Proceedings of ICFTBA 2026 Symposium: Driving Organizational Strategic Change for Financial Gains
DOI: 10.54254/2754-1169/2026.BH39170

Reorder Prediction and Advertising Strategy Design for Retail Media Network Product Ranking

Yuxuan Dong

School of Journalism and Communication, Shaanxi Normal University, Xi'an, China
yuxuan212026@outlook.com

Abstract. Reorder prediction has been widely studied as a classification problem in e-commerce, but little research has explored how to use the predicted reorder probabilities for downstream product ranking. Extend reordering prediction to Retail Media Network (RMN) advertising in this paper and consider sponsored products as competitors to organic recommendations for display slots. A LightGBM model trained on user-level, product-level, and user-product interaction features achieved a mean AUC of 0.6771 in five-fold cross-validation and 0.6968 on a held-out test set, and consistently outperformed logistic regression. Personalised Top-10 recommendation lists were generated for the organic baseline. A simulated RMN environment compared the four advertising strategies in terms of product selection, bid generation, and advertisement quantity. Each strategy was evaluated according to the quality of recommendations (Precision@10, Recall@10, NDCG@10) and advertising results (exposure rate, average ad rank, simulated revenue). As a result, the strategy that preserved the recommendation quality better (NDCG@10 = 0.8904) had a lower revenue (109.19), whereas the revenue-maximising strategy (542.40) resulted in more disruption (NDCG@10 = 0.7108). It was not due to the selection of products or the quantity of advertisements; random bidding unexpectedly had a lower average ad rank (2.9956) than conversion-based bidding (around 4.97). Users with weaker organic recommendation signals were more likely to be shown the ads. Based on the above results, an empirical foundation for the design of RMN strategy has been provided; platform operators should choose a position on the recommendation-revenue trade-off curve explicitly, and per-user monitoring needs to be introduced to prevent excessive harm to vulnerable users.

Keywords: reorder prediction, Retail Media Network, personalized recommendation, advertising strategy, LightGBM

1. Introduction

Repeat purchases drive the revenue of e-commerce platforms. Grocery platforms, household retailers, and subscription services all have a large number of transactions for regularly purchased items by customers. Predict when and what a customer will purchase to take corresponding actions, such as sending a reminder before the customer runs out of a product, offering a discount to a hesitant buyer, or ensuring that the correct product is at the top of the recommendation list. Most

© 2026 The Authors. This is an open access article distributed under the terms of the Creative Commons Attribution License 4.0 (https://creativecommons.org/licenses/by/4.0/).

DOI: 10.54254/2754-1169/2026.BH39170

research on reorder prediction has focused on improving the accuracy of classification and has not considered the problem to be solved. The predictions are relatively accurate; now we need to find a way to use them in a real-life system.

This paper studies the problem of reorder prediction and its application in Retail Media Network advertising. RMNs are the advertising systems in e-commerce platforms that place sponsored products alongside organically recommended ones for a limited number of slots. Platforms face a continuous conflict: to increase advertising income, they need to increase the exposure of sponsored products, but this may harm the user experience. Research on reorder prediction has rarely been extended to ranking applications, and most studies of RMN advertising focus on the design of the auction mechanism without connecting back to the underlying prediction model that determines what is considered a relevant recommendation.

Therefore, there are three problems. First, can the reorder prediction model be trained reliably enough to provide a ranking basis for personalised recommendations? Second, how do different advertising strategies, such as choosing different products, adjusting bids, changing the number of advertisements, etc., affect the performance of the campaign?, affect the trade-off between recommendation quality and advertising revenue? Third, does the effect of advertising differ among users, and are some groups of users systematically more vulnerable to ranking displacement?

The three are as follows: First, a complete pipeline for reorder prediction and RMN ranking simulation based on actual transaction data will be built instead of artificial construction. Secondly, four sets of advertising strategies are offered at the three levels of design, and a bidding mechanism is employed to quantify the recommendation-revenue trade-off. Third, it shows that the effect of advertising varies by individual; for example, those with weaker organic recommendation signals are displaced more frequently and are not included in the aggregate data. The rest of this paper is as follows. Section 2 reviews related work in reorder prediction, product ranking, and RMN advertising. Section 3 is the method. Section 4 shows the results. Section 5 discusses implications and deficiencies. Section 6 ends.

