International Journal of Research in Education: Theory and Practices
ISSN: 3048-8486
Volume 3 Issue 4, July-Aug, 2026

An Intelligent Machine Learning Framework for Predicting Customer Purchase Decisions Using Classification and Regression Techniques in E-Commerce

S.Srinivas[1], Thumu Pradeepthi[2]
[1]Assistant Professor, Department of MCA, Megha Institute of Engineering and Technology for Women, Edulabad (Village), Ghatkesar (Mandal), Medchal District, Telangana – 501301, India.
[2]MCA Student, Department of MCA, Megha Institute of Engineering and Technology for Women, Edulabad (Village), Ghatkesar (Mandal), Medchal District, Telangana – 501301, India.

Abstract. The rapid expansion of e-commerce platforms has generated enormous volumes of consumer interaction data, creating new opportunities for understanding purchasing behavior through intelligent data analytics. Accurate prediction of customer purchase decisions enables online retailers to improve personalized recommendations, optimize promotional campaigns, and enhance customer satisfaction. This study proposes a machine learning framework for predicting whether a customer will purchase a product after browsing it on an e-commerce platform. The research utilizes a large-scale transaction dataset obtained from JD.com containing customer information, product characteristics, browsing history, pricing details, and promotional activities. Before model construction, comprehensive data preprocessing and feature engineering techniques are performed to improve data quality and create informative predictive variables. The proposed framework evaluates three supervised machine learning algorithms, namely Decision Tree, Random Forest, and Logistic Regression, for binary purchase classification. Hyperparameter optimization is conducted using the Random Search strategy to maximize prediction performance while reducing overfitting. Model effectiveness is evaluated using Accuracy, Precision, Recall, and Area Under the ROC Curve (AUC). Experimental results demonstrate that the optimized Random Forest classifier achieves the highest prediction performance with a testing accuracy of 0.999871 and an AUC of 0.9998, outperforming the remaining models. Feature importance analysis further indicates that coupon discount level and quantity discount level contribute most significantly to customer purchasing decisions. The proposed framework offers a reliable and computationally efficient solution for customer behavior prediction and provides valuable support for personalized recommendation systems, pricing optimization, inventory planning, and intelligent marketing strategies in modern e-commerce platforms.

keywords: Customer Purchase Prediction, Machine Learning, Random Forest, Decision Tree, Logistic Regression, Feature Engineering, Classification, E-Commerce Analytics, Consumer Behavior, Predictive Modeling

I. Introduction

The continuous advancement of digital technologies has transformed the global retail industry by enabling businesses to conduct commercial activities through electronic commerce platforms. Modern e-commerce websites provide consumers with convenient access to a wide variety of products while allowing transactions to be completed anytime and from any location. As online shopping continues to expand, platforms such as JD.com, Amazon, Flipkart, and Alibaba generate massive volumes of customer interaction data, including browsing behavior, purchase history, product preferences, pricing information, and promotional responses. These large-scale datasets provide valuable opportunities for applying machine learning techniques to understand consumer purchasing behavior and support intelligent business decision-making. Predicting whether a customer is likely to purchase a product has therefore become one of the most important research problems in e-commerce analytics.

Customer purchasing decisions are influenced by numerous interconnected factors, including product characteristics, pricing strategies, discount offers, customer demographics, browsing history, previous purchasing experience, seasonal demand, and promotional campaigns. Unlike traditional retail environments, online shopping platforms continuously collect detailed records of every customer interaction, enabling researchers to analyze consumer behavior with greater precision. However, the large volume and diversity of available information also increase the complexity of accurately identifying purchasing intentions. Since multiple variables simultaneously influence buying decisions, conventional statistical approaches often fail to model these complex nonlinear relationships effectively.

Recent developments in artificial intelligence and machine learning have significantly improved the capability of predictive systems used in e-commerce. Supervised learning algorithms can automatically identify hidden relationships within historical customer data and generate highly accurate purchase predictions. These predictive models assist online retailers in improving recommendation systems, optimizing product placement, planning promotional campaigns, reducing inventory costs, and enhancing customer satisfaction. As businesses increasingly rely on data-driven decision-making, machine learning has become an essential component of modern e-commerce platforms.

