

Prediction of E-Commerce Customer Purchase Intention Using a Hybrid SMOTE-XGBoost Method with Model Comparison Evaluation on an Imbalanced Dataset

Rafli Rahmandani¹, Fajri Rakhmat Umbara², Ridwan Ilyas³

^{1,2,3}Universitas Jenderal Achmad Yani, Cimahi, Indonesia

Email: raflirahmandani11@gmail.com

Abstract

Class imbalance in e-commerce customer behavior data caused classification models to be biased toward the non-purchase class and made accuracy misleading. This study aimed to build and evaluate a hybrid model that combined the Synthetic Minority Oversampling Technique (SMOTE) with Extreme Gradient Boosting (XGBoost) to predict customer purchase intention on the Online Shoppers Purchasing Intention Dataset, which contained 12,330 sessions of which only 15.47% ended in a purchase. The method comprised data quality auditing, one-hot encoding, an 80:20 stratified train-test split, and standardization and oversampling placed inside the pipeline to prevent data leakage. Hyperparameters were tuned separately for the baseline and hybrid models using randomized search with stratified 5-fold cross-validation, and seven metrics were used for evaluation. The proposed model was compared with Decision Tree, Random Forest, and Support Vector Machine. On the testing data, the hybrid model achieved the highest F1-Score (0.6625) and Recall (0.6963), whereas XGBoost without oversampling obtained the best Accuracy (0.9019), Matthews Correlation Coefficient (0.5983), auROC (0.9322), and auPR (0.7493). Oversampling raised Recall by 0.1073 but lowered Precision. The hybrid model was therefore more suitable when maximizing the detection of potential buyers was the main priority, although it was not superior on every metric.

Keywords: *E-commerce, Imbalanced Dataset, Purchase Intention, SMOTE, XGBoost.*

-----◆-----

A. INTRODUCTION

The rapid growth of the e-commerce industry has shifted consumer shopping patterns from conventional transactions toward online transactions. This shift increases the need for systems capable of accurately predicting customers' purchase intention, because such information can be used to design more targeted promotions, improve product recommendations, and increase sales conversion (Rita et al., 2019; Suchacka & Chodak, 2017).

One of the main challenges in predicting purchase intention is class imbalance: the number of sessions that end in a purchase is far smaller than the number of sessions that only involve browsing. In the Online Shoppers Purchasing Intention Dataset introduced by Sakar et al. (2019), only about 15% of sessions belong to the purchase category. This imbalance causes classification models to be biased toward the majority class and to fail to optimally recognize patterns in the minority class (Mondal et al., 2021; Rayhan et al., 2017).

Previous studies on online purchase intention can be grouped into several streams. Sakar et al. (2019) introduced the benchmark dataset and tested MLP and LSTM models for real-time prediction, while Fernandes and Teixeira (2015) analyzed clickstream data from an online retailer to predict purchase intention and visitor engagement in real time. Other studies analyzed purchase intention through association rules (Suchacka & Chodak, 2017), online consumer behavior factors (Law et al., 2016), digital service quality (Rita et al., 2019), purchase intention as part of cognitive consumer behavior (Zarrad & Debabi, 2012), the influence of perceived risk on online purchase intention (Ariffin et al., 2018), and the influence of input device type (touch versus mouse) on consumer engagement and purchase decisions (Chung et al., 2018), but these studies remain focused on statistical analysis or behavioral patterns. Machine learning has been applied more broadly through SVM (Suchacka et al., 2015), Random Forest (Baati & Mohsil, 2020), XGBoost (Song & Liu, 2020), and the testing of various classification algorithms using behavioral features (Kabir et al., 2019). In terms of feature engineering, Prayogo and Karimah (2021) showed that feature selection can improve purchase intention prediction, although it has not been integrated with modern boosting algorithms. As an alternative to oversampling, Saha et al. (2022) proposed cluster-based instance selection to reduce redundancy in the training data while preserving the representation of each class, although this approach is not applied in the present study.

Most of these studies still focus on single classification algorithms such as Decision Tree, Random Forest, Support Vector Machine (SVM), or XGBoost without explicit treatment of class imbalance. On the same dataset used in this study, Hamami and Muzakki (2021) built a machine learning pipeline for purchase intention classification and reported Random Forest as the best algorithm with an accuracy of 90%, but without considering class imbalance or metrics other than Accuracy. Reliance on Accuracy is less representative on imbalanced datasets, whereas metrics such as the Matthews Correlation Coefficient (MCC) and the area under the precision-recall curve (auPR) are considered to better reflect actual performance under such conditions (Rayhan et al., 2017).