2. Literature review

The three types of research related to this study are reorder prediction, product ranking, and the recommendation-advertising trade-off. Review them one by one and find the deficiency in this study.

2.1. Reorder prediction

The first type of prediction for customer purchasing behaviour in e-commerce analytics was inventory-oriented. Suryono and others [1] have applied RFID sensor data and fuzzy time series to predict the reorder point of a warehouse and found that consumption patterns contain signals of future demand. Later work focused on customer-level behavioural prediction using machine learning. Droomer and Bekker [2] compared neural networks, XGBoost, and sequence models for predicting the next purchase time of individual customers and found that models trained on pooled user-product data outperformed single-pair models. Manikandan et al. [3] combined information-gain feature selection with fuzzy clustering and multilayer perceptrons for customer segmentation and purchase prediction. Kim and others have used several supervised algorithms on retail transaction data to find that behavioural characteristics, such as the frequency and timeliness of purchases, are more predictive than demographic factors [4]. Related work has extended prediction

DOI: 10.54254/2754-1169/2026.BH39170

to various forms of e-commerce, such as live-streaming [5], and has also addressed related problems such as customer lifetime value and information demand prediction [6].

Collectively, these studies have shown that machine learning can extract useful purchase signals from behavioural data. The other is that the prediction is taken to be the end goal. How to incorporate prediction outputs in downstream ranking or advertising systems has not been studied much.

2.2. Product ranking

The step from prediction to ranking is related to recommendation systems research. Kim et al. [7] introduce the general architecture of personalization and recommendation, and explain how user behaviour and preference data are used to suggest content. Long and others [8] provided a specific index by putting forward a ranking system for the predicted best-selling volume of product searches. Rather than ranking solely based on query relevance, their model elevated items with higher predicted sales to improve the placement of products users were likely to purchase without reducing relevance. Prediction signals can help improve ranking, as shown in this paper; however, no advertisements are added to the search environment.

The difference between the organic and advertising-influenced rankings is the root of the RMN problem. Traditional Recommender Systems Only Optimize for User Relevance. In an RMN environment, the paying advertisers need to be displayed in the same small area; therefore, the single-objective ranking problem has been converted into a multi-objective ranking problem.

2.3. Recommendation–advertising trade-off

Research on RMN advertising mechanisms has grown alongside the commercial importance of retail media. Liu and others [9] have directly tackled the multi-objective problem of e-commerce advertising by putting forward a Deep Neural Auction that jointly optimizes user experience, advertiser utility, and platform revenue. According to the above research, a traditional auction model that considers only a single indicator performs poorly in a multi-objective environment, and the addition of organic ranking signals and bidding rules can improve allocation efficiency.

Most RMN studies have not yet identified the underlying reasons. The empirical trade-off between recommendation quality and advertising revenue under different strategy designs is not well-studied at the level of individual users.

2.4. Positioning of this study

This paper connects the two flows of research that have been developing in parallel. Most research on reordering prediction has stopped at model evaluation, and most RMN research assumes an underlying relevance model. Build an operating prediction model, generate organic rankings according to this model, and then evaluate the advertising strategy based on both revenue and recommendation preservation to provide empirical support for the RMN strategy. In addition, the addition of individual-level analysis in this study is not present in previous research.

DOI: 10.54254/2754-1169/2026.BH39170

3. Methodology

3.1. Problem formulation

This paper studies product reorder prediction in e-commerce and its application to Retail Media Network advertising. Based on the user's past purchase history and product attributes, predict whether the user will buy the same product again in the next shopping session. The outcome is a binary label. Instead of seeking a high-accuracy prediction, the predicted probabilities are used as a ranking signal to build personalised recommendation lists, and thus, a simulated advertising system is created where sponsored products compete with organic recommendations.

