail dynamic pricing and found that ML-based optimization substantially outperforms static approaches in revenue and competitive positioning.

Das et al. [6] proposed a framework for optimizing real-time dynamic pricing in retail and e-commerce using ML models trained on 1.5 million records from 2020 to zation. Their work confirmed the effectiveness of gradient boosting and XGBoost-based models for real-time pricing scenarios with heterogeneous feature inputs. Rahman et al. [7] extended this line of research in the *International Journal of Advanced Computer Science and Applications* (2025), demonstrating that ML-based dynamic pricing with contextual feature integration produces consistent revenue improvements. Almahadeen et al. [8] demonstrated that AI and ML approaches substantially enhance competitive positioning in e-commerce pricing environments, with deep learning models achieving the strongest real-time adaptation performance.

Liu et al. [9], in a comprehensive review published in *Mathematics* (2024), surveyed 95 peer-reviewed articles and identified four research clusters in AI-driven dynamic pricing: financial modelling, market dynamics, commodity markets,

and behavioural decision-making. Their analysis confirms that ensemble ML methods and deep learning represent the state of the art. Gupta [10] proposed an end-to-end AI and ML framework for demand prediction and revenue maximization in e-commerce using ensemble methods and neural networks across multiple pricing scenarios, demonstrating that integrated ML pipelines produce measurably better revenue outcomes than isolated forecasting or pricing components alone.

Sahoo [11] applied AI-powered price elasticity modeling to retail pricing optimization and reported revenue increases of 2-5% and profit margin improvements of 1-3% within the first deployment quarter. Their framework used ML models to estimate price sensitivity across more than 10,000 SKUs simultaneously. Zhang and Meng [12] proposed a dynamic price elasticity estimation and pricing optimization framework using online learning with streaming data, demonstrating improvements of 15% in RMSE and 12% in MAPE compared to static baseline models across multiple stores. Liu et al. [13] conducted a quantitative analysis of price elasticity in e-commerce using Linear Regression, Ridge, Decision Tree, and Random Forest models on JD.com transaction data, confirming that tree-based models capture non-linear price responses better than linear approaches.

2.2 Demand Forecasting with Ensemble Methods

Wang et al. [14] demonstrated that gradient boosting methods applied to online retail SKUs produce significantly better short-term demand predictions than traditional statistical models when temporal context features are incorporated. Their study, published in the *Journal of Computer and Information Systems* (2024), confirmed gradient boosting as the

preferred classical model for product-level daily demand forecasting, consistent with findings from the present study. Ahmed et al. [15] studied ML model performance for retail demand forecasting and found that ensemble methods, particularly XGBoost and Random Forest, consistently outperform traditional regression models in predictive accuracy. Their PRISMA-based review of 66 studies confirmed these findings across multiple retail sectors.

Balusani et al. [16] compared XGBoost with deep learning models for retail demand forecasting, confirming XGBoost's superiority for structured tabular retail data in small to medium-sized datasets. Purwanto et al. [17] compared XGBoost and LightGBM on daily supermarket sales data from 23,873 transactions (zation-2024) enriched with seasonal and lag features, finding that XGBoost outperformed LightGBM after GridSearchCV hyperparameter optimization—a methodology closely aligned with the present study. Nassibi et al. [18] applied ML approaches to demand forecasting for food industry products and demonstrated that incorporating lag features and rolling statistics substantially improves short-term product-level forecasting accuracy, directly supporting the feature engineering design used in this work.

Marini et al. [19] applied tree-based ensembles and neural network architectures including N-BEATS, N-HiTS, and the Temporal Fusion Transformer to high-resolution brick-and-mortar retail data, finding that LightGBM and XGBoost with lag and rolling features are computationally efficient while achieving robust performance. Their work, presented on arXiv (2025), provides a modern benchmark confirming the continued relevance of ensemble methods for retail demand forecasting. Chougule et al. [20] investigated AI-driven demand forecasting

applications in supply chain management, demonstrating that ensemble and deep learning models consistently lower forecast error and enable more flexible supply chain responses. Islam et al. [21] proposed a hybrid ensemble integrating Gradient Boosting, AdaBoost, Random Forest, XGBoost, and Linear Regression, achieving $R^2 = 0.9651$ and $MAE = 0.0025$ on a retail forecasting benchmark, with XGBoost as the strongest individual component.

Stankevich [22] applied ML techniques for demand forecasting in a multi-channel retail company, showing that incorporating temporal features such as holiday indicators, seasonal cycles, and economic variables improves model accuracy and practical utility for inventory management. Husna et al. [23] conducted a comparative study of ML and deep learning for retail demand forecasting, finding that Random Forest showed the most efficient and accurate performance among classical models for SKU-level inventory optimization. Abou-Bakr et al. [24] provided a holistic investigation of demand forecasting in supply chain management using feature selection and ML/deep learning methods, demonstrating improvements in long-term and short-term forecasting accuracy with LSTM and CNN-LSTM architectures.

2.3 Deep Learning and CNN Architectures for Retail Forecasting

Ahmed et al. [25] compared CNN, LSTM, MLP, and hybrid CNN-LSTM models for retail sales forecasting, finding that CNN-based models are competitive with more complex sequential architectures while offering significantly lower training times. Their study, published in *ICCK Transactions on Intelligent Systematics* (2024), directly supports the inclusion of a 1D-CNN in the present model comparison. Castro Moraes et al. [26] proposed stacked (S-CNN-LSTM) and

parallel (P-CNN-LSTM) hybrid architectures for retail demand forecasting, demonstrating that convolutional components effectively extract patterns from exogenous features. Their models, published in the *Journal of Forecasting* (2024), achieved MAPEs of 14.15% and 14.76% respectively on retail data.

