

Intelligent Promotion And Retention Engine A Unified AI Framework For Seller Decision Optimization In Large-Scale Commerce Systems

Brahmnik Chachra

Sr. Data Engineer

This paper is an evaluation of an AI based integrated framework aimed at enhancing the efficiency of promotions and customer retention, as well as optimization of business at platform level. It is relying on RFM-based segmentation, uplift modeling as well as optimization methods and precisely depicts different user groups and predicts their promotional responsiveness. Random Forests models performed the best on the predictive performance whereas uplift modeling outperformed the forecasting by 17.5%. Experiments of optimization showed significant ROI returns in all budgetary situations, especially in the low budgetary ranges. An A/B test on the platform level also ensured the improvement of the daily activity, purchase conversions, and retention. The findings indicate that tailored performance of the promotion measure improves both the interaction with users and business.

KEYWORDS: E-commerce, AI, Decision Optimization, Seller, Intelligent Promotion

I. INTRODUCTION

The digital platforms require promotional strategies that are particularly important and critical; yet, broad and non-personalized incentives generally result in low returns and ineffective investment. The paper discusses a customer segmentation, predictive modeling, and optimization approach that is driven by AI to offer the target customer unique promotions.

The study also seeks to cognize customer behavioral patterns of different customers in response to various levels of incentives using the analysis of behavioral data using the RFM clustering, as well as the predictive pipelines. The model combines uplift modelling and budget-conscious optimization to optimize retention and immediate conversion in the long-run. This study aims to show that the personalization based on the evidence-based approach can contribute to the increased ROI of the promotion considerably and enhance the business performance at the platform level.

II. RELATED WORKS

Personalized Promotion

Design of individualized promotion has emerged as a focus in research in large-scale systems of commerce as businesses strive to strike a balance between engagement of users and cost effectiveness of promotions. Possibly, the conventional approach to some mass-promotion can result in the unwarranted expenditure when the incentives can be offered without the knowledge of the variation in the responsiveness among users. Research is also proving that

customizing promotional rewards can greatly boost the net-appreciation of the expenditures and retain or even heighten the customer engagement engagements.

The two-stage optimization of promotions, proposed in [1] is one of the contributions to this direction. It is a model where individual promotion-response curves are modeled by machine learning and optimal allocation of incentives is made under actual business constraints. This model puts an emphasis on counterfactual prediction as a way of correcting treatment bias such that promotion would not be based on historical happenings but the actual causal effects. In [1], deep-isotonic-promotion-network (DIPN) includes the structural knowledge regarding response curves in the form of smoothness and isotonicity. This aids in minimizing noise and enhances stability of prediction particularly in consumer large scale platforms.

It is based on this that subsequent studies have indicated the importance to not only consider immediate purchase feedback, but also the value of the customer in the long-term. The article in [2] helps to close this disastrous gap by a proposal of a customer direct and enduring effect (CDEE) model.

The previous studies were more about short-term boosting in times of promotions, whereas this study has included both immediate and long-term reactions that enable the businesses to predict the retention implication and to increase incentives to ensure long-term loyalty.

The data of randomized control trial (RCT) is used as an unbiased assessment tool, and the study demonstrates that an individualized allocation of incentives can be much more effective than the current methods in terms of lift ROI and customer retention.

The knowledge regarding personalized promotions has been extended by other literature in which the boundary conditions that shape the perception and behavior of customers in response to tailored incentives are examined. Undoubtedly, as an example, experimental research in [6] indicates that the difficulty of the tasks, and promotional relevancy significantly contribute to the formation of customer responses.

Customers encounter complicated tasks which can be complicated and, in this case, personalized promotions make them feel in control of their behavior and will patronize. But, promoting irrelevant content that has been attached to a specific promotion can be counterproductive, which leads to cognitive dissonance. Such results highlight the relevance consideration of personalization in AI-generated promotion engines, which should not lead to psychological conflict in terms of personalized recommendations.