3.2. Features and data

The three groups of features are: user-level, product-level, and user-product interaction features. This division reflects the nature of purchase decisions. The frequency of use of a user on the platform generally affects how likely a product is to be repurchased, and whether that specific user has a history of buying that product previously. Only user-level attributes cannot differentiate between various product preferences; only product-level attributes apply to all users uniformly. Interaction features address this problem by learning the personalised association between users and items. All eight variables are shown in Table 1.

Table 1. Full feature list used for model training

Feature Category	Feature Name	Definition
User-level	total_orders	Total historical orders generated by the user on the platform
User-level	order_dow	The day of the week when the user placed the order
User-level	order_hour_of_day	The hour of the day when the user placed the order
User-level	days_since_prior_order	Time interval (days) between the current order and the user's last historical order
Product-level	global_reorder_rate	Overall reorder proportion of the commodity among all purchasing users
Product-level	total_purchase_volume	Cumulative total purchase quantity of the product in the full dataset
User-product interaction	user_buy_times	Accumulated purchase times of the target product by the specific user
User-product interaction	personal_reorder_ratio	The personal reorder proportion of the target product for the individual user

Missing values, which occur because a user has never purchased a specific product, are replaced by the median of that feature for all users or products. For purchase-related attributes that are skewed to the right, such as days_since_prior_order, a median is used instead of a mean to avoid the

DOI: 10.54254/2754-1169/2026.BH39170

influence of outliers on the mean. The median is representative of the user group. As shown in Figure 1, the distribution of days_since_prior_order is as follows.

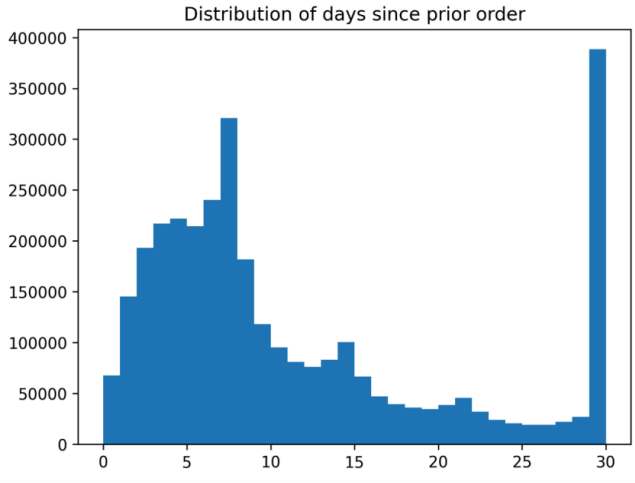


Figure 1. Distribution of days since previous order

Delete duplicate records and ensure the uniformity of the timestamp column prior to model training. As the target variable is imbalanced, stratified sampling was used to divide the dataset into a training set (80%) and a test set (20%), and the ratio of reorders to non-reorders in both subsets was maintained. The Test Set is used for the final test.

3.3. Models and evaluation

Two models are shown with different learning modes. Logistic Regression with L2 Regularisation is a simple linear model that uses L2 regularisation to prevent overfitting by penalising large coefficients. LightGBM is a gradient boosting algorithm that builds an ensemble of decision trees by iteratively splitting nodes to learn non-linear relationships in the data. Together, the two models can determine whether non-linear modelling is more suitable than linear modelling.

Stratified five-fold cross-validation of the training set is used to select a model based on accuracy and AUC. AUC is the model's capacity to assign a higher rank to the positive reordering cases than to the negative ones, and it is relatively insensitive to class imbalance compared with accuracy. Retrain the selected model using all the training data and then evaluate it once on the held-out test set.

3.4. Recommendation list construction

The trained model outputs a predicted repurchase probability for each candidate item in a user's candidate set. The two types of products in the candidate pool are items the user has purchased in the past and items in the same category that the user has never ordered. The first type meets the current demand, and the second adds new recommendations and covers the replenishment needs of familiar items. Products that are off-the-shelf or have had no total sales are excluded. Sort the items according to their predicted probabilities in descending order and select the top 10 as the organic recommendation list for that user. Table 2 is a single-user example list.

