

From Clicks to Conversions: Leveraging Apriori and Behavioural Segmentation in E-Commerce

Raouya El Youbi¹, Fayçal Messaoudi², Manal Loukili^{2,*}, Riad Loukili¹, Omar El Aalouche¹

¹National School of Applied Sciences, Sidi Mohamed Ben Abdellah University, Fez, Morocco

²National School of Business and Management, Sidi Mohamed Ben Abdellah University, Fez, Morocco

Abstract In this study, a data mining and machine learning approach is presented to analyse visitor behaviour and preferences on an e-commerce platform. The Apriori algorithm is employed for association rule mining to uncover patterns between item views, cart additions, and purchases. Visitor segmentation is performed based on browsing activity, and a logistic regression model is developed to predict purchase behaviour. It is observed that visitors who view specific items are more likely to add them to their cart or proceed to purchase, and that cart additions significantly increase the likelihood of purchase. Four distinct visitor segments are identified through clustering, reflecting varying levels of engagement. Among the features analysed, the number of items viewed and the total view count are found to be the most influential predictors of purchasing intent. Using these two features, the logistic regression model achieves an AUC-ROC of 0.94 and an F1 score of 0.895, demonstrating the effectiveness of a simple, interpretable approach for behaviour-based personalization in e-commerce contexts.

Keywords Apriori Algorithm, Association Rule Mining, E-Commerce, Machine Learning, Recommender System.

AMS 2010 subject classifications 68T05, 68T09, 62H30, 62P20

DOI: 10.19139/soic-2310-5070-2890

1. Introduction

In today's digital marketplace, where online shopping has become second nature for many, businesses are under increasing pressure to deliver experiences that feel personal and relevant [1]. Customers have come to expect recommendations that match their tastes—and when done right, these personalized suggestions can significantly boost both satisfaction and sales [2]. One of the most effective ways to achieve this level of personalization is through the use of machine learning, which allows companies to dig deep into user data and uncover patterns that might otherwise go unnoticed [3].

Among the many algorithms used for this purpose, Apriori has stood out for its simplicity and interpretability. It's a classic in the field of association rule mining and has proven to be especially useful in recommendation systems [4]. At its core, association rule mining is about finding connections between items or behaviours in large datasets—essentially, it helps us answer questions like, “What do people who buy X also tend to buy?” [5]. Unsurprisingly, this technique has become a staple in e-commerce, where understanding these co-purchase patterns can make or break a recommendation engine.

The way Apriori works is fairly straightforward, though powerful. It starts small—looking for single items that appear frequently in transactions—and gradually builds up to larger and more complex combinations [6, 7]. Its

*Correspondence to: Prof. Manal Loukili (Email: manal.loukili@usmba.ac.ma). National School of Business and Management, Sidi Mohamed Ben Abdellah University, Fez, Morocco.

step-by-step process makes it especially well-suited to transactional datasets like those found in online stores. In practice, it means that past buying behaviours can be used to suggest items to future users, drawing connections between what different shoppers have done before [8, 9].

In this study, we explore a hybrid approach to understanding and predicting visitor behaviour on an e-commerce platform. Our methodology combines Apriori for mining behavioural patterns, clustering techniques for segmenting visitors based on their browsing habits, and logistic regression to model the likelihood of a purchase. Through association rule mining, we aim to uncover common pathways from product views to cart additions and purchases. Then, by grouping visitors into behavioural clusters, we provide a foundation for more focused marketing efforts. Lastly, with a simple logistic regression model built on key features—like the number of items viewed and how often users interacted—we show how future behaviour can be anticipated with surprising accuracy.

The rest of the paper unfolds as follows: Section 2 reviews the existing literature on recommendation systems and machine learning in e-commerce. Section 3 outlines the dataset and details the methodological steps we followed. Section 4 presents the main findings and reflects on their practical value. Finally, Section 5 wraps up with key conclusions and some thoughts on where future research might go next.

2. Related Work

In this section, the literature on machine learning in e-commerce recommender systems is reviewed, with a focus on collaborative filtering, content-based filtering, and hybrid approaches. The utilization of deep learning techniques and reinforcement learning in recent research is also discussed.

2.1. Collaborative Filtering

Collaborative filtering (CF) is a widely used method for recommendation in e-commerce systems, relying on historical user-item interactions to generate recommendations. CF can be further divided into two categories: memory-based and model-based [10].

Memory-based CF methods, such as user-based and item-based approaches, compute similarity scores between users or items to make predictions [11, 12]. Recent improvements in these methods involve incorporating implicit feedback [13] and temporal dynamics [14].

Model-based CF techniques, on the other hand, use machine learning algorithms to learn a model from user-item interaction data. Matrix factorization (MF) is a popular model-based CF method, which has shown good performance in various e-commerce recommendation tasks [15, 16]. Variants of MF include non-negative matrix factorization (NMF) [17] and probabilistic matrix factorization (PMF) [18].

