Article

ARTIFICIAL INTELLIGENCE IN MARKETING ANALYSIS AND PREDICTIVE CONSUMER BEHAVIOUR MODELLING: AN EMPIRICAL MACHINE-LEARNING CASE STUDY

Inteligência artificial na análise de marketing e na modelagem preditiva do comportamento do consumidor: um estudo de caso empírico com aprendizado de máquina

Inteligencia artificial en el análisis de marketing y en el modelado predictivo del comportamiento del consumidor: un estudio de caso empírico con aprendizaje automático

Nupur Chakraborty
Doctorate in Business Administration (DBA) Student, Western State University (WSU), California, EUA.
https://orcid.org/0009-0002-7028-7986
E-mail: nupurchakraborty2@gmail.com

Histórico do artigo
Recebido em: 12 de junho de 2026
Aceito em: 13 de agosto de 2026
Revisado em: 20 de agosto de 2026
Publicado em: 3 de setembro de 2026

Editor responsável: Dr. Tarsis Barreto Oliveira
Editor assistente: Dr. Leonardo de Andrade Carneiro

ABSTRACT

Artificial Intelligence (AI) has expanded the analytical capabilities of marketing by enabling organisations to process large volumes of consumer data, identify behavioural patterns, and support predictive decision-making. This study examines the application of AI and machine-learning techniques to marketing analysis through a critical review of the literature and an empirical case study based on transactional consumer data. The empirical analysis compares Logistic Regression, Random Forest, Support Vector Machine (SVM), and XGBoost using accuracy, precision, recall, F1-score, and ROC-AUC as evaluation metrics. Among the evaluated models, XGBoost achieved the strongest overall performance, with an accuracy of 83.87%, precision of 81.79%, F1-score of 79.19%, and ROC-AUC of 0.9126, while Logistic Regression achieved the highest recall, at 80.33%. The findings indicate that machine-learning methods can identify predictive associations in transactional data and support consumer classification, although such associations should not be interpreted as

| ISSN eletrônico: 2447-9896 | ISSN impresso: 2177-0360

causal effects. The study also discusses methodological limitations related to the behavioural proxy adopted, the available predictors, and the generalisability of the findings. Overall, the results indicate the potential contribution of AI-based predictive analytics to marketing decision-making while emphasising the need for transparent, interpretable, and ethically responsible applications.

Keywords: Artificial Intelligence; Marketing Analytics; Machine Learning; Consumer Behaviour; Predictive Modelling.

RESUMO

A Inteligência Artificial (IA) ampliou as capacidades analíticas do marketing ao permitir que as organizações processem grandes volumes de dados de consumidores, identifiquem padrões comportamentais e apoiem a tomada de decisões preditivas. Este estudo examina a aplicação da IA e de técnicas de aprendizado de máquina à análise de marketing por meio de uma revisão crítica da literatura e de um estudo de caso empírico baseado em dados transacionais de consumidores. A análise empírica compara Regressão Logística, Random Forest, Support Vector Machine (SVM) e XGBoost, utilizando acurácia, precisão, recall, F1-score e ROC-AUC como métricas de avaliação. Entre os modelos avaliados, o XGBoost apresentou o melhor desempenho geral, com acurácia de 83,87%, precisão de 81,79%, F1-score de 79,19% e ROC-AUC de 0,9126, enquanto a Regressão Logística apresentou o maior recall, de 80,33%. Os resultados indicam que os métodos de aprendizado de máquina podem identificar associações preditivas em dados transacionais e apoiar a classificação de consumidores, embora tais associações não devam ser interpretadas como efeitos causais. O estudo também discute limitações metodológicas relacionadas à proxy comportamental adotada, aos preditores disponíveis e à generalização dos resultados. De modo geral, os resultados indicam a contribuição potencial da análise preditiva baseada em IA para a tomada de decisões de marketing, ao mesmo tempo que ressaltam a necessidade de aplicações transparentes, interpretáveis e eticamente responsáveis.

Palavras-chave: Inteligência Artificial; Análise de Marketing; Aprendizado de Máquina; Comportamento do Consumidor; Modelagem Preditiva.

1 INTRODUCTION

Artificial Intelligence (AI) has become increasingly integrated into marketing analysis as organisations seek to transform large volumes of consumer data into information capable of supporting decision-making. The expansion of digital environments has increased the availability of transactional, behavioural, and interaction data, creating conditions for more

| ISSN eletrônico: 2447-9896 | ISSN impresso: 2177-0360

granular forms of consumer analysis and for the development of data-driven marketing strategies (Wedel; Kannan, 2016).

Within marketing, AI encompasses computational techniques capable of supporting activities such as consumer segmentation, personalisation, recommendation, classification, demand analysis, and predictive modelling. Machine-learning methods are particularly relevant because they allow models to identify statistical patterns from historical observations and apply those patterns to previously unseen data. This analytical capacity has contributed to the transition from predominantly descriptive approaches towards predictive forms of marketing analytics (Davenport *et al*., 2020; Huang; Rust, 2021).

The increasing use of predictive analytics also changes the role of consumer data in marketing decision-making. Rather than relying exclusively on aggregated historical indicators, organisations can use transactional and behavioural information to identify patterns associated with consumer outcomes and support decisions concerning targeting, customer relationships, product recommendations, and resource allocation (Wedel; Kannan, 2016). In this context, AI-based analytics can expand the capacity to process heterogeneous data and identify relationships that may be difficult to represent through conventional analytical procedures.

AI-driven personalisation constitutes another relevant application of these technologies. By analysing consumer interactions and behavioural histories, algorithmic systems can differentiate recommendations, communications, and offers according to patterns identified in available data. However, the increasing use of personal information also creates tensions involving privacy, trust, transparency, and consumer perceptions regarding the collection and use of their data (Martin; Murphy, 2017).

