

# Latent Customer Segmentation and Value-Based Recommendation Leveraging a Two-Stage Model with Missing Labels

Keerthi Gopalakrishnan\*  
Walmart Global Tech  
Sunnyvale, California, USA  
k.gopalakrishnan@walmart.com

Yokila Arora  
Walmart Global Tech  
Sunnyvale, California, USA  
yokila.arora@walmart.com

Sushant Kumar  
Walmart Global Tech  
Sunnyvale, California, USA  
sushant.kumar@walmart.com

Tianning Dong\*  
Walmart Global Tech  
Sunnyvale, California, USA  
Tianning.Dong@walmart.com

Topojoy Biswas  
Walmart Global Tech  
Sunnyvale, California, USA  
topojoy.biswas@walmart.com

Kannan Achan  
Walmart Global Tech  
Sunnyvale, California, USA  
kannan.achan@walmart.com

Chia-Yen Ho  
Walmart Global Tech  
Sunnyvale, California, USA  
chiayen.ho@walmart.com

Jason Cho  
Walmart Global Tech  
Sunnyvale, California, USA  
jason.cho@walmart.com

## Abstract

1

The success of businesses is highly dependent on their ability to convert consumers into loyal customers. The value proposition of a customer is considered the primary determinant in this conversion process, and establishing intrinsic product worth at a specified value point is essential. While affordability remains a significant factor for online shoppers, brands are transitioning away from broad-scale promotions and campaigns that reduce the perceived brand worth. Implementing such widespread marketing campaigns can inadvertently erode brand equity and diminish the probability of achieving the desired return on marketing investment. Existing solutions have fallen short in addressing the twin challenges of finding the optimal balance between the value proposition and the long-term brand value. At the same time, existing dynamic economic algorithms often misidentify highly engaging customers as ideal campaign targets, leading to suboptimal engagement and conversion rates, thereby diminishing customer loyalty.

This article introduces a two-stage multi-model architecture that employs Self-Paced Loss to enhance customer categorization. The first layer is a Multi-Class Neural Network that differentiates high engaging customers stimulated by a campaign, high-engagement customers not influenced by campaign nudges, and

low-engagement customers. The Binary Label Correction Model forms the second layer, further refining the classification of highly engaged customers by differentiating between those responding to a campaign and those engaging organically. The actual customer's intent can be determined using a missing label framework. This level gathers data from customers displaying variable value proposition behavior, rectifying their true state labels during training for more efficient customer segmentation. By distinguishing customer engagement intent (prompted versus organic), the suggested solution allows businesses to enhance their marketing campaign strategies and target prompted engaged segment, reducing exposure rates while boosting conversion rates. When testing this solution using an A/B test framework, we noticed an increase of more than 100 basis points in the success metric.

## CCS Concepts

• Information systems → Recommender systems; Personalization; • Computing methodologies → Machine learning.

## Keywords

Recommender System, Personalized Advertising, Machine Learning, Customer Segmentation, Price Sensitivity, Label Correction

## ACM Reference Format:

Keerthi Gopalakrishnan, Tianning Dong, Chia-Yen Ho, Yokila Arora, Topojoy Biswas, Jason Cho, Sushant Kumar, and Kannan Achan. 2025. Latent Customer Segmentation and Value-Based Recommendation Leveraging a Two-Stage Model with Missing Labels . In *Companion Proceedings of the ACM Web Conference 2025 (WWW Companion '25)*, April 28-May 2, 2025, Sydney, NSW, Australia. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3701716.3715243>

## 1 Introduction

In highly competitive business landscape nowadays, understanding and catering to customer behavior and preferences is paramount for businesses to maintain and expand market share [1]. Traditional

\*Both authors contributed equally to this work.

<sup>1</sup>This is the author's version of the work. The definitive published version is available at ACM Digital Library: DOI 10.1145/3701716.3715243.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WWW Companion '25, Sydney, NSW, Australia

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-1331-6/25/04

<https://doi.org/10.1145/3701716.3715243>

customer segmentation methods, which often rely on static demographic information, are increasingly being replaced by more personalized approaches [42, 21]. The advent of big data and advancements in machine learning have enabled businesses to analyze vast amounts of data in real-time [28], providing deeper insights into customer needs and behaviors [12, 42]. This shift towards dynamic customer segmentation allows for more personalized marketing strategies, improved customer satisfaction, and increased profitability [33, 41].

While these personalized marketing campaigns and offers are primary influences for online buyers [22], in today's competitive market, simply launching site-wide campaigns and offers may not be enough to attract and retain customers. Mental accounting perspective suggests that customers evaluate products based on their perceived value [40]. In the context of personalized marketing campaigns and offers, understanding the customer's perceived value of a product can help retailers tailor their pricing strategies to individual customers.