Several machine learning algorithms have been successfully applied to customer purchase prediction problems. Decision Tree classifiers provide transparent decision-making structures that are easy to interpret while efficiently handling both categorical and numerical variables. Random Forest, an ensemble learning technique, combines multiple decision trees to improve prediction accuracy, reduce variance, and minimize overfitting. Logistic Regression remains one of the most widely used binary classification algorithms because of its computational simplicity and probabilistic interpretation of customer purchasing behavior. Each algorithm offers unique advantages depending on dataset characteristics and business requirements, making

II. Literature Survey

The rapid expansion of e-commerce has transformed the way businesses interact with customers, creating enormous volumes of transactional and behavioral data that can be analyzed to understand purchasing decisions. Online shopping platforms continuously record customer interactions such as product browsing, purchase history, pricing preferences, promotional responses, and demographic information. These data provide an excellent foundation for applying machine learning techniques to predict customer purchasing behavior. Accurate prediction models help e-commerce companies optimize personalized recommendations, improve inventory planning, design targeted promotional campaigns, and enhance customer satisfaction. Consequently, numerous researchers have investigated statistical learning, classification algorithms, and ensemble machine learning techniques to develop intelligent customer purchase prediction systems.

One of the important studies in this area was conducted by Cao, who examined the competitive advantages and operational challenges of modern e-commerce platforms. The research highlighted that the rapid growth of online retail has significantly increased the availability of customer interaction data, enabling organizations to make data-driven business decisions. The study emphasized that understanding customer purchasing behavior has become essential for improving marketing strategies, customer retention, and recommendation systems. Although the research primarily focused on business perspectives rather than predictive modeling, it established the importance of applying advanced analytical techniques to improve commercial decision-making in digital marketplaces.

A comprehensive investigation into consumer purchasing behavior was presented by Yong et al., who analyzed the various factors influencing online shopping decisions. Their research considered demographic characteristics, product quality, pricing strategies, promotional offers, website usability, and consumer trust as major determinants affecting purchasing intention. Statistical analysis revealed that promotional activities and product pricing significantly influence customer decisions when multiple alternatives are available. The authors concluded that integrating multiple customer-related factors provides a more comprehensive understanding of purchasing behavior than analyzing individual variables separately. However, the study relied mainly on traditional statistical methods and did not employ advanced machine learning models capable of handling large-scale transaction data.

To improve customer behavior analysis, Zhang proposed the application of machine learning algorithms for consumer purchasing prediction using historical transaction records. The study demonstrated that supervised learning techniques effectively identify hidden patterns within customer behavior by analyzing browsing history, purchasing frequency, and promotional responses. Experimental evaluation indicated that machine learning algorithms outperform conventional statistical models in capturing nonlinear relationships among multiple variables. The research further suggested that intelligent predictive systems could significantly improve personalized recommendation services and customer segmentation strategies within modern e-commerce platforms.

Large-scale transaction datasets have become increasingly important for developing reliable predictive models. Shen et al. introduced the publicly available JD.com transaction-level dataset for the MSOM Data-Driven Research Challenge. The dataset contains detailed information regarding customer profiles, product characteristics, browsing activities, purchase records, inventory, and logistics information collected during March 2018. This comprehensive dataset has become an important benchmark for evaluating machine learning algorithms in e-commerce research because it represents realistic customer purchasing behavior under various promotional and operational conditions. The availability of such extensive transactional information has enabled researchers to construct highly accurate predictive models for customer purchase classification.

III. Research Methodology

The proposed research methodology presents a structured machine learning framework for predicting customer purchasing decisions using historical transaction data collected from an e-commerce platform. The framework integrates comprehensive data preprocessing, feature engineering, supervised machine learning, hyperparameter optimization, and performance evaluation to develop an accurate purchase prediction model. The primary objective is to determine whether a customer will purchase a product after viewing it by analyzing product characteristics, customer information, browsing behavior, and promotional activities. The complete methodology consists of six major stages: data collection, data preprocessing, feature engineering, model development, hyperparameter optimization, and performance evaluation. Each stage contributes to improving prediction accuracy, reducing model complexity, and enhancing the overall reliability of the classification system.

