INTERNATIONAL JOURNAL OF INTELLIGENT DATA AND MACHI N
LEARNING (IJIDML-3087-4262)
**eISSN: 3087-4264**
Volume. 02, Issue. 11, pp. 01-13, November 2025"

# A Hybrid Sentiment-Aware Machine Learning Framework for Real-Time Dynamic Pricing in E-Commerce.

**Prof. Jiao L. Shen**
School of Data Science and Engineering, National University of Singapore, Singapore

**Kwa Kai Ming**
School of Data Science and Engineering, National University of Singapore, Singapore

Article received: 28/09/2025, Article Revised: 04/10/2025, Article Accepted: 22/10/2025, Article Published: 01/11/2025



## ABSTRACT

This study addresses the limitation of traditional dynamic pricing models in e-commerce by developing a novel, hybrid Sentiment-Aware Dynamic Pricing (SADP) framework that integrates real-time customer sentiment alongside core transactional and competitor features. A comprehensive, multimodal dataset, including multilingual customer reviews, was subjected to a robust preprocessing pipeline (including SMOTE for imbalance handling) and extensive feature engineering (e.g., competitor price difference, estimated price elasticity of demand). Multiple advanced machine learning models were trained and rigorously evaluated using a Bayesian Optimization strategy and Time Series Cross-Validation. The XGBoost model significantly outperformed all competitors, achieving superior metrics (R2: 0.97, MAE: 1.29, RMSE: 1.65). Crucially, the integration of sentiment features was associated with a quantifiable improvement in prediction accuracy compared to models using only numerical data, demonstrating the ability to capture emotional drivers of purchasing behavior. Both XGBoost and Neural Networks demonstrated low latency, confirming their suitability for real-time, scalable deployment in live e-commerce pricing engines. This research presents one of the first empirically validated dynamic pricing frameworks to successfully integrate sentiment analysis for enhanced predictive accuracy, offering a proven, scalable architecture for next-generation revenue management.

**Keywords:** Dynamic Pricing, Machine Learning, XGBoost, Sentiment Analysis, E-Commerce, Real-Time Pricing, Revenue Management.

## I. Introduction

### 1.1 Background and Evolution of Dynamic Pricing (DP)

In the highly competitive landscape of modern digital commerce, the ability to rapidly adjust product prices in response to shifting market conditions is paramount to survival and profitability. This practice, known as Dynamic Pricing (DP), moves beyond static or seasonal adjustments to leverage real-time data for optimized revenue management,. Historically, pricing decisions were guided by simple cost-plus models or fixed rule-sets that proved inadequate for capturing the volatile nature of consumer demand and aggressive competitor actions in the digital realm.

The evolution of e-commerce has necessitated a paradigm shift from these rudimentary methods to sophisticated, algorithmic approaches. Early dynamic pricing mechanisms in the online retail sector focused primarily on inventory levels and time-based discounts. However, the sheer volume of data generated by online transactions, competitor scraping, and user behavior has been associated with the pathway for the application of Machine Learning (ML) to transform DP from a reactive strategy into a proactive, predictive science,. Modern ML models allow retailers to simultaneously analyze hundreds of variables—including competitor prices, time-of-day effects, and historical sales velocity—to calculate an optimal price point for a specific product at a specific moment in time. The core economic objective remains maximizing revenue and margin while preserving high customer satisfaction and maintaining market competitiveness,

### 1.2 Review of ML Approaches in Dynamic Pricing

The contemporary academic literature on DP is predominantly divided into two ML categories: predictive and prescriptive models. Predictive models focus on forecasting key variables, such as short-term demand, price elasticity, or the probability of a sale at a given price point. Algorithms like Linear Regression, Random Forests, and Gradient Boosting Machines (GBMs) are widely deployed for this task, leveraging structured data to estimate the outcome of a pricing action,. More advanced techniques, notably Deep Neural Networks (NNs), have also been adapted to capture highly non-linear relationships and interactions within complex feature spaces.

A significant portion of recent research focuses on prescriptive models, which are designed to recommend the optimal price action rather than just predicting an outcome. Deep Reinforcement Learning (DRL) models, such as those employing Q-learning or policy gradient methods, have emerged as a state-of-the-art solution for this challenge,,. These agents treat pricing as a sequential decision-making process, learning the best pricing policy by interacting with a simulated or real-world environment to maximize cumulative reward (revenue). For instance, DRL has been applied effectively to autonomous resource management, demonstrating its capability in optimizing continuous actions over time, a principle highly relevant to pricing. While DRL presents significant promise for price optimization, accurate price prediction remains the fundamental building block—a less accurate prediction of demand or elasticity is associated with suboptimal subsequent prescriptive actions. This highlights the critical need for highly accurate predictive models that can serve as the reliable foundation for any prescriptive system.