In other domains with similar imbalance patterns, combining the Synthetic Minority Oversampling Technique (SMOTE) with boosting-based ensemble algorithms has proven effective: Andespa et al. (2025) showed improvements in Recall and F1-Score for bank customer churn prediction using SMOTE with Random Forest and XGBoost, while Meng et al. (2020) demonstrated the effectiveness of combining SMOTE and XGBoost for credit card fraud detection. However, the application of this combination to e-commerce purchase intention prediction remains limited and has not been comprehensively evaluated using metrics designed for imbalanced data.

Based on this gap, this study proposes a Hybrid SMOTE-XGBoost pipeline that places standardization and SMOTE only on the training data, including inside each validation fold, to prevent the testing data from influencing the generation of synthetic samples. The novelty of this study is not a new SMOTE or XGBoost

algorithm, but rather the design of an anti-leakage pipeline, hyperparameter tuning performed separately for the baseline and hybrid models, a comprehensive evaluation using seven metrics, and the interpretation of feature importance and SHAP in the e-commerce purchase intention domain. This study thus extends the application of Hybrid SMOTE-XGBoost from the churn and fraud detection domains (Andespa et al., 2025; Meng et al., 2020) to e-commerce purchase intention prediction, so that its novelty lies in the application, the evaluation procedure, and the interpretation of results.

Accordingly, the main contributions of this study are (1) the implementation and analysis of Hybrid SMOTE-XGBoost on the Online Shoppers Purchasing Intention Dataset using a pipeline that prevents data leakage, (2) an evaluation using seven metrics relevant to imbalanced class distributions, and (3) a systematic comparison with Decision Tree, Random Forest, and SVM under the same experimental setup, complemented by an interpretation of feature contributions.

B. METHOD

This study employed a quantitative-experimental approach through computational experiments. The research stages follow the machine learning pipeline workflow from data collection to drawing conclusions, with `random_state` fixed at 42 so that the entire process is reproducible. Figure 1 shows the overall research method flow:




Figure 1. Research Method Flow

The dataset used is the Online Shoppers Purchasing Intention Dataset (Sakar et al., 2019) from the UCI Machine Learning Repository, consisting of 12,330 visit sessions and 18 attributes (17 predictor attributes and one target attribute, Revenue). The predictor attributes are a combination of numerical, categorical, and boolean data describing the number of pages visited, visit duration, page value (PageValues), bounce rate, exit rate, technical access characteristics, visit time, and visitor type. The proportion of the purchase class is only 15.47% (a majority-to-minority ratio of

approximately 5.46:1), a condition that can push a model to favor the majority class if training and evaluation are not specifically designed for it (Mondal et al., 2021; Rayhan et al., 2017).

A data quality audit was conducted before exploratory data analysis (EDA) to examine the structure, data types, missing values, duplicates, and target consistency. The audit results and the class distribution are presented in Table 1.

Table 1. Summary of Dataset Audit and Class Distribution

Component	Result
Dataset size	12,330 rows × 18 attributes
Missing values	0 (none)
Duplicate data	125 rows (1.01%), retained
Class 0 (non-purchase)	10,422 (84.53%)
Class 1 (purchase)	1,908 (15.47%)

Duplicate rows were retained because the dataset does not provide a unique session identifier, so the recurrence of identical attribute patterns is not necessarily a recording error. The audit and EDA results confirm the need for a Stratified Train-Test Split, the application of SMOTE only to the training data, and the use of metrics that do not rely solely on Accuracy.

Initial preprocessing was performed before dataset splitting and included data type conversion, one-hot encoding of the Month, VisitorType, and Weekend attributes, transformation of the target label Revenue into 0/1, and the assignment of the features (X) and the target (y). One-hot encoding increased the number of columns from 18 to 27 (26 predictor features). The dataset was then divided using an 80:20 Stratified Train-Test Split into 9,864 training samples and 2,466 testing samples, with the proportion of the purchase class preserved in both subsets (15.47% and 15.49%). Standardization of numerical features using StandardScaler was placed inside the pipeline of each model and was fitted only on the training data in every training process or validation fold, so that the validation and testing data did not influence the standardization parameters.