The first stage involves data collection, where transaction-level data are obtained from the JD.com e-commerce platform through the INFORMS 2020 MSOM Data-Driven Research Challenge dataset. The dataset contains comprehensive information related to customer transactions recorded during March 2018. It consists of several interconnected tables describing product information, customer profiles, browsing activities, purchase records, delivery details, inventory status, and network characteristics. For this research, the most relevant information is extracted from the SKU, Users, Clicks, and Orders tables. These tables provide valuable attributes including product specifications, customer demographics, browsing history, pricing information, and promotional discounts that collectively influence purchasing decisions. Using a large-scale real-world dataset ensures that the proposed prediction model accurately represents customer behavior in practical e-commerce environments.

Following data acquisition, the framework performs data preprocessing to improve dataset quality before model development. Initially, information collected from different tables is integrated using common identifiers such as user_ID and sku_ID to generate a unified transaction dataset. During this process, duplicate records, inconsistent entries, and incomplete observations are identified and appropriately handled. Missing values are treated using different strategies depending on their occurrence. Attributes containing only a small number of missing observations are removed, whereas variables with substantial missing information are estimated using a Random Forest regression-based imputation approach. This strategy predicts missing price and discount values by considering related product characteristics and customer information, thereby preserving important relationships within the dataset. After missing value treatment, categorical variables are converted into numerical representations using one-hot encoding, allowing the machine learning algorithms to process non-numeric information efficiently. Finally, redundant records resulting from repeated browsing activities are eliminated to reduce unnecessary computational complexity and improve model performance.

The next stage focuses on feature engineering, which plays a crucial role in improving prediction accuracy. Instead of relying solely on the original variables, several new features are generated to better represent customer purchasing behavior. Continuous price information is transformed into discrete price levels based on predefined price intervals. Similarly, four promotional variables—direct discount, quantity discount, bundle discount, and coupon discount—are converted into corresponding discount levels according to the percentage discount relative to the original product price. An additional binary variable identifies whether customer browsing occurred on a weekday or weekend. Furthermore, a new dependent variable named "purchased" is constructed by examining customer browsing history and purchase records. If a customer places an order within thirty minutes after viewing a product, the browsing session is labeled as a successful purchase; otherwise, it is classified as a non-purchase. These engineered variables improve the predictive capability of the classification models while reducing the influence of noise and extreme values present in the original dataset.

IV. Existing System

Existing customer purchase prediction systems primarily rely on conventional statistical analysis and supervised machine learning algorithms to estimate whether a customer is likely to purchase a product after browsing it on an e-commerce platform. These systems generally utilize historical transaction records, customer profiles, product information, and browsing behavior to construct predictive models. Commonly employed algorithms include Logistic Regression, Decision Tree, Random Forest, Support

Vector Machine (SVM), Naïve Bayes, and other traditional classification techniques. Although these approaches have demonstrated satisfactory performance for customer behavior analysis, they often face challenges when processing large-scale e-commerce datasets containing complex customer interactions and numerous influencing variables.

Most existing systems begin by collecting customer transaction data from online shopping platforms and performing basic preprocessing operations such as removing duplicate records, handling missing values, and converting categorical variables into numerical representations. After preprocessing, machine learning algorithms are trained using customer demographic information, browsing history, pricing details, and purchase records to classify whether a customer will complete a transaction. While this workflow provides acceptable prediction accuracy, many existing models depend heavily on the quality of manually selected features and often fail to capture the complex relationships among promotional activities, customer preferences, and purchasing behavior.

Traditional classification models such as Logistic Regression assume relatively simple linear relationships between predictor variables and customer purchase decisions. Although Logistic Regression offers fast computation and straightforward interpretation, its predictive capability decreases when customer behavior exhibits highly nonlinear patterns. Similarly, Decision Tree classifiers generate easy-to-understand classification rules but are highly sensitive to variations in training data and frequently suffer from overfitting, resulting in unstable prediction performance when applied to unseen customer transactions. These limitations reduce their effectiveness for large-scale e-commerce applications where customer behavior continuously changes over time.