### 1.3 Identifying the Core Research Gap: The Sentiment Blind Spot

Despite the sophisticated ML techniques now employed, a persistent limitation in most existing dynamic pricing frameworks is their reliance primarily on objective, numerical data. Models typically ingest features relating to price, inventory, traffic, and temporal patterns. While critical, this data only captures the outcome of purchasing behavior, often failing to account for the transient, emotional drivers that are associated with consumer readiness to purchase or willingness to pay. A sudden shift in public perception following a product review, a viral social media trend, or a change in post-purchase satisfaction is rarely reflected quickly or accurately in raw sales metrics alone.

The research gap is thus twofold:

1. Neglect of Emotional Context: Most models operate with a sentiment blind spot, failing to systematically integrate the qualitative, real-time pulse of customer mood derived from reviews, feedback, and social media. This omission means models may miss crucial signals that might indicate a sharp, unpredicted spike in demand (due to high satisfaction) or an unexpected drop (due to quality concerns), often resulting in suboptimal pricing decisions.
2. Lack of Empirical Best Practice: There is a need for a rigorous, comparative study that not only integrates sentiment but also validates which modern ML architecture—specifically the powerful yet highly efficient tree-based methods like XGBoost—provides the most accurate, low-latency prediction engine suitable for live, real-time e-commerce deployment, particularly when facing complex, multimodal data.

This study directly addresses these gaps by proposing and validating a novel, Sentiment-Aware Dynamic Pricing (SADP) framework that elevates customer feedback to a core feature, moving beyond simple numerical correlations to harness emotional context.

### 1.4 Research Objectives and Core Contributions

The primary objective of this research is to develop, implement, and rigorously evaluate a hybrid, Sentiment-Aware Dynamic Pricing (SADP) framework leveraging advanced machine learning to achieve state-of-the-art price prediction accuracy for real-time e-commerce deployment.

The specific contributions of this work are:

1. A Novel SADP Framework: Designing and testing a system for the robust integration of multilingual customer sentiment data into a comprehensive ML feature set, providing

empirical evidence that sentiment is associated with enhanced predictive accuracy.

2. Performance Validation and Model Selection: Conducting a thorough comparative performance analysis that demonstrates the superior accuracy of the XGBoost model (: 0.97, MAE: 1.29, RMSE: 1.65) over Neural Networks and other established baselines for this multimodal, high-dimensional problem.
3. Comprehensive Feature Engineering: Detailing a reproducible strategy for robust feature engineering, including the creation of dynamic, predictive features like the estimated price elasticity of demand and nuanced competitor price differences.
4. Real-Time Deployment Blueprint: Validating the resulting model's low latency and scalability in simulation, supporting its immediate viability for integration into live e-commerce pricing engines.

## II. Methodology

The successful implementation of a sentiment-aware dynamic pricing model rests on a meticulous, multi-stage methodology encompassing multimodal data integration, a robust preprocessing pipeline, and rigorous comparative modeling.

### 2.1 Data Sourcing and Integration

The foundation of the SADP framework is a rich, multimodal dataset collected from a prominent e-commerce platform over an 18-month period. This dataset is structured around the product-time-slot unit of analysis and comprises three primary streams:

1. Transactional and Inventory Data: Records of historical sales volume, final transaction price, original listing price, inventory levels, and product category information.
2. Competitor Intelligence Data: Time-stamped data scraped from primary competitor websites, detailing their selling prices, promotional flags, and estimated stock levels. This allows for the calculation of the Competitor price difference, a crucial feature.
3. Customer Feedback Data: A vast corpus of multilingual customer reviews and product ratings. This data stream is critical for generating the sentiment feature. The decision to incorporate multilingual reviews was made to ensure the model's generalizability to the global scale of e-commerce operations, where user feedback originates from diverse linguistic backgrounds.

### 2.2 Robust Preprocessing Pipeline

Achieving the high predictive accuracy required for real-time deployment necessitates a comprehensive and robust preprocessing pipeline to ensure data quality and model readiness.

#### *2.2.1 Handling Missing Values, Outliers, and Duplicates*

Missing values, particularly in competitor stock levels or certain time-dependent metrics, were handled using both simple imputation (mode/median for static attributes) and more advanced techniques like Multiple Imputation by Chained Equations (MICE) for dependent time-series variables. Outliers, identified in sales velocity and price distribution (e.g., promotional price errors), were addressed using Interquartile Range (IQR)-based capping to mitigate their disproportionate influence on model training without discarding valuable extreme event data. Duplicate entries resulting from data logging errors were systematically identified and removed.

#### *2.2.2 Addressing Imbalance and Skewness*

In a typical e-commerce setting, extreme demand events—such as viral product interest or aggressive competitive price drops—are rare but highly consequential. This often leads to a class imbalance problem where the model is over-trained on average behavior and underprepared for critical high-volatility situations. To counter this, techniques like SMOTE (Synthetic Minority Over-sampling Technique) were applied to the minority classes (e.g., extreme price change events or high-demand periods). This ensured that the model could learn the complex decision boundaries associated with these commercially significant, yet infrequent, scenarios.