In the context of current economic uncertainties, a significant number of consumers have developed a more discerning approach to seeking value for their money [38]. This shift in consumer behavior has led to a keener focus on gaining maximum value from each dollar spent. At one end of the spectrum, there are certain dynamic individuals and consumers from Generation Z who have exhibited a noticeable change in their product value preferences [9]. They are more inclined to purchase items in higher-priced categories and demonstrate a readiness to pay the full price or even more for certain products [43]. These customers do not place a high reliance on marketing campaigns or special offers to drive their purchasing decisions.

On the other end of the spectrum, there is another faction of consumers who tend to respond more to marketing campaigns and offers [14]. This group of consumers often require some form of prompt from a campaign before they decide to make a purchase. However, this group of consumers can be transformed into loyal customers if they are presented with competitive pricing and promotions that align with their budgetary constraints. By acknowledging and understanding the diverse needs and preferences of these two distinct consumer groups, retailers can develop targeted marketing strategies and flexible pricing models to cater to them [29]. This approach can enable retailers to appeal to a wider range of consumers, thereby potentially enhancing their sales figures and fostering greater customer loyalty.

Understanding the dynamics between consumers who respond to marketing campaigns and those who transact organically is essential for businesses aiming to craft effective strategies [17, 26]. Each group exhibits distinct behaviors and characteristics, impacting customer acquisition costs, loyalty, and long-term value [25]. Thus, the value proposition of a customer plays a predominant role in the success of these marketing campaigns. The value proposition of a customer represents their perceived worth to a business, balancing the potential profitability they bring against the cost of engaging them. It encompasses their likelihood to respond to campaigns, transact organically, and contribute to long-term brand equity [4].

**Consumers Responding to Marketing Campaigns:** Consumers who respond to marketing campaigns are primarily influenced by targeted promotional efforts such as advertisements, email campaigns, social media ads, and special offers [2]. These consumers are often drawn to businesses due to compelling discounts, limited-time deals, or eye-catching advertisements. They exhibit high price sensitivity and are more likely to make impulse purchases driven by immediate incentives [5].

The engagement level of these consumers with the brand is typically higher during the campaign period, as they interact with various marketing channels [23]. However, their loyalty may be short-lived, often contingent upon the continuation of promotions [8]. This group tends to have a higher Customer Acquisition Cost (CAC) since businesses need to invest in marketing efforts to attract them [25, 44]. While marketing campaigns can provide a significant sales boost in the short term, businesses must implement effective retention strategies to convert these consumers into repeat customers. Importantly, marketing campaigns generate valuable data on consumer preferences and behaviors, which can be leveraged to refine future marketing efforts [19].

**Consumers Transacting Organically:** In contrast, consumers who transact organically discover and purchase products without the direct influence of marketing campaigns. Their decision-making process is driven by personal needs, brand reputation, word-of-mouth recommendations, and organic search results [32]. These consumers often exhibit strong brand loyalty and trust, having conducted thorough research and read reviews before making a purchase.

Organic consumers have a lower CAC since they are not acquired through paid marketing efforts [31]. They tend to provide a more stable and predictable revenue stream due to their repeat purchases and higher Customer Lifetime Value (CLV) [39]. This group's purchasing decisions are influenced by the brand's reputation and the quality of the customer experience, making it crucial for businesses to maintain a positive brand image and deliver exceptional service [11]. Investments in search engine optimization and content marketing are pivotal in enhancing organic reach and attracting these consumers [37].

**Comparing and Balancing Both Groups:** Marketing campaign responders offer the advantage of a quick sales boost and measurable return on investment (ROI), but they require continuous promotional efforts to maintain engagement and loyalty [15]. On the other hand, organic consumers provide sustainable revenue and higher loyalty, relying on the strength of the brand and organic discovery channels.

To achieve a balanced approach, businesses should integrate marketing efforts to attract new customers while leveraging organic strategies [10] to retain them [24]. Personalization based on consumer data can enhance engagement and long-term loyalty [30]. Emphasizing an exceptional customer experience is critical in converting campaign responders into loyal customers. The continuous analysis of data from both groups allows businesses to optimize their strategies and improve overall performance [35].

By understanding and leveraging the strengths of both consumer groups, businesses can develop more comprehensive and effective marketing strategies, ensuring both immediate sales boosts and long-term growth.

## 2 Related Work

In marketing research, identifying target audiences with precise value propositions is critical to personalizing marketing campaigns and optimizing their effectiveness [4]. The application of machine learning models to personalized marketing is a rapidly evolving field, showing significant potential to revolutionize marketing strategies for e-commerce companies. Existing works primarily focus on predictive modeling of customer behaviors using supervised approaches such as tree-based models [13, 20] and neural networks [6, 16]. Meanwhile, unsupervised methods, including k-means clustering [18, 27] and its integration with traditional RFM (Recency, Frequency, Monetary) marketing techniques [36], have also gained considerable traction.