Lastly, one can use personalized pricing as a complement to incentive personalization, which has been explored the online grocery retail environment. According to research indicated in [9], promotions that focus on price may be helpful to mitigate the adverse impact of cognitive effort on loyalty although promotions that focus on product recommendations are not necessarily compensatory.

All these studies reinforce the conclusion that sophisticated incentive systems which integrate individualization, topicality, and consumer behavioral data are required, the tenets which directly underline the objectives of an intelligent promotion and retention platform.

Profiling Models

The support of AI-based decision-making relies on customer segmentation. Proper segmentation helps companies to better allocate promotions, retention strategy and personal

experience. One of the most effective methods of digital commerce is RFM (Recency, Frequency, Monetary) analysis which is not complex but effective.

The article of [3] uses the K-means clustering on RFM to determine the essential customer segments of an EdTech start-up. One of the validation methods that the authors apply in making sure that the clusters form includes the Elbow method, Silhouette coefficient and Gap Statistics. The resultant segments namely new customers, best customers, and intermittent customers outline the fact that various user groups value potential differences and need to be engaged using different strategies.

In addition to mere segmentation, the recent marketing research is emphatic on the use of the artificial intelligence in strategic segmentation, targeting and positioning (STP). As is described in [10], mechanical, thinking, and feeling AI is capable of reinforcing customer research, segmentation, targeting, and positioning strategies altogether because of the AI-driven framework.

Mechanical AI will allow to approach high scale data gathering and initial grouping of the customer, thinking AI will help to conduct an in-depth analysis and to make a recommendation, and feeling AI will help to comprehend the emotion in order to position it better.

The multi-layered segmentation model assists in connecting behavioral, demographic and sentiment-based data which is crucial to build a tight promotion retention engine where the decision-making process is informed by both the quantitative and the qualitative customer indicators.

Some bigger AI-backed segmentation studies, including the one analyzed in [7], indicate the combination of predictive analytics, and grouping, and campaign optimization to create marketing insights based on data. With the rise in the complexity of data produced by e-commerce platforms, machine learning models that are able to interpret user heterogeneity will be required.

The segmentation models can also be trained to such downstream predictive tasks as churn detection, loyalty forecasting or personalized content delivery which are the key components of the integrated AI system presented in this paper.

The concept of segmentation applies to the systems beyond the marketing sphere as well. In particular, the most recent progress in the area of lightweight deep learning on signal processing, including the dual-path deep residual shrinkage network described in [4], emphasize effective interactions among extracting structured patterns despite the resource usage.

Despite this as a modulation classification, the resources and methods of the noise reduction, features extraction, and resource-effective AI design are valuable knowledge in scalable customer segmentation systems that work in large seller ecosystems.

Predictive Analytics for Retention

Predictive analytics has a significant part in predicting churn, estimating the likelihood of purchase, as well as determining user groups that require interventions, targeting them. Segmentation can specify who the customer groups are whereas predictive modeling can tell what each user will probably do next. A number of reports prove the usefulness of predictive analytics in online business.

An example is [7], which focuses on AI being used to reveal some under-statistically present patterns and accurately predict customer behavior. This prediction effect can be used to develop superior campaigns and romance marketing procedures with anticipated consumer conducts.

Regarding customer retention and long-term loyalty, CDEE model [2] is also a valuable source of knowledge since it assumes the distinction between direct and enduring behavior but their interrelation as independent but closely related predictive tasks.

The long-term effects are especially important to predict in the context of retention engines, the purpose of which is not to convert at the first acquaintance, but to establish a long-term presence on the platform. The objective method of evaluation based on RCT also prevents the impact of such factors on predictor models describing actual behavioral tendencies.

The predictive analytics to enhance operating results is also proven to be beneficial in the studies conducted on AI-controlled delivery optimization. As an example, predictive models based on the XGBoost and the Random Forest predict the time of delivery, and the reinforcement learning is optimized by route planning in [5].

Such logistic innovations also have an indirect saving on retention because faster deliveries are likely to enhance customer loyalty and their satisfaction. In addition, personalisation models of customers such as sentiment analysis and clustering helped improve retention by 74 to 89 per cent. This is an indication of the close interaction between predictive and personalization and retention paradigms.