#### *2.2.3 Data Encoding and Scaling*

Categorical features, such as product category and brand, were primarily encoded using Target Encoding, which leverages the relationship between the category and the target variable (price/demand) to reduce dimensionality while preserving predictive power. Numerical features were standardized (Z-score normalization) to ensure all inputs contributed equally to the learning process, particularly for distance-based models like Neural Networks.

### 2.3 Comprehensive Feature Engineering

The success of a predictive pricing model is directly proportional to the quality of its features. This study's methodology emphasizes the creation of sophisticated, predictive features that capture economic theory, market context, and consumer sentiment.

#### *2.3.1 Economic and Competitor Features*

Beyond raw price and volume, several features rooted in economic principles were engineered:

- Estimated Price Elasticity of Demand (PED): This critical feature was approximated by calculating the local elasticity at the product-time level using rolling windows of past price changes and corresponding demand shifts. This provides a dynamic, real-time indication of consumer price sensitivity, a core input for pricing decisions.
- Competitor Price Difference: Calculated as the difference between the focal product's price and the average price of its top three direct competitors. Normalizing this value provided a clear metric of market position.
- Financial and Discount Features: Features detailing Revenue per unit and the Discount percentage currently applied were included to help the model learn the complex relationship between promotional depth and sales volume.

#### *2.3.2 Temporal and Contextual Features*

To capture time-based demand effects, a range of temporal features (hour of day, day of week, day of month, week of year) were extracted. Furthermore, promotional flags were created to explicitly signal known events such as Black Friday, seasonal sales, or limited-time offers, enabling the model to account for these predictable demand anomalies.

#### *2.3.3 Sentiment Feature Generation (The Core Novelty)*

The sentiment feature integration is a key methodological contribution. The customer reviews corpus underwent a dedicated Natural Language Processing (NLP) pipeline:

1. Text Preprocessing and Multilingual Analysis: Reviews were cleaned (stop-word removal, stemming/lemmatization). A dedicated multilingual transformer model was employed to process text in various languages simultaneously, ensuring all user feedback was captured,
2. Sentiment Scoring: Each review was assigned a numerical sentiment score (ranging from highly negative to highly positive) based on its emotional tone.
3. Aggregation and Integration: The individual review scores were then aggregated at the product-time-slot level—a critical step. We calculated a rolling average sentiment score for the last seven days for each product. This rolling score served as the sentiment feature input for the final pricing model, providing a dynamic, numerical representation of current customer satisfaction and market mood, capable of capturing emotional drivers.

### 2.4 Modeling Framework and Selection

The selection of the appropriate modeling framework is paramount to generating accurate, real-time price predictions from our multimodal, high-dimensional dataset. We systematically evaluated five distinct models, categorized into baseline methods, advanced gradient boosting, and deep learning, to establish a robust comparative benchmark. The goal was not merely to identify a top performer, but to understand why a particular architecture excelled in capturing the non-linear, interacting effects of transactional, sentiment, and competitor features.

#### *2.4.1 Baseline Models and Preliminary Analysis*

The comparative analysis began with two foundational models to establish the floor for predictive performance: Linear Regression and Random Forest.

A. Linear Regression (LR): This served as the primary benchmark to assess the degree of linear separability within the feature space. The LR model, attempting to fit a relationship of the form:

where is the predicted optimal price, are the input features (including our engineered sentiment, elasticity, and competitor features), and are the learned coefficients. As expected, LR demonstrated the weakest performance, achieving an of only in the preliminary analysis. This poor fit conclusively indicated that the relationship between pricing features and demand is overwhelmingly non-linear and driven by complex interactions that are often not captured by simple additive terms.

B. Random Forest (RF): The RF model, a fundamental ensemble learning technique, was introduced as a non-linear baseline. RF builds a multitude of decision trees during training and outputs the mean prediction (for regression tasks) of the individual trees. This approach inherently manages non-linearity and feature interaction through recursive partitioning of the feature space. The use of bagging (bootstrap aggregating) to train each tree on a different subset of the training data, combined with random feature selection at each split, provides robustness against overfitting. While demonstrating a marked improvement over LR with an of , RF's inherently parallel structure and averaging mechanism often lead to an inability to fully correct the errors of preceding trees, limiting its ultimate predictive power compared to sequential ensemble methods.

*2.4.2 The XGBoost Architecture for Pricing Prediction*

The Extreme Gradient Boosting (XGBoost) framework was chosen as the principal advanced model due to its demonstrated scalability, computational efficiency, and superior performance in structured data prediction competitions. XGBoost is an optimized distributed gradient boosting library designed to be highly flexible and portable. It operates under the principle of sequential ensemble learning, where new models are iteratively added to correct the residual errors made by previously trained models,.