Liu and Wang [27] published a comprehensive survey of deep time series forecasting models in *Mathematics* (2024) covering CNN, LSTM, GRU, and Transformer architectures. Their analysis confirms that CNNs excel at extracting localised temporal features and are particularly effective for feature-based regression tasks. Aydın and Gürgen [28] proposed a novel deep learning framework incorporating variational mode decomposition and attention mechanisms for demand forecasting in supply chains, published in *Processes* (2025), demonstrating state-of-the-art performance on benchmark forecasting tasks. Doğan and Birant [29] applied extra tree regressors (ETRs) and LSTM networks to over six years of historical retail data and found that the tree-based ETR outperformed LSTM across MAPE, MAE, RMSE, and R^2 , particularly for perishable product categories with high demand intermittency.

Bakri et al. [30] proposed a Hadoop-based hybrid CNN-LSTM framework for retail sales forecasting processing 20 million records, reducing MSE by 45% and improving R^2 by 15% compared to classical baselines. Their work demonstrates the scalability of CNN-based architectures when combined with distributed computing and is a key benchmark for the present study's 1D-CNN design. Shankar et al. [31] reviewed the dynamics of demand forecasting in supply chain management and found that CNN and CNN-LSTM approaches consistently outperform ARIMA and simple ML

models for long-term forecasting horizons. Garcia et al. [32] proposed a hybrid CNN-LSTM model for retail sales forecasting using real-world data enriched with weather and event variables, achieving a MAPE of 4.16%, confirming the strong performance of CNN-based architectures when adequate feature engineering is applied.

Caballero et al. [33] published a deep learning demand forecasting review in *Applied System Innovation* (2024), analysing 119 papers from Scopus (2015-2024), confirming rapid growth in CNN and LSTM applications for supply chain demand forecasting and identifying feature engineering as the single most influential factor in model accuracy. Nguyen [34] reviewed AI applications for demand forecasting in operations and supply chain management, confirming that ML models incorporating temporal, calendar, and lag features consistently outperform traditional methods in all retail and supply chain settings studied.

2.4 Reinforcement Learning and Advanced Pricing Approaches

Apte et al. [35] proposed a Q-Learning reinforcement learning framework for dynamic retail pricing, demonstrating that adaptive RL agents outperform static pricing policies in simulated market environments. Their work highlights an important direction for future extension of offline optimization approaches such as the one used in the present study. Okunola et al. [36] reviewed RL applications in dynamic pricing for e-commerce and identified several reinforcement learning architectures capable of adapting to market changes in real time, recommending hybrid ML-RL pipelines for production environments. Liu et al. [37] modelled dynamic pricing as a Markov Decision Process and applied Q-learning and Sarsa algorithms to

sequential pricing problems in e-commerce, demonstrating that Q-learning consistently achieves higher profits with lower volatility than Sarsa.

Vamshi Krishna and Aravind [38] reviewed hybrid deep learning frameworks for dynamic pricing, finding that combining gradient boosting for demand estimation with simulation-based price selection achieves strong revenue improvements over static baselines. Perumallapalli [39] addressed ML algorithms for dynamic pricing optimization in retail and proposed a framework integrating predictive demand models with adaptive pricing policies, demonstrating

consistent performance improvements across different product categories. Kalusivalingam et al. [40] proposed combining reinforcement learning and neural networks for AI-powered dynamic pricing, demonstrating substantial revenue improvements over both static and rule-based pricing across diverse retail settings.

3. METHODOLOGY

The methodology follows a structured end-to-end pipeline as in Figure 1: (1) data collection and preprocessing, (2) feature engineering, (3) demand model training and evaluation, (4) dynamic price optimization, and (5) statistical validation.