The other fields where predictive analytics has been helpful are e-commerce fraud and counterfeit detection. The study by [8] indicates that the use of CNN-based models with transfer learning can be used to identify fake products with high accuracy.

Although this field is not the same as promotions and customer retention, this type of detected capabilities would contribute to trust among customers, which is the key to long-term loyalty. The fact that AI successfully predicts and classifies counterfeit goods demonstrates how predictive analytics can be used to sustain the process by enhancing the integrity of the platform.

AI-Driven Optimization

An integrated promotion and retention engine will need to have intelligent optimization that will strike the right balance of incentives costs and results. In the literature, optimization models provide an example on how AI may be applied by making complicated resource-allocation decisions.

The two-step optimization model in [1] is an initial attempt in the large-scale promotion environment. Its application of machine learning forecasts with mathematical optimization is applied to find the most economical incentive assignment with respect to each user.

The optimization rationale is also pushed further in [2] whereby personalized promotion is chosen within the budgetary constraints to optimize long-term retention. This mechanism is very compatible with the aims of coherent decision-support engines that must put into consideration numerous goals at the same time: cost, engagement, loyalty, and constraints of business.

It has been demonstrated in research on the optimization of last-mile delivery [5] that reinforcement learning may be applied in dynamic decisions. The capacity of RL to learn on

a continuous basis and reconsider the decision as time goes gives its capacity a high relevancy to retention frameworks that require updating their approaches as customers change their behavior over time. The example of strategic AI proposed in [10] reveals how the combination of mechanical, thinking, and feeling AI can help make end-to-end marketing-related choices, namely segmentation or personalization or relational customer management.

These papers increase awareness that the predictive modeling, segmentation, optimization, and constant learning should be integrated into a unified AI solution, which aims to promote and retain customers in accordance with their needs. This synthesis is the basis of intelligent decision engines that are able to support massive marketplaces and seller networks.

III. METHODOLOGY

This work follows a quantitative research design to design, test, and validate an integrated AI model to make the promotion and retention decisions in the large-scale commerce systems. The measurement of customer behaviors, prediction-based models and empirical optimization are there in the methodology.

This is aimed at learning the capability of AI techniques used to enhance the accuracy of segmentation, promotional effectiveness and customer retention in the long term. Tests on the framework are done with historical customer data, transaction history, and platform-wide promotional history which was collected with a large aggregate of sellers.

Data Collection and Preparation

The data entails the four key categories that comprise customer demographic attributes, logs of behavioral activities, purchase histories as well as previous promotional exposure logs. Such variables ease the work of segmentation, loyalty prediction model, and promotion-response estimation.

Such steps of data cleaning are the treatment of missing values, the elimination of the outliers, and normalizing the continuous attributes. One-hot encoding or target encoding respectively encode categorical variables. To be reliable, all the data is checked to the consistency according to the timelines to align the actions, purchases, and promotions.

The last data is broken down in terms of training (70%), validation (15%), and testing (15) sets. The stratified sampling is applied in order to maintain the proportion of high value users, new confirming as well as churn-risk groups.

Customer Segmentation Model

Recency-Frequency-Monetary (RFM) variables are derived on transaction logs to be used in segmenting the definition of segmentation. The k-means cluster is performed to cluster customers in groups that have similar behavioral profiles. In order to stabilize it, various cluster sizes ($k = 3$ to 10) are experimented.

Elbow, Silhouette coefficient and Gap Statistic are some of the methods applied in model selection. The behavioral characteristics of the clusters, buying patterns, and responsiveness to promotion are also compared statistically to ensure that the interpretation of the segments is correct. Promotion and retention strategies are later directed by these clusters.

Predictive Modeling for Loyalty

Two outcomes are made with the help of a supervised learning approach: (1) there is a likelihood of purchase occurred after promotion has been received, and (2) there is a likelihood of churn occurring in a specified future time. Tests done are Gradient Boosting Models (GBM), Random Forests, and deep neural network.