The analytical capabilities of AI should nevertheless be distinguished from causal explanation. Machine-learning algorithms can identify associations between observed characteristics and specified outcomes, but predictive performance does not establish that a particular variable or marketing action causes a change in consumer behaviour. Prediction and causal inference address different analytical questions and require different research designs and assumptions. Consequently, high classification performance should be interpreted as

| ISSN eletrônico: 2447-9896 | ISSN impresso: 2177-0360

evidence of predictive association unless the empirical design supports causal identification.

The adoption of AI in marketing also raises methodological and ethical questions concerning data quality, model interpretability, privacy, algorithmic bias, and transparency. As increasingly complex algorithms are incorporated into organisational decision processes, predictive accuracy must be considered alongside the capacity to understand how models generate their outputs and under which conditions those outputs can be appropriately interpreted. The literature on interpretable machine learning has emphasised that predictive performance alone is insufficient when algorithmic results are expected to support human decision-making (Doshi-Velez; Kim, 2017).

These issues are particularly relevant in consumer-behaviour modelling because computational systems generally do not directly observe latent psychological constructs such as intention, motivation, trust, or preference. Instead, models commonly operate with observable behavioural indicators and proxies derived from transactions, ratings, interactions, and other digital traces. The relationship between the theoretical construct and its empirical operationalisation therefore requires explicit justification so that predictive performance is not interpreted as evidence for a construct that the available data do not directly measure. Against this background, this study examines how Artificial Intelligence and machine-learning techniques can contribute to marketing analysis and predictive consumer-behaviour modelling. It combines a critical examination of the literature with an empirical case study based on transactional consumer data. The empirical component compares four supervised machine-learning approaches: Logistic Regression, Random Forest, Support Vector Machine (SVM), and XGBoost. Model performance is evaluated using accuracy, precision, recall, F1-score, and the area under the Receiver Operating Characteristic curve (ROC-AUC).

The comparative design allows the study to examine whether more flexible machine-learning approaches provide improvements over a conventional statistical baseline when applied to the same consumer-classification problem. At the same time, the analysis explicitly distinguishes predictive association from causal inference and considers the methodological limitations associated with the construction of the behavioural outcome used in the empirical model. Accordingly, the study addresses the following research question: *To*

| ISSN eletrônico: 2447-9896 | ISSN impresso: 2177-0360

what extent can machine-learning models identify predictive patterns in transactional consumer data that support consumer-behaviour classification for marketing analysis?

To answer this question, the study pursues three objectives: first, to examine the theoretical and analytical contributions of AI to contemporary marketing analysis; second, to compare the predictive performance of four supervised-learning algorithms using a common empirical dataset and evaluation framework; and third, to discuss the managerial, methodological, and ethical implications of applying predictive models to consumer data.

The contribution of the study lies in connecting the conceptual discussion of AI-based marketing analytics with an empirical comparison of predictive models. Rather than treating AI solely as an abstract technological capability, the analysis evaluates how different algorithms perform under the same classification framework while explicitly considering the distinction between predictive performance and causal explanation. This approach enables a more circumscribed assessment of the applicability and limitations of machine-learning techniques in consumer analytics.

The remainder of the article is organised as follows. Section 2 reviews the literature on Artificial Intelligence, marketing analytics, consumer behaviour, predictive modelling, and responsible AI. Section 3 describes the research design, dataset, preprocessing procedures, classification criterion, machine-learning models, and evaluation metrics. Section 4 presents the empirical results. Section 5 discusses the findings in relation to the literature. Section 6 examines theoretical, managerial, and ethical implications. Section 7 presents the limitations and directions for future research, and Section 8 concludes the study.

2 LITERATURE REVIEW

2.1 Artificial Intelligence and Marketing Analytics

Artificial Intelligence has expanded the analytical capabilities of marketing by enabling organisations to process large and heterogeneous datasets and identify patterns associated with consumer behaviour. The increasing availability of data generated through digital commerce, social media, online retailing, and multichannel interactions has contributed to the transition from predominantly descriptive analytics towards predictive and

| ISSN eletrônico: 2447-9896 | ISSN impresso: 2177-0360

decision-oriented approaches (Wedel; Kannan, 2016). In addition to structured transactional data, computational approaches can extract information from unstructured consumer-generated content through techniques such as opinion mining and sentiment analysis (Pang; Lee, 2008; Cambria *et al*., 2013).

In this context, machine learning allows historical observations to be used to identify statistical relationships and support activities such as consumer classification, segmentation, recommendation, and predictive modelling. These applications extend conventional marketing analytics by enabling the simultaneous analysis of multiple consumer and transactional characteristics, including relationships that may not be adequately represented through simpler analytical procedures (Wedel; Kannan, 2016).

However, the analytical capacity of AI depends on the quality, representativeness, and theoretical relevance of the data used. Predictive models learn from observed historical patterns and therefore remain subject to limitations arising from incomplete information, changes in consumer behaviour, inappropriate operationalisation of variables, and the specific context from which the data were obtained. Accordingly, predictive performance should be interpreted within the empirical boundaries of the dataset and research design.

2.2 AI-Driven Personalisation and Consumer Behaviour

AI-driven personalisation is an important application of artificial intelligence in digital marketing. Personalisation systems use consumer and behavioural data to differentiate recommendations, communications, and offers according to patterns identified in previous interactions. Transaction history, product preferences, ratings, purchase frequency, and other observable characteristics can consequently provide information for data-driven consumer analysis.