SMOTE generates synthetic samples for the minority class through linear interpolation between neighboring minority samples in the feature space rather than by duplicating data, as applied to class imbalance problems in other domains (Andespa et al., 2025; Meng et al., 2020). Mathematically, a new sample is formed through:

$$x_{\text{new}} = x_i + \delta \times (x_{\text{nn}} - x_i) \quad (1)$$

where x_i is a minority sample, x_{nn} is its nearest neighbor, and δ is a random number in the range 0–1.

Extreme Gradient Boosting (XGBoost) is an ensemble learning algorithm based on gradient boosting that sequentially combines a number of decision trees, with each new tree focused on the errors not yet corrected by the previous model (Song & Liu, 2020). Optimization is performed by minimizing the following objective function (Rayhan et al., 2017; Song & Liu, 2020):

$$L = \sum_i l(y_i, \hat{y}_i) + \sum_k \Omega(f_k) \quad (2)$$

where $l(y_i, \hat{y}_i)$ is the loss function and $\Omega(f_k)$ is the tree-complexity regularization component. XGBoost is widely used for customer behavior modeling because of its ability to handle many features and complex relationships (Song & Liu, 2020), but it is less sensitive to the minority class on imbalanced data and is therefore often combined with SMOTE (Andespa et al., 2025; Meng et al., 2020).

As comparison models, Decision Tree builds classification rules in the form of a tree that is easy to interpret but prone to overfitting (Kabir et al., 2019); Random Forest combines many decision trees to reduce variance and has been applied to online customer behavior prediction (Baati & Mohsil, 2020); and Support Vector Machine (SVM) separates classes through the best hyperplane, although its performance may decline under highly imbalanced class distributions (Suchacka et al., 2015).

In the Hybrid SMOTE-XGBoost pipeline, SMOTE was applied only to the training data, after StandardScaler and before XGBoost, so that synthetic samples never enter the validation or testing data. The `sampling_strategy` and `k_neighbors` parameters were searched jointly with the XGBoost hyperparameters through `RandomizedSearchCV`; the best combination obtained was a `sampling_strategy` of 0.5 and `k_neighbors` of 7, which means that the resampled minority class was made approximately 50% of the majority class (a ratio of roughly 1:2), rather than a full 50:50 balance, to avoid excessive False Positives.

The hyperparameters of the baseline XGBoost and the Hybrid SMOTE-XGBoost were optimized separately using `RandomizedSearchCV` with Stratified 5-Fold Cross Validation, 20 random parameter combinations per pipeline (100 fitting processes), and F1-Score as the scoring function, because this study emphasizes the Precision–Recall balance on the minority class. The hyperparameter search space is shown in Table 2.

Table 2. Hyperparameter Search Space

Component	Hyperparameter	Search Space
XGBoost	<code>n_estimators</code>	100; 200; 300; 500
	<code>max_depth</code>	3; 4; 5; 6
	<code>learning_rate</code>	0.01; 0.03; 0.05; 0.10
	<code>subsample</code>	0.7; 0.8; 0.9; 1.0
	<code>colsample_bytree</code>	0.7; 0.8; 0.9; 1.0
	<code>min_child_weight</code>	1; 3; 5
	<code>gamma</code>	0; 0.1; 0.3
	<code>reg_alpha</code>	0; 0.01; 0.1
	<code>reg_lambda</code>	1; 2; 5
SMOTE	<code>sampling_strategy</code>	0.5; 0.75; 1.0
	<code>k_neighbors</code>	3; 5; 7

The best configuration resulting from the search is shown in Table 3. The baseline and hybrid models required different configurations, so separate tuning improves the validity of the comparison between pipelines.

Table 3. Best Hyperparameter Configuration

Parameter	XGBoost	SMOTE-XGBoost
n_estimators	200	100
max_depth	5	6
learning_rate	0.03	0.03
subsample	1.0	0.9
colsample_bytree	0.9	0.8
min_child_weight	1	1
gamma	0	0
reg_alpha	0.1	0
reg_lambda	2	2
sampling_strategy	–	0.5
k_neighbors	–	7
Best CV F1-Score	0.6664	0.6950

Five experimental scenarios were used to compare the main model with the baseline and the comparison algorithms, as shown in Table 4. All scenarios used the same dataset split and the same classification threshold of 0.5.

Table 4. Experimental Scenarios

Scenario	Model	Balancing
S1	XGBoost	Without SMOTE
S2	SMOTE + XGBoost	SMOTE inside the pipeline
S3	Decision Tree	None
S4	Random Forest	None
S5	SVM	None

The evaluation used seven metrics: Accuracy, Precision, Recall, F1-Score (the primary metric), Matthews Correlation Coefficient (MCC), auROC, and auPR, with predicted labels determined at a classification threshold of 0.5. MCC is a more informative metric on imbalanced datasets (Rayhan et al., 2017). The equations of the main metrics are as follows:

$$Accuracy = (TP + TN) / (TP + TN + FP + FN) \quad (3)$$

$$Precision = TP / (TP + FP) \quad (4)$$

$$Recall = TP / (TP + FN) \quad (5)$$

$$F1 = 2 \times (Precision \times Recall) / (Precision + Recall) \quad (6)$$

$$MCC = (TP \cdot TN - FP \cdot FN) / \sqrt{[(TP+FP)(TP+FN)(TN+FP)(TN+FN)]} \quad (7)$$

auROC measures the model's ability to distinguish the positive and negative classes across various threshold values, whereas auPR is more sensitive to the quality of minority-class predictions and is recommended for imbalanced datasets (Mondal et al., 2021). The best model was determined using F1-Score as the primary

reference and was then verified together with Precision, Recall, MCC, auROC, auPR, the confusion matrix, and the cross-validation results.

C. RESULTS AND DISCUSSION

1. Experimental Environment and Cross-Validation Results

This study was implemented in Python 3.x using pandas, NumPy, scikit-learn, imbalanced-learn, XGBoost, Matplotlib, Seaborn, SciPy, SHAP, and joblib in a Jupyter Notebook run through Google Colaboratory.

Stratified 5-Fold Cross Validation on the training data showed that the Hybrid SMOTE-XGBoost obtained the highest mean Recall and F1-Score, whereas the baseline XGBoost excelled in Accuracy, auROC, and auPR, and SVM had the highest Precision. The complete results are presented in Table 5.

Table 5. Results of Stratified 5-Fold Cross-Validation

Model	Accuracy	Precision	Recall	F1-Score	MCC	auROC	auPR
XGBoost	0.9062	0.7423	0.6048	0.6664	0.6169	0.9342	0.7531
SMOTE-XGBoost	0.9008	0.6639	0.7300	0.6950	0.6372	0.9325	0.7453
Decision Tree	0.8788	0.6198	0.5629	0.5894	0.5197	0.8158	0.5158
Random Forest	0.9040	0.7476	0.5734	0.6489	0.6015	0.9288	0.7445
SVM	0.8934	0.7625	0.4522	0.5673	0.5339	0.8800	0.6864

Note. Bold values indicate the best result for each metric.

2. Final Evaluation on the Testing Data

The final evaluation on the testing data (Table 6) confirms a similar pattern: SMOTE-XGBoost ranked first in F1-Score (0.6625) and Recall (0.6963), whereas the baseline XGBoost remained superior in Accuracy (0.9019), MCC (0.5983), auROC (0.9322), and auPR (0.7493).

Table 6. Final Evaluation Results on the Testing Data

Model	Accuracy	Precision	Recall	F1-Score	MCC	auROC	auPR
XGBoost	0.9019	0.7258	0.5890	0.6503	0.5983	0.9322	0.7493
SMOTE-XGBoost	0.8901	0.6318	0.6963	0.6625	0.5981	0.9287	0.7303
Decision Tree	0.8832	0.6478	0.5393	0.5886	0.5242	0.8478	0.5361
Random Forest	0.9002	0.7361	0.5550	0.6328	0.5841	0.9181	0.7248
SVM	0.8844	0.7137	0.4241	0.5320	0.4917	0.8736	0.6505

Note. Bold values indicate the best result for each metric.

Figure 2 shows the performance profile of the five models across the seven metrics simultaneously. The baseline XGBoost and Random Forest are relatively stronger in Accuracy and Precision, whereas the Hybrid SMOTE-XGBoost stands out in Recall and F1-Score; in MCC, auROC, and auPR, the baseline XGBoost remains slightly higher. No single model dominates all metrics, so model selection needs to be adjusted to the priorities among the precision of positive predictions, the coverage of buyer detection, and the quality of probability discrimination.




Figure 2. Comparison of the Seven Evaluation Metrics on the Testing Data for All Models

3. Effect of SMOTE on XGBoost Performance

Table 7 summarizes the changes in performance and in the confusion matrix components of XGBoost before and after the application of SMOTE on the testing data.

Table 7. Effect of SMOTE on XGBoost Performance

Metric	XGBoost	SMOTE-XGBoost	Difference
Accuracy	0.9019	0.8901	-0.0118
Precision	0.7258	0.6318	-0.0940
Recall	0.5890	0.6963	+0.1073
F1-Score	0.6503	0.6625	+0.0122
MCC	0.5983	0.5981	-0.0002
auROC	0.9322	0.9287	-0.0035
auPR	0.7493	0.7303	-0.0190
TN	1,999	1,929	-70
FP	85	155	+70
FN	157	116	-41
TP	225	266	+41

Applying SMOTE increased Recall by 10.73 percentage points and F1-Score by 1.22 points, but decreased Accuracy, Precision, auROC, and auPR and yielded a practically identical MCC. In other words, the hybrid model found 41 additional buyers (True Positives rose from 225 to 266, and False Negatives fell from 157 to 116), but also produced 70 additional False Positives (rising from 85 to 155). The hypothesis that SMOTE improves performance in general is therefore not fully supported; SMOTE provides a partial improvement in sensitivity and in the Precision–Recall balance rather than a universal improvement across all metrics. This is in line with findings in the churn and fraud detection domains (Andespa et

al., 2025; Meng et al., 2020), although in this study MCC, auROC, and auPR did not improve as well.

From the perspective of operational trade-offs, the Hybrid SMOTE-XGBoost is more suitable when the cost of losing a potential buyer (False Negative) is greater than the cost of reaching a non-buyer (False Positive), for example in digital campaigns with low contact costs. Conversely, the baseline XGBoost is more suitable when the cost of mis-targeting is high or follow-up capacity is limited, because its Precision and auPR are better. The choice of the best model should therefore be adjusted to business priorities, not merely to the ranking of F1-Score values.

4. Feature Importance and Interpretation

Correlation analysis showed that PageValues has the most prominent positive relationship with Revenue, whereas ExitRates and BounceRates are negatively oriented. The feature importance of the Hybrid SMOTE-XGBoost model (Figure 3) confirms PageValues as the most dominant feature (importance 0.6135), followed by Month_Nov, Month_Sep, Administrative, Month_May, Month_Mar, OperatingSystems, ExitRates, TrafficType, and VisitorType_Returning_Visitor. The SHAP interpretation reinforces this finding, with contribution directions that are consistent with the feature importance and correlation results.




Figure 3. Top Ten Feature Importances of the Hybrid SMOTE-XGBoost Model
Model Selection

With F1-Score as the primary reference, verified together with Recall, the cross-validation results, the confusion matrix, and the consequences of False Positives and False Negatives, the Hybrid SMOTE-XGBoost was determined to be the best model for research objectives that prioritize the coverage of potential-buyer detection. This decision is contextual: the baseline XGBoost remains the more appropriate choice when Accuracy, Precision, MCC, auROC, and auPR are the priority. The improvement in Recall and F1-Score after oversampling in this study is in line with findings in the banking churn domain (Andespa et al., 2025) and in fraud detection (Meng et al., 2020); however, MCC, auROC, and auPR did not improve, which indicates that the benefit of balancing is specific to the data domain,

the oversampling ratio, and the threshold used, and which also underscores the novelty of this study in the e-commerce purchase intention domain and in its evaluation procedure.

D. CONCLUSION

This study successfully built and evaluated an anti-leakage Hybrid SMOTE-XGBoost pipeline for predicting e-commerce customer purchase intention on the Online Shoppers Purchasing Intention Dataset, which has a class imbalance of 84.53% versus 15.47%. On the testing data, SMOTE-XGBoost obtained the highest F1-Score (0.6625) and the highest Recall (0.6963) among the five models tested, whereas the baseline XGBoost remained superior in Accuracy (0.9019), MCC (0.5983), auROC (0.9322), and auPR (0.7493). Applying SMOTE increased the Recall of XGBoost by 0.1073 and its F1-Score by 0.0122, but decreased Precision, Accuracy, auROC, and auPR, so SMOTE was shown to provide a partial improvement in minority-class sensitivity rather than a universal performance improvement. The Hybrid SMOTE-XGBoost is determined to be the best model when the main priority is maximizing the coverage of potential-buyer detection, whereas the baseline XGBoost is more suitable when the precision of positive predictions is the priority.

Future research is recommended to apply nested cross-validation and external validation on other e-commerce datasets, to compare SMOTE with other balancing techniques such as Borderline-SMOTE, ADASYN, or SMOTE-Tomek, to optimize the classification threshold and conduct cost-sensitive analysis according to business needs, and to explore additional algorithms such as LightGBM or CatBoost while strictly maintaining the separation of the training and testing data to prevent data leakage.

REFERENCES

- Andespa, R., Sadik, K., Suhaeni, C., & Soleh, A. M. (2025). Evaluating random forest and XGBoost for bank customer churn prediction on imbalanced data using SMOTE and SMOTE-ENN. *Media Statistika*, 18(1), 25–36.
- Ariffin, S. K., Mohan, T., & Goh, Y.-N. (2018). Influence of consumers' perceived risk on consumers' online purchase intention. *Journal of Research in Interactive Marketing*, 12(3), 309–327.
- Baati, K., & Mohsil, M. (2020). Online shopper purchase prediction using random forest. In *Proceedings of the IFIP International Conference on Artificial Intelligence Applications and Innovations* (pp. 43–51).
- Chung, S., Kramer, T., & Wong, E. M. (2018). Do touch interface users feel more engaged? The impact of input device type on online shoppers' engagement, affect, and purchase decisions. *Psychology & Marketing*, 35(11), 795–806.
- Fernandes, R. F., & Teixeira, C. M. (2015). *Using clickstream data to analyze online purchase intentions* [Master's thesis, University of Porto].
- Hamami, F., & Muzakki, A. (2021). Machine learning pipeline for online shopper intention classification. *AIP Conference Proceedings*, 2329(1), 050014.

7. Kabir, M. R., Ashraf, F. B., & Ajwad, R. (2019). Predicting purchase intention from empirical data. In *Proceedings of the IEEE ICCIT* (pp. 1–6).
8. Law, M., Kwok, R. C.-W., & Ng, M. (2016). An extended online purchase intention model for middle-aged online users. *Electronic Commerce Research and Applications*, 20, 132–146.
9. Meng, C., Zhou, L., & Liu, B. (2020). A case study in credit fraud detection with SMOTE and XGBoost. *Journal of Physics: Conference Series*, 1601(5).
10. Mondal, I. A., Haque, M. E., Hassan, A.-M., & Shatabda, S. (2021). Handling imbalanced data for credit card fraud detection. In *Proceedings of the IEEE ICCIT*.
11. Prayogo, R. D., & Karimah, S. A. (2021). Feature selection for purchase intent prediction. In *Proceedings of the ICAICTA* (pp. 1–5).
12. Rayhan, F., Ahmed, S., Shatabda, S., Farid, D. M., Mousavian, Z., Dehzangi, A., & Rahman, M. S. (2017). iDTI-ESBoost: Identification of drug target interaction using evolutionary and structural features with boosting. *Scientific Reports*, 7(1), 1–18.
13. Rita, P., Oliveira, T., & Farisa, A. (2019). The impact of e-service quality and satisfaction on online shopping behavior. *Heliyon*, 5(10).
14. Saha, S., Sarker, P. S., Al Saud, A., Shatabda, S., & Newton, M. A. H. (2022). Cluster-oriented instance selection for classification problems. *Information Sciences*, 602, 143–158.
15. Sakar, C. O., Polat, S. O., Katircioglu, M., & Kastro, Y. (2019). Real-time prediction of online shoppers' purchasing intention using multilayer perceptron and LSTM recurrent neural networks. *Neural Computing and Applications*, 31(10), 6893–6908.
16. Song, P., & Liu, Y. (2020). XGBoost algorithm for predicting behavior on e-commerce. *Tehnički Vjesnik*, 27(5), 1467–1471.
17. Suchacka, G., & Chodak, G. (2017). Using association rules to assess purchase probability in online stores. *Information Systems and E-Business Management*, 15(3), 751–780.
18. Suchacka, G., Skolimowska-Kulig, M., & Potempa, A. (2015). Classification of e-customer sessions using SVM. In *Proceedings of the ECMS* (pp. 594–600).
19. Zarrad, H., & Debabi, M. (2012). Online purchasing intention: Factors and effects. *International Business Management*, 4(1), 37–47.