Traditional marketing approaches like RFM have been widely adopted due to their simplicity and effectiveness in targeting high-value transactional customers. However, these methods often rely on generalized cohort-level analyses based purely on high transactional metrics (i.e., high recency, frequency, and monetary value) without adequately personalizing campaigns or pricing for individual customers. Such approaches fail to account for nuanced customer behaviors, particularly the varying sensitivities to value propositions and pricing strategies.

To address these limitations, companies like Amazon have implemented advanced dynamic pricing methods such as the "Intelligent Pricing" initiative [7]. This approach uses a two-pronged strategy: analyzing both historical and real-time market data to understand evolving consumer price sensitivities, and recommending personalized price adjustments at critical decision-making moments. While effective in some scenarios, these methods still lack granularity, particularly in distinguishing between price-sensitive customers and those willing to pay full price organically. Consequently, dynamic pricing systems often result in blanket price adjustments that overlook customer-specific behaviors, leading to suboptimal outcomes.

Another prominent technique in the domain of campaign optimization is uplift modeling [3, 34], which measures the incremental impact of marketing actions on individual behaviors. Although uplift modeling is effective in evaluating campaign effectiveness, its reliance on high-quality experimental data with well-designed control and treatment groups presents a significant challenge in real-world scenarios. The resulting data sparsity and missing labels, particularly for customers who have never been subjected to campaigns, further limit its applicability.

Our proposed method addresses several of these limitations. Unlike traditional RFM techniques, it incorporates value proposition features to create a more holistic customer segmentation framework. By distinguishing between price-sensitive and non-price-sensitive users, our approach avoids the pitfalls of generalized dynamic pricing systems, ensuring personalized interventions for each customer segment. Additionally, our method overcomes the missing label issue inherent in uplift modeling by employing robust label correction techniques within a two-stage model architecture. This enables precise targeting of both highly transactional and value-driven customers, significantly enhancing campaign effectiveness and customer engagement.

## 3 Methodology

### 3.1 Latent Customer Segmentation

Based on transaction history and engagement with marketing initiatives, we identify three customer segments:

- Engaged segment consists of customers who have previously engaged with our marketing campaigns and have made a subsequent purchase.
- Unengaged segment consists of customers who have not engaged with our marketing campaigns and have made purchases without any known marketing influence.
- Inactive segment consists of customers who have not engaged with our marketing campaigns and have not made any recent purchases.

Understanding these customer segments can help tailor our marketing strategies and increase the effectiveness of our campaigns. However, the segmentation only considers customer behaviors while overlooking their underlying intents, particularly for the engaged group which could have two different motivations. Some engaged customers are prompted, and due to lower perceived value, they could potentially shift their loyalty to competing products or brands in the absence of such campaigns. Conversely, other organically engaged customers, who have similar user pattern as the unengaged segment, might continue their purchases irrespective of any marketing influence.

To reduce marketing costs and boost marginal profits, Walmart and other e-commerce companies aim to focus on the first scenario, targeting customers who are undecided and can be nudged by marketing campaigns. This requires model to not only predict customer responses to campaigns precisely, but also comprehends the underlying intent, which is often not directly observable and results in a missing label issue.

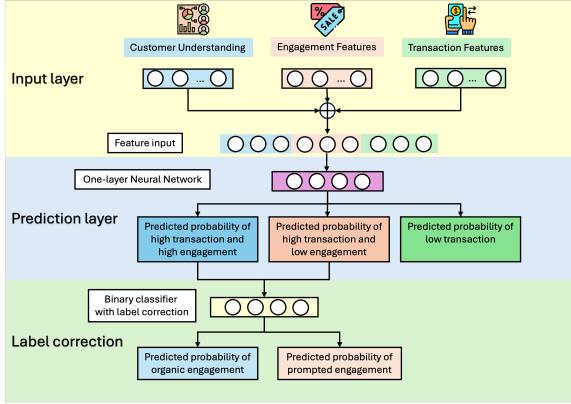
### 3.2 Overall Architecture

To solve the missing label issue and capture customer intent behind their behaviors, we propose a two-stage architecture as shown in Figure 1. The proposed architecture predicts customer behaviors and excludes inactive customers in the first stage. In the second stage, the proposed architecture further predicts customer intent with label correction, and identifies the true incremental gains achievable through marketing promotions, i.e. customers who will respond but would not have made a purchase without exposure to campaigns.

In the first stage, we train a single-layer neural network for multi-class classification, and three segments defined in Section 3.1 are labeled as three classes. Considering the imbalanced sample sizes of three classes, we adopt a weighted categorical cross entropy loss given by

$$\mathcal{L}_{\text{CCE}} = - \sum_{i=1}^N \sum_{c=1}^C w_c y_{i,c} \log(p_{i,c}), \quad (1)$$

where  $N$  denotes the number of samples,  $C$  denotes the number of classes,  $w_c$  is the weight for class  $c$ ,  $y_{i,c}$  is a binary indicator (i.e. 0 or 1) if class  $c$  is the correct classification for the  $i$ -th sample, and  $p_{i,c}$  is the predicted probability that the  $i$ -th sample is of class  $c$ . Here each class  $c$  is given a weight  $w_c$ , which adjusts the contribution of



**Figure 1: The proposed model architecture for predicting customer behaviors and intent, including label correction and incremental gain identification.**

that class to the total loss and helps to address the class imbalance by giving more importance to underrepresented classes.