The increasing use of consumer information, however, creates a tension between the analytical benefits of personalisation and concerns regarding privacy, transparency, and trust. Data-driven marketing practices depend not only on the capacity to extract information from consumer data but also on how such information is collected, processed, and used. Responsible data practices therefore affect the relationship between organisations and

| ISSN eletrônico: 2447-9896 | ISSN impresso: 2177-0360

consumers and should be considered alongside analytical performance (Martin; Murphy, 2017).

A further methodological issue concerns the distinction between observable behaviour and latent psychological constructs. Transactional datasets provide evidence about realised actions and outcomes, but they generally do not directly measure motivations, attitudes, preferences, or intentions. Consequently, behavioural indicators used as proxies for psychological constructs require explicit conceptual justification, and the interpretation of empirical results should remain consistent with what the available data actually measure.

This distinction is particularly relevant when analysing purchase intention. Purchase intention ordinarily refers to a consumer's willingness or likelihood to purchase a product or service and may be measured directly through stated preferences or intention scales. When secondary transactional data are used instead, observed outcomes can provide behavioural indicators related to purchasing behaviour, but they should not automatically be considered equivalent to psychological intention. Nevertheless, historical transactional behaviour can provide relevant information for customer-base analysis and the identification of changes in behavioural loyalty within non-contractual retail settings (Buckinx; Van den Poel, 2005).

2.3 Predictive Consumer Behaviour Modelling

Predictive consumer-behaviour modelling uses historical observations to identify patterns associated with specified consumer outcomes. Supervised machine-learning methods are suitable for this purpose because models can be trained on observations with previously defined outcomes and subsequently evaluated on data not used during model estimation.

Different modelling approaches may capture different aspects of the underlying data structure. Relatively interpretable statistical models can provide useful baselines, whereas more flexible machine-learning methods can represent non-linear relationships and interactions among predictors. Comparing alternative models under the same analytical conditions therefore provides a means of assessing whether increased model flexibility is accompanied by improved predictive performance.

| ISSN eletrônico: 2447-9896 | ISSN impresso: 2177-0360

Recommender systems represent a specific application of data-driven consumer analytics, using information about users, items, and previous interactions to support personalised recommendations (Ricci; Rokach; Shapira, 2015). Nevertheless, predictive performance and causal explanation address different analytical questions. A model may accurately classify consumer outcomes without demonstrating that the variables used to generate the prediction caused those outcomes. Predictive associations should therefore not be interpreted as causal effects unless the research design and analytical strategy provide a basis for causal identification.

This distinction also affects model interpretation. Complex machine-learning methods may improve predictive performance while making the relationship between predictors and outcomes more difficult to explain. Interpretability is therefore relevant when algorithmic outputs are used to support managerial decisions, particularly when decision-makers need to understand the basis and limitations of model predictions (Doshi-Velez; Kim, 2017).

Model assessment should consequently consider more than overall classification accuracy. The evaluation of predictive performance requires complementary measures capable of capturing different types of classification error and discriminatory capacity. The specific metrics and algorithms adopted in this study are presented in the methodological section rather than in the literature review.

2.4 Methodological and Ethical Boundaries of AI-Based Consumer Analytics

The application of AI to consumer analytics involves methodological and ethical constraints that affect the interpretation and practical use of predictive models. One limitation concerns dependence on historical data. Relationships identified from past transactions may not remain stable when consumer preferences, market conditions, product characteristics, or digital environments change. Predictive models therefore require evaluation beyond the data on which they were originally developed (Wedel; Kannan, 2016).

Another concern involves explainability. Complex algorithms can identify predictive structures that are difficult to translate into substantively interpretable relationships. This creates a distinction between knowing that a model discriminates effectively between classes

| ISSN eletrônico: 2447-9896 | ISSN impresso: 2177-0360

and understanding why particular observations receive particular predictions. Interpretability becomes especially relevant when predictive outputs inform organisational decisions affecting consumers (Doshi-Velez; Kim, 2017).

Privacy and responsible data use constitute an additional constraint. Extensive collection and analysis of behavioural information can generate concerns regarding surveillance, intrusive personalisation, discrimination, and consumer trust. Responsible AI-based marketing therefore requires attention to legitimate data use, transparency, security, and the proportionality of data-driven interventions (Martin; Murphy, 2017).

These methodological and ethical considerations establish the interpretive boundaries for empirical predictive analysis. Machine-learning models can identify useful statistical patterns in transactional data, but the validity of their conclusions depends on the correspondence between theoretical constructs and observable variables, the quality of the data, the adequacy of model evaluation, and the distinction between prediction and causal explanation. The empirical analysis developed in the following sections is examined within these boundaries.

3 RESEARCH METHODOLOGY

3.1 Research Design

This study adopts a quantitative, exploratory case-study design to examine the application of AI-enabled predictive analytics to consumer behaviour modelling. The research combines a literature-based conceptual foundation with an empirical analysis of transactional retail data.

The case-study design enables the predictive modelling process to be examined through a defined analytical sequence comprising data preparation, variable construction, consumer classification, model training, and performance evaluation. The empirical analysis focuses on predictive association rather than causal inference. Accordingly, the results indicate the extent to which observed transactional characteristics distinguish High Purchase Intention from Low Purchase Intention observations, without assuming that these characteristics cause the observed outcome.

| ISSN eletrônico: 2447-9896 | ISSN impresso: 2177-0360

3.2 Dataset and Case Context

The empirical analysis uses the Amazon India Sales 2025 dataset, which contains transaction-level retail records. The available attributes include Product Category, Quantity, Unit Price (INR), Total Sales (INR), Payment Method, Delivery Status, Review Rating, State, and other transaction-related information.

The dataset provides an empirical basis for examining predictive consumer behaviour because observable transactional characteristics can be transformed into variables suitable for machine-learning classification. However, the transactional nature of the data establishes limits on interpretation. The dataset represents observed purchasing outcomes and does not directly measure psychological motivations, attitudes, consumer satisfaction, or individual exposure to proprietary AI-driven personalisation systems.