DOI: 10.54254/2754-1169/2026.BH39170

Table 2. Example recommendation list for a sampled user

Rank	Product ID
1	24852
2	49302
3	13176
4	27845
5	32403
6	43394
7	35914
	20082
9	47209
10	21903

3.5. Advertising simulation

Add the sponsored products to the organic Top-10 list in the simulated RMN environment and then re-rank all items according to the combined score. Four strategies are compared based on the following three attributes: product selection (long-tail vs. popular), bid generation (global_reorder_<rate-based vs. random), and advertisement quantity (two vs. three per user). Table 3 shows the configurations.

Table 3. Strategy configurations

Strategy	Product Selection	Bid Generation	Ad Count
A (Baseline)	Long-tail (low total_purchase_volume)	global_reorder_rate ±20%	2
B	Popular (high total_purchase_volume)	global_reorder_rate ±20%	2
C	Long-tail (low total_purchase_volume)	Random bid	2
D	Long-tail (low total_purchase_volume)	global_reorder_rate ±20%	3

Long-tail products are chosen as the subjects of advertising because high-volume products have already appeared many times in users' transaction records and recommendation lists, and advertising them offers little additional exposure. Long-tail products, on the other hand, lack natural exposure and are not suitable for advertising to change behaviour. Bid prices are based on the global_reorder_rate of each product, and a ±20% floating range is used to simulate mild merchant competition without significantly reducing the intrinsic conversion capacity. Strategy C randomly generates bids and does not link them to product-level conversion indicators. Natural recommendation items have a bid of 0, and sponsored items have a simulated bid value. Based on

DOI: 10.54254/2754-1169/2026.BH39170

the combined scores, all items are re-ranked, and a new Top-10 list is generated for each user. Table 4 shows a representative case of the before-and-after ranking change.

Table 4. Example ranking change after advertising intervention

Rank	Before Intervention	After Intervention
1	24852	37620 (Ad)
2	49302	24852
3	13176	49302
4	27845	13176
5	32403	27845
6	43394	32403
7	35914	43394
8	20082	35914
9	47209	20082
10	21903	9896 (Ad)

The simulation has the following simplifying assumptions. All of the selected sponsored items are considered displayable; there is no quality score filter, budget constraint, or minimum bid threshold. It is assumed that users will click on organic and paid items in the same way. The above assumptions rule out ranking displacement due to changes in platforms or behaviour.

4. Results

4.1. Model performance

The mean accuracy and mean AUC of LightGBM in five-fold cross-validation are 0.6402 and 0.6771, respectively, and 0.6334 and 0.6700 for Logistic Regression, respectively. Both LightGBM improve the indicators, but the difference is small, about 0.007 in mean AUC. The fold-wise AUC values of LightGBM are 0.6597 and 0.7031. Both models have their lowest AUC at Fold 4 (Logistic Regression: 0.6490; LightGBM: 0.6590), and this is not due to a defect in the model; rather, this fold contains more difficult samples.

LightGBM had an AUC of 0.6968 on the independent test set and was only 0.0002 higher than that of Logistic Regression (0.6966). The two ROC curves are almost the same, and both are much higher than the random baseline. Figure 2 shows the ROC curves of both models. LightGBM performed slightly better than all the other models in all the evaluation scenarios, so it was chosen as the final model. The AUC of the test set is also higher than the cross-validation mean; therefore, retraining with the full training set did not result in overfitting. Table 5 summarizes the cross-validation and test set performance for both models.