V. Proposed System

The proposed system presents an intelligent machine learning framework for predicting customer purchasing decisions by integrating comprehensive data preprocessing, advanced feature engineering, hyperparameter optimization, and supervised classification algorithms. Unlike conventional prediction approaches that rely solely on basic customer information or default machine learning configurations, the proposed framework is designed to improve classification accuracy by effectively utilizing customer browsing behavior, product characteristics, pricing information, and promotional activities. The system evaluates three supervised machine learning algorithms—Decision Tree, Random Forest, and Logistic Regression—under a unified experimental environment to determine the most effective model for customer purchase prediction. This integrated framework provides accurate and reliable purchase classification while supporting intelligent decision-making for modern e-commerce platforms.

The proposed methodology begins with data acquisition, where customer transaction records are collected from the JD.com e-commerce dataset released through the INFORMS 2020 MSOM Data-Driven Research Challenge. The dataset contains detailed information related to customer profiles, product specifications, browsing history, purchase records, promotional offers, and pricing information. These data represent real customer interactions occurring on an online shopping platform and provide sufficient information for developing reliable purchase prediction models. Collecting large-scale real-world transaction data improves the robustness of the proposed framework and enables accurate analysis of customer purchasing behavior under different promotional and business conditions.

After collecting the data, the framework performs an extensive data preprocessing stage to improve dataset quality before model development. Initially, information from multiple database tables is merged using common customer and product identifiers to generate a unified transaction dataset. Duplicate records and inconsistent observations are removed to eliminate redundancy and improve data consistency. Missing values occurring within pricing and promotional variables are estimated using a Random Forest-based imputation approach, preserving relationships among important customer and product attributes. Furthermore, categorical variables are transformed into numerical representations using one-hot encoding, enabling the classification algorithms to process non-numeric information efficiently. This preprocessing stage significantly improves data quality and establishes a reliable foundation for subsequent machine learning operations.

Following preprocessing, the proposed framework performs feature engineering to construct meaningful variables that better represent customer purchasing behavior. Continuous product prices are categorized into different price levels, while promotional variables including direct discount, quantity discount, bundle discount, and coupon discount are transformed into corresponding discount levels according to the percentage reduction in product price. An additional feature is created to distinguish customer browsing activities occurring on weekdays from those occurring during weekends. Finally, the target variable representing customer purchasing decisions is generated by examining the relationship between browsing records and completed purchase transactions. A browsing session is labeled as a successful purchase when the customer completes an order within a predefined time interval after viewing the product. These engineered features improve the discriminative capability of the classification models while reducing the influence of noisy and redundant variables.

The processed dataset is subsequently divided into training and testing subsets to evaluate the prediction models objectively. Approximately 60% of the available records are used for training the classifiers, while the remaining 40% are reserved for testing their performance using previously unseen customer transactions.

VI. System Architecture

he proposed system architecture presents a comprehensive machine learning framework for predicting customer purchasing decisions on an e-commerce platform. The architecture integrates multiple processing stages, beginning with data acquisition and ending with purchase prediction and business intelligence generation. Each stage is designed to systematically process customer and product information while improving data quality, enhancing feature representation, optimizing model performance, and generating highly accurate purchase predictions. The modular design enables efficient handling of large-scale transaction datasets and provides a reliable decision-support system for intelligent e-commerce applications.

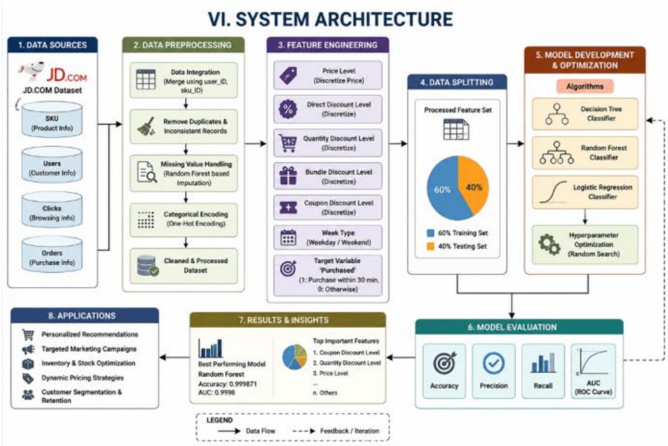

The first stage of the architecture is data collection, where historical transaction data are obtained from the JD.com e-commerce dataset. The collected information includes multiple relational tables containing product details, customer profiles, browsing history, and purchase records. Specifically, the SKU table provides product characteristics, the Users table contains customer demographic information, the Clicks table records browsing activities, and the Orders table stores transaction details such as product prices and promotional discounts. These heterogeneous data sources collectively provide a complete representation of customer shopping behavior and establish the foundation for accurate purchase prediction.