3.3 Data Preparation and Preprocessing

The analytical process began with data preparation and preprocessing to ensure consistent representation of the variables across the machine-learning models. Six predictor variables were retained in the final analysis:

1. Product Category;
2. Quantity;
3. Unit Price (INR);
4. Total Sales (INR);
5. Payment Method;
6. State.

The categorical variables, Product Category, Payment Method, and State, were transformed using one-hot encoding. The numerical variables, Quantity, Unit Price (INR), and Total Sales (INR), were standardised using StandardScaler. These procedures generated a consistent numerical feature matrix for the classification models and reduced the influence of differences in measurement scale on scale-sensitive algorithms.

3.4 Consumer Classification Criterion

The target variable, Purchase_Intention_Label, was derived from observable Delivery Status and Review Rating. An observation was classified as:

| ISSN eletrônico: 2447-9896 | ISSN impresso: 2177-0360

1 (High Purchase Intention): when the order status was *Delivered* and the review rating was 4 or 5.

0 (Low Purchase Intention): when the order was *Delivered* with a review rating between 1 and 3, or when the order status was *Pending* or *Returned*.

This operationalisation converts observable transactional outcomes into a binary behavioural proxy suitable for supervised classification. The resulting distribution comprised approximately 40% High Purchase Intention and 60% Low Purchase Intention observations. Because the imbalance was moderate rather than severe, the original class distribution was retained without artificial resampling.

The classification criterion requires a specific interpretative qualification. Purchase_Intention_Label represents a behavioural proxy and not a direct psychological measurement of purchase intention. Delivery status and review ratings represent observed transactional outcomes and therefore do not directly measure the cognitive or motivational processes underlying consumer intention. This distinction is maintained throughout the interpretation of the empirical findings.

3.5 Predictive Modelling Framework

A supervised binary-classification framework was employed to assess whether the selected transactional variables could distinguish between High Purchase Intention and Low Purchase Intention observations. Four algorithms were evaluated:

1. Logistic Regression;
2. Random Forest;
3. Support Vector Machine (SVM);
4. Extreme Gradient Boosting (XGBoost).

The models represent different classification approaches. Logistic Regression provides a comparatively interpretable linear baseline. Random Forest combines multiple decision trees and can capture non-linear relationships and interactions among predictors. SVM constructs decision boundaries between classes and can represent non-linear relationships through its kernel-based formulation. XGBoost uses sequentially constructed gradient-boosted decision trees to model complex relationships and interactions in structured data.

All four models were trained using the same predictor variables and evaluated under the same training-testing framework. This common analytical structure permits direct

| ISSN eletrônico: 2447-9896 | ISSN impresso: 2177-0360

comparison of predictive performance while limiting differences attributable to preprocessing or evaluation procedures. The modelling exercise was designed to assess the predictive information contained in the selected transactional characteristics. Consequently, model performance is interpreted as evidence of predictive association, not as evidence that any individual predictor or AI-based personalisation mechanism causes purchase intention.

3.6 Training and Testing Procedure

The dataset was divided into training and testing subsets using an 80:20 train-test split. Stratified sampling was applied to preserve the relative distribution of High Purchase Intention and Low Purchase Intention observations in both subsets. A fixed random state of 42 was used to support reproducibility. The training subset was used to fit the four classification models, whereas the independent testing subset was reserved for evaluating predictive performance on observations not used during model fitting.

This separation provides a basis for assessing model generalisation beyond the observations used for training and ensures that all four algorithms are compared using the same testing conditions.

3.7 Evaluation Metrics

Model performance was evaluated using five complementary classification metrics: Accuracy, Precision, Recall, F1-score, and ROC-AUC. Accuracy represents the proportion of observations correctly classified by the model. Precision represents the proportion of observations predicted as High Purchase Intention that were actually classified as High Purchase Intention.

Recall represents the proportion of actual High Purchase Intention observations correctly identified by the model.

F1-score combines precision and recall and provides a balanced measure when both false-positive and false-negative classifications are relevant.

ROC-AUC evaluates the model's ability to discriminate between the two classes across different classification thresholds.

| ISSN eletrônico: 2447-9896 | ISSN impresso: 2177-0360

The use of multiple metrics is appropriate because the target classes are not perfectly balanced. Accuracy alone could conceal differences in the models' ability to identify High Purchase Intention observations. Precision, recall, F1-score, and ROC-AUC therefore provide complementary evidence concerning classification performance.

4 RESULTS

4.1 Comparative Model Performance

The comparative performance of the four machine-learning models is presented in Table 1.

Table 1 – Comparative performance of the machine-learning models

Model	*Accuracy*	*Precision*	*Recall*	*F1-score*	*ROC-AUC*
Logistic Regression	80.43%	73.31%	**80.33%**	76.66%	0.8934
Random Forest	80.27%	73.14%	80.08%	76.45%	0.8786
SVM	82.20%	76.73%	79.67%	78.17%	0.8986
XGBoost	**83.87%**	**81.79%**	76.75%	**79.19%**	**0.9126**

Source: Research data.

All four models achieved accuracy above 80%, indicating that the selected transactional variables contain predictive information for distinguishing High Purchase Intention from Low Purchase Intention observations within the adopted classification framework. XGBoost achieved the strongest overall performance, with 83.87% accuracy, 81.79% precision, 76.75% recall, 79.19% F1-score, and a ROC-AUC of 0.9126. It obtained the highest accuracy, precision, F1-score, and ROC-AUC among the four evaluated models.