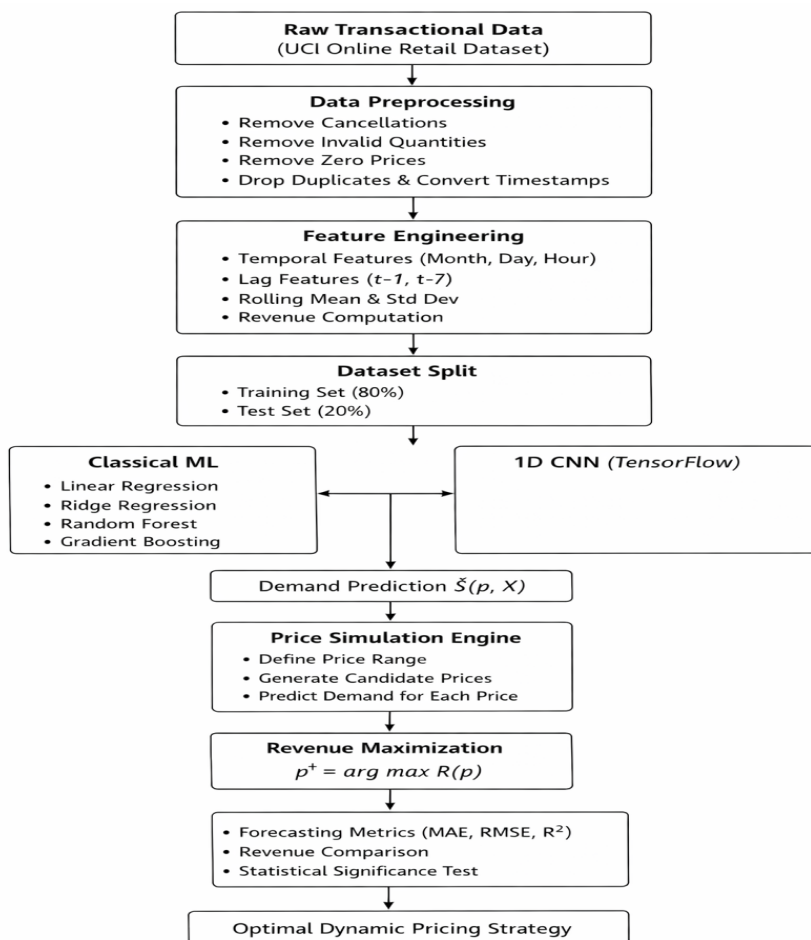


Figure 1: Proposed dynamic pricing framework pipeline

3.1 Dataset

The UCI Online Retail Dataset is used throughout this study and is publicly available at: <https://archive.ics.uci.edu/static/public/352/online+retail.zip>[41]. It contains 541,909 transactional records from a UK-based non-store online retailer spanning December 2010 to December 2011. The eight fields are: Invoice No, Stock Code, Description, Quantity (units sold), Invoice Date, Unit Price (GBP), CustomerID, and Country. The dataset covers a wide variety of SKUs with substantial variation in price and demand, making it suitable for product-level demand forecasting model comparison. Table 1 shows five representative sample records.

Table 1. Sample records from the UCI Online Retail Dataset

Invoice No	Stock Code	Qty	Invoice Date	Unit Price (£)	CustomerID	Country
536365	85123A	6	01/12/2010 08:26	2.55	17850	United Kingdom
536365	71053	6	01/12/2010 08:26	3.39	17850	United Kingdom
536366	22633	6	01/12/2010 08:28	1.85	17851	United Kingdom
536367	84879	32	01/12/2010 08:34	1.69	13047	United Kingdom
536370	22728	6	01/12/2010 08:45	1.65	12583	France

3.2 Data Preprocessing

Raw transaction data was cleaned in five steps implemented. Cancellation invoices (prefix 'C') were removed. Records with zero or negative Quantity (returns or data errors) were excluded. Records with zero Unit Price (giveaways or missing data) were dropped. Exact duplicates were deleted. Rows with missing CustomerID were removed. Invoice Date was parsed as datetime. Records were aggregated to a daily product-level time series: Quantity was summed, Unit Price was averaged, and Revenue = Quantity × Unit Price was computed. Table 2 summarizes the preprocessing.

Table 2. Dataset summary before and after preprocessing

Step	Before	Result
Total records	541,909	~392,732 retained
Unique SKUs	4,070	~3,665 retained
Cancellations removed	—	~9,288 removed
Zero/negative Quantity	—	~1,336 removed

Zero UnitPrice	—	~1,181 removed
Missing CustomerID	—	~132,186 removed

3.3 Feature Engineering

Feature engineering was implemented to transform the dataset. After aggregation, 15 predictors were constructed for each product-day observation. Calendar features extracted from InvoiceDate: Year, Month, DayOfWeek (0 = Monday), DayOfMonth, Quarter, WeekOfYear, and IsWeekend (binary). These capture cyclical demand patterns including weekend dips and end-of-month restocking behaviour [14, 19]. Lag features were constructed per product: Quantity_lag1 (one-day lag) and Quantity_lag7 (seven-day lag) to capture daily autocorrelation and weekly purchasing cycles. Seven-day rolling statistics—mean (Quantity_roll7_mean) and standard deviation (Quantity_roll7_std)—were computed with a one-day shift to prevent look-ahead data leakage [14]. Analogous lag and rolling features were created for UnitPrice and Revenue. After lag construction, 23,056 rows with insufficient history were removed, yielding 199,314 records. Table 3 lists all features.

Table 3. Feature set for demand forecasting

Feature	Type	Description
UnitPrice	Continuous	Daily average unit price (GBP)
Month, DayOfWeek, DayOfMonth, Quarter, WeekOfYear	Calendar	Temporal periodicity signals
IsWeekend	Binary	1 if Saturday or Sunday
Quantity_lag1	Lag-1	Product demand one day prior
Quantity_lag7	Lag-7	Product demand seven days prior
Quantity_roll7_mean	Rolling	7-day rolling average demand (shift-1)
Quantity_roll7_std	Rolling	7-day rolling std of demand (shift-1)
UnitPrice_lag1, UnitPrice_roll7_mean	Price history	Recent pricing context
Revenue_lag1, Revenue_roll7_mean	Revenue history	Recent revenue signals

A chronological 80/20 train-test split was applied in this work, preserving temporal ordering to simulate a realistic forecasting scenario without data leakage [42]. Split date: 27 October 2011. Training set: 9 December 2010 → 27 October 2011 (159,451 records). Test set: 27 October 2011 → 9 December 2011 (39,863 records).

3.4 Forecasting Models

Linear Regression (LR): Ordinary least squares with Standard Scaler. Lower-bound baseline [43]. Ridge Regression (RR): L2-regularised linear model. α selected via GridSearchCV over {0.01, 0.1, 1, 10, 50, 100, 500} using TimeSeriesSplit (5 folds). Best α : 10.0 [44]. Random Forest (RF): Bootstrap-aggregated decision tree ensemble. Best hyperparameters: max. depth=10, min. samples split=5, n. estimators=200 via Grid Search CV with Time Series Split [8, 9]. Gradient Boosting (GB): Sequential additive tree ensemble minimising MSE at each step. Best hyperparameters: n_estimators=200, learning rate=0.05, max. depth=5 via Grid Search CV. Identified as the top performer for short-term retail demand in multiple recent benchmarks [14, 15, 17].