Table 5. Five-fold cross-validation and test set performance

Model	CV Mean Accuracy	CV Mean AUC	Test Accuracy	Test AUC
Logistic Regression	0.6334	0.6700	0.6506	0.6966

DOI: 10.54254/2754-1169/2026.BH39170

Table 5. (continued)

LightGBM	0.6402	0.6771	0.6556	0.6968

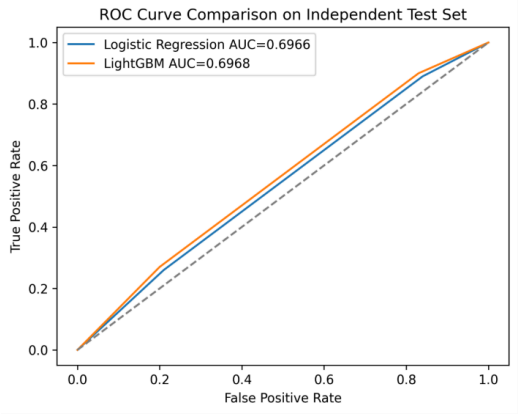


Figure 2. ROC curves on the independent test set

4.2. Baseline performance: organic recommendations

Before advertising is introduced, the organic Top-10 recommendation lists produce a Precision@10 of 0.0826, Recall@10 of 0.0831, and NDCG@10 of 0.0991. A Precision@10 of about 0.083 means that fewer than one in ten of the items are repurchased. Based on the above data, the tasks are relatively difficult; the pool of candidates for each user is large, and most users purchase a small number of products at any given time. The above are the reference bases for measuring the effects of advertising.

4.3. Strategy comparison

Four Advertising Strategies Are Compared. Table 3 in the Methodology section shows the configurations. Table 6 shows the recommendation quality and advertising outcome indicators for Strategies B, C, and D. Strategy A is used as the basic reference, and its own recommendation quality metric is set to 1.0.

Table 6. Strategy comparison results

Metric	Strategy B	Strategy C	Strategy D
P@10	0.8009	0.9000	0.7018
R@10	0.8009	0.9000	0.7018
NDCG@10	0.8210	0.8904	0.7108
Ad Exposure Rate	1.0	1.0	1.0
Avg. Ad Rank	4.9757	2.9956	4.9720
Simulated Revenue	361.60	109.19	542.40

The three patterns are: First, the quality of recommendations is consistently ranked as C > B > D. Strategy C is closest to the original recommendation list (NDCG@10 = 0.8904), and Strategy D causes the most disruption (NDCG@10 = 0.7108). Strategy B falls in between (NDCG@10 =

DOI: 10.54254/2754-1169/2026.BH39170

0.8210). Second, the spread of advertising is uneven. Strategy C has the best average ad rank (2.9956) and places sponsored products near the top three positions on average. Strategies B and D produce nearly identical average ranks (4.9757 and 4.9720), clustering around the fifth position. The Bidding Mechanism Determines the Location of Ads, Not Which Products or How Many Ads. Third, simulated revenue is negatively correlated with recommendation quality. Strategy D has the highest revenue (542.40), followed by Strategy B (361.60), and Strategy C is the lowest (109.19). No strategy can achieve both goals.

4.4. Individual-level analysis

Aggregate indicators mask the differences among users. Two cases are shown above. Strategy A has minimal disruption for user 173665; only one sponsored product appears in the list at Rank 3, and the top two organic items are still at the top. User 68590 has been modified considerably; two sponsored products have been added at positions 3 and 7, only the top two organic items remain in their original places, and the two lowest-ranked organic items have been deleted. Table 7 presents the before-and-after ranking comparison for both users.

Table 7. Individual ranking comparison of two contrasting users

Rank	User 173665 (Before)	User 173665 (After)	User 68590 (Before)	User 68590 (After)
1	47766	47766	27845	27845
2	27845	27845	22035	22035
3	4799	48364 (Ad)	30489	39928 (Ad)
4	2078	4799	2108	30489
5	1862	2078	11422	2108
6	18531	1862	42828	11422
7	23029	18531	30561	13176 (Ad)
8	21461	23029	5450	42828
9	9092	21461	48628	30561
10	19895	—	8518	5450

User 173665 shows little change again, and user 68590 has lost several organic items again. The consistency of the strategies shows that the structure of the user's original recommendation list, specifically the distribution and density of predicted repurchase probabilities for the top 10 items, is the main reason for advertising displacement vulnerability. Users with relatively high and concentrated organic prediction scores are buffered; those with weak or dispersed scores are gradually exposed.

5. Discussion

This paper studies whether reorder prediction can be used as the foundation for product ranking in a Retail Media Network. The results show that it can be achieved; however, the design of the advertising layer introduces trade-offs that are measurable, consequential, and unequally distributed among users.

DOI: 10.54254/2754-1169/2026.BH39170

The Prediction Model Worked as Expected. LightGBM performed marginally better than Logistic Regression, and as indicated by the near-overlap of their ROC curves, the main predictors in this dataset are predominantly linear. This is a good negative result: for this specific task and set of features, the increased complexity of a tree-based model does not lead to substantial improvements. Future work on richer feature spaces or larger datasets may discover more obvious non-linear structures, but within the scope of this study, a simpler model would have been sufficient.

The strategy comparison showed that the quality of recommendations and advertising revenue pull were in opposition. Strategy C randomly bid for long-tail products to keep the recommendations relatively accurate but resulted in the lowest revenue. Strategy D added three conversion-oriented ads per user to maximise revenue but caused significant disruption to organic recommendations. Strategy B fell in between. This trade-off is not an artifact of the simulation. It is a structural feature of RMN systems that advertising occupies display slots that would otherwise be used for organic recommendations, and the two goals cannot be achieved simultaneously under a fixed list length.

What was not expected was the function of the bidding mechanism. According to traditional logic, linking bids to product-level conversion indicators will enable more efficient advertisement placement by reflecting the actual value of the product in the bids. However, according to the data, random bidding has a higher average ad rank than conversion-based bidding. The other is a duplicate. If the bid price is related to the global reorder rate, the composite ranking score will enhance this signal twice: once in the prediction model and once in the bid, thus clustering sponsored products in a narrow rank range. Decoupling bids from organic signals distributes the ad display across the list and can show some paid products at the top. This result is in line with previous research on the interaction between organic ranking and bidding rules in sponsored search [9], and indicates that bid design should receive more attention than is currently given to discussions of the RMN strategy.

The pattern of the general index for individuals was different. Under the same advertising plan, some users are barely affected; others have lost several organic recommendations. It seems that the reason is the structure of each user's original recommendation list: users with a high concentration of prediction scores are suppressed, and those with weak or dispersed scores are systematically more exposed. This has implications for platform operations. Uniform advertising policies that show the same number of advertisements for all users are regressive and place a greater burden on users with weaker organic recommendations. These users may be the ones who will gain the most from good recommendations because the system has not learned their preferences well yet. An adaptive mechanism that lowers the advertising load according to the signal strength of each user can be employed to solve this problem, but it is more complex to implement.

The three applications are as follows: First, platform operators should consider the recommendation-revenue trade-off as an intentional strategic choice by the system, rather than the outcome of the system. No such arrangement can achieve both objectives simultaneously. Second, the design of the bid mechanism requires more engineering attention than the rules for the quantity or type of goods. Third, the scope of monitoring should be expanded to track the distribution of individual-level effects; otherwise, a relatively stable average may not show significant disruptions in the experience of a particular group of users.

Several Deficiencies Limit the Results. To reduce the complexity of the simulation, quality score filtering, budget limits, and different user responses to labelled advertisements have been excluded. Only one dataset and a single model architecture were used for the analysis; therefore, the specific numerical values should not be considered predictions of the live platform's performance. The indicators of evaluation only show the coverage of recommendations, not user satisfaction or

DOI: 10.54254/2754-1169/2026.BH39170

continued use. The reduction was done deliberately, but only the ranking displacement mechanism was isolated; therefore, the simulation results did not align with those in actual operation.

The following will be the future work. According to the per-user signal strength, adjust the insertion parameter to solve the problem of heterogeneity directly. Build a multi-round simulation framework to add user feedback (clicks, purchases, drop-offs) and improve ecological validity. Live A/B testing will be carried out to determine whether the above trade-offs are still valid in practice.

6. Conclusion

This study determines whether reorder prediction can serve as a basis for product ranking in a Retail Media Network. A LightGBM model was trained on multi-level transaction features, and then, using the predicted probabilities, personalised recommendation lists were generated. A simulated advertising environment then tested the four strategies of product selection, bid generation, and advertisement quantity.

Three results are shown. First, the quality of recommendations and advertising income cannot be increased simultaneously; the strategy that best preserved organic recommendations generated the least revenue, and the revenue-maximising strategy caused the most disruption. Secondly, the bidding mechanism was the reason for the result; by separating bids from product-level conversion indicators, the average number of ad placements unexpectedly increased. Third, the impact of advertising varied by user; those with weak or more scattered organic recommendation signals were consistently ranked lower, and this discrepancy was not reflected in the aggregated indicators.

Therefore, the platform operator should be aware that there is a trade-off between recommendations and revenue, and it should not be expected that advertising will be added for free. The design of bids needs to be optimised, and per-user monitoring should be implemented to prevent overly heavy burdens on vulnerable groups of users.

Only one qualification is used, and the results are obtained from a simulation that abstracts away real-world RMN features, such as quality score filtering, budget constraints, and user behaviour in response to advertisements. Therefore, the patterns found here should be considered structural changes rather than changes in operation.

References

[1] Suryono, S., Suseno, J. E., Mashuri, C., Sabila, A. D., Nugraha, J. A. M., & Primasiwi, M. H. (2017). RFID sensor for automated prediction of reorder point (ROP) values in a vendor management inventory (VMI) system using fuzzy time series. *Advanced Science Letters, 23*(3), 2398-2400. https://doi.org/10.1166/asl.2017.8774.

[2] Droomer, M., & Bekker, J. (2020). Using machine learning to predict the next purchase date for an individual retail customer. *South African Journal of Industrial Engineering, 31*(3), 69-82. https://doi.org/10.7166/31-3-2419.

[3] Somasundara Manikandan, R., Santhosh Kumar, S., & Jegathesh Amalraj, J. (2026). Leveraging advanced machine learning with feature selection approach for robust customer segmentation with purchase behavior prediction. *International Journal of Information Technology*, (prepublish), 1-8. https://doi.org/10.1007/S41870-025-03061-3.

[4] Kim, S., Lee, D., Shanto, T. T. I., & Kim, S. (2026). Data-driven prediction of customer purchase behavior in retail and e-commerce using machine learning models. *International Journal of Grid Computing & Applications, 17*(2).

[5] Huang, Y., & Rojniruttikul, N. (2026). Structural equation model of e-commerce live broadcasting influencing customers purchase intention prediction using machine learning. *International Journal of Information Technology*. https://doi.org/10.1007/s41870-025-02851-z

[6] Long, H., Chen, D., & Li, Y. (2026). User information demand prediction and intelligent communication strategies based on customer lifetime value. *Engineering Reports, 8*(4), e70704–e70704. https://doi.org/10.1002/ENG2.70704.

[7] Kim, S., Lee, S., Biswas, D., Shahnawaj, M. D., Mahmood, K. N., & Bhardwaj, P. (2025). Personalization and recommendation systems: Leveraging machine learning algorithms to offer personalized product recommendations

DOI: 10.54254/2754-1169/2026.BH39170

and content to customers based on their behavior, preferences, and purchasing history. *International Journal of Grid Computing & Applications, 16*(2), 1–4. https://doi.org/10.5121/IJGCA.2025.16201.

[8] Long, B., Bian, J., Dong, A., & Chang, Y. (2012). Enhancing product search by best-selling prediction in e-commerce. In *Proceedings of the 21st ACM International Conference on Information and Knowledge Management*. ACM. https://doi.org/10.1145/2396761.2398671

[9] Liu, X., Yu, C., Zhang, Z., Zheng, Z., Rong, Y., Lv, H., Huo, D., Wang, Y., Chen, D., Xu, J., Wu, F., Chen, G., & Zhu, X. (2021). Neural auction: End-to-end learning of auction mechanisms for e-commerce advertising. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. ACM. https://doi.org/10.1145/3447548.3467103