SVM achieved the second-highest accuracy (82.20%) and ROC-AUC (0.8986), followed by Logistic Regression, with 80.43% accuracy and 0.8934 ROC-AUC. Random Forest presented similar accuracy (80.27%), although its ROC-AUC of 0.8786 was the lowest among the evaluated models. The results also reveal a trade-off between precision and recall. Although XGBoost achieved substantially higher precision (81.79%), Logistic Regression obtained the highest recall (80.33%), followed by Random Forest (80.08%), SVM (79.67%), and XGBoost (76.75%). Thus, no single metric should be considered independently when assessing model suitability.

| ISSN eletrônico: 2447-9896 | ISSN impresso: 2177-0360

4.2 Comparative Interpretation of the Models

The performance of Logistic Regression indicates that a considerable portion of the predictive structure can be captured through a relatively simple linear classification approach. Its accuracy of 80.43% and ROC-AUC of 0.8934 suggest that the selected predictors contain substantial discriminatory information even without a more flexible modelling architecture.

Random Forest produced a comparable accuracy of 80.27%, but its ROC-AUC was lower than that of the other models. This result indicates weaker discrimination between the two classes across classification thresholds, despite an overall accuracy similar to that obtained by Logistic Regression.

SVM improved upon both Logistic Regression and Random Forest in terms of overall accuracy and F1-score, reaching 82.20% and 78.17%, respectively. Its ROC-AUC of 0.8986 also indicates strong discriminatory capacity. XGBoost produced the strongest combined performance. Its ROC-AUC of 0.9126 indicates the highest discriminatory capacity among the evaluated models, while its precision of 81.79% suggests a lower proportion of false-positive classifications among observations identified as High Purchase Intention. Its lower recall, however, shows that this improvement involves a trade-off: XGBoost identifies fewer actual High Purchase Intention observations than the other models.

This distinction has practical implications for model selection. Where the analytical objective prioritises the reduction of false-positive targeting, higher precision may be preferred. Conversely, where the objective is to identify the largest possible proportion of consumers belonging to the positive class, recall becomes comparatively more relevant. Model selection should therefore reflect the analytical objective rather than relying exclusively on overall accuracy.

4.3 Relationship Between the Empirical Findings and the Theoretical Framework

The empirical findings indicate that observable transactional characteristics contain predictive information associated with the behavioural outcome operationalised as Purchase_Intention_Label. The fact that all four models achieved accuracy above 80%

| ISSN eletrônico: 2447-9896 | ISSN impresso: 2177-0360

suggests that the classification task contains identifiable structure rather than random variation.

The comparatively strong performance of Logistic Regression indicates that part of this structure can be represented through a relatively simple classification model. At the same time, the higher overall performance of XGBoost suggests that more flexible modelling may capture additional relationships among the selected variables. This interpretation should remain cautious, however, because superior predictive performance alone does not identify the specific non-linear relationships or interactions responsible for the difference between models.

The results are consistent with the broader literature on data-rich marketing environments, which identifies predictive analytics as a means of extracting decision-relevant patterns from consumer and transactional data (Wedel; Kannan, 2016). In this study, the empirical comparison demonstrates that different modelling approaches can extract predictive information from the same set of observed transactional characteristics.

Nevertheless, the findings do not demonstrate that the selected predictors cause High Purchase Intention. The models estimate predictive associations within the available dataset and classification framework. Establishing causal relationships would require a different research design, such as experimental, quasi-experimental, or longitudinal approaches capable of isolating the effects of specific marketing interventions.

This qualification is particularly relevant because the study does not directly observe consumers' exposure to AI-generated recommendations or proprietary personalisation mechanisms. Consequently, the results should not be interpreted as evidence that AI-based personalisation caused the classified behavioural outcomes. Rather, the machine-learning analysis demonstrates the capacity of the selected models to identify patterns in observable transactional data within a digital-commerce context.

4.4 Analytical Contribution of the Empirical Case Study

The empirical component connects the conceptual discussion of AI-enabled marketing analytics with an operational predictive-modelling exercise. Its contribution lies in comparing

| ISSN eletrônico: 2447-9896 | ISSN impresso: 2177-0360

a relatively interpretable statistical baseline with alternative machine-learning approaches under the same classification framework, predictors, training-testing procedure, and evaluation metrics.

The results also demonstrate why model evaluation should not rely on accuracy alone. XGBoost achieved the highest overall accuracy, precision, F1-score, and ROC-AUC, whereas Logistic Regression achieved the highest recall. The comparison therefore shows that the relative suitability of a predictive model depends partly on the type of classification error considered more consequential for the intended application.

The analysis further reinforces the distinction between prediction and explanation. The models can identify statistical patterns associated with the operationalised behavioural outcome, but predictive performance does not establish the psychological mechanisms underlying consumer decisions. Accordingly, the empirical contribution should be understood as evidence concerning classification performance rather than as a causal model of consumer behaviour.

4.5 Managerial Implications

The findings indicate that transactional data can support data-driven consumer classification and provide information potentially useful for marketing decision-making. Predictive models may assist organisations in identifying observations associated with higher probabilities of the target outcome and, consequently, support activities such as targeting, prioritisation, customer analysis, and allocation of marketing resources.

The comparative results also show that model choice should be aligned with managerial objectives. XGBoost may be preferable when higher precision and overall discriminatory performance are prioritised. By contrast, a model with higher recall may be preferable when the cost of failing to identify a positive observation is comparatively high.

These applications should nevertheless be treated as decision-support mechanisms rather than autonomous decision rules. The study does not empirically evaluate the effects of campaign targeting, retention interventions, recommendation systems, or budget allocation.

| ISSN eletrônico: 2447-9896 | ISSN impresso: 2177-0360

These represent potential applications of the predictive framework rather than outcomes directly demonstrated by the present analysis.