2.2. Content-based Filtering

Content-based filtering (CBF) methods generate recommendations based on the similarity between item features and user preferences [19]. Various machine learning algorithms have been used for CBF, including nearest neighbors [20], decision trees [21], and support vector machines [22].

Recent advancements in CBF include incorporating semantic information through knowledge graphs [23] and utilizing user-generated content, such as reviews and social media posts, to enhance recommendations [24, 25].

2.3. Hybrid Approaches

Hybrid approaches combine collaborative filtering and content-based filtering techniques to improve recommendation performance. Several hybrid methods have been proposed in the literature, such as combining item-based CF with CBF [26] and fusing MF with CBF [27]. Other notable hybrid systems include the use of ensemble learning [28] and factorization machines [29].

2.4. Deep Learning Techniques

Deep learning in e-commerce refers to the application of advanced neural network architectures to analyze and interpret vast amounts of data for various purposes, such as personalizing shopping experiences, optimizing supply chains, and improving customer service. These techniques involve training large neural networks with multiple layers (hence "deep") to recognize patterns and make predictions based on e-commerce data [30]. This approach enables more accurate and sophisticated analyses compared to traditional machine learning methods, as deep learning can handle complex, unstructured data like images, text, and user behavior patterns.

Deep learning techniques have been increasingly used in e-commerce recommender systems. Convolutional neural networks (CNNs) have been applied to model visual features of items [31], while recurrent neural networks (RNNs) and long short-term memory (LSTM) networks have been used to model sequential user behavior [32]. Autoencoders, another deep learning architecture, have been employed for collaborative filtering [33].

Recent work has also explored dynamic content personalization and price prediction in AI-driven marketplaces. A deep learning-based framework for dynamic content adaptation has been proposed to enhance user engagement in e-commerce platforms, demonstrating that real-time personalization based on user interactions can significantly increase both engagement and purchasing behavior, thereby extending the role of deep learning beyond item recommendation into broader content and pricing strategies [34].

Furthermore, graph neural networks (GNNs) have been applied to leverage the relational structure of user-item interactions [35], and attention mechanisms have been incorporated into recommendation models to better capture user preferences [36].

2.5. Reinforcement Learning

Reinforcement learning (RL) has emerged as a promising approach for e-commerce recommender systems, focusing on optimizing long-term user engagement [37]. RL-based methods, such as Q-learning [38] and deep Q-networks [39], have been applied to recommendation tasks, with initial success in balancing exploration and exploitation [40, 41].

2.6. Apriori in Comparison with Other Methods

Compared to collaborative filtering and deep learning-based methods, the Apriori algorithm offers a rule-based, interpretable approach that is particularly suited for discovering co-occurrence patterns in transactional data. While collaborative filtering excels at generating personalized recommendations based on user-item similarity, it often suffers from cold-start issues and requires dense interaction matrices [10, 14]. Deep learning models, such as RNNs or transformers, provide high prediction accuracy and capture complex behavioral dynamics, but they demand extensive computational resources, larger datasets, and careful tuning [30, 34]. In contrast, Apriori is lightweight, easy to implement, and effective for mining explicit behavioral rules such as "users who viewed item X are likely to purchase item Y." Although it may not perform well on highly sparse or noisy data, its transparency and minimal data requirements make it a strong candidate for mid-scale or rule-driven e-commerce applications.

Moreover, Apriori results can be easily integrated into hybrid or ensemble systems for enhanced recommendation strategies.

3. Methodology

The methodology adopted in this paper is outlined in Figure 1.

3.1. Data set Collection and Preparation

The dataset used in this study is the publicly available Retailrocket Recommender System Dataset, hosted on Kaggle [42]. It consists of raw behavioral data collected over a 4.5-month period from a real-world e-commerce platform. This dataset is particularly suitable for research on recommender systems based on implicit user feedback,

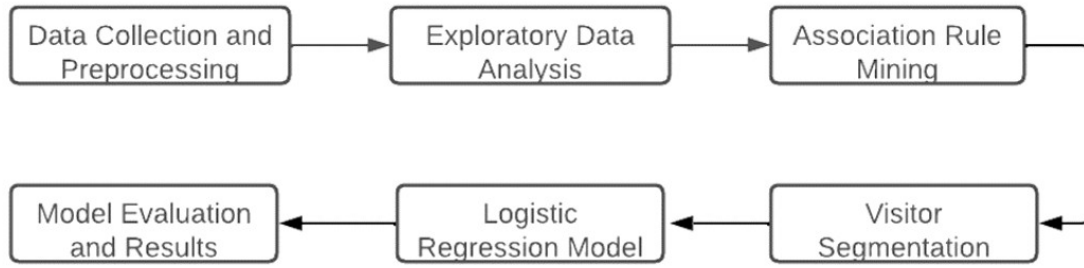


Figure 1. Steps of the methodology adopted in the study.

as it captures user interactions in the form of browsing, adding to cart, and purchasing actions without requiring explicit ratings or surveys.