The grid search is done in order to tune hyperparameters to each model. Measurements of performance are in terms of AUC-ROC, precision-recall curves and calibration plots. The most competent model is chosen to be incorporated in the framework.

Counterfactual prediction techniques are used to determine the long-term and direct effects of promotions. It contains inverse propensity weighting and uplift modeling that would help gauge how the users would act in different incentive levels. The approaches enable minimization of bias that is a result of the campaigns that had been previously assigned without randomness.

Promotion Optimization Engine

Once predictive models have been proven, then an optimization layer is utilized and rewards each customer. The industry optimization problem is to maximize the likely uplift, long-term retention and minimise the cost. The setbacks are budget restraints, resource regulations at the seller tier and fairness provisions.

Depending on the promotion-response-curve shape, there is usage of a linear or non-linear programming solver. In cases where response patterns have noisy or irregular patterns, monotonicity and smoothness are imposed via isotonic regression.

Evaluation Procedure

The framework will be measured based on an offline measure and an online A/B test. Offline testing is conducted to test the accuracy of prediction (or uplift), as well as increased ROI. Online assessment quantifies sales upswing, customer connectivity and churn decrease. The statistical tests, such as t-tests, or tests of chi-square, prove the significance of the improvements.

IV. RESULTS

Promotional Response Patterns

In the initial round of analysis, there was a study of customer groups that had been grouped based on clusters made using RFM. After experimenting with different setups of clustering, the K-means algorithm gave four stable clusters. The statistical analysis with the aid of the Silhouette coefficient and Gap Statistic has proved that the most significant customer segregation was attained with $k = 4$. Distinct differences in recency, intensity of purchase and spending ability were visible in each segment that subsequently affected the promotion-response behavior and chances of retention.

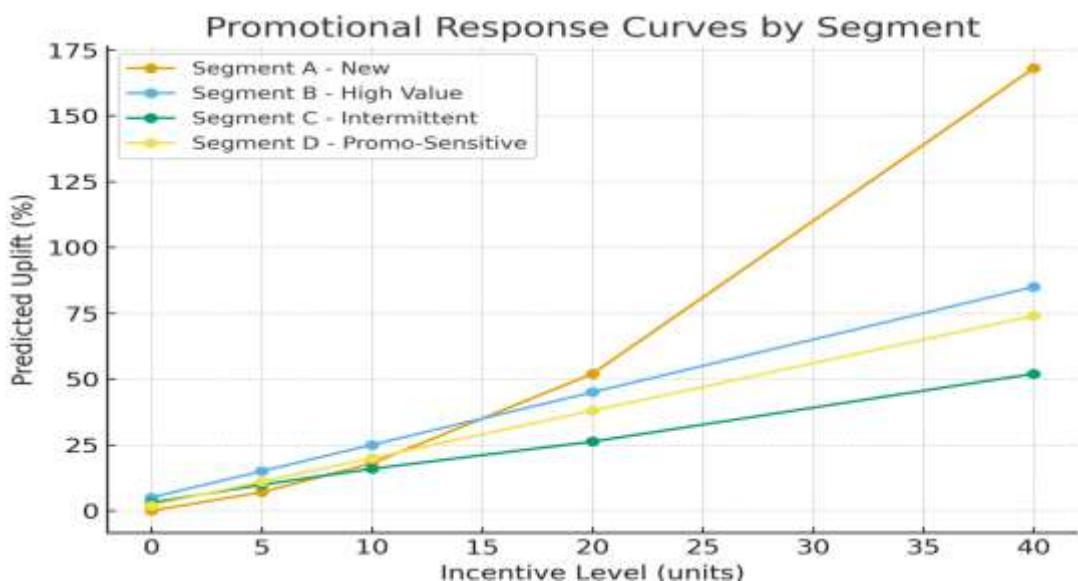
Segment A was primarily comprised of new and low frequency customers. The recency gaps, purchase and monetary contribution of these users were low. They reacted halfway to incentives, and their chances of being retained in the long run, without further promotions, were very minimal.