4.6 Ethical and Responsible Use of Predictive Analytics

The use of predictive consumer analytics requires attention to privacy, transparency, interpretability, and responsible data governance. Predictive accuracy does not by itself establish that a model should be deployed in all marketing contexts. The collection and processing of consumer information may affect trust and raise concerns regarding intrusive personalisation and the appropriate use of behavioural data (Martin; Murphy, 2017).

Interpretability is also relevant when model outputs inform organisational decisions. More complex algorithms may provide improved predictive performance while reducing the transparency of the relationships between inputs and outputs. This creates a need to balance predictive performance with the ability to understand and communicate the basis and limitations of algorithmic decisions (Doshi-Velez; Kim, 2017).

Accordingly, organisations using predictive analytics should consider model performance alongside data quality, privacy safeguards, human oversight, monitoring for potential bias, and transparency regarding the role of automated systems in decision-making.

4.7 Limitations of the Empirical Analysis

The empirical findings are subject to several limitations. First, the analysis is based on historical transactional data. Consumer behaviour may change over time as market conditions, preferences, technologies, and competitive environments evolve. The predictive relationships identified in the dataset therefore cannot be assumed to remain stable indefinitely.

Second, *Purchase_Intention_Label* is a behavioural proxy derived from Delivery Status and Review Rating rather than a direct psychological measure of purchase intention. The results consequently concern the operationalised classification outcome and should not be interpreted as direct measurement of consumers' underlying intentions.

Third, the predictor set is limited to Product Category, Quantity, Unit Price, Total Sales, Payment Method, and State. Other potentially relevant factors, including demographic

| ISSN eletrônico: 2447-9896 | ISSN impresso: 2177-0360

characteristics, browsing behaviour, promotional exposure, prior purchase history, and direct measures of attitudes or preferences, are not included.

Fourth, the analysis does not directly measure exposure to AI-driven personalisation or recommendation systems. It therefore cannot establish whether algorithmic personalisation influenced the observed consumer outcomes.

Finally, the results derive from a specific dataset and modelling framework. External validation using other datasets, periods, markets, and consumer populations would be required before the findings could be generalised more broadly.

5 DISCUSSION AND IMPLICATIONS

The empirical findings indicate that machine-learning techniques can identify predictive patterns in transactional consumer data, although their performance varies according to the modelling approach and evaluation criterion adopted. The results are consistent with the broader development of data-driven marketing analytics, in which increasingly detailed consumer and transactional information can support predictive forms of analysis and decision-making (Wedel; Kannan, 2016).

The comparative performance of the models provides two relevant findings. First, XGBoost achieved the strongest overall discriminatory performance, outperforming the other algorithms in accuracy, precision, F1-score, and ROC-AUC. Second, this advantage was not consistent across all metrics, since Logistic Regression achieved the highest recall. Model performance should therefore be interpreted multidimensionally rather than through accuracy alone.

The comparatively strong performance of Logistic Regression is also noteworthy. Despite its simpler specification, its results were relatively close to those obtained by the more flexible machine-learning approaches. This indicates that part of the predictive information contained in the selected variables can be captured without relying exclusively on complex modelling architectures. Conversely, the higher overall performance of XGBoost suggests that greater model flexibility may capture additional predictive structure within the data.

| ISSN eletrônico: 2447-9896 | ISSN impresso: 2177-0360

This interpretation requires caution. The superior performance of XGBoost does not, by itself, demonstrate that specific non-linear relationships or interactions determine consumer behaviour. Establishing the substantive contribution of individual predictors would require additional interpretability analyses. The present comparison therefore demonstrates differences in predictive performance rather than the mechanisms responsible for those differences.

5.1 Predictive Consumer Behaviour and the Limits of the Behavioural Proxy

The findings demonstrate that observable transactional characteristics contain information associated with the behavioural outcome operationalised in the study. This result is consistent with the use of consumer and transactional data as inputs for predictive marketing analytics (Wedel; Kannan, 2016).

However, the interpretation of the findings depends on the distinction between observable transactional outcomes and psychological purchase intention. The study operationalises Purchase_Intention_Label through a behavioural proxy derived from Delivery Status and Review Rating. Consequently, the models predict this operationalised outcome rather than directly measuring the psychological construct of purchase intention.

Transactional behaviour can provide information for consumer classification, but it does not directly reveal the motivations, attitudes, preferences, or cognitive processes underlying purchasing decisions. The empirical findings should therefore be interpreted as evidence of predictive association within the adopted operational framework.

The same restriction applies to AI-driven personalisation. Although predictive analytics may support personalised marketing strategies, the dataset does not directly measure consumer exposure to recommendation algorithms or other AI-based marketing interventions. The findings consequently do not demonstrate that AI-driven personalisation caused the observed outcomes.

This distinction narrows the inferential scope of the study but does not invalidate its predictive contribution. Rather, it establishes the appropriate boundary between what the

| ISSN eletrônico: 2447-9896 | ISSN impresso: 2177-0360

empirical analysis demonstrates and broader claims about consumer intention or the effects of AI-mediated marketing interventions.

5.2 Prediction and Causal Interpretation

Predictive modelling and causal inference address different analytical questions. The models evaluated in this study were designed to classify observations based on statistical patterns identified in historical transactional data. Their performance therefore indicates how effectively the selected variables discriminate between the two classes in observations not used for model training.

The findings do not establish that changes in Product Category, Quantity, Unit Price, Total Sales, Payment Method, or State would cause corresponding changes in purchase intention. Such a conclusion would require a research design capable of identifying causal effects and addressing alternative explanations.

This distinction has direct implications for the use of predictive analytics in marketing. A model may identify consumers associated with a particular outcome without identifying which intervention would change that outcome. Predictive classification can therefore support the identification and prioritisation of observations, whereas decisions concerning the effectiveness of specific marketing interventions require additional empirical evidence.