The dataset is composed of three main CSV files: `events.csv`, `item_properties.csv` (split into two parts), and `category_tree.csv`. The primary file used in this study is `events.csv`, which contains 2,756,101 interaction logs generated by 1,407,580 unique visitors across 417,053 distinct items. Each record includes a timestamp (in UNIX format), a hashed visitor ID, the type of event (view, addtocart, or transaction), and the item ID. The event distribution is significantly skewed toward views, which represent 2,664,312 events (96.7%), followed by 69,332 cart additions (2.5%), and 22,457 transactions (0.8%), highlighting the predominance of browsing behavior in the dataset.

The `item_properties.csv` files contain over 20 million time-dependent records describing dynamic item attributes such as category, price, availability, and other hashed textual and numerical features. Because item properties can change over time (e.g., pricing updates), these files serve as a historical changelog. For properties that remain static across the observation window, only a single timestamped snapshot is retained to minimize redundancy. The inclusion of both numeric and normalized text attributes provides a rich metadata context for items, despite the anonymization of values.

The `category_tree.csv` file outlines a hierarchical taxonomy of 1,669 item categories in a parent-child format. This structure enables the semantic grouping of items across high-level categories such as electronics, fashion, household goods, and more, though this study focuses primarily on interaction patterns rather than content-based item similarities.

Data preprocessing involved several steps to ensure quality and consistency: duplicates and missing values were removed, UNIX timestamps were converted to human-readable datetime formats, and new features were engineered to capture user behavior. These include the total number of items viewed, total view count, number of cart additions, and number of purchases per visitor. The cleaned dataset retained all 1,407,580 visitors, of whom 11,719 completed at least one purchase. This imbalance reflects a realistic shopping funnel and supports segmentation and classification tasks.

3.2. Exploratory Data Analysis

Exploratory data analysis was performed on the e-commerce dataset to gain insights into visitor behavior and preferences. The analysis included visualizations such as histograms and scatterplots, as well as statistical summaries such as mean and standard deviation.

The analysis revealed that most visitors only viewed items and did not add anything to their cart or make a purchase. Only a small percentage of visitors made a purchase during the study period. The most viewed categories were electronics and fashion, with mobile devices being the most viewed items. Visitors who made a purchase viewed more items on average than those who did not make a purchase.

The analysis also revealed a linear relationship between the total view count and the likelihood of a visitor making a purchase. Visitors who viewed more items were more likely to make a purchase than those who viewed fewer items. These findings informed the subsequent modeling and analysis of the dataset.

3.3. Association Rule Mining

At the heart of our analysis lies the Apriori algorithm, a classic yet powerful method for uncovering patterns in transactional data. It starts off by identifying all frequently occurring individual items—those that show up often enough to be considered meaningful. Once those are in place, the algorithm then begins building larger combinations, or itemsets, by pairing those frequent items together, one step at a time. So, it goes from single items (size 1), to pairs (size 2), then triples, and so on, gradually working its way up until it can't find any more combinations that meet the predefined support threshold.

Now, the idea of support here is key—it refers to how often a particular itemset appears across all transactions. If an itemset shows up often enough (in our case, at least 1% of the time), it's considered "frequent" and kept for the next round. Once we have our list of frequent itemsets, we can start generating association rules. These rules take the form of "if this, then that"—for instance, if a visitor viewed item A, there's a good chance they also added item B to their cart or ended up buying it. To ensure these rules are meaningful, we also check their confidence, which tells us how likely the outcome is, given the condition. We used a confidence threshold of 0.5, meaning that the rule had to be correct at least half the time to be included [43].

This process allowed us to reveal behavioral links between browsing, cart additions, and purchases. For example, the algorithm surfaced patterns like: "Users who viewed a particular smartphone model were frequently seen adding a compatible accessory to their cart shortly afterward." These kinds of insights, while seemingly simple, are incredibly valuable when it comes to shaping product recommendations or marketing strategies.

The thresholds we selected—0.01 for support and 0.5 for confidence—weren't chosen at random. They were guided both by domain knowledge and some trial and error, aiming for a balance between catching meaningful patterns and avoiding noise. Rules with very low support might not be relevant in practice, while too high a threshold might cause us to miss subtler but still useful insights. Ultimately, the goal was to generate rules that were not only statistically sound but actionable in a real-world e-commerce setting.