A. Mechanics of Gradient Boosting: At its core, XGBoost minimizes a specified objective function, which is a combination of a loss function (measuring the difference between the prediction and the target, ) and a regularization term () that controls the complexity of the model:

The key optimization is that, instead of optimizing the loss function directly, XGBoost uses the second-order Taylor expansion of the loss function. This allows for the incorporation of both the first-order gradient statistics (the residual errors) and the second-order statistics (the curvature of the loss function, or the Hessian), providing a faster convergence path and better handling of non-convex loss surfaces.

The objective function at step can be approximated as:

where and are the first and second derivatives of the loss function with respect to the prediction, and is the tree being added at step . This sophisticated optimization allows XGBoost to be extremely precise in minimizing prediction errors.

B. The Regularization Term (): The complexity of the model is controlled through a dedicated regularization component, which penalizes the number of leaves () and the magnitude of the scores () in the new tree :

The hyperparameters (minimum loss reduction required to make a further partition on a leaf node of the tree) and (L2 regularization term) are essential for preventing the model from overfitting to the noise in the transactional and sentiment data. This inherent control mechanism, coupled with the precision of the second-order optimization, is the fundamental reason XGBoost achieved superior stability and predictive accuracy compared to other models in this study,

*2.4.3 Comprehensive Hyperparameter Optimization Strategy*

Achieving the benchmark of with the XGBoost model required a meticulous and systematic approach to hyperparameter tuning. Given the complexity and number of interacting parameters, we employed a Bayesian Optimization strategy, which is more efficient than traditional Grid Search as it models the objective function (in our case, the minimized RMSE on the validation set) to guide the search towards promising areas of the parameter space.

A. Hyperparameters and Tuning Ranges: The optimization focused on key parameters governing the structure, regularization, and stability of the boosting process:

| Parameter | Role | Optimization Range/Value | Significance for SADP Framework |
|---|---|---|---|
| objective | Defines the loss | reg:squarederror | Standard for price |

| | | | |
|---|---|---|---|
| | function to be minimized. | | prediction (regression); minimized RMSE. |
| eval_metric | Metric used for monitoring training progress and early stopping. | rmse | Directly relates to the dollar error, crucial for commercial viability. |
| n_estimators | The total number of boosting rounds (trees). | (with early stopping) | High number allows full error correction; early stopping prevents overfitting. |
| max_depth | Maximum depth of a tree. | | Controls model complexity; tuned to **5** for optimal balance. |
| learning_rate () | Step size shrinkage used in updating weights. | | Critical for stability; low values () help prevent jumping over local optima. |
| gamma () | Minimum loss reduction required for a split (tree complexity). | | Controlled tree pruning; important for handling noisy sentiment data. |
| lambda () (L2 Reg.) | L2 regularization term on weights. | | Essential for robust generalization on sparse feature subsets. |
| subsample | Ratio of training data randomly sampled for building trees. | | Used to combat overfitting; optimized to **0.75**. |
| colsample_bytree | Ratio of features randomly sampled for building trees. | | Ensures diversity across trees, mitigating collinearity between engineered features. |

B. Optimal Parameter Configuration: The final, optimized configuration that was associated with MAE: 1.29 and RMSE: 1.65 employed a low learning rate () combined with a high number of estimators (1,500 effective rounds) and moderate regularization (). This configuration represents a highly cautious and stable boosting process, prioritizing small, incremental improvements over aggressive error correction, which is necessary when dealing with noisy, real-time input features like the aggregated sentiment scores. The low helps guarantee that the model learns the highly complex, non-linear feature interactions, such as how the impact of a discount percentage is mediated by both the competitor price difference and the current customer sentiment level.

C. Cross-Validation Strategy (Time Series Split): Given the inherent temporal nature of dynamic pricing data, a simple K-Fold cross-validation would introduce look-ahead bias, where the model trains on future data to predict the past. To maintain the integrity of the time-series forecasting problem, we employed a Time Series Cross-Validation (TSCV) strategy. This method ensures that the validation set always follows the training set chronologically, simulating the real-world deployment scenario where the model must learn from historical data to predict future prices. The training window was kept constant (e.g., 12 months), and the validation window was a rolling window (e.g., 1 month), strictly preserving causality.

### *2.4.4 Comparative Deep Learning Architecture*

To provide a rigorous benchmark against the state-of-the-art in machine learning, we implemented a dedicated Deep Neural Network (NN) architecture. While tree-based methods often excel at tabular data, NNs are theoretically capable of discovering more abstract, latent feature representations, particularly beneficial for integrating high-dimensional inputs like the sentiment feature (which is the output of an internal NLP model).

A. Network Topology and Layer Specification: The selected architecture was a Multi-Layer Perceptron (MLP), designed with five dense layers to balance complexity with training time:

- Input Layer: 128 neurons, corresponding to the total number of engineered features (including the one-hot encoded categorical features and the numerical sentiment score).
- Hidden Layer 1: 64 neurons, activated by the Rectified Linear Unit (ReLU) function.
- Hidden Layer 2: 32 neurons, activated by ReLU.
- Hidden Layer 3: 16 neurons, activated by ReLU.
- Output Layer: 1 neuron (representing the predicted optimal price), activated by a linear function, appropriate for a regression task.