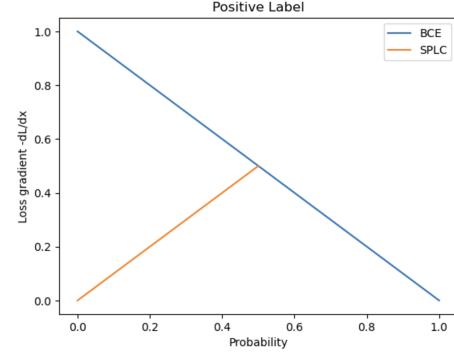
For the  $i$ -th sample, the model predicts a probability vector  $\hat{p}_i = (\hat{p}_{i,0}, \hat{p}_{i,1}, \hat{p}_{i,2})$ , where each component represents the predicted probability of the sample customer belonging to the corresponding segment, respectively. Given the model prediction on customer behaviors, we can distinguish those customers with high predicted probabilities of belonging to the inactive segment, and focus on customers more likely to have engagement with marketing campaigns.

In the second stage, we aim to distinguish value proposition customers and predict if customers will make purchases without any marketing stimuli. But for the engaged group who have been exposed and engaged with marketing campaigns, whether they would have purchased the product without exposure is unknown. So we start with initial labels generated by customer behaviors and make label correction at further iterations. Initial labels are assigned as labeling the unengaged segment as the negative class that will make purchases without any marketing incentives, while labeling the engaged segment as the positive class. True positive labels are users who promptly engaged and motivated by marketing incentives, while true negative labels are organically engaged without marketing. In the initial labels, we could have mislabeled some customers in the engaged group who would have made a purchase even without campaigns.

Thus for label correction, we need a robust loss function of binary classification, which can flip negative labels to positive if predicted probability is lower than a given threshold. So we start with a binary cross entropy (BCE) loss given by

$$\mathcal{L} = -\frac{1}{N} \sum_{i=1}^N [y_i L_i^+ + (1 - y_i) L_i^-], \quad (2)$$

where  $N$  denotes the number of users,  $y_i$  denotes the true label for the  $i$ -th user, and  $L_i^+$  and  $L_i^-$  denote the positive and negative losses



**Figure 2: Loss gradient by probability, illustrating the relationship between predicted probability and loss gradient for SPLC.**

of the  $i$ -th user. The positive and negative losses are given by

$$\begin{cases} L_i^+ = \log p_i \\ L_i^- = \log(1 - p_i) \end{cases}, \quad (3)$$

where  $p_i$  denotes the predicted probability of the  $i$ -th user being in the positive class.

To make the BCE loss adaptive to correct false positives, we want  $L_i^+$  to fall back to  $L_i^-$  when model prediction suggests the  $i$ -th user is highly likely to engage organically. Following [45], we adopt the self-paced loss correction (SPLC) and change the positive loss to

$$L_i^+ = \begin{cases} \log p_i, & \text{if } p_i > \tau \\ \log(1 - p_i), & \text{if } p_i \leq \tau \end{cases}, \quad (4)$$

where  $\tau \in (0, 1)$  denotes a threshold to identify whether it is a true positive or false positive based on the model prediction probability  $p_i$ . Specifically, for a given positive label, it would be corrected with low probability  $p_i > \tau$ .

To make the robust loss function more flexible to incorporate business insights, we further define the threshold  $\tau$  as a tunable parameter. Then finalized loss function is given by

$$\begin{cases} L_i^+(\theta) = \mathcal{I}(p_i(\theta) \leq \tau(\theta)) \log(1 - p_i(\theta)) \\ \quad + \mathcal{I}(p_i(\theta) > \tau(\theta)) \log p_i(\theta) \\ L_i^-(\theta) = \log(1 - p_i(\theta)) \end{cases}, \quad (5)$$

where  $\mathcal{I}$  denotes an indicator function and  $\theta$  denotes parameters of binary classifier. For the  $i$ -th sample, the model generates a score as  $p_i$ , which is the predicted probability of the customer only converting if exposed to marketing campaigns. In the tailored marketing strategy, we tend to prioritize those customers with high scores due to marketing costs.

Figure 2 demonstrates how the SPLC adapts to false positives with an example of  $\tau(\theta) = 0.5$ . Here, the x axis is the predicted probability  $p_i$ , and the y axis is the loss gradient  $-\frac{dL}{dx}$  where  $x$  is the predict logit and  $y = (1 + e^{-x})^{-1}$ . For a positive label, the smaller the probability, the more likely it is a false positive. Compared to the BCE, for which the loss increases as the predict probability decreases, the SPLC loss function assigns a lower weight to the false positives and reduces the effect of missing labels.

## 4 Experiment

In this section, we will discuss the offline model evaluation, and the online experiments and deployment in detail. We first introduce the dataset and the evaluation metrics used in this paper. Then, we illustrate the effectiveness of model by comparing to a strong baseline model widely used in online advertising. Finally, we present the successful online performance of the model on the Walmart promotion engine.

### 4.1 Dataset

We collected a sample of marketing response data and utilized a variety of features that help us gain insight into customer behavior and preference. User features were collected from user data before marketing started and labels were collected from user response after marketing initiatives were launched.

User labels for the multi-class neural network at the first stage are generated based on user behaviors including engagement with campaign and transactions after that. For the binary label correction stage, true user labels are generated on users who were initially not included in marketing events but got exposure later, so we could identify if their engagement is prompted or organic.

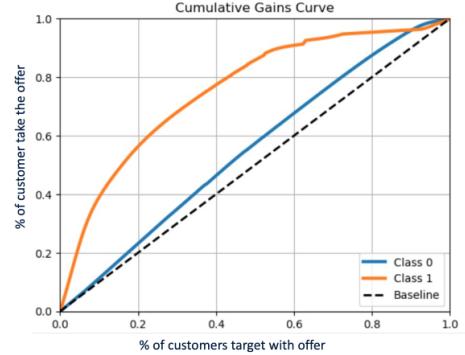
User features are primarily divided into three categories: customer understanding, user engagement, and transactions. Customer understanding features include their family type, price affinity, fulfillment preferences, and personas. User engagement metrics include all impressions, views and click interactions on banners, campaigns, and splash pages. Data was collected from logs of all Walmart's platform experiences: desktop web, mobile web, and IOS and Android native apps. Transaction features include the customer's transaction history with Walmart.

### 4.2 Baseline and Metrics

To illustrate that the optimal campaign audience are not highly engaging customers but high value proposition customers, we use a transaction-based user segmentation called RFM segment as the baseline model, compared to our proposed value proposition based user segment. RFM is a marketing technique used to quantitatively rank and group customers based on the recency, frequency and monetary total of their recent transactions. We selected the top 3 RFM segments of the most active transacting customer segments (i.e. champions, loyal customers and prospects) as our baseline model.

For comparison of classification models, we report the precision, recall, balanced accuracy and weighted F1-score of the models as the metrics to evaluate their performance. The definitions of these metrics are listed as follows:

- Precision measures the proportion of positive predictions that are actually correct.
- Recall measures the proportion of true positive cases that the model correctly identifies.
- Balanced accuracy is the average of sensitivity (recall for the positive class) and specificity (recall for the negative class).
- Weighted F1-score is a weighted average of the F1 scores for each class, where the weights are proportional to the number of samples in each class.



**Figure 3: Cumulative gain curve.** The curve illustrates how many customers to target and what returns to expect from a marketing campaign.

### 4.3 Offline Evaluation

We evaluated the offline model performance compared to the top 3 RFM segments that are most active in transactions.

Table 1 shows the proposed model outperforms the baseline on the common evaluation metrics for a classification model such as precision, recall, balanced accuracy and weighted F1 score.

Figure 3 shows the cumulative gain curve to further illustrate that the proposed model aids in resource allocation by focusing on the most promising customers. X axis represents the percentage of customers that have been targeted, and y axis represents the cumulative percentage of actual promptly engaged customers that can be captured in target audience. Figure 3 reveals that by targeting 40% of model-selected customers, 78% of the actual value proposition customers can be reached. That is, despite a reduction in the overall campaign exposure rate, we can achieve a higher conversion rate. The curve helps us make informed decisions regarding the number of customers required to be targeted for achieving marketing objectives.

### 4.4 Online evaluation

To assess the efficacy of our model in real-world scenarios, we conducted an online A/B testing experiment. The marketing campaign targets a considerable user base of 10.8 million, with 6.6 million users who have never been part of Walmart campaigns, and 4.2 million users who used to engage with campaigns but have been churned. This test was carefully structured to incorporate two concurrently running setups: a control group and a treatment group.

To ensure balanced sample sizes and unbiased results, we randomly split the users into a control group and a treatment group in a stratified manner to make sure the population sizes of churned and never users are the same in the control and treatment group, respectively. For the treatment group, we ranked all the users based on their model scores and selected the top users to show the offers. For the control group, we random sampling a subset to get an equal amount of users. Both the control and the treatment group got identical exposure to the marketing campaign.

Table 2 and Table 3 were obtained by comparing the results of random group and model group. On both user groups, there was

**Table 1: Offline evaluation results.**

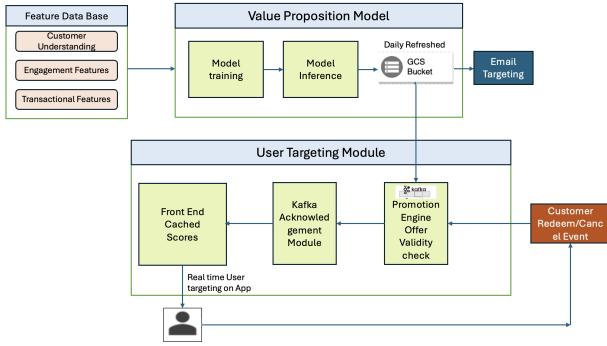
|                          | precision    | recall       | Balanced accuracy | Weighted F1 score |
|--------------------------|--------------|--------------|-------------------|-------------------|
| Champion (RFM)           | 0.371        | 0.433        | 0.624             | 0.746             |
| Loyalist (RFM)           | 0.280        | 0.207        | 0.537             | 0.718             |
| Potential Loyalist (RFM) | 0.273        | 0.132        | 0.522             | 0.720             |
| Model segment            | <b>0.448</b> | <b>0.600</b> | <b>0.724</b>      | <b>0.817</b>      |

**Table 2: Online evaluation results for new customers.**

| group         | User count | %Site visit   | Conversion rate |
|---------------|------------|---------------|-----------------|
| Random Sample | 3.3M       | 79%           | 0.17%           |
| Model Driven  | 3.3M       | 96%           | 0.34%           |
| Lift          |            | <b>21.52%</b> | <b>99.85%</b>   |

**Table 3: Online evaluation results for churned customers.**

|               | User count | %Site visit  | Conversion rate |
|---------------|------------|--------------|-----------------|
| Random Sample | 2.1M       | 87%          | 0.68%           |
| Model Driven  | 2.1M       | 90%          | 0.76%           |
| Lift          |            | <b>3.45%</b> | <b>10.78%</b>   |

**Figure 4: Overall framework of system deployment.**

a statistically significant increase in conversion rates. Specifically, there was a remarkable 99.85% uplift among never customers, and a 10.78% lift in previously churned customers. This data demonstrates the effectiveness of our model in targeting potential customers, resulting in a significant improvement in their overall engagement and interaction with the brand.

#### 4.5 Online Deployment

The proposed model has been deployed to support campaign targeting by emails and promotion engine for real-time user advertising since April 2024.

Figure 4 illustrates the role of the value proposition model in the Walmart ecosystem. The user features are collected and saved in the back end, and passed for model inference on a daily basis. The model refreshes user scores in a batched manner and updates

in cloud-based storage. For email targeting, the optimal user list is directly retrieved from Google cloud storage. For online advertising triggered by real-time site visit, the model-driven user list will be checked by promotion engine for offer validity and sent to front end via Kafka acknowledge module. Customer response data keeps being feed back to promotion engine for logging and model re-training purpose.

#### 5 Conclusion

The paper introduced a two-stage multi-model architecture that utilizes Self-Paced Loss to improve customer categorization, addressing the issue of striking a balance between value proposition and long-term brand value. This system allows for a more efficient customer segmentation by discerning between customers who engage due to marketing campaigns and those who engage naturally.

The best trained model was deployed to our production environment and evaluated in an online A/B test experiment. Based on model predictions, we achieved a nearly 100% lift on conversion rate compared to the baseline model. Therefore, the proposed solution presents a novel approach to enhancing marketing strategies, reducing exposure rates, and increasing conversion rates. This paper highlights the need for dynamic customer engagement, paving the way for more complex, customer-centric strategies in the future.

## References

- [1] Mohammad Akbar. 2024. Customer-centric strategies: navigating the dynamics of marketing management for competitive advantage. *Advances in Business & Industrial Marketing Research*, 2, (May 2024). doi:10.60079/abim.v2i2.288.
- [2] Christy Ashley and Tracy Tuten. 2015. Creative strategies in social media marketing: an exploratory study of branded social content and consumer engagement. *Psychology & Marketing*, 32, (Jan. 2015). doi:10.1002/mar.20761.
- [3] Daniel Baier and Björn Stöcker. 2022. Profit uplift modeling for direct marketing campaigns: approaches and applications for online shops. *Journal of Business Economics*, 92, 4, 645–673.
- [4] Alexander Bleier and Maik Eisenbeiss. 2015. Personalized online advertising effectiveness: the interplay of what, when, and where. *Marketing Science*, forthcoming, (July 2015). doi:10.1287/mksc.2015.0930.
- [5] Haipeng Chen. 2012. When more is less: the impact of base value neglect on consumers' preferences for bonus packs over price discounts. *Journal of Marketing*, 76, (Jan. 2012), 64–77.
- [6] Shiuann-Shuoh Chen, Bhaskar Choubey, and Vinay Singh. 2021. A neural network-based price sensitive recommender model to predict customer choices based on price effect. *Journal of Retailing and Consumer Services*, 61, 102573. doi:https://doi.org/10.1016/j.jretconser.2021.102573.
- [7] Joe Cooprider and Shima Nassiri. 2023. Science of price experimentation at amazon. *Business Economics*, 58, (Jan. 2023), 1–8. doi:10.1057/s11369-023-00303-9.
- [8] Omkar Dastane and Andy Chin Fook. 2021. Effectiveness of loyalty programs in customer retention: a multiple mediation analysis. *Jindal Journal of Business Research*, 10, (June 2021), 1–27. doi:10.1177/22786821211000182.
- [9] Elmira Djafarova and Tamar Bowes. 2020. 'Instagram made me buy it': generation z impulse purchases in fashion industry. *Journal of Retailing and Consumer Services*, 59, (Oct. 2020), 102345. doi:10.1016/j.jretconser.2020.102345.
- [10] Sahar F. 2018. Machine-learning techniques for customer retention: a comparative study. *International Journal of Advanced Computer Science and Applications*, 9, (Jan. 2018). doi:10.14569/IJACSA.2018.090238.
- [11] Muhammad Fawad, Muhammad Sheikh, Muhammad Rizwan, and Qasim Maqsood. 2014. The role of brand and company image in building brand loyalty through service quality and brand trust. *Journal of Sociological Research*, 5, (May 2014). doi:10.5296/jsr.v5i1.6571.
- [12] Amir Gandomi and Murtaza Haider. 2015. Beyond the hype: big data concepts, methods, and analytics. *International Journal of Information Management*, 35, (Apr. 2015), 137–144. doi:10.1016/j.ijinfomgt.2014.10.007.
- [13] Leo Guelman, Montserrat Guillén, and Ana M. Pérez-Marín. 2015. A decision support framework to implement optimal personalized marketing interventions. *Decision Support Systems*, 72, 24–32. doi:https://doi.org/10.1016/j.dss.2015.01.010.
- [14] Sunil Gupta and Donald Lehmann. 2003. Customers as assets. *Journal of Interactive Marketing*, 17, (Dec. 2003), 9–24. doi:10.1002/dir.10045.
- [15] Tarun Gupta and Supriya Bansal. 2021. Impact of paid advertising on organic sales. *International Journal of Science and Research (IJSR)*, 10, (Nov. 2021), 1495. doi:10.21275/SR24213021840.
- [16] Kaizhu Huang, Shufei Zhang, Rui Zhang, and Amir Hussain. 2020. Novel deep neural network based pattern rui classification architectures. *Neural Networks*, 127, 82–95. doi:https://doi.org/10.1016/j.neunet.2020.03.011.
- [17] Christina T. Fong Jeffrey Pfeffer. 2005. Building organization theory from first principles: the self-enhancement motive and understanding power and influence. *Organization Science - ORGAN SCI*, 16, (Aug. 2005), 372–388. doi:10.1287/orsc.1050.0132.
- [18] Gayathri K and Arunodhaya R. 2021. Customer segmentation and personalized marketing using k-means and apriori algorithm. In *Proceedings of the First International Conference on Combinatorial and Optimization, ICCAP 2021, December 7-8 2021, Chennai, India*. EAI, (Dec. 2021). doi:10.4108/eai.7-12-2021.2314561.
- [19] Shirsendu Kar. 2023. Impact of artificial intelligence on digital marketing. *INTERANATIONAL JOURNAL OF SCIENTIFIC RESEARCH IN ENGINEERING AND MANAGEMENT*, 07, (July 2023). doi:10.55041/IJSREM25001.
- [20] Jong Woo Kim, Byung Lee, Michael Shaw, Hsin-Lu Chang, and Matthew Nelson. 2001. Application of decision-tree induction techniques to personalized advertisements on internet storefronts. *International Journal of Electronic Commerce / Spring*, 5, (Mar. 2001), 45–62. doi:10.1080/10864415.2001.11044215.
- [21] V. Kumar. 2018. Transformative marketing: the next 20 years. *Journal of Marketing*, 82, (July 2018), 1–12. doi:10.1509/jm.82.41.
- [22] Mui Kung, Kent Monroe, and Jennifer Cox. 2002. Pricing on the internet. *Journal of Product & Brand Management*, 11, (Sept. 2002), 274–288. doi:10.1108/10610420210442201.