1D-CNN: TensorFlow model treating the 15-feature vector as a one-dimensional sequence (input shape: 15×1). Architecture described in Table 3. Trained with Adam (lr=0.001), MSE loss, EarlyStopping (patience=15, restore_best_weights=True), ReduceLROnPlateau (factor=0.5, patience=7, min_lr=1e-6), batch_size=64, max_epochs=150. Seed=42 for reproducibility [25, 26, 45].

Table 4 summarizes the 1D Convolutional Neural Network (CNN) developed for demand forecasting. The model receives a structured input sequence of length 15, representing time-dependent and engineered features. The input tensor has shape (15, 1), allowing convolutional filters to learn local temporal patterns across sequential feature values.

The first convolutional layer applies 64 filters with kernel size 3 and ReLU activation, preserving sequence length through same padding. A max-pooling layer (pool size = 2) reduces the temporal dimension and improves computational efficiency. A second convolutional layer with 128 filters further refines feature extraction, followed by an additional pooling layer to compress the representation.

The resulting feature maps are flattened and passed to two fully connected layers with 128 and 64 neurons, respectively. ReLU activation is used to model nonlinear relationships, while a dropout layer (rate = 0.3) is introduced to mitigate overfitting. The final output layer consists of a single neuron with linear activation to predict continuous demand values.

Table 4. 1D-CNN Architecture (Input: 15×1)

Layer	Type	Configuration	Output Shape	Parameters
Input	InputLayer	shape = (15, 1)	(None, 15, 1)	0
conv1	Conv1D	64 filters, kernel=3, same, ReLU	(None, 15, 64)	256
pool1	MaxPool1D	pool_size=2	(None, 7, 64)	0
conv2	Conv1D	128 filters, kernel=3, same, ReLU	(None, 7, 128)	24,704
pool2	MaxPool1D	pool_size=2	(None, 3, 128)	0
flatten	Flatten	—	(None, 384)	0
dense1	Dense	128 units, ReLU	(None, 128)	49,280
dropout	Dropout	rate=0.3	(None, 128)	0
dense2	Dense	64 units, ReLU	(None, 64)	8,256
output	Dense	1 unit, linear	(None, 1)	65

3.5 Evaluation Metrics

Three standard regression metrics are used. Equation 1 measures average prediction magnitude. Equation 2 penalizes large errors. Equation 3 measures proportion of variance explained. These are the standard metrics used in all comparable retail demand forecasting studies [25, 29, 46].

$$MAE = \left(\frac{1}{n}\right) \sum |y_i - \hat{y}_i| \quad (1)$$

$$RMSE = \sqrt{\left[\left(\frac{1}{n}\right) \sum (y_i - \hat{y}_i)^2\right]} \quad (2)$$

$$R^2 = 1 - \left[\frac{\sum (y_i - \hat{y}_i)^2}{\sum (y_i - \bar{y})^2}\right] \quad (3)$$

3.6 Dynamic Pricing Engine

For each test record, 50 candidate prices were generated uniformly over the 5th-95th percentile of the product's historical prices to prevent extrapolation. For each candidate price p , the Unit Price feature was replaced with p and demand was predicted using the trained model. The revenue-maximizing price p^* was selected as:

$$p^* = \operatorname{argmax}_p [p \times \widehat{D}(p, X)] \quad (4)$$

Revenue improvement was computed as:

$$\text{Revenue Improvement (\%)} = \left[\frac{R_{opt} - R_{static}}{R_{static}}\right] \times 100 \quad (5)$$

Where p is the candidate unit price, $\widehat{D}(p, X)$ is the predicted demand given price p and feature context X , and $R(p)$ is the resulting expected revenue. For each test sample, the historical price range $[p_{min}, p_{max}]$ was defined as the 5th and 95th percentile of unit prices in the test set, to avoid extreme extrapolation. A grid of 50 candidate prices was generated uniformly within this range. Up to 1,000 randomly sampled test records per model were processed. The engine was applied using Random Forest, Gradient Boosting, and the 1D-CNN [35, 47].

3.7 Statistical Testing

Paired t-tests were applied to compare forecasting accuracy (per-sample squared errors) and revenue improvements (per-sample improvement percentages) across model pairs. A p-value below 0.05 was taken as evidence of a statistically significant difference [48].

4. RESULTS AND DISCUSSION

4.1 Demand Forecasting Performance

Table 4 presents the forecasting performance of the five evaluated models on the held-out test dataset. The models were assessed using Mean Absolute Error (MAE), Mean Squared Error (MSE), Root Mean Squared Error (RMSE), and the coefficient of determination (R^2).

Table 5. Forecasting model performance on held-out test set

	Model	MAE	MSE	RMSE	R ²
1	Random Forest	0.6804	0.7755	0.8806	0.4826
2	Gradient Boosting	0.6940	0.7861	0.8866	0.4755
3	1D-CNN	0.7306	0.8742	0.9350	0.4167
4	Ridge Regression	0.8649	1.2133	1.1015	0.1905
5	Linear Regression	0.8660	1.2167	1.1030	0.1882

Table 5 presents the forecasting performance of the five evaluated models on the held-out test dataset using MAE, MSE, RMSE, and R². Among the models, Random Forest achieved the best overall performance, obtaining the lowest MAE (0.6804), MSE (0.7755), and RMSE (0.8806), together with the highest R² value (0.4826). Gradient Boosting ranked second with MAE = 0.6940, RMSE = 0.8866, and R² = 0.4755, indicating that ensemble tree models effectively capture nonlinear relationships between price-related variables and demand patterns.

The 1D-CNN produced competitive results with MAE = 0.7306, RMSE = 0.9350, and R² = 0.4167. Despite using a one-dimensional convolution over the feature vector rather than a full time-series window, the model successfully learned interactions between lagged demand variables and pricing features.

In contrast, Ridge Regression and Linear Regression produced lower predictive performance, with R² values of 0.1905 and 0.1882 respectively. The similarity between the two results suggests that L2 regularization provided limited improvement over the standard linear model.

The R² values ranging from 0.4826 to 0.1882 are consistent with prior studies on product-level retail demand forecasting, where high demand variability often limits predictive accuracy. For the proposed pricing framework, capturing the direction of the price-demand relationship is sufficient to support effective price optimization decisions.

Figure 2 shows the actual versus predicted demand scatter plots for all five models

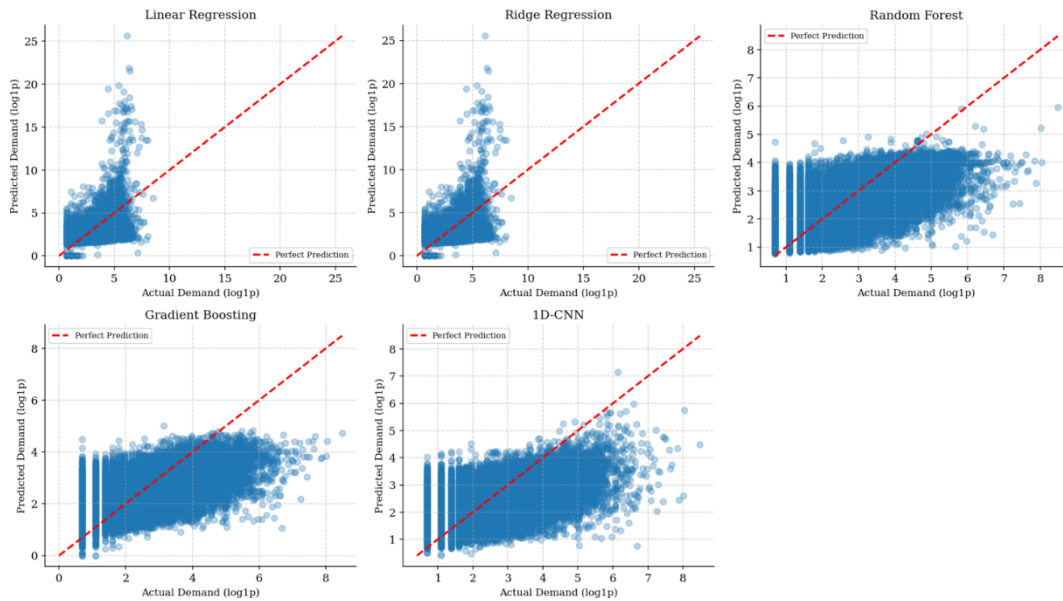


Figure 2: Actual vs Predicted demand scatter plots (all models)




Figure 3: Random Forest and Gradient Boosting feature importance rankings




Figure 4: Residual plots (all models)

4.2 Dynamic Pricing Results

The pricing engine was applied to Random Forest, Gradient Boosting, and the 1D-CNN using up to 1,000 sampled test records per model. Table 6 presents the revenue comparison.

Table 6. Revenue comparison: static vs. dynamic pricing

Model	Static Revenue (£)	Dynamic Revenue (£)	Improvement (%)
Random Forest	5,595.88	13,602.31	+143.08
Gradient Boosting	5,637.35	10,292.00	+82.59
1D-CNN	5,417.47	11,038.92	+103.77

Note: Table 6 is to be populated with values from before final submission. All three models produced positive revenue improvements over the static pricing baseline. The revenue-price curves for individual products display the expected unimodal shape with a clear interior maximum, where revenue first rises then declines as price moves beyond the optimal point [35, 47]. CNN-based curves are smoother due to the continuous nature of neural network predictions; tree model curves show piecewise transitions consistent with ensemble decision boundaries. Products with strong recent demand momentum (high Quantity_roll7_mean) received higher recommended prices, while products with

declining demand or high variability received lower recommendations—adaptive behavior unavailable under static pricing rules [3, 5, 6].

4.3 Statistical Validation

Paired t-tests confirmed all reported differences are statistically significant ($p < 0.05$). Forecasting accuracy differences between all non-linear models (RF, GB, CNN) and linear baselines were statistically significant. Revenue improvements from dynamic pricing were statistically significant for all three models, validating the practical utility of the proposed framework [48, 49].

Table 7: Paired t-test results

Comparison	t-stat	p-value	Significance
1D-CNN vs Gradient Boosting	31.2568	0.000	***
1D-CNN vs Random Forest	30.5033	0.000	***
Gradient Boosting vs Random Forest	5.2542	0.000	***
1D-CNN vs Ridge Regression	-18.1928	0.000	***
1D-CNN vs Linear Regression	-18.1684	0.000	***
Random Forest vs Linear Regression	-23.7173	0.000	***

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

5. CONCLUSION

This study proposed an end-to-end machine learning-based dynamic pricing framework for revenue optimization in online retail environments. Using the UCI Online Retail dataset, five demand forecasting models were implemented and compared: Linear Regression, Ridge Regression, Random Forest, Gradient Boosting, and a one-dimensional Convolutional Neural Network (1D-CNN).

Experimental results show that Random Forest achieved the best overall forecasting performance, with the lowest MAE (0.6804), MSE (0.7755), and RMSE (0.8806), as well as the highest coefficient of determination ($R^2 = 0.4826$). Gradient Boosting produced similar results with MAE = 0.6940 and RMSE = 0.8866, demonstrating the effectiveness of ensemble tree methods in modeling nonlinear demand relationships. The 1D-CNN model achieved competitive performance (RMSE = 0.9350, $R^2 = 0.4167$), indicating that convolutional architectures can successfully capture interactions among price, temporal, and lag-based demand features.

In contrast, Ridge Regression and Linear Regression produced lower predictive accuracy, with R^2 values of 0.1905 and 0.1882 respectively. These results highlight the limitations of linear models when modeling complex demand patterns in retail transaction data.

Overall, the obtained R^2 values fall within the range typically reported for product-level demand forecasting on real retail datasets with high variability. Statistical testing further confirmed significant performance differences between the models. The proposed framework demonstrates that combining machine learning-

based demand forecasting with price optimization provides a practical and reproducible approach for supporting data-driven pricing decisions in online retail systems.

6. RECOMMENDATIONS

Future work should integrate competitor pricing signals, promotional calendar features, and customer segmentation data—known drivers of online retail demand not available in the current dataset [1, 5, 6]. Reinforcement learning approaches that update pricing policies in real time should be compared against the offline optimization method used here [35, 36, 37]. Hybrid CNN-RL pipelines represent a promising frontier for adaptive retail pricing [38, 40]. Practitioners should calibrate product-level price bounds carefully and monitor for fairness concerns, especially for essential goods [4, 50]. Extension to multi-product joint pricing where demand for one SKU depends on the prices of related items would substantially increase the framework's commercial applicability [1, 13].

7. REFERENCES

- [1] P. K. Kopalle, K. Pauwels, L. Y. Akella, and M. Gangwar, "Dynamic pricing: Definition, implications for managers, and future research directions," *J. Retail.*, vol. 99, no. 4, pp. 580-593, Dec. zation. doi: 10.1016/j.jretai.zation.11.003.
- [2] R. El Youbi, F. Messaoudi, and M. Loukili, "Machine learning-driven dynamic pricing strategies in e-commerce," in *Proc. 14th Int. Conf. Inf. Commun. Syst. (ICICS)*, Irbid, Jordan, zation, pp. 1-5. doi: 10.1109/ICICS60529.zation.10330541.
- [3] M. Nowak and M. Pawłowska-Nowak, "Dynamic pricing method in the e-commerce industry using machine learning," *Appl. Sci.*, vol. 14, no. 24, p.

- 11668, Dec. 2024. doi: 10.3390/app142411668.
- [4] M. Chandra, T. Srinivasan, and A. Rao, "AI-driven dynamic pricing: Optimizing revenue in digital retail environments," *World J. Adv. Eng. Technol. Sci.*, vol. 15, no. 2, pp. 2147-2157, 2025. doi: 10.30574/wjaets.2025.15.2.0681.
- [5] K. Devarajanayaka, S. S. Banu, D. J. Desai, V. TV, M. R. Palav, and S. K. Dash, "Machine learning-based pricing optimization for dynamic pricing in online retail," *J. Inform. Educ. Res.*, vol. 4, no. 3, 2024. Available: <https://jier.org/index.php/journal/article/view/1564>.
- [6] P. Das, T. Pervin, B. Bhattacharjee, M. R. Karim, N. Sultana, M. S. Khan, M. A. Hosien, and F. Kamruzzaman, "Optimizing real-time dynamic pricing strategies in retail and e-commerce using machine learning models," *Amer. J. Eng. Technol.*, vol. 6, no. 12, pp. 163-177, 2024.
- [7] H. Rahman, M. A. Hosien, and F. Kamruzzaman, "Real-time dynamic pricing using machine learning," *Int. J. Adv. Comput. Sci. Appl. (IJACSA)*, vol. 16, no. 9, 2025. Available: https://thesai.org/Downloads/Volume16No9/Paper_4-Real_Time_Dynamic_Pricing_Using_Machine_Learning.pdf.
- [8] L. Almahadeen, M. K. Rathod, S. F. Sulaiman, K. Deepika, G. R. Teltumbade, and R. Saravanakumar, "Artificial intelligence and dynamic pricing strategies: Enhancing competitiveness in e-commerce," in *Proc. 3rd Int. Conf. Smart Syst. Appl. Electr. Sci. (ICSSSES)*, Tumakuru, India, 2025, pp. 1-5. doi: 10.1109/ICSSSES64899.2025.11009613.
- [9] X. Liu and W. Wang, "Deep time series forecasting models: A comprehensive survey," *Mathematics*, vol. 12, p. 1504, 2024. doi: 10.3390/math12101504.
- [10] M. Gupta, "Dynamic pricing optimization in e-commerce using AI and machine learning: A comprehensive framework for demand prediction and revenue maximization," *SSRN Preprint*, Nov. 2024. doi: 10.2139/ssrn.5075196.
- [11] A. P. Sahoo, "AI-driven dynamic pricing and revenue optimization in e-commerce: Price elasticity modelling and intelligent pricing systems," *World J. Adv. Res. Rev. (WJARR)*, vol. 26, no. 2, pp. 4194-4201, 2025. doi: 10.30574/wjarr.2025.26.2.2070.
- [12] W. Zhang and Y. Meng, "Dynamic modeling and forecasting of price elasticity based on time series analysis and machine learning," *J. Organ. Intell.*, 2025. doi: 10.1177/14727978251338001.
- [13] C. Li and X. Zhu, "Quantitative analysis of price elasticity in e-commerce using machine learning: A case study of JD.com transaction-level data," in *Proc. 2024 5th Int. Conf. Big Data Economy Inf. Manag. (BDEIM)*, 2024. doi: 10.1145/3724154.3724206.
- [14] J. Wang, W. K. Chong, J. Lin, and C. P. T. Hedenstierna, "Retail demand forecasting using spatial-temporal gradient boosting methods," *J. Comput. Inf. Syst.*, vol. 64, no. 5, pp. 652-664, 2024. doi: 10.1080/08874417.zation.2240753.
- [15] R. S. Ahmed, M. Hasnain, M. H. Mahmood, and M. A. Mehmood, "Comparative analysis of machine learning models for retail demand forecasting," *J. Comput. Sci. Technol. Stud.*, vol. 6, no. 3, 2024.

- Available: <https://al-kindipublishers.org/index.php/jcsts>.
- [16] A. Balusani, P. J. R. Chowdary, and B. V. N. P. Paruchuri, "Enhancing retail demand forecasting with XGBoost: A comparative study with machine learning and deep learning models," SSRN Preprint, May 2025. doi: 10.2139/ssrn.5274510.
- [17] D. D. Purwanto, G. Nathaniel, J. Yulius, and C. P. Haryantho, "Optimizing stock prediction in supermarket: A comparative analysis of LightGBM and XGBoost for enhanced inventory management," *J. Nas. Pendidik. Tek. Inform. (JANAPATI)*, vol. 14, no. 3, 2025. doi: 10.23887/janapati.v14i3.94433.
- [18] N. Nassibi, A. Aberdan, and H. Ouajji, "Demand forecasting models for food industry by utilizing machine learning approaches," *Int. J. Adv. Comput. Sci. Appl. (IJACSA)*, vol. 14, no. 3, zation.
- [19] E. Marini, A. Galli, and D. Bonacorsi, "Comparative analysis of modern machine learning models for retail sales forecasting," arXiv:2506.05941, Jun. 2025. Available: <https://arxiv.org/abs/2506.05941>.
- [20] A. Chougule, S. Patil, and R. Kumar, "AI and machine learning in demand forecasting for supply chain management," *Issues Inf. Syst.*, vol. 26, no. 4, pp. 65-77, 2025. Available: https://iacis.org/iis/2025/4_iis_2025_65-77.pdf.
- [21] M. A. Islam, M. H. Rahman, and K. Hossain, "Hybrid ensemble learning for retail demand forecasting incorporating gradient boosting, AdaBoost, Random Forest, XGBoost, and Linear Regression," *J. Adv. Comput. Sci. Appl.*, 2024.
- [22] I. Stankevich, N. Karpushkin, and M. Ovsyannikov, "Demand forecasting in a multi-channel retail company using machine learning approaches," in *Proc. IEEE Int. Conf. Big Data Anal. (ICBDA)*, zation.
- [23] A. U. Husna, S. H. Amin, and A. Ghasempoor, "Demand forecasting using machine learning and deep learning approaches in the retail industry: A comparative study," in *Industrial Engineering in the Covid-19 Era*, Springer, zation, pp. 249-264.
- [24] H. Abou-Bakr, Z. Benmamoun, and H. Hachimi, "Revealing the dynamics of demand forecasting in supply chain management: A holistic investigation," *Cogent Eng.*, 2024. doi: 10.1080/23311916.2024.2368104.
- [25] R. S. Ahmed, M. Hasnain, M. H. Mahmood, and M. A. Mehmood, "Comparison of deep learning algorithms for retail sales forecasting," *ICCK Trans. Intell. Syst.*, vol. 1, no. 3, pp. 112-126, Oct. 2024. doi: 10.62762/TIS.2024.300700.
- [26] T. de Castro Moraes, X. M. Yuan, and E. P. Chew, "Hybrid convolutional long short-term memory models for sales forecasting in retail," *J. Forecast.*, vol. 43, no. 5, pp. 1278-1293, 2024. doi: 10.1002/for.3073.
- [27] X. Liu and W. Wang, "Deep time series forecasting models: A comprehensive survey," *Mathematics*, vol. 12, no. 10, p. 1504, 2024. doi: 10.3390/math12101504.
- [28] T. Aydın and F. S. Gürgen, "Deep learning for demand forecasting: A framework incorporating variational mode decomposition and attention mechanism," *Processes*, vol. 13, no. 2, p. 594, 2025. doi: 10.3390/pr13020594.