Accordingly, the present findings should be interpreted as predictive rather than causal evidence. Maintaining this distinction avoids attributing explanatory power to machine-learning models beyond that supported by the research design.

5.3 Model Selection, Interpretability and Marketing Decision-Making

The comparison between the four algorithms illustrates the relationship between predictive performance, classification objectives, and interpretability. Logistic Regression provides a comparatively transparent baseline, whereas models such as XGBoost can represent more complex predictive structures but are less directly interpretable.

The higher overall performance of XGBoost may favour its adoption when discriminatory capacity and precision are prioritised. However, Logistic Regression achieved

| ISSN eletrônico: 2447-9896 | ISSN impresso: 2177-0360

higher recall, demonstrating that the choice of model depends partly on the consequences associated with different classification errors. For example, a marketing application seeking to reduce unnecessary targeting may place greater value on precision because false-positive classifications could result in inefficient allocation of resources. Conversely, an application in which failing to identify potentially relevant consumers is more costly may place greater weight on recall. Model selection should therefore correspond to the decision problem rather than being based solely on the highest overall accuracy.

Interpretability introduces another consideration. When predictive outputs support managerial decisions, understanding the basis and limitations of model predictions may be as relevant as marginal improvements in predictive performance. This is particularly important when algorithmic classifications influence decisions concerning individual consumers (Doshi-Velez; Kim, 2017).

Thus, the results do not support the unconditional selection of the most complex or highest-performing algorithm. Model choice should consider predictive performance, classification costs, interpretability requirements, available analytical capacity, and the organisational context in which predictions will be used.

5.4 Managerial Implications

The empirical results indicate that transactional information can support consumer classification and complement conventional marketing analysis. Predictive models may assist organisations in identifying observations associated with the target outcome and prioritising analytical or managerial attention accordingly.

The comparative results also demonstrate that the managerial usefulness of a model depends on the objective for which it is employed. Models optimised for precision may be more appropriate when organisations seek to minimise inefficient targeting, whereas higher-recall models may be preferable when identifying the greatest possible proportion of potentially relevant consumers is the primary objective.

Predictive analytics can consequently contribute to activities such as customer targeting, campaign prioritisation, segmentation, and resource allocation. However, these

| ISSN eletrônico: 2447-9896 | ISSN impresso: 2177-0360

should be understood as potential applications of the analytical framework rather than outcomes directly tested in the present study. The empirical analysis does not measure the effects of campaigns, recommendation systems, retention programmes, or changes in marketing expenditure.

For this reason, predictive outputs are more appropriately treated as decision-support information than as autonomous decision rules. Their managerial value depends on integration with contextual knowledge, operational constraints, additional consumer information, and evidence concerning the effectiveness of specific marketing interventions.

5.5 Ethical and Responsible AI Implications

The use of consumer data for predictive marketing requires attention to privacy, transparency, interpretability, and responsible data governance. The ability to infer behavioural patterns from transactional information can support marketing decisions while simultaneously creating concerns regarding how consumer information is collected, processed, and used (Martin; Murphy, 2017).

Historical data may also contain patterns reflecting existing inequalities, selection processes, or biases. When such data are used to train predictive models, these patterns may be reproduced in subsequent classifications. Predictive performance alone is therefore insufficient to establish that a model is appropriate for organisational deployment.

Responsible implementation requires adequate data governance, privacy safeguards, monitoring of model behaviour, evaluation of potential bias, and human oversight. Organisations should also consider whether the level of data collection and personalisation employed is proportionate to the intended marketing purpose.

Transparency and interpretability are particularly relevant when algorithmic outputs affect individual consumers. Organisations should be capable of explaining the analytical purpose of predictive systems, their principal limitations, and the role assigned to automated outputs within the decision process (Doshi-Velez; Kim, 2017).

Accordingly, responsible AI in marketing requires balancing predictive performance with privacy protection, accountability, transparency, and appropriate human supervision.

| ISSN eletrônico: 2447-9896 | ISSN impresso: 2177-0360

6 LIMITATIONS AND FUTURE RESEARCH

The findings should be interpreted in light of several methodological limitations.

First, the analysis is based on historical transactional data. Consumer behaviour may change over time in response to market conditions, preferences, technologies, competitive dynamics, and other contextual factors. The predictive relationships identified in the dataset therefore cannot be assumed to remain stable across different periods or environments.

Second, the dependent variable represents a behavioural proxy rather than a direct psychological measurement of purchase intention. Purchase_Intention_Label is derived from Delivery Status and Review Rating and consequently captures an operationalised transactional outcome. Although this construction enables supervised classification, it does not directly measure the motivational or cognitive dimensions conventionally associated with purchase intention.

Third, the predictor set is restricted to Product Category, Quantity, Unit Price, Total Sales, Payment Method, and State. Consumer behaviour may also be associated with demographic characteristics, browsing histories, previous purchases, promotional exposure, price sensitivity, product attributes, attitudes, and other variables not represented in the dataset.

Fourth, the empirical analysis does not directly measure exposure to AI-driven personalisation, recommendation algorithms, or other automated marketing interventions. The results therefore cannot establish whether such mechanisms influenced the observed behavioural outcomes.

Fifth, the analysis is based on a single dataset and empirical context. External validation using data from other markets, periods, platforms, and consumer populations would be necessary to assess the generalisability of the predictive relationships identified.

Future research should address these limitations through longitudinal and multi-source datasets, broader predictor sets, and direct measures of consumer attitudes and purchase intentions. Studies could also incorporate information concerning actual exposure to recommendation systems and personalised marketing interventions, enabling a clearer

| ISSN eletrônico: 2447-9896 | ISSN impresso: 2177-0360

distinction between predictive consumer classification and the effects of AI-mediated marketing actions.