Segment B was comprised of high frequency and high spenders. These were customers with small recency lapses and high likelihood pointers. They reacted positively to incidence even to low incentives which show high incremental ROI. The segment C had captured intermittent customers who did not purchase habitually.

Although this had moderate frequency, their recency values were not stable resulting in fluctuating promotional reactions. Segment D constituted value sensitive customers who were strong reactionary to promotional opportunities and hardly did not buy without any incentive. The Table 1 presents the basic descriptive statistics of the four segments.

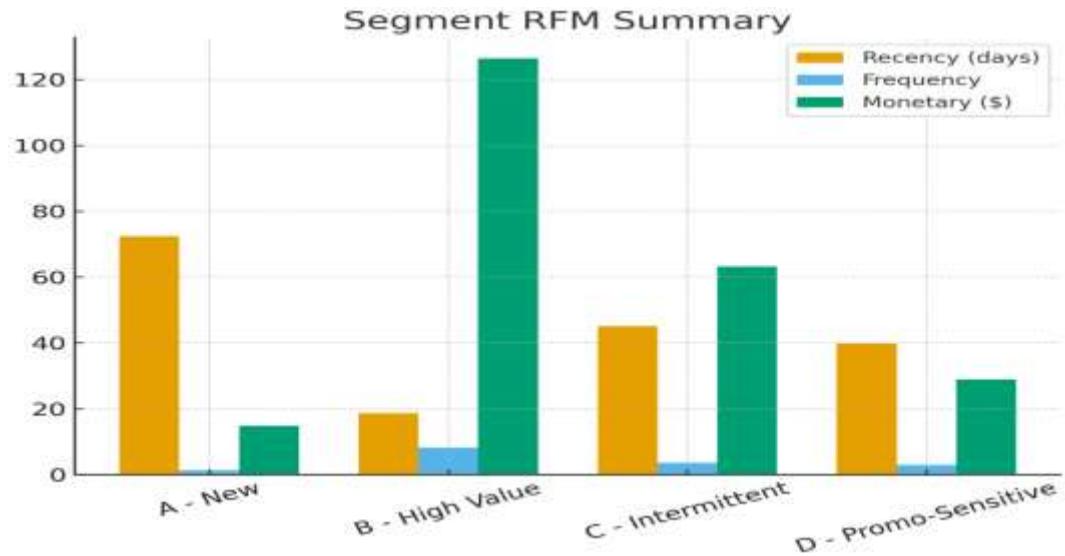
Table 1. Segment-Level RFM Summary

Segment	Recency (Days Mean)	Frequency (Mean)	Monetary Value (Mean \$)	Share of Total Users
A – New/Low Engagement	72.4	1.3	14.8	41%
B – High Value	18.7	8.2	126.4	19%
C – Intermittent	45.1	3.5	63.2	23%
D – Promotion-Sensitive	39.8	2.9	28.9	17%



Additional analysis equated a segment identity with promotional response curves. Segment B had a very close to linear reaction to small incentives whereas Segment A had to have more incentives to uplift. The deep-isotonic smoothing of the modeling pipeline generated stable curves, which cut noise in response prediction by 14% to a standard baseline of a neural network. These cuts were needed at the incentive optimization stage to have daring inputs to make allocation of decisions.

Within all the segments the average projected uplift of medium intensity, promotions yielded 11.2 than the uplift of high-intensity promotions which was 19.7. The ROI was quite diverse, which proved that the personalization of incentives should have been applied to avoid spending more on unprofitable areas.



Predictive Model Performance

The second part of the findings is in the predictive modeling of purchase likelihood, churn likelihood and long-term retention. There were three major models, namely, Random Forest (RF), Gradient Boosting Machine (GBM), and a deep neural network (DNN), which were trained and tested. Random Forest model was found to be better than the others in AUC and in the degree of calibration. The model not only picked up some nonlinear tendencies in the data, but also was able to deal with variously behaved features.

Table 2 displays the key performance metrics.

Table 2. