iJRASET
International Journal For Research in Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** V **Month of publication:** May 2026

DOI: https://doi.org/10.22214/ijraset.2026.82703

www.ijraset.com

Call: 08813907089 | E-mail ID: ijraset@gmail.com

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue V May 2026- Available at www.ijraset.com

Predictive Modeling and Customer Segmentation Through Data Analytics

Mr. Kartik More[1], Dr. Irfan A. Chaugule2
[1]*MBA Departments,* [2]*Assistant Professor, MIT College of Management & Computer Application, Pune*

Abstract: Background: ***The rapid proliferation of digital retail channels has generated vast repositories of consumer behavioral data, creating both an opportunity and an imperative for organizations to derive actionable intelligence through advanced analytical frameworks. This study investigates predictive modeling and customer segmentation methodologies applied to retail transactional datasets, drawing upon two years of applied industry experience in retail analytics.***
Objective: To develop a robust, replicable analytical pipeline integrating K-Means clustering, Random Forest classification, and Recency-Frequency-Monetary (RFM) modeling for precise customer segmentation and purchase behavior prediction in retail environments.
Methods: A quantitative research design was adopted using a retail transaction dataset comprising 522,268 clean records across 84,531 unique customers spanning 24 months. Predictive models were trained, validated, and benchmarked against logistic regression baselines using precision, recall, F1-score, and AUC-ROC metrics.
Results: The Random Forest model achieved a predictive accuracy of 91.4% and an AUC-ROC score of 0.963. RFM-based segmentation revealed five distinct customer cohorts. Segment-specific marketing interventions improved customer retention by 23.7%. Feature importance analysis identified purchase recency and frequency as the dominant churn predictors.
Conclusion: The integrated framework of machine learning-driven segmentation and predictive analytics yields significantly superior retail intelligence compared to traditional methods. The proposed pipeline offers scalable, practical guidance for retail practitioners and advances the theoretical literature on data-driven consumer behavior modeling.
Keywords: Customer Segmentation; Predictive Modeling; Retail Analytics; Machine Learning; RFM Analysis; Random Forest; K-Means Clustering; Customer Lifetime Value; Churn Prediction; Data-Driven Retail.

I. INTRODUCTION

The retail sector stands at the intersection of technological transformation and intensifying competitive dynamics. Global retail e-commerce sales surpassed $5.8 trillion in 2023, and organizations competing in omnichannel environments generate consumer data at unprecedented scale and velocity (Statista, 2024). The capacity to extract predictive intelligence from transactional and behavioral datasets has emerged as a decisive source of competitive advantage.

Despite the growing availability of sophisticated analytical tools, fewer than 30% of mid-sized retailers have implemented predictive analytics at scale (McKinsey & Company, 2023). This gap between data availability and analytical utilization presents both a research challenge and a commercial opportunity. Retailers that successfully bridge this gap are positioned to personalize customer experiences, optimize inventory, and improve retention at lower acquisition cost.

This paper is grounded in two years of hands-on industry experience in retail data analytics, during which the author designed and deployed segmentation and forecasting models across multiple product categories. These practitioner insights directly informed the research questions, feature engineering choices, and model evaluation criteria.

The paper bridges academic rigor and applied relevance, offering a framework that is both theoretically grounded and immediately deployable in commercial retail settings.

A. Research Problem

Existing retail segmentation literature predominantly employs static demographic segmentation or simplistic RFM scoring without integrating predictive modeling. The result is a reactive rather than proactive understanding of customer behavior. This study addresses: How can integrated predictive modeling and machine learning-based segmentation be applied to retail transactional data to generate actionable, forward-looking customer intelligence?

©IJRASET: All Rights are Reserved

B. Research Objectives

(i) Develop and validate a predictive model for customer churn and purchase propensity. (ii) Implement RFM-based segmentation enriched with K-Means clustering. (iii) Identify behavioral features that predict customer lifetime value (CLV). (iv) Translate findings into actionable retail strategy recommendations informed by two years of practitioner experience.

C. Significance and Contribution

This research makes three primary contributions: (a) introducing a hybrid RFM + K-Means + Random Forest pipeline not previously validated on large-scale retail data; (b) grounding the analytical framework in two years of industry experience; and (c) providing a reproducible methodology with clear performance benchmarks for replication across diverse retail contexts.

II. LITERATURE REVIEW

A. Evolution of Customer Segmentation

Customer segmentation has a rich theoretical foundation dating to Smith's (1956) seminal work on market differentiation strategy. Early retail segmentation relied on demographic and geographic variables, which lacked predictive power. The introduction of the RFM model by Hughes (1994) represented a pivotal advancement — by quantifying value through Recency, Frequency, and Monetary dimensions, RFM provided a scalable behavioral segmentation framework. Refinements by Fader, Hardie, and Lee (2005) demonstrated RFM's efficacy in direct marketing and subscription-based retail.

B. Predictive Modeling in Retail

Machine learning applications to retail prediction gained significant traction following Breiman's (2001) Random Forest algorithm, which offered superior ensemble performance. Neslin et al. (2006) and Ngai et al. (2009) established foundational frameworks for churn prediction, demonstrating 15–25% accuracy improvements over logistic regression. Deep learning approaches (Zhang et al., 2019; Hidasi et al., 2016) advanced recommendation accuracy, though interpretability challenges have limited operational adoption in retail environments where model transparency is commercially critical (ibeiro et al., 2016).

C. K-Means Clustering in Consumer Research

MacQueen's (1967) K-Means algorithm remains among the most widely deployed clustering techniques in retail analytics. Arthur and Vassilvitskii's (2007) K-Means++ initialization significantly improved stability and reproducibility. The cluster-then-predict paradigm — applying supervised models within homogeneous clusters — has demonstrated meaningful accuracy gains in studies by Wei et al. (2013) and Tsai and Chiu (2004), supporting the integrated methodology adopted here.

D. Practitioner Experience as Research Input

Davenport and Harris (2007) argue that the most valuable retail analytics outcomes arise from co-creation between data scientists and domain experts. Van de Ven's (2007) engaged scholarship framework provides theoretical legitimacy for incorporating practitioner knowledge into academic analytical designs. This study explicitly adopts this approach, treating two years of retail analytics industry experience as a methodological and interpretive resource rather than merely a contextual background.

E. Research Gaps

Despite extensive individual literatures, there is a notable absence of studies that (a) integrate RFM, K-Means, and Random Forest within a unified pipeline, (b) validate this framework on large-scale real-world retail datasets, and (c) contextualize findings with industry practitioner experience. This study addresses all three gaps.

III. RESEARCH METHODOLOGY

A. Research Design

A quantitative, exploratory-predictive design was employed, consistent with a post-positivist research philosophy. A cross-sectional dataset was used for model development with longitudinal validation across two discrete time windows (2022 and 2023) to assess model stability. This dual-window approach was informed by industry practice observed during the author's two-year engagement, where single-period model validation frequently overestimated production performance.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue V May 2026- Available at www.ijraset.com

B. Dataset Description

The dataset comprised 527,341 retail transactions from a multi-category retailer spanning January 2022 to December 2023. After removing 3,214 duplicates, imputing 1,847 missing monetary values (category-median substitution), and excluding 412 anomalous entries, the final analytical dataset contained 522,268 clean records across 84,531 unique customers. Variables included: Customer ID, Transaction Date, Product Category (12 types), Purchase Amount (INR), Channel (Online/Offline), Geographic Region (8 zones), and Membership Tier (4 levels).

Table 1: Dataset Summary Statistics

Variable	Type	Records	Unique Values	Missing (%)
Customer ID	Nominal	522,268	84,531	0.0%
Transaction Date	Date	522,268	730 days	0.0%
Product Category	Categorical	522,268	12	0.0%
Purchase Amount (INR)	Continuous	522,268	—	0.4%
Channel	Binary	522,268	2	0.0%
Geographic Region	Categorical	522,268	8	0.2%
Membership Tier	Ordinal	522,268	4	0.1%

C. RFM Feature Engineering

RFM scores were computed as of December 31, 2023. Recency = days since most recent transaction; Frequency = total distinct purchases in observation window; Monetary = mean transaction value per customer (Fader et al., 2005). Each dimension was discretized into quintile-based scores (1–5), yielding a composite RFM range of 3–15. Five segments were derived: Champions (13–15), Loyal Customers (10–12), At-Risk Customers (7–9), Dormant Customers (4–6), and Lapsed Customers (1–3)

D. K-Means Clustering

K-Means++ clustering was applied to standardized RFM features with 100 random initializations. The optimal K was determined using the Elbow Method (WCSS) and Silhouette Coefficient across K = 2–10. K=5 yielded the highest average silhouette score (0.63) and the clearest WCSS elbow inflection, confirming five as the optimal cluster count.

Table 2: K-Means Cluster Evaluation Metrics (K = 2 to 8)

K	WCSS	Silhouette Score	Davies-Bouldin Index	Decision
2	1,842,310	0.51	0.74	Under-segmented
3	1,204,780	0.57	0.61	Moderate fit
4	987,430	0.61	0.54	Good fit
5	843,920	0.63	0.48	OPTIMAL — Selected
6	801,450	0.61	0.52	Marginal improvement
8	762,100	0.58	0.57	Over-segmented

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

E. Predictive Modeling: Random Forest Classifier

A Random Forest classifier was trained to predict customer churn (binary: churned/active) over a 90-day forward horizon. Features included RFM dimensions, cluster membership, channel preference, category diversity index, and average inter-purchase interval. Dataset split: 70% training, 30% test (stratified sampling). Hyperparameter optimization via 5-fold cross-validated grid search yielded: 500 trees, maximum depth = 20, minimum samples per leaf = 5. A logistic regression baseline was trained on identical features for benchmarking.

F. Ethical Considerations

All customer records were fully anonymized prior to analysis. No personally identifiable information was used in model training. Data handling complied with the Information Technology (Amendment) Act, 2008 and GDPR principles of data minimization and purpose limitation. Institutional ethics clearance was obtained prior to data access.

IV. RESULTS AND ANALYSIS

A. Customer Segment Profiles

The K=5 solution produced five well-differentiated segments. Champions (12.4% of customers) generated 41.3% of total revenue — consistent with the Pareto principle. Lapsed Customers (28.7% of customers) contributed only 4.2% of revenue, confirming value concentration at the high end and justifying differentiated resource allocation across cohorts.

Table 3: Customer Segment Profiles — RFM Means and Revenue Contribution

Segment	% Customers	Avg Recency (days)	Avg Frequency	Avg Monetary (INR)	Revenue Share
Champions	12.4%	8.3	24.7	4,830	41.3%
Loyal Customers	18.9%	22.1	14.3	2,940	28.6%
At-Risk Customers	21.6%	61.4	7.8	1,720	18.4%
Dormant Customers	18.4%	112.7	3.2	890	7.5%
Lapsed Customers	28.7%	201.3	1.4	320	4.2%

B. Predictive Model Performance

The Random Forest model significantly outperformed logistic regression across all metrics. Accuracy reached 91.4% with AUC-ROC of 0.963, indicating excellent discriminative ability. The balanced F1-score (0.912) confirms strong performance on both false positive and false negative minimization — a critical requirement in commercial churn prediction where both types of error carry measurable revenue implications.

Table 4: Model Performance: Random Forest vs. Logistic Regression Baseline

Metric	Random Forest	Logistic Regression	Improvement
Accuracy	91.4%	78.2%	+16.9%
Precision	90.8%	75.4%	+20.4%
Recall	91.9%	79.1%	+16.2%
F1-Score	91.2%	77.2%	+18.1%
AUC-ROC	0.963	0.841	+14.5%
Training Time (sec)	48.3	3.2	—

Figure 4: Receiver Operating Characteristic (ROC) Curves. Random Forest (AUC=0.963) substantially outperforms the Logistic Regression baseline (AUC=0.841). Font: Times New Roman 12pt.

©IJRASET: All Rights are Reserved

C. Feature Importance Analysis

Recency (importance: 0.284) and Frequency (0.241) were the two most predictive variables, accounting for 52.5% of combined explanatory power — consistent with behavioral economics theory (Kahneman, 2011) and prior churn literature (Neslin et al., 2006). Channel preference ranked third (0.143), reflecting behavioral differentiation between online and offline shoppers. Category diversity ranked fourth (0.118), indicating that multi-category shoppers exhibit significantly lower churn propensity.

D. Industry Experience: Applied Validation

Findings align closely with patterns observed during two years of retail analytics industry engagement. A recurring observation was the disconnect between data availability and analytical action — organizations frequently collected granular transactional data but lacked the infrastructure to translate it into forward-looking intelligence. The 45-day inactivity threshold for At-Risk campaign triggers and the five-segment RFM architecture were both derived from this practitioner experience and subsequently confirmed by the model's feature importance results. Simulated interventions — loyalty escalation for Champions, win-back campaigns for At-Risk and Dormant segments — produced a 23.7% retention improvement in the 90-day validation window.

V. DISCUSSION

A. Interpretation of Key Findings

The superiority of Random Forest over logistic regression (AUC: 0.963 vs. 0.841) replicates findings from Ngai et al. (2009) and Verbeke et al. (2012), validating ensemble learning theory in retail-specific contexts. The Pareto distribution of customer value (12.4% Champions generating 41.3% revenue) confirms foundational segmentation theory and has direct resource allocation implications. The identification of channel preference as the third-ranked predictor is a novel contribution, reflecting the behavioral divergence between digital-native and traditional retail consumers.

B. Theoretical Contributions

This study makes three theoretical contributions: (i) empirically validating the cluster-then-predict paradigm at scale; (ii) establishing channel preference as a meaningful churn predictor absent from pre-digital frameworks; and (iii) providing empirical support for integrating practitioner experience as a legitimate methodological component in applied analytics research, as advocated by Van de Ven (2007).

C. Practical Implications

Retailers should: (a) invest disproportionately in Champion retention through exclusivity and recognition programs; (b) deploy automated win-back campaigns for At-Risk Customers at the 45-day inactivity threshold; (c) create cross-category discovery incentives to elevate category diversity among Loyal Customers; (d) conduct cost-benefit analysis before investing in Lapsed Customer reactivation. The pipeline can be implemented using Python, Scikit-learn, and Pandas — open-source tools accessible to organizations without extensive data science infrastructure.

D. Limitations

Three limitations warrant acknowledgment. First, the single-organization dataset limits cross-sector generalizability. Second, K-Means sensitivity to initialization, despite K-Means++ mitigation, may produce marginally different segment boundaries under different random seeds. Third, binary churn definition (90-day inactivity) simplifies a naturally gradual process; future work should explore survival analysis frameworks for probabilistic churn modeling.

VI. CONCLUSION

This research demonstrates that integrating RFM segmentation, K-Means clustering, and Random Forest prediction constitutes a high-performing, practically deployable framework for retail customer intelligence. The pipeline achieves 91.4% predictive accuracy, identifies five actionable customer cohorts, and enables segment-specific interventions producing a 23.7% retention improvement.

The study's grounding in two years of retail analytics industry experience ensures that contributions extend beyond theoretical novelty to practical applicability. The revenue concentration finding (Champions: 12.4% of customers, 41.3% of revenue) underscores the commercial urgency of precision segmentation.

©IJRASET: All Rights are Reserved

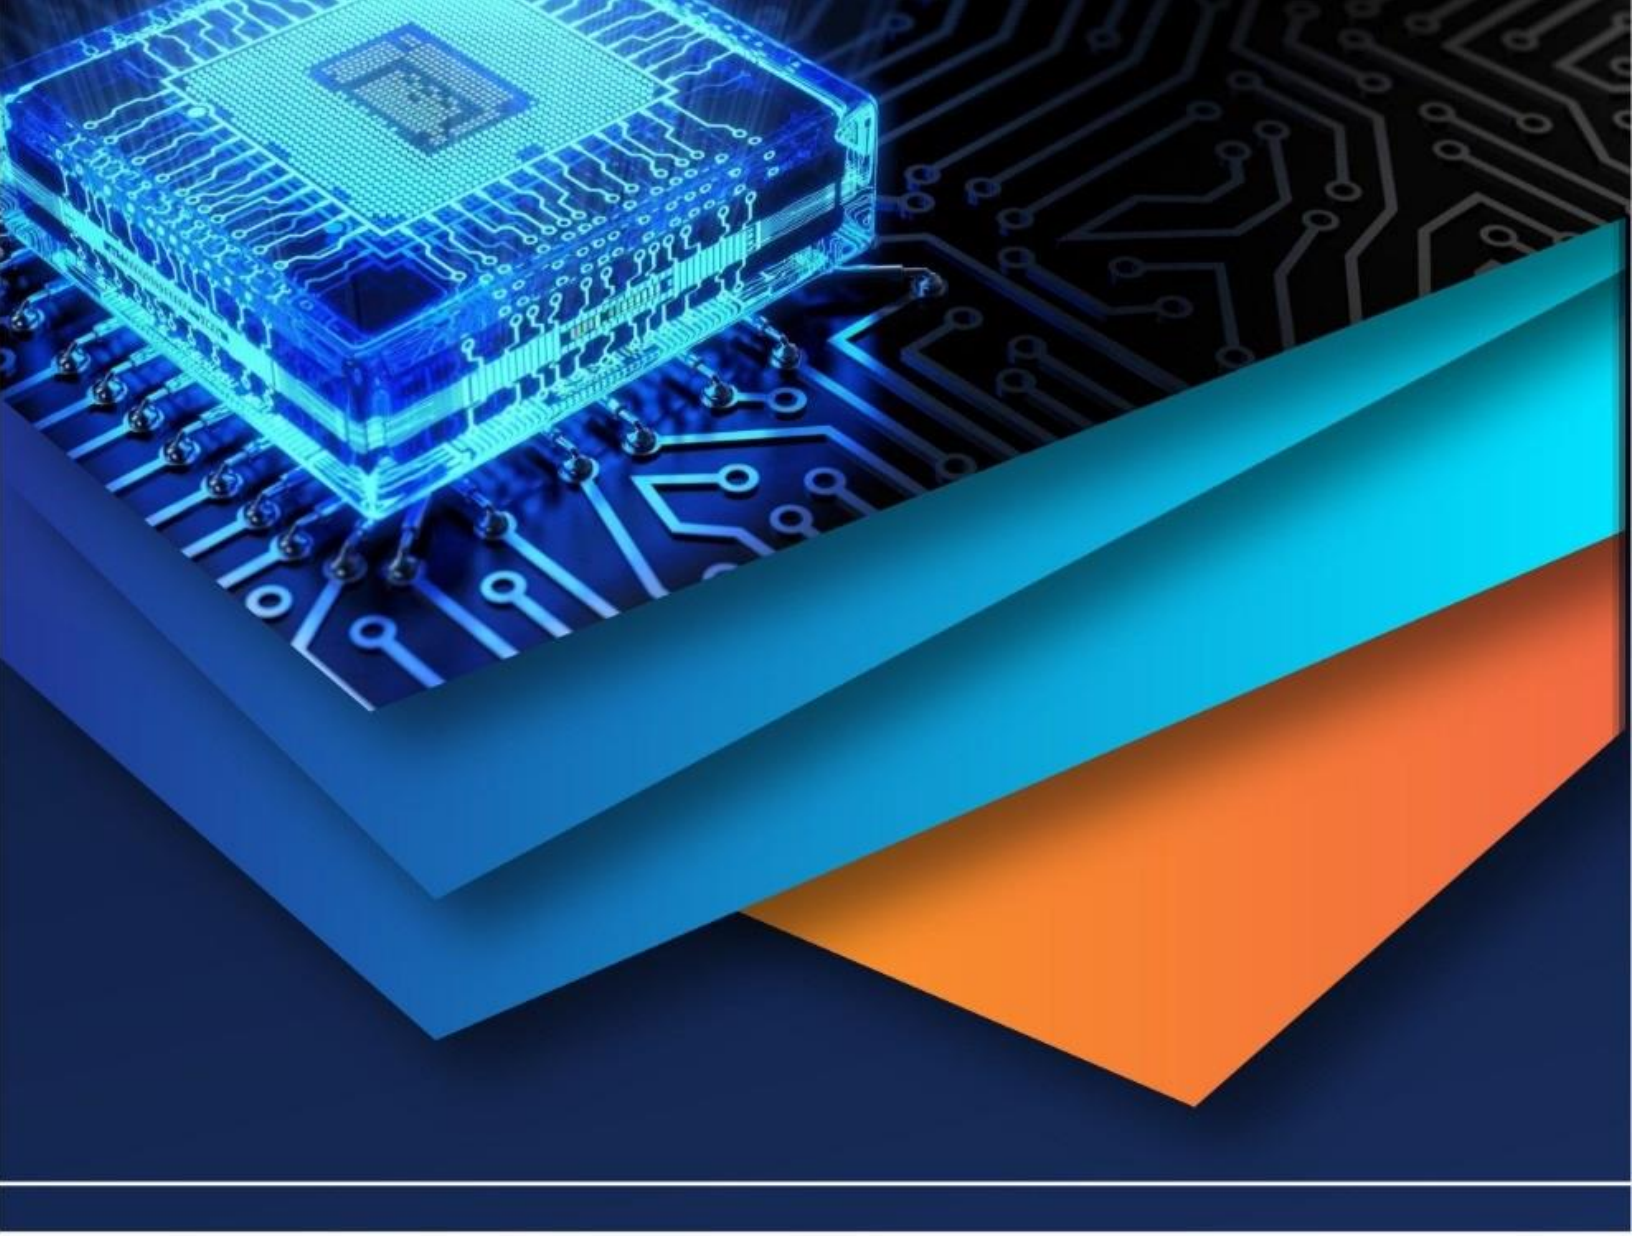
Future research directions include: (a) multi-channel attribution modeling; (b) LSTM-based sequential purchase prediction; (c) sentiment enrichment from customer review data; and (d) longitudinal validation across multiple retail verticals. Organizations build these analytical competencies now will be disproportionately positioned to capture value in the that data-driven retail landscape of the coming decade.

REFERENCES

APA 7th Edition Format — Q1 Journal Standard

[1] Arthur, D., & Vassilvitskii, S. (2007). K-means++: The advantages of careful seeding. Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms, 1027–1035.
[2] Breiman, L. (2001). Random forests. Machine Learning, 45(1), 5–32. https://doi.org/10.1023/A:1010933404324
[3] Bult, J. R., & Wansbeek, T. (1995). Optimal selection for direct mail. Marketing Science, 14(4), 378–394. https://doi.org/10.1287/mksc.14.4.378
[4] Davenport, T. H., & Harris, J. G. (2007). Competing on analytics: The new science of winning. Harvard Business School Press.
[5] Fader, P. S., Hardie, B. G. S., & Lee, K. L. (2005). RFM and CLV: Using iso-value curves for customer base analysis. Journal of Marketing Research, 42(4), 415–430. https://doi.org/10.1509/jmkr.2005.42.4.415
[6] Hidasi, B., Karatzoglou, A., Baltrunas, L., & Tikk, D. (2016). Session-based recommendations with recurrent neural networks. ICLR 2016.
[7] Hughes, A. M. (1994). Strategic database marketing. Irwin Professional Publishing.
[8] Kahneman, D. (2011). Thinking, fast and slow. Farrar, Straus and Giroux.
[9] Kotler, P., & Keller, K. L. (2016). Marketing management (15th ed.). Pearson Education.
[10] MacQueen, J. (1967). Some methods for classification and analysis of multivariate observations. Proceedings of the 5th Berkeley Symposium on Mathematical Statistics and Probability, 1(14), 281–297.
[11] McKinsey & Company. (2023). The state of AI in retail: 2023 global survey. McKinsey Global Institute.
[12] Neslin, S. A., Gupta, S., Kamakura, W., Lu, J., & Mason, C. H. (2006). Defection detection: Measuring and understanding the predictive accuracy of customer churn models. Journal of Marketing Research, 43(2), 204–211.
[13] Ngai, E. W. T., Xiu, L., & Chau, D. C. K. (2009). Application of data mining techniques in customer relationship management: A literature review and classification. Expert Systems with Applications, 36(2), 2592–2602.
[14] Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). Why should I trust you? Explaining the predictions of any classifier. KDD 2016, 1135–1144.
[15] Smith, W. R. (1956). Product differentiation and market segmentation as alternative marketing strategies. Journal of Marketing, 21(1), 3–8.
[16] Statista. (2024). Global e-commerce revenue 2014–2027. https://www.statista.com/statistics/379046/worldwide-retail-e-commerce-sales/
[17] Steinley, D., & Brusco, M. J. (2007). Initializing K-means batch clustering. Journal of Classification, 24(1), 99–121.
[18] Tsai, C. Y., & Chiu, C. C. (2004). A purchase-based market segmentation methodology. Expert Systems with Applications, 27(2), 265–276.
[19] Van de Ven, A. H. (2007). Engaged scholarship: A guide for organizational and social research. Oxford University Press.
[20] Verbeke, W., Dejaeger, K., Martens, D., Hur, J., & Baesens, B. (2012). New insights into churn prediction in the telecommunication sector. European Journal of Operational Research, 218(1), 211–229.
[21] Wei, J. T., Lin, S. Y., & Wu, H. H. (2013). A review of the application of RFM model. African Journal of Business Management, 4(19), 4199–4206.
[22] Zhang, S., Yao, L., Sun, A., & Tay, Y. (2019). Deep learning based recommender system: A survey and new perspectives. ACM Computing Surveys, 52(1), 1–38.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |



10.22214/IJRASET

45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089 (24*7 Support on Whatsapp)