- [23] Michel Laroche, Isar Kiani, Nectarios Economakis, and Marie-Odile Richard. 2014. Effects of multi-channel marketing on consumers' online search behavior: the power of multiple points of connection. *Journal of Advertising Research*, 53, (Jan. 2014), 431. doi:10.2501/JAR-53-4-431-443.
- [24] Katherine Lemon and Peter Verhoef. 2016. Understanding customer experience throughout the customer journey. *Journal of Marketing*, 80, (June 2016). doi:10.1509/jm.15.0420.
- [25] Marcin Majka. 2024. Understanding customer acquisition cost, (July 2024).
- [26] Wendy Moes and Peter Fader. 2004. Dynamic conversion behavior at e-commerce sites. *Management Science*, 50, (Mar. 2004), 326–335. doi:10.1287/mnsc.1040.0153.
- [27] Georgios P. Papamichail and Dimitrios P. Papamichail. 2007. The k-means range algorithm for personalized data clustering in e-commerce. *European Journal of Operational Research*, 177, 3, 1400–1408. doi:https://doi.org/10.1016/j.ejor.2005.04.011.
- [28] Jegatheeswari Perumalsamy, Bhavani Krothapalli, and Chandrashekhar Althati. 2022. Machine learning algorithms for customer segmentation and personalized marketing in life insurance: a comprehensive analysis, (Dec. 2022).
- [29] Chandra Potluri, G. Rao, L.V.R Kumar, Kedir Allo, Yaregal Tarekegn, and Abdilkerim Seman. 2024. Machine learning-based customer segmentation and personalised marketing in financial services, (May 2024), 1570–1574. doi:10.1109/IC3SE62002.2024.10593143.
- [30] Kaledio Potter, Favour Olaoye, and Lucas Doris. 2024. Analyzing the effectiveness of personalized marketing strategies in improving customer engagement and loyalty. *Advances in consumer research. Association for Consumer Research (U.S.)*, (Oct. 2024).
- [31] João Ramos and Fabrício Silva. 2023. Customer lifetime value prediction: a machine learning approach. *Anais do Encontro Nacional de Inteligência Artificial e Computacional (ENIAC)*, (Sept. 2023), 486–500. doi:10.5753/eniac.2023.234262.
- [32] F. Reichheld and P. Schefter. 2000. E-loyalty: your secret weapon on the web. *Harvard Business Review*, 78, (July 2000).
- [33] Roland Rust, Katherine Lemon, and Valarie Zeithaml. 2004. Return on marketing: using customer equity to focus marketing strategy. *Journal of Marketing - J MARKETING*, 68, (Jan. 2004), 109–127. doi:10.1509/jmkg.68.1.109.24030.
- [34] Piotr Rzepakowski and Szymon Jaroszewicz. 2012. Uplift modeling in direct marketing. *Journal of telecommunications and information technology*, 43–50.
- [35] Krishnamoorthi S and Gopal Shyam. 2024. Drl-hifa: a dynamic recommendation system with deep reinforcement learning based hidden markov weight update and factor analysis. *Multimedia Tools and Applications*, 83, (Mar. 2024), 1–25. doi:10.1007/s11042-024-18296-8.
- [36] Malay Sarkar, Aisharya Roy Puja, and Faiaz Rahat Chowdhury. 2024. Optimizing marketing strategies with rfm method and k-means clustering-based ai customer segmentation analysis. *Journal of Business and Management Studies*, 6, 2, 54–60.
- [37] Ralph Shad, Kaledio Potter, and Abram Gracias. 2024. Natural language processing (nlp) for sentiment analysis: a comparative study of machine learning algorithms. (Oct. 2024). doi:10.20944/preprints202410.2338.v1.
- [38] Jagdish Sheth. 2020. Impact of covid-19 on consumer behavior: will the old habits return or die? *Journal of Business Research*, 117, (June 2020). doi:10.1016/j.jbusres.2020.05.059.
- [39] Siddharth Singh and Dipak Jain. 2010. Measuring customer lifetime value: models and analysis. *Review of Marketing Research*, 6, (Nov. 2010), 37–62. ISBN: 978-0-7656-2127-6. doi:10.1108/S1548-6435(2009)0000006006.
- [40] Richard Thaler. 1985. Mental accounting and consumer choice. *Marketing Science*, 4, (Aug. 1985), 199–214. doi:10.1287/mksc.4.3.199.
- [41] Jari Vesanen. 2007. What is personalization? a conceptual framework. *European Journal of Marketing*, 41, (June 2007), 409–418. doi:10.1108/03090560710737534.
- [42] Michel Wedel and P. K. Kannan. 2016. Marketing analytics for data-rich environments. *Journal of Marketing*, 80, (May 2016). doi:10.1509/jm.15.0413.
- [43] Bradley Wilson. 2019. Book review: marketing to gen z: the rules for reaching this vast and very different generation of influencers. *Journal of Macromarketing*, 39, (Feb. 2019), 027614671983015. doi:10.1177/0276146719830156.
- [44] Dennis Zhang, Hengchen Dai, Lingxiu Dong, Fangfang Qi, Nannan Zhang, Xiaofei Liu, Zhongyi Liu, and Jiang Yang. 2018. How do price promotions affect customer behavior on retailing platforms? evidence from a large randomized experiment on alibaba. *Production and Operations Management*, 27, (Dec. 2018), 2343–2345. doi:10.1111/poms.12964.
- [45] Youcai Zhang, Yuhao Cheng, Xinyu Huang, Fei Wen, Rui Feng, Yaqian Li, and Yandong Guo. 2021. Simple and robust loss design for multi-label learning with missing labels. (2021). <https://arxiv.org/abs/2112.07368> arXiv: 2112.07368 [cs.LG].