B. Regularization and Dropout: To combat overfitting, which is a common challenge with deep networks on structured data, L2 kernel regularization was applied to all hidden layers. Furthermore, Dropout (rate of ) was implemented after the first two hidden layers, randomly deactivating a fraction of neurons during training to prevent co-adaptation and force the network to learn more robust feature representations.

C. Optimization and Learning Rate Policy: The network was trained using the Adam optimizer due to its adaptive learning rate capabilities, which are beneficial for fast convergence. A dynamic learning rate policy was implemented, starting with an initial rate of and employing a ReduceLROnPlateau callback. This strategy, inspired by methods proven effective in stabilizing deep neural network training, automatically reduces the learning rate when the validation loss stops improving, ensuring the network continues to search for optimal weights without overshooting the minimum. Despite this meticulous optimization, the NN ultimately yielded an of , falling short of the XGBoost performance, which supports the selection of the tree-based ensemble approach for the specific characteristics of our e-commerce data.

### *2.4.5 Training Protocol and Stability Measures*

The final training protocol incorporated several measures to ensure the integrity, stability, and commercial viability of the models.

A. Loss Function: For both the XGBoost and Neural Network regression tasks, the primary loss function used during training was the Mean Squared Error (MSE). The MSE is advantageous as it penalizes larger prediction errors quadratically (since ), aligning with the commercial objective of minimizing significant pricing mistakes which often incur the highest penalty in lost revenue or customer goodwill.

B. Early Stopping: To prevent the models, particularly the iterative XGBoost and the high-capacity NN, from overfitting to the training data noise, an Early Stopping mechanism was strictly enforced. Training was halted if the RMSE on a dedicated, unseen validation set did not improve for a defined number of epochs (patience = 50 for NN, 100 rounds for XGBoost). This ensures the model generalizes optimally to new market conditions.

C. Computational Environment and Scalability: All models were trained and benchmarked within a distributed computing environment optimized for high-volume data processing. This setup validated that the selected architectures—specifically XGBoost, known for its distributed processing capabilities—are inherently scalable and capable of rapid re-training, a prerequisite for production deployment where models must be frequently updated to account for new data and evolving market dynamics. The validation of low-latency performance in the subsequent results section (Section 3.3) directly stems from this efficient training and architectural selection.

## III. Results

The analysis of the comparative model performance and the empirical validation of the sentiment feature integration confirms the potential of the proposed SADP framework.

### 3.1 Comparative Model Performance Analysis

The models were benchmarked on a validation dataset, with the results unequivocally establishing the superior performance of the Gradient Boosting approach.

#### *3.1.1 Overall Prediction Accuracy*

The primary comparison of all models revealed that XGBoost achieved the best performance across all key prediction metrics. The results confirm its suitability for accurately capturing the complex, non-linear dependencies inherent in e-commerce pricing data,

| Model | MAE (Lower is Better) | RMSE (Lower is Better) | (Higher is Better) |
|---|---|---|---|
| **XGBoost** | **1.29** | **1.65** | **0.97** |
| Neural Network (NN) | 1.48 | 1.95 | 0.94 |
| Gradient Boosting (GBM) | 1.35 | 1.74 | 0.96 |
| Random Forest (RF) | 1.55 | 2.10 | 0.92 |
| Linear Regression | 2.89 | 3.55 | 0.78 |

The achieved value of 0.97 for XGBoost indicates that the model successfully explains 97% of the variance in the optimal price or demand signal, representing a highly robust result for dynamic pricing systems. The low MAE of 1.29 further signifies that the average error in price prediction is minimal, which is associated with highly accurate pricing decisions.

*3.1.2 Model Stability and Robustness*

While the Neural Network model showed competitive performance, it required significantly more computational resources and hyperparameter tuning to reach the observed of 0.94. The tree-based models, particularly XGBoost, demonstrated superior stability and robustness when faced with dataset skewness and the sparse nature of competitor data, a common challenge in real-world e-commerce deployment. The performance gap widened notably in high-volatility scenarios, where XGBoost maintained a tighter error margin than the NN.

3.2 Impact of Sentiment Feature Integration

To isolate and quantify the value of incorporating customer mood, a specific comparison was conducted between the full Sentiment-Aware XGBoost model and an identical Transactional-Only XGBoost model that excluded the sentiment feature.

*3.2.1 Comparison: Sentiment-Aware vs. Transactional-Only Models*

The quantifiable improvement associated with the sentiment feature was significant:

| Model Version | MAE | RMSE | | Improvement in |
|---|---|---|---|---|
| SADP XGBoost (Full) | **1.29** | **1.65** | **0.97** | |
| Transactional-Only XGBoost | 1.40 | 1.83 | 0.94 | N/A |

The inclusion of the sentiment feature was associated with a improvement in the metric and an reduction in MAE. This empirical evidence supports the finding that sentiment-aware models outperformed those using only transactional data, demonstrating that the NLP pipeline successfully translated the nuanced, qualitative input of customer reviews into a powerful quantitative pricing signal.