After data collection, the framework proceeds to the data preprocessing stage, where the raw transaction records are prepared for machine learning analysis. Initially, information from different database tables is integrated using common identifiers such as user_ID and sku_ID to generate a unified dataset. Duplicate entries, inconsistent records, and irrelevant observations are removed to improve data consistency. Missing values associated with pricing and promotional variables are estimated using a Random Forest-based imputation technique, allowing incomplete information to be recovered while preserving relationships among product and customer attributes. Furthermore, categorical variables are transformed into numerical representations through one-hot encoding, enabling the classification algorithms to process non-numeric information efficiently. These preprocessing operations significantly improve dataset quality and provide reliable input for subsequent learning stages.

The processed dataset is then supplied to the feature engineering module, where several informative variables are constructed to better represent customer purchasing behavior. Continuous product prices are converted into discrete price levels according to predefined price intervals. Similarly, promotional variables including direct discount, quantity discount, bundle discount, and coupon discount are transformed into corresponding discount levels based on their percentage reduction relative to the original product price. An additional binary variable distinguishes customer browsing activities occurring on weekdays from those taking place during weekends. Finally, the target variable "Purchased" is generated by analyzing customer browsing sessions and completed purchase transactions, enabling the prediction models to learn whether a browsing activity ultimately results in a successful purchase. These engineered variables reduce data complexity while improving the discriminative capability of the machine learning algorithms.

Following feature engineering, the prepared dataset enters the data partitioning stage, where it is divided into training and testing subsets using a 60:40 ratio. Approximately sixty percent of the available records are allocated for model training, while the remaining forty percent are reserved for evaluating model performance on previously unseen customer transactions. This partitioning strategy enables objective assessment of prediction capability and minimizes the possibility of model overfitting by ensuring that testing data remain independent from the training process.

The model development stage constitutes the central component of the proposed architecture. Three supervised machine learning algorithms—Decision Tree, Random Forest, and Logistic Regression—are trained using the prepared transaction dataset. Decision Tree generates hierarchical classification rules based on customer and product characteristics, providing interpretable decision-making structures. Random Forest constructs multiple decision trees using randomly selected samples and feature subsets before combining their predictions through majority voting, thereby improving prediction stability and reducing overfitting. Logistic Regression estimates the probability of customer purchase by applying a logistic function to the weighted combination of input variables. To maximize prediction performance, the framework incorporates Random Search hyperparameter optimization, which automatically evaluates different parameter combinations and selects the configuration that produces the highest classification accuracy while maintaining strong generalization capability..

VII. Conclusion

This study presented an intelligent machine learning framework for predicting customer purchasing decisions using historical transaction data collected from the JD.com e-commerce platform. The proposed methodology integrated comprehensive data preprocessing, advanced feature engineering, supervised machine learning algorithms, and Random Search hyperparameter optimization to develop an accurate and reliable purchase prediction system. By utilizing customer demographic information, product characteristics, browsing history, pricing details, and promotional variables, the framework effectively modeled customer purchasing behavior and generated highly accurate classification results. The comparative evaluation of Decision Tree, Random Forest, and Logistic Regression classifiers demonstrated the effectiveness of applying machine learning techniques for intelligent customer behavior prediction in modern e-commerce environments.

Experimental analysis confirmed that the optimized Random Forest classifier achieved the highest predictive performance among all evaluated models, producing a testing accuracy of 0.999871 and an Area Under the ROC Curve (AUC) of 0.9998. These results indicate that ensemble learning techniques are highly effective in capturing complex relationships among customer characteristics, product attributes, and promotional activities. Furthermore, feature importance analysis identified coupon discount level and quantity discount level as the most influential factors affecting purchasing decisions, followed by product-related attributes and discount information. These findings provide valuable insights into consumer behavior and demonstrate the significant role of promotional strategies in influencing online purchasing decisions.