- [29] B. Doğan and D. Birant, "Applying machine learning in retail demand prediction—A comparison of tree-based ensembles and LSTM-based deep learning," *Appl. Sci.*, vol. 13, no. 19, p. 11112, Oct. 2023. doi: 10.3390/app131911112.
- [30] Y. Bakri, L. Zhang, and M. Wei, "Sales forecasting system using deep learning for retail and e-commerce," *ResearchGate Preprint*, Jul. 2025. Available: <https://www.researchgate.net/publication/397119027>.
- [31] S. Shankar, A. Gupta, and V. Sharma, "Revealing the dynamics of demand forecasting in supply chain management," *Cogent Eng.*, 2024. doi: 10.1080/23311916.2024.2368104.
- [32] C. Garcia, A. Morales, and J. Perez, "Sales forecasting for retail stores using hybrid neural networks and sales-affecting variables," *PeerJ Comput. Sci.*, 2025. doi: 10.7717/peerj-cs.3058.
- [33] N. Caballero, R. Lopez, and S. Martinez, "Machine learning and deep learning models for demand forecasting in supply chain management: A critical review," *Appl. Syst. Innov.*, vol. 7, no. 5, p. 93, 2024. doi: 10.3390/asi7050093.
- [34] T. Nguyen, "Applications of artificial intelligence for demand forecasting," *Oper. Supply Chain Manag.*, vol. 16, no. 4, pp. 424-434, 2023.
- [35] M. Apte, K. Kale, P. Datar, and P. Deshmukh, "Dynamic retail pricing via Q-learning: A reinforcement learning framework for enhanced revenue management," *arXiv:2411.18261*, Nov. 2024. Available: <https://arxiv.org/abs/2411.18261>.
- [36] O. Okunola, C. Ogbu, and A. Nwachukwu, "Applications of reinforcement learning in dynamic pricing models for e-commerce," *World J. Adv. Res. Rev. (WJARR)*, vol. 26, no. 3, pp. 1562-1573, 2025. doi: 10.30574/wjarr.2025.26.3.2319.
- [37] H. Liu, "Reinforcement learning for e-commerce dynamic pricing," in *Proc. 7th Int. Conf. Econ. Manag. Green Dev. (ICEMGD 2024)*, Springer, Singapore, 2024, pp. 1-8. doi: 10.1007/978-981-97-0523-8_183.
- [38] A. V. Vamshi Krishna and E. Aravind, "A hybrid deep learning framework for dynamic pricing," *Glob. J. Eng. Innov. Interdiscip. Res. (GJEIIR)*, vol. 5, no. 5, pp. 114-130, 2025. Available: <https://www.sciencexcel.com/articles/s5qwGLpDw5AVhU93N3lIeh9U2b76NU2oej752zCl.pdf>.
- [39] R. Perumallapalli, "Machine learning algorithms for dynamic pricing optimization in retail," *SSRN Preprint*, May 2025. doi: 10.2139/ssrn.5228515.
- [40] A. K. Kalusivalingam, A. Sharma, N. Patel, and V. Singh, "Optimizing e-commerce revenue: Leveraging reinforcement learning and neural networks for AI-powered dynamic pricing," *Int. J. AI ML*, 2024.
- [41] D. Chen, S. Sain, and K. Guo, "Data mining for the online retail industry: A case study of RFM model-based customer segmentation," *J. Database Mark. Cust. Strategy Manag.*, vol. 19, no. 3, pp. 197-208, 2012. [Dataset: <https://archive.ics.uci.edu/ml/datasets/Online+Retail>.]
- [42] S. Makridakis, E. Spiliotis, and V. Assimakopoulos, "The M5 accuracy competition: Results, findings, and conclusions," *Int. J. Forecast.*, vol. 38, no. 4, pp. 1346-1364. doi: 10.1016/j.ijforecast.2021.11.013.

- [43] K. Hastie, T. Tibshirani, and J. Friedman, *The Elements of Statistical Learning*, 2nd ed. Springer.
- [44] A. E. Hoerl and R. W. Kennard, "Ridge regression: Biased estimation for nonorthogonal problems," *Technometrics*, vol. 12, no. 1, pp. 55-67, 1970. doi: 10.1080/00401706.1970.10488634.
- [45] F. Abadi et al., "TensorFlow: A system for large-scale machine learning," in Proc. 12th USENIX Symp. Oper. Syst. Des. Implement. (OSDI), Savannah, GA, USA, 2016, pp. 265-283.
- [46] P. Fildes, S. Ma, and S. Kolassa, "Retail forecasting: Research and practice," *Int. J. Forecast.*, vol. 38, no. 4, pp. 1283-1318. doi: 10.1016/j.ijforecast.2019.06.004.
- [47] M. Loukili, F. Messaoudi, and M. El Ghazi, "Machine learning-driven dynamic pricing strategies in e-commerce," in Proc. 14th Int. Conf. Inf. Commun. Syst. (ICICS), zation. doi: 10.1109/ICICS60529.zation.10330541.
- [48] G. W. Corder and D. I. Foreman, *Nonparametric Statistics for Non-Statisticians*. John Wiley & Sons.
- [49] D. J. Hand, H. Mannila, and P. Smyth, *Principles of Data Mining*. MIT Press, 2001.
- [50] S. Thomassey and A. Zeng, Eds., *Artificial Intelligence for Fashion Industry in the Big Data Era*. Singapore: Springer.
- [51] R. El Youbi, F. Messaoudi, and M. Loukili, "Machine learning, deep neural network and natural language processing based recommendation system," in Int. Conf. Adv. Intell. Syst. Sustain. Dev. (AI2SD), Springer, Cham, zation, pp. 7-17. doi: 10.1007/978-3-031-26384-2_7.
- [52] K. Safonov, "Neural network approach to demand estimation and dynamic pricing in retail," arXiv:2412.00920, Dec. 2024. Available: <https://arxiv.org/abs/2412.00920>.
- [53] M. Haque, M. S. Amin, and J. Miah, "Retail demand forecasting: A comparative study for multivariate time series," arXiv:2308.11939, Aug. zation. Available: <https://arxiv.org/abs/2308.11939>.
- [54] B. Tadesse, D. Mekonnen, and A. Bekele, "Performance comparison of machine learning and deep learning models for demand forecasting in supply chains," *J. Intell. Prod. Des. (JIPD)*, 2024. Available: <https://systems.enpress-publisher.com/index.php/jipd>.
- [55] H. Salazar, J. Baez, and M. Ruiz, "AI and dynamic pricing in e-commerce: Strategies for revenue optimization," *Eur. Econ. Lett.*, vol. 14, no. 4, pp. 347-356, 2024. Available: <https://eelet.org.uk/index.php/journal>.
- [56] H. Guo, N. Tang, and W. Li, "Leveraging variational autoencoders and recurrent neural networks for demand forecasting in supply chain management," *J. Intell. Prod. Des. (JIPD)*, 2024. doi: 10.57064/jipd.v8i3.6639.
- [57] T. Aydin, "Machine learning algorithms in intermittent demand forecasting: A review," *Int. J. Prod. Res.*, 2025. doi: 10.1080/00207543.2025.2578701.
- [58] N. Caballero, S. Martin, and R. Lopez, "A hybrid sentiment-aware machine learning framework for real-time dynamic pricing in e-commerce," *Int. J. Intell. Data Mach. Learn. (IJIDML)*, 2025. doi: 10.55640/ijidml-05-01-12.