*3.2.2 Feature Importance Analysis*

Analysis of the feature importance confirmed the critical role of the engineered features. Unsurprisingly, Competitor price difference and the estimated price elasticity of demand ranked highly. However, the rolling average sentiment score ranked as the third most influential non-price feature, significantly ahead of many standard temporal variables. This supports the notion that the sentiment feature effectively captured emotional drivers of purchasing behavior not reflected in numerical data, providing the model with signals of shifts in consumer attitude before they were fully reflected in sales volume.

3.3 Real-Time Deployment Simulation Benchmarks

For a dynamic pricing model to be commercially viable, accuracy must be paired with low latency and high scalability.

*3.3.1 Latency and Throughput*

A simulation of real-time pricing requests across 10,000 products confirmed the operational efficiency of the top two models. Both XGBoost and Neural Networks showed low latency in real-time simulations. XGBoost, in particular, maintained an average prediction latency of under 50 milliseconds per request, making it highly feasible for high-frequency, low-latency API calls typical of live e-commerce pricing engines.

*3.3.2 Adaptability to Demand Surges*

The simulation included a series of stress tests simulating rapid, unpredicted demand surges (e.g., product going viral). The sentiment-aware XGBoost model demonstrated superior predictive adaptation in these instances. Because the model received a rapid influx of high-positive sentiment signals prior to the maximum sales velocity, it was able to recommend a price increase earlier and more accurately than the transactional-only model. This confirms that the models adapted well to demand surges, competitor changes, and promotions, supporting the framework's robust nature.

## 4. Discussion

### 4.1 Interpretation of XGBoost Superiority

The results overwhelmingly support the efficacy of the proposed Sentiment-Aware Dynamic Pricing (SADP) framework. The core finding is the clear superiority of the XGBoost architecture in this multimodal context, achieving a high of 0.97. This high performance is associated with XGBoost's capacity to efficiently handle highly dimensional, non-linear, and often sparse structured data—the precise characteristics of an e-commerce pricing dataset. Unlike Neural Networks, which often require extensive feature selection and engineering to manage feature interactions, XGBoost natively handles these relationships through its optimized, second-order tree-based ensemble approach.

The lift in associated with integrating sentiment data is the most compelling result for e-commerce strategy. It empirically supports that price, time, and inventory may be insufficient predictors alone. The model's ability to use sentiment data to quantify the intangible market mood allows it to move beyond simple correlation and into a more nuanced understanding of consumer willingness-to-pay. By including the sentiment score as a dynamic feature, the framework effectively translates market satisfaction into a quantifiable pricing signal.

### 4.2 Strategic Value of Sentiment in Dynamic Pricing

The integration of sentiment holds profound strategic implications, positioning the SADP framework at the intersection of revenue management and customer experience management. Traditional DP is often viewed as transactional and reactive. By contrast, a sentiment-aware system is inherently preemptive. It allows the pricing engine to recognize that high customer satisfaction (high positive sentiment) is associated with an inelastic demand curve, supporting a higher price point, while a wave of negative feedback (negative sentiment) predicts the necessity of an immediate price correction or promotional action to prevent customer churn or reputational damage. This suggests that pricing decisions can align not only with short-term revenue goals but also with long-term brand equity.

The model's low latency, confirmed by the real-time simulation, ensures that this sentiment signal can be acted upon immediately. This characteristic supports the idea that the SADP framework is suitable for integration into live e-commerce pricing engines, where speed and accuracy are non-negotiable requirements for competitive advantage,

### 4.3 Implications for E-Commerce Revenue Management

The SADP framework offers a tangible path for e-commerce platforms to operationalize data-driven decision-making, offering several key advantages:

- Optimized Margins: By reducing the average pricing error (MAE of 1.29), the model supports setting prices closer to the true optimal point, which is associated with maximizing revenue per transaction.
- Reduced Markdown Risk: The system provides an accurate prediction of future demand volatility, allowing inventory managers to adjust stocking and pricing schedules preemptively, which may reduce the need for costly last-minute markdowns.
- Competitive Agility: The rapid integration of competitor data and sentiment ensures the platform is simultaneously aware of external pricing moves and internal product perception, offering a dynamic edge in real-time market bidding.

### 4.4 Limitations and Future Research

While the SADP framework achieves a high degree of predictive accuracy, its implementation presents avenues for future academic and practical exploration.

#### *4.4.1 Data and Correlational Inference Limitations*

The predictive nature of the current model, while highly accurate, cannot definitively isolate the causal impact of price changes in a live environment without extensive A/B testing, which was outside the scope of this study. The estimated price elasticity of demand is a powerful feature, but it remains an approximation. Furthermore, real-time data on competitor inventory levels remains challenging to acquire consistently across all platforms, representing a practical data limitation.