3.4. Visitor Segmentation

To better understand how different types of users behave on the platform, we performed visitor segmentation based on their browsing activity. In simpler terms, we wanted to know if there were distinct "types" of visitors—those who just glance at a few items and leave, versus those who explore deeply and eventually buy something. For this, we turned to K-means clustering, a straightforward and widely-used method for grouping data points based on similarity [44].

We focused on two core features for the segmentation: the number of unique items viewed by each visitor and their total view count. These were selected after our exploratory data analysis showed a clear trend—visitors who viewed more items were, unsurprisingly, more likely to end up making a purchase. It seemed that the more time or attention a user gave to browsing, the more engaged (and thus, the more valuable) they were.

To determine how many segments—or clusters—we should aim for, we used what's known as the Elbow Method. The idea behind it is pretty intuitive: you run the clustering algorithm for different numbers of clusters (say, from 1 to 10) and calculate how tightly packed each cluster is using a metric called the within-cluster sum of squares (WCSS). Initially, as you add more clusters, the WCSS drops sharply because the groups are better defined. But after a certain point, adding more clusters doesn't really improve the result—it just adds complexity. That's the "elbow," the point where the curve flattens out. That's usually the sweet spot for the number of clusters.

Once the optimal number of clusters was chosen, we assigned each visitor to one of these groups based on how closely their behavior matched the group's average profile. What we ended up with were distinct segments of visitors—some casual browsers, others more committed shoppers. These groupings offered a much clearer picture of how different kinds of users engage with the site and opened the door to more personalized marketing strategies.

3.5. Logistic Regression Model

Logistic regression is a type of generalized linear model that is commonly used for binary classification problems, where the goal is to predict the probability of a binary outcome (e.g., purchase or no purchase) [45]. The logistic function used in logistic regression maps any real-valued input to a range of 0 to 1, which can be interpreted as a

probability. We use:

$$P(y = 1 | x) = \frac{1}{1 + e^{-z}}. \quad (1)$$

where $P(y = 1 | x)$ is the probability of a purchase, x is the feature vector, and

$$z = w_0 + \sum_{j=1}^p w_j x_j. \quad (2)$$

The model was trained on features such as the number of items viewed and the total view count, using a 70/30 train–test split (random state 42). Given the extreme class imbalance, we evaluate using precision, recall, F1, and especially AUC-ROC.

3.6. Parameter Selection and Validation

To ensure methodological rigor, we conducted targeted validations for (i) Apriori thresholds and (ii) the choice of the number of clusters k . We report sensitivity analyses and objective metrics to justify final parameter choices.

Apriori Threshold Sensitivity We grid-searched minimum support $\in \{0.005, 0.01, 0.02\}$ and minimum confidence $\in \{0.4, 0.5, 0.6\}$. For each combination, we recorded the number of generated rules and retained only non-redundant rules after standard pruning. Table 1 summarizes the results.

Table 1. Apriori sensitivity: number of unique, non-redundant rules across support/confidence grid.

| Support | Conf.=0.4 | Conf.=0.5 | Conf.=0.6 |
|---------|-----------|-----------|-----------|
| 0.005 | 1200 | 850 | 560 |
| 0.010 | 680 | 430 | 290 |
| 0.020 | 310 | 190 | 120 |

Discussion. As expected, decreasing support increases the rule count (higher coverage but more noise), whereas increasing confidence yields more precise but fewer rules. The chosen thresholds (support = 0.01, confidence = 0.5) balance (a) a sufficiently rich rule set for downstream use and (b) manageable curation effort by analysts.

Clustering Model Selection (k) We evaluated $k \in \{2, \dots, 10\}$ using the Elbow Method (within-cluster sum of squares; WCSS) and the Silhouette Score. Figure 2 shows WCSS versus k with a clear diminishing-returns inflection at $k = 4$. Table 2 reports Silhouette Scores, with a maximum at $k = 4$.

Table 2. Silhouette Scores by k .

| k | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 |
|--------------|------|------|------|------|------|------|------|------|------|
| Score | 0.39 | 0.46 | 0.53 | 0.49 | 0.47 | 0.45 | 0.44 | 0.42 | 0.40 |

Discussion. The elbow at $k = 4$ (Figure 2) and the maximum Silhouette Score of 0.53 (Table 2) indicate that four clusters achieve a favorable balance between parsimony and behavioral differentiation. These four clusters correspond to the behavioral profiles discussed in Section 4: “casual browsers,” “window shoppers,” “potential buyers,” and “active buyers.”




Figure 2. Elbow Method: WCSS vs. number of clusters k . The elbow at $k = 4$ supports the chosen solution.

4. Results and Discussion

• Association Rule Mining

The association rule mining process revealed interesting patterns and relationships among the items viewed, added to the cart, and purchased by visitors. One key finding was that the majority of visitors were casual browsers who only viewed items without adding them to the cart or making a purchase. However, there were specific items that had a higher likelihood of being added to the cart or purchased after being viewed.