Further research could compare predictive performance across different temporal periods and market contexts and examine model stability when consumer behaviour changes. External validation would provide stronger evidence regarding the transferability of the models beyond the dataset used in this study.

Another direction concerns model interpretability. Techniques capable of identifying the contribution of individual variables to model predictions could clarify why more flexible algorithms achieve different levels of performance and provide greater substantive insight into the relationships captured by the models.

Finally, experimental, quasi-experimental, or appropriately designed longitudinal studies would be required to move beyond predictive association and investigate whether particular marketing interventions produce changes in consumer behaviour.

7 FINAL CONSIDERATIONS

This study examined the application of Artificial Intelligence and machine-learning techniques to marketing analysis and predictive consumer-behaviour modelling, combining a literature-based discussion with an empirical comparison of four supervised-learning algorithms. The findings indicate that transactional consumer data contain predictive information capable of supporting classification within the analytical framework adopted in the study.

The comparative analysis showed that all four models achieved relevant predictive performance, although their results varied across the evaluation metrics. XGBoost presented the strongest overall performance, particularly in accuracy, precision, F1-score, and ROC-AUC, whereas Logistic Regression achieved the highest recall. These differences reinforce that model selection should not depend on a single performance metric but should consider the analytical objective and the relative consequences of different classification errors.

| ISSN eletrônico: 2447-9896 | ISSN impresso: 2177-0360

The results also require a clear distinction between predictive association and causal explanation. The models identify statistical patterns associated with the operationalised behavioural outcome, but they do not demonstrate that the selected transactional characteristics cause changes in consumer behaviour. Likewise, because Purchase_Intention_Label is constructed from Delivery Status and Review Rating, it should be interpreted as a behavioural proxy rather than as a direct psychological measure of purchase intention.

From an applied perspective, the findings indicate that machine-learning models can complement conventional marketing analytics by supporting consumer classification and data-driven decision-making. Their practical use, however, depends not only on predictive performance but also on interpretability, data quality, privacy protection, transparency, and appropriate human oversight. Predictive outputs should therefore support managerial judgement rather than function as autonomous decision rules.

Overall, the study contributes to the discussion of AI-enabled marketing analytics by empirically comparing different predictive approaches while explicitly delimiting the scope of the conclusions that can be drawn from transactional data. Future applications of these methods should combine predictive performance with stronger construct measurement, external validation, interpretability analysis, and research designs capable of distinguishing prediction from causal effects.

Revista Esmat, Palmas, v. 18, n. 2, 2026, p. 01-26.

| ISSN eletrônico: 2447-9896 | ISSN impresso: 2177-0360

REFERENCES

BUCKINX, Wouter; VAN DEN POEL, Dirk. Customer base analysis: partial defection of behaviourally loyal clients in a non-contractual FMCG retail setting. *Journal of Marketing Research*, v. 42, n. 3, p. 252-268, 2005. DOI: 10.1509/jmkr.42.3.252. https://doi.org/10.1509/jmkr.42.3.252.

CAMBRIA, Erik; SCHULLER, Björn; XIA, Yunqing; HAVASI, Catherine. New avenues in opinion mining and sentiment analysis. *IEEE Intelligent Systems*, v. 28, n. 2, p. 15-21, **2013**. DOI: 10.1109/MIS.2013.30. https://doi.org/10.1109/MIS.2013.30.

DOSHI-VELEZ, Finale; KIM, Been. *Towards a rigorous science of interpretable machine learning*. 2017. arXiv:1702.08608. arXiv:1702.08608.

HUANG, Ming-Hui; RUST, Roland T. Artificial intelligence in service. *Journal of Service Research*, v. 21, n. 2, p. 155-172, **2018**. DOI: 10.1177/1094670517752459. https://doi.org/10.1177/1094670517752459.

KUMAR, V.; RAMACHANDRAN, Divya; KUMAR, Binay. Influence of new-age technologies on marketing: a research agenda. *Journal of Business Research*, v. 125, p. 864-877, **2021**. DOI: 10.1016/j.jbusres.2020.01.007. https://doi.org/10.1016/j.jbusres.2020.01.007.

MARTIN, Kelly D.; MURPHY, Patrick E. The role of data privacy in marketing. *Journal of the Academy of Marketing Science*, v. 45, n. 2, p. 135-155, 2017. DOI: 10.1007/s11747-016-0495-4. https://doi.org/10.1007/s11747-016-0495-4.

PANG, Bo; LEE, Lillian. Opinion mining and sentiment analysis. *Foundations and Trends in Information Retrieval*, v. 2, n. 1-2, p. 1-135, 2008. DOI: 10.1561/1500000011. https://doi.org/10.1561/1500000011.

RICCI, Francesco; ROKACH, Lior; SHAPIRA, Bracha (ed.). *Recommender systems handbook*. 2. ed. New York: Springer, 2015. DOI: 10.1007/978-1-4899-7637-6. https://doi.org/10.1007/978-1-4899-7637-6.

VERBEKE, Wouter; MARTENS, David; BAESENS, Bart. Social network analysis for customer churn prediction. *Applied Soft Computing*, v. 14, p. 431-446, 2014. DOI: 10.1016/j.asoc.2013.09.020. https://doi.org/10.1016/j.asoc.2013.09.020.

WEDEL, Michel; KANNAN, P. K. Marketing analytics for data-rich environments. *Journal of Marketing*, v. 80, n. 6, p. 97-121, 2016. DOI: 10.1509/jm.15.0413. https://doi.org/10.1509/jm.15.0413.

| ISSN eletrônico: 2447-9896 | ISSN impresso: 2177-0360