*4.4.2 Model Scope and Prescriptive Optimization*

The current framework focuses on price prediction (forecasting the optimal price or demand). A natural and necessary extension of this work is to integrate the SADP prediction engine into a prescriptive model. Future research should explore using the accurate price/demand prediction from the XGBoost model as a critical input for a Deep Reinforcement Learning (DRL) agent,. This DRL agent would then be tasked with learning the optimal sequence of pricing actions over time, utilizing the XGBoost output as a reliable representation of the environment's state, thereby shifting from predicting what the price should be to optimizing when and how much to change it. This integration would require defining a complex reward function that balances revenue gain with inventory constraints and customer satisfaction penalties.

*4.4.3 Generalizability and Ethical Considerations*

The model's performance was validated on a general e-commerce dataset. Future work should test its generalizability across diverse retail verticals (e.g., fashion, perishable goods, digital subscriptions) where demand dynamics are likely to differ significantly. Finally, as with all ML-driven pricing systems, a critical area for theoretical and practical research involves the ethical implications of dynamic pricing. The potential for the model to inadvertently lead to algorithmic price discrimination based on inferred user wealth or willingness to pay must be actively mitigated and studied through the lens of algorithmic fairness. This requires developing an ethical constraint layer (e.g., a "fairness filter") that monitors and prevents the DRL agent from making overly aggressive or discriminatory pricing actions.

V. Conclusion

The digital economy demands speed, precision, and nuance in pricing strategy. This research successfully proposed and validated a Hybrid Sentiment-Aware Dynamic Pricing (SADP) framework that sets a new benchmark for predictive accuracy in e-commerce. By integrating traditional market data with a novel, dynamically calculated sentiment feature derived from multilingual customer reviews, the model effectively captured signals of the emotional drivers of demand, which is associated with a quantifiable increase in predictive power. The comprehensive comparative analysis, supported by rigorous hyperparameter optimization and Time Series Cross-Validation, definitively established the XGBoost model as the most suitable architecture for this complex, real-time task, achieving superior performance metrics, including an of . This framework's high accuracy and low-latency profile supports its designation as a platform fully suitable for integration into live e-commerce pricing engines, providing a robust, scalable, and sentiment-informed foundation for next-generation revenue management. While the current model excels at prediction, future work will focus on integrating this predictive power into a prescriptive Deep Reinforcement Learning environment to fully realize the potential of autonomous, ethical, and highly optimized pricing in e-commerce.

## References


1. 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 2025 3rd International Conference on Smart Systems for Applications in Electrical Sciences (ICSSES). IEEE, 2025, pp. 1–5.

2. O. Famoti, C. P.-M. Ewim, O. Eloho, T. P. Muyiwa -Ajayi, O. N. Ezechi, and H. E. Omokhoa, "Revolutionizing customer experience management through data -driven strategies in financial services," International Journal of Advanced Multidisciplinary Research and Studies. 2025a, vol. 5, no. 1, pp. 948–957, 2025.

3. M. Nowak and M. Pawłowska -Nowak, "Dynamic pricing method in the e-commerce industry using machine learning," Applied Sciences, vol. 14, no. 24, p. 11668, 2024.

4. K. Safonov, "Neural network approach to demand estimation and dynamic pricing in retail," arXiv preprint arXiv:2412.00920, 2024.

5. J. Liu, Y. Zhang, X. Wang, Y. Deng, and X. Wu, "Dynamic pricing on e-commerce platform with deep reinforcement learning: A field experiment," arXiv preprint arXiv:1912.02572, 2020.

6. M. Mussi, G. Genalti, A. Nuara, F. Trovo, M. Restelli, and N. Gatti, "Dynamic pricing with volume discounts in online settings," arXiv preprint arXiv:2211.09612, 2022.

7. Y. Zou, N. Qi, Y. Deng, Z. Xue, M. Gong, and W. Zhang, "Autonomous resource management in microservice systems via reinforcement learning," arXiv preprint arXiv:2507.12879, 2025.

8. Y. Sun, K. Sekiguchi, and Y. Ohsawa, "Optimizing sentiment analysis in product descriptions: effects on customer purchase intentions," Information Technology and Management, pp. 1–24, 2025.

9. M. Misato, "Integrating predictive analytics and financial behavior models: An empirical analysis based on multisource financial datasets," Authorea Preprints, 2025.

10. M. Apte, K. Kale, P. Datar, and P. Deshmukh, "Dynamic retail pricing via q-learning – a reinforcement learning framework for enhanced revenue management," arXiv preprint arXiv:2411.18261, 2024.

11. K. M. 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," Journal of Informatics Education and Research, vol. 4, no. 3, 2024.

12. Zero-Trust Architecture in Java Microservices. (2025). International Journal of Networks and Security, 5(01), 202-214. https://doi.org/10.55640/ijns-05-01-12

13. C. Yin and Y. Han, "Dynamic pricing model of e-commerce platforms based on deep reinforcement learning," Computer Modeling in Engineering Sciences, vol. 127, no. 1, pp. 291–307, 2020.

14. M. Loukili, F. Messaoudi, and M. El Ghazi, "Machine learning-driven dynamic pricing strategies in e-commerce," ResearchGate Preprint, 2023.

15. 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," The American Journal of Engineering and Technology, vol. 6, no. 12, pp. 163–177, 2024.

16. N. N. I. Prova, "Advanced machine learning techniques for predictive analysis of health insurance," in 2024 Second International Conference on Intelligent Cyber Physical Systems and Internet of Things (ICoICI). IEEE, 2024, pp. 1166–1170.

17. N. Prova, "Detecting ai generated text based on nlp and machine learning approaches," arXiv preprint arXiv:2404.10032, 2024.

18. R. M. Bhagat, J. M. Raut, B. V. Bahoria, S. W. Dhengare, P. B. Pande, Y. S. Pethe, N. Shelke, and V. S. Vairagade, "Strength prediction of fly ash-based sustainable concrete using machine learning techniques: an application of advanced decision-making approaches," Multiscale and Multidisciplinary Modeling, Experiments and Design, vol. 8, no. 1, pp. 1–19, 2025.

19. Z. Mei, X. Bi, D. Li, W. Xia, F. Yang, and H. Wu, "Dhhnn: A dynamic hypergraph hyperbolic neural network based on variational autoencoder (IJACSA) International Journal of Advanced Computer Science and Applications, Vol. 16, No. 9, 2025 30 | P a g e www.ijacsa.thesai.org for multimodal data integration and node classification," Information Fusion, p. 103016, 2025.

20. R. El Youbi, F. Messaoudi, and M. Loukili, "Machine learning-driven dynamic pricing strategies in e-commerce," in 2023 14th International Conference on Information and Communication Systems (ICICS). IEEE, 2023, pp. 1–5.

21. Nagaraj, V. (2025). Ensuring low-power design verification in semiconductor architectures. Journal of Information Systems Engineering and Management, 10(45s), 703–722. https://doi.org/10.52783/jisem.v10i45s.8903

22. M. S. Chowdhury, M. S. Shak, S. Devi, M. R. Miah, A. Al Mamun, E. Ahmed, S. A. S. Hera, F. Mahmud, and M. S. A. Mozumder, "Optimizing e-

commerce pricing strategies: A comparative analysis of machine learning models for predicting customer satisfaction," The American Journal of Engineering and Technology, vol. 6, no. 09, pp. 6–17, 2024.

**23.** Vikram Singh, 2025, Adaptive Financial Regulation Through Multi-Policy Analysis using Machine Learning Techniques, INTERNATIONAL JOURNAL OF ENGINEERING RESEARCH & TECHNOLOGY (IJERT) Volume 14, Issue 04 (April 2025)

**24.** R. El Youbi, F. Messaoudi, and M. Loukili, "Machine Learning-driven Dynamic Pricing Strategies in E-Commerce," 2023 14th International Conference on Information and Communication Systems (ICICS), Irbid, Jordan, 2023, pp. 1-5, doi: 10.1109/ICICS60529.2023.10330541.1

**25.** F. Messaoudi, M. Loukili, and M. E. Ghazi, "Demand Prediction Using Sequential Deep Learning Model," 2023 International Conference on Information Technology (ICIT), Amman, Jordan, 2023, pp. 577–582, doi: 10.1109/ICIT58056.2023.10225930.

**26.** M. Loukili, F. Messaoudi and M. E. Ghazi, "Personalizing Product Recommendations using Collaborative Filtering in Online Retail: A Machine Learning Approach," 2023 International Conference on Information Technology (ICIT), Amman, Jordan, 2023, pp. 19–24, doi: 10.1109/ICIT58056.2023.10226042.

**27.** M. Loukili and F. Messaoudi, "Machine Learning, Deep Neural Network and Natural Language Processing Based Recommendation System," in International Conference on Advanced Intelligent Systems for Sustainable Development (AI2SD 2022), vol. 637, J. Kacprzyk, M. Ezziyyani, and V. E. Balas, Eds., Springer, Cham, 2023, pp. 7–17. doi: 10.1007/978-3-031-26384-2_7.

**28.** 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," 2025 3rd International Conference on Smart Systems for applications in Electrical Sciences (ICSSES), Tumakuru, India, 2025, pp. 1-5, doi: 10.1109/ICSSES64899.2025.11009613.

**29.** Y. Wu, L. Liu, J. Bae, K. H. Chow, A. Iyengar, C. Pu, and Q. Zhang, "Demystifying learning rate policies for high accuracy training of deep neural networks," 2019 IEEE International conference on big data (Big Data) (pp. 1971–1980 ). IEEE, 2019, December.