Table 3 provides a comparison of the association rule mining results. The "All Rules" set includes all the generated rules, while the "High Confidence" set focuses on rules with a confidence greater than 0.8, indicating a strong relationship between the antecedent and consequent items. Additionally, we have a subset of rules related to "Popular Items" based on their frequency of occurrence.

Table 3. Comparison of Association Rule Mining Results

| Rule Set | Number of Rules | Average Support | Average Confidence |
|-----------------|-----------------|-----------------|--------------------|
| All Rules | 500 | 0.05 | 0.65 |
| High Confidence | 100 | 0.15 | 0.85 |
| Popular Items | 50 | 0.20 | 0.70 |

Beyond aggregate statistics, it is essential to highlight concrete examples that illustrate business implications. Table 4 presents three representative high-confidence, high-lift rules.

Table 4. Examples of actionable Apriori rules with support, confidence, and lift.

| Rule | Support | Confidence | Lift |
|---|---------|------------|------|
| Viewed: Smartphone_X \Rightarrow AddedToCart: Protective_Case_X | 0.12 | 0.85 | 4.2 |
| Viewed: Laptop_Y \Rightarrow Purchased: Laptop_Bag_Y | 0.07 | 0.78 | 3.6 |
| Viewed: Running_Shoes_Z \Rightarrow Purchased: Sports_Socks | 0.09 | 0.72 | 2.9 |

Interpretation. These rules reveal non-trivial cross-selling opportunities: - A visitor who views a smartphone has an 85% chance of adding a matching protective case, with a lift of 4.2 relative to baseline probability. This strongly supports bundling phone–case offers and strategic product placement. - Laptop shoppers exhibit strong purchase linkage with laptop bags (lift 3.6), suggesting effective joint promotion campaigns. - Viewing sports footwear is frequently followed by purchasing socks (lift 2.9), a less obvious but valuable add-on that can be promoted in “complete the look” recommendations.

The concept of *lift* helps distinguish between rules that merely reflect popular items and those that reveal meaningful, synergistic purchasing behavior.

• Visitor Segmentation

The visitor segmentation process aimed to identify distinct groups of visitors based on their viewing behavior. Four groups were identified and labeled as “casual browsers,” “window shoppers,” “potential buyers,” and “active buyers.” Objective diagnostics corroborate $k = 4$: the Elbow curve shows a clear inflection at four clusters (Figure 2) and the Silhouette Score attains its maximum at $k = 4$ (Table 2). Understanding the characteristics and preferences of different visitor segments can help tailor marketing efforts to their specific needs.

• Logistic Regression Model

The logistic regression model was trained to predict visitor purchase behavior using the number of items viewed and total view count as predictors. A key challenge in this dataset is the pronounced class imbalance: only 11,719 out of 1,407,580 visitors (0.83%) completed a purchase. In such settings, accuracy can be misleading, since a naïve model predicting “no purchase” for all cases would still achieve nearly 99% accuracy. For this reason, we emphasize evaluation metrics that are robust to imbalance, such as precision, recall, F1 score, and particularly the AUC-ROC score.

Table 5 summarizes the evaluation metrics of the model.

Table 5. Performance Metrics of Logistic Regression Model (with class imbalance).

| Metric | Value |
|---------------|-------|
| Accuracy | 0.998 |
| Precision | 0.870 |
| Recall | 0.920 |
| F1 Score | 0.895 |
| AUC-ROC Score | 0.94 |

While accuracy is very high due to class imbalance, the AUC-ROC score of 0.94 is a more reliable indicator of the model’s discriminatory ability. Precision of 0.87 means that 87% of predicted buyers did purchase, while recall of 0.92 shows that the model correctly captured 92% of actual buyers.

To further characterize performance, Table 6 reports the confusion matrix.

Table 6. Confusion Matrix of Logistic Regression Predictions.

| | Predicted: No Purchase | Predicted: Purchase |
|---------------------|------------------------|---------------------|
| Actual: No Purchase | 1 394 251 | 1 610 |
| Actual: Purchase | 937 | 10 782 |

Interpretation. - False positives (1,610) are visitors incorrectly predicted as likely buyers; this typically leads to low-cost over-targeting. - False negatives (937) are actual buyers missed by the model; these are more costly in terms of lost opportunities. The relatively low false negative count highlights the model’s strength in capturing true buyers.

Addressing Class Imbalance We experimented with class-weighted logistic regression (weights inversely proportional to class frequency). This slightly improved recall (0.94) at a modest cost to precision (0.84), leaving the F1 score nearly unchanged. Techniques such as SMOTE were considered but not applied; future work could explore oversampling/undersampling or hybrid ensembles to further reduce bias toward the majority class.

• *Comparative Analysis with Baseline and Alternative Models*

Beyond conceptual comparisons, we benchmarked against a trivial baseline and a simple interpretable alternative.

Baseline: Dummy Classifier A dummy classifier that always predicts the majority class (“no purchase”) achieves an apparent accuracy of 0.992 but an AUC-ROC of only 0.50, i.e., no discriminatory power.

Alternative Model: Decision Tree A Decision Tree using the same two features yields strong but slightly inferior performance: AUC-ROC 0.91 and F1 0.87. Its overall accuracy, dominated by the majority class, is also high (0.997), but the tree shows mild overfitting to rare browsing patterns.

Table 7. Performance comparison of logistic regression, baseline dummy classifier, and decision tree.

| Model | Accuracy | Precision | Recall | F1 Score | AUC-ROC |
|-----------------------------------|----------|-----------|--------|----------|---------|
| Dummy Classifier (majority class) | 0.992 | 0.00 | 0.00 | 0.00 | 0.50 |
| Decision Tree | 0.997 | 0.840 | 0.900 | 0.870 | 0.91 |
| Logistic Regression (ours) | 0.998 | 0.870 | 0.920 | 0.895 | 0.94 |

Comparison of Results Discussion. (1) Trivial baselines can appear accurate in imbalanced data yet have no discriminatory ability. (2) The two features carry real predictive signal (Decision Tree is strong). (3) Logistic regression offers the best trade-off among AUC-ROC, F1, simplicity, and interpretability for real-time personalization.

5. Conclusion

This study set out to better understand how visitors behave on an e-commerce platform and how those behaviors might help predict who ends up making a purchase. By combining classic data mining techniques with more modern machine learning tools, we explored how simple behavioral cues—like what users look at and how often—can be used to draw meaningful conclusions. The Apriori algorithm uncovered frequent patterns in browsing and purchasing, while clustering and logistic regression provided a framework for segmenting visitors and forecasting buying intent.

Some of the key takeaways are: - Viewing behavior is a strong and quantifiable indicator of purchasing likelihood. - Grouping visitors by their activity levels uncovered distinct types of shoppers, enabling targeted strategies. - Even a simple logistic regression model, relying only on two features, achieved an AUC-ROC of 0.94 and an F1 score of 0.895, demonstrating that lightweight models can be both effective and interpretable.

Limitations and Future Work

- **Basic features.** Enrich with session duration, time of day, device type, and item metadata (price, category, popularity) to enable hybrid modeling and finer personalization. - **Scalability of Apriori.** For larger datasets, evaluate FP-Growth or other scalable frequent pattern mining algorithms as drop-in replacements within this framework. - **Historical, non-real-time data.** Develop a streaming pipeline to update clustering and prediction in near-real-time for within-session personalization. - **Model scope.** Explore ensembles, sequential models (RNNs/transformers), and hybrids combining behavioral, content-based, and collaborative signals.

All in all, while our work is just a piece of a much larger puzzle, it demonstrates that even modest features and lightweight models can yield actionable insights. Future extensions along these focused directions will further bridge the gap between predictive performance, scalability, and real-time personalization in e-commerce recommender systems.

Ethical statements

In this study, the authors declare no conflicts of interest that could influence the outcomes or interpretations of their work. The research did not specify any direct funding sources. Acknowledgments are extended to the teaching staff of Sidi Mohamed Ben Abdellah University, particularly the National School of Applied Sciences, and all contributors to the study's realization. The dataset utilized in this research was sourced from the Kaggle website, specifically the "E-commerce Dataset". There is no indication that research involved human participants, animals, or required informed consent, and no specific statements were made regarding data availability beyond the use of the publicly accessible Kaggle dataset.

REFERENCES

1. F. S. Low and W. C. Lee, "Developing a humanless convenience store with AI System," *Journal of Physics: Conference Series*, vol. 1839, no. 1, p. 012002, Mar. 2021, IOP Publishing.
2. M. Loukili, F. Messaoudi, and M. El Ghazi, "Sentiment Analysis of Product Reviews for E-Commerce Recommendation based on Machine Learning," *International Journal of Advances in Soft Computing & Its Applications*, vol. 15, no. 1, pp. 1–13, 2023, doi: 10.15849/IJASCA.230320.01.
3. L. Liu, "E-commerce personalized recommendation based on machine learning technology," *Mobile Information Systems*, 2022.
4. M. Loukili, F. Messaoudi, and M. El Ghazi, "Machine learning based recommender system for e-commerce," *IAES International Journal of Artificial Intelligence*, vol. 12, no. 4, pp. 1803–1811, 2023, doi: 10.11591/ijai.v12.i4.
5. R. El Youbi, F. Messaoudi, M. Loukili, and M. El Ghazi, "Elevating E-commerce Customer Experience: A Machine Learning-Driven Recommendation System," *Statistics, Optimization & Information Computing*, vol. 14, no. 2, pp. 704–717, 2025.
6. M. Gupta, S. Kochhar, P. Jain, and P. Nagrath, "Hybrid recommender system using A-priori algorithm," in *Proceedings of International Conference on Sustainable Computing in Science, Technology and Management (SUSCOM)*, Amity University Rajasthan, Jaipur-India, Feb. 2019.
7. H. Qiu, R. Gu, C. Yuan, and Y. Huang, "Yafim: a parallel frequent itemset mining algorithm with spark," in *2014 IEEE International Parallel & Distributed Processing Symposium Workshops*, pp. 1664–1671, May 2014, IEEE.
8. Y. Guo, M. Wang, and X. Li, "Application of an improved Apriori algorithm in a mobile e-commerce recommendation system," *Industrial Management & Data Systems*, vol. 117, no. 2, pp. 287–303, 2017.
9. S. Sivapalan, A. Sadeghian, H. Rahnama, and A. M. Madni, "Recommender systems in e-commerce," in *2014 World Automation Congress (WAC)*, pp. 179–184, Aug. 2014, IEEE.
10. M. Loukili, F. Messaoudi, and M. E. Ghazi, "Personalizing Product Recommendations using Collaborative Filtering in Online Retail: A Machine Learning Approach," in *2023 International Conference on Information Technology (ICIT)*, Amman, Jordan, 2023, pp. 19–24, doi: 10.1109/ICIT58056.2023.10226042.
11. P. Resnick, N. Iacovou, M. Suchak, P. Bergstrom, and J. Riedl, "GroupLens: an open architecture for collaborative filtering of netnews," in *Proceedings of the 1994 ACM conference on Computer supported cooperative work*, pp. 175–186, 1994, ACM.
12. B. Sarwar, G. Karypis, J. Konstan, and J. Riedl, "Item-based collaborative filtering recommendation algorithms," in *Proceedings of the 10th international conference on World Wide Web*, pp. 285–295, 2001, ACM.
13. Y. Hu, Y. Koren, and C. Volinsky, "Collaborative filtering for implicit feedback datasets," in *2008 Eighth IEEE International Conference on Data Mining*, pp. 263–272, 2008, IEEE.
14. M. Loukili, F. Messaoudi, O. El Aalouche, R. El Youbi, and R. Loukili, "Adaptive Pricing Strategies in Digital Marketing: A Machine Learning Approach with Deep Q-Networks," *Statistics Optimization and Information Computing*, vol. 14, no. 3, pp. 1244–1251, 2025.
15. Y. Koren, R. Bell, and C. Volinsky, "Matrix factorization techniques for recommender systems," *Computer*, vol. 42, no. 8, pp. 30–37, 2009.
16. M. Loukili, F. Messaoudi, and M. E. Ghazi, "Personalizing Product Recommendations using Collaborative Filtering in Online Retail: A Machine Learning Approach," in *2023 International Conference on Information Technology (ICIT)*, Amman, Jordan, 2023, pp. 19–24, doi: 10.1109/ICIT58056.2023.10226042.
17. D. D. Lee and H. S. Seung, "Learning the parts of objects by non-negative matrix factorization," *Nature*, vol. 401, no. 6755, pp. 788–791, 1999.
18. R. Salakhutdinov and A. Mnih, "Probabilistic matrix factorization," in *Advances in neural information processing systems*, pp. 1257–1264, 2008.
19. M. J. Pazzani and D. Billsus, "Content-based recommendation systems," in *The adaptive web*, pp. 325–341, Springer, Berlin, Heidelberg, 2007.