Overall, the proposed framework provides a scalable, accurate, and computationally efficient solution for customer purchase prediction. The integration of robust preprocessing techniques, informative feature engineering, optimized machine learning models, and systematic performance evaluation enables the framework to generate reliable purchase predictions while maintaining excellent generalization capability. The proposed system can assist e-commerce organizations in implementing personalized recommendation systems, targeted promotional campaigns, customer segmentation, inventory optimization, demand forecasting, dynamic pricing strategies, and intelligent business decision-making. In addition, the modular architecture allows future integration with real-time customer interaction data, deep learning techniques, and cloud-based analytics platforms to further enhance prediction performance and support next-generation intelligent e-commerce applications.

References

1. M. Cao, "Competitive Advantages and Challenges of E-Commerce Platforms," Advances in Economics, Management and Political Sciences, vol. 47, pp. 114–119, 2023.
2. S. C. S. C. Yong, R. T. Huan, W. S. Poh, M. Osman, and D. C. W. Ng, "Assessing the Factors Influencing Consumer Behaviour in E-Commerce Platforms," International Journal of Multidisciplinary Applied Business and Education Research, vol. 4, no. 10, pp. 3725–3735, 2023.
3. Y. Zhang, "Machine Learning Algorithms for Consumer Behaviour Prediction in E-Commerce," Applied Computing and Engineering, vol. 49, pp. 213–219, 2024.
4. M. Shen, C. S. Tang, D. Wu, R. Yuan, and W. Zhou, "JD.com Transaction-Level Dataset for the 2020 MSOM Data-Driven Research Challenge," Manufacturing & Service Operations Management, vol. 26, no. 1, pp. 2–10, 2024.

5. H. Blockeel, L. Devos, B. Frénay, G. Nanfack, and S. Nijssen, "Decision Trees: From Efficient Prediction to Responsible Artificial Intelligence," Frontiers in Artificial Intelligence, vol. 6, pp. 1–18, 2023.
6. L. Breiman, "Random Forests," Machine Learning, vol. 45, no. 1, pp. 5–32, 2001.
7. J. Brownlee, Machine Learning with Logistic Regression. Melbourne, Australia: Machine Learning Mastery, 2023.
8. A. Biswas, "Feature Engineering Techniques for Machine Learning Using One-Hot Encoding," International Journal of Data Science and Analytics, vol. 12, no. 3, pp. 185–196, 2024.
9. M. H. Hassan, "Random Search-Based Hyperparameter Optimization for Machine Learning Models," Journal of Artificial Intelligence Research and Applications, vol. 8, no. 2, pp. 88–101, 2023.
10. F. Pedregosa, G. Varoquaux, A. Gramfort et al., "Scikit-learn: Machine Learning in Python," Journal of Machine Learning Research, vol. 12, pp. 2825–2830, 2011.
11. T. Chen and C. Guestrin, "XGBoost: A Scalable Tree Boosting System," Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 785–794, 2016.
12. G. James, D. Witten, T. Hastie, and R. Tibshirani, An Introduction to Statistical Learning, 2nd ed. New York, NY, USA: Springer, 2021.
13. T. Hastie, R. Tibshirani, and J. Friedman, The Elements of Statistical Learning, 2nd ed. New York, NY, USA: Springer, 2009.
14. C. M. Bishop, Pattern Recognition and Machine Learning. New York, NY, USA: Springer, 2006.
15. I. Goodfellow, Y. Bengio, and A. Courville, Deep Learning. Cambridge, MA, USA: MIT Press, 2016.
16. C. C. Aggarwal, Data Mining: The Textbook. Cham, Switzerland: Springer, 2015.
17. J. Han, M. Kamber, and J. Pei, Data Mining: Concepts and Techniques, 4th ed. Burlington, MA, USA: Morgan Kaufmann, 2023.
18. D. Dua and C. Graff, "UCI Machine Learning Repository," University of California, Irvine, CA, USA, 2019.
19. K. P. Murphy, Machine Learning: A Probabilistic Perspective. Cambridge, MA, USA: MIT Press, 2012.
20. S. Russell and P. Norvig, Artificial Intelligence: A Modern Approach, 4th ed. Hoboken, NJ, USA: Pearson, 2021.