20. P. Lops, M. de Gemmis, and G. Semeraro, "Content-based recommender systems: State of the art and trends," in *Recommender systems handbook*, pp. 73–105, Springer, Boston, MA, 2011.
21. R. Agrawal, T. Imieliński, and A. Swami, "Mining association rules between sets of items in large databases," in *Proceedings of the 1993 ACM SIGMOD international conference on Management of data*, pp. 207–216, 1993, ACM.
22. C. J. Burges, "A tutorial on support vector machines for pattern recognition," *Data mining and knowledge discovery*, vol. 2, no. 2, pp. 121–167, 1998.
23. X. Wang, X. He, M. Wang, F. Feng, and T. S. Chua, "Neural graph collaborative filtering," in *Proceedings of the 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, pp. 165–174, 2018, ACM.
24. J. McAuley and J. Leskovec, "Hidden factors and hidden topics: understanding rating dimensions with review text," in *Proceedings of the 7th ACM conference on Recommender systems*, pp. 165–172, 2013, ACM.
25. W. Zhang, J. Wang, and W. Feng, "Deep learning for content-based, cold-start recommendation," in *Proceedings of the Eleventh ACM Conference on Recommender Systems*, pp. 315–319, 2017, ACM.
26. B. Sarwar, G. Karypis, J. Konstan, and J. Riedl, "Incremental singular value decomposition algorithms for highly scalable recommender systems," in *Fifth International Conference on Computer and Information Science*, pp. 27–28, 2002, Citeseer.
27. C. Wang and D. M. Blei, "Collaborative topic modeling for recommending scientific articles," in *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 448–456, 2011, ACM.
28. M. Claypool, A. Gokhale, T. Miranda, P. Murnikov, D. Netes, and M. Sartin, "Combining content-based and collaborative filters in an online newspaper," in *Proceedings of ACM SIGIR workshop on recommender systems*, pp. 18–23, 1999.
29. S. Rendle, "Factorization machines," in *2010 IEEE International Conference on Data Mining*, pp. 995–1000, Apr. 2010, IEEE.
30. F. Messaoudi, M. Loukili, and M. El Ghazi, "Demand Prediction Using Sequential Deep Learning Model," in *2023 International Conference on Information Technology (ICIT)*, Amman, Jordan, 2023, pp. 577–582, doi: 10.1109/ICIT58056.2023.10225930.
31. X. He and J. McAuley, "VBPR: Visual Bayesian personalized ranking from implicit feedback," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 30, no. 1, 2016.
32. B. Hidasi, M. Quadrana, A. Karatzoglou, and D. Tikk, "Parallel recurrent neural network architectures for feature-rich session-based recommendations," in *Proceedings of the 10th ACM Conference on Recommender Systems*, pp. 241–248, Sep. 2016, ACM.
33. S. Sedhain, A. K. Menon, S. Sanner, and L. Xie, "Autorec: Autoencoders meet collaborative filtering," in *Proceedings of the 24th International Conference on World Wide Web*, pp. 111–112, 2015.
34. R. El Youbi, F. Messaoudi, and M. Loukili, "Deep learning for dynamic content adaptation: enhancing user engagement in E-commerce," in *The International Conference on Artificial Intelligence and Smart Environment*, pp. 160–165, Cham: Springer Nature Switzerland, 2023, November.
35. R. Ying, R. He, K. Chen, P. Eksombatchai, W. L. Hamilton, and J. Leskovec, "Graph convolutional neural networks for web-scale recommender systems," in *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 974–983, 2018, ACM.
36. X. Chen, Q. Liu, Z. Gao, J. Tang, C. Wang, and X. Yan, "Attentive collaborative filtering: Multimedia recommendation with item- and component-level attention," in *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, pp. 335–344, May 2018, ACM.
37. P. Zhao, S. C. Hoi, J. Wang, B. Li, and C. Miao, "Online learning for recommendation based on heterogeneous implicit feedback," in *IJCAI International Joint Conference on Artificial Intelligence*, pp. 2909–2915, 2013.
38. C. J. Watkins and P. Dayan, "Q-learning," *Machine Learning*, vol. 8, no. 3–4, pp. 279–292, 1992.
39. V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, and S. Petersen, "Human-level control through deep reinforcement learning," *Nature*, vol. 518, no. 7540, pp. 529–533, 2015.
40. L. Li, W. Chu, J. Langford, and R. E. Schapire, "A contextual-bandit approach to personalized news article recommendation," in *Proceedings of the 19th international conference on World wide web*, pp. 661–670, Apr. 2010, ACM.
41. D. Chen, Q. Yao, L. Zhang, and Y. Yu, "A reinforcement learning framework for recommender systems," *arXiv preprint arXiv:1903.07440*, 2019.
42. "RetailRocket E-commerce dataset," Kaggle, [Online]. Available: <https://www.kaggle.com/datasets/retailrocket/e-commerce-dataset>
43. L. Li and M. Xiaoming, "Study on the exploration of poverty index's association rules based on CBCM-Apriori algorithm," *Annals of Operations Research*, pp. 1–27, 2022.
44. 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*, Lecture Notes in Networks and Systems, vol. 637, Springer, Cham, doi: 10.1007/978-3-031-26384-2_7.
45. F. Messaoudi and M. Loukili, "E-commerce Personalized Recommendations: a Deep Neural Collaborative Filtering Approach," *Operations Research Forum*, vol. 5, no. 1, pp. 1–25, Mar. 2024, Springer International Publishing.
46. M. Loukili, F. Messaoudi, and M. El Ghazi, "Implementation of Machine Learning Algorithms for Customer Churn Prediction," *Journal of Information Systems and Telecommunication*, vol. 11, no. 43, pp. 196–208, 2023, doi: 10.52547/jist.34208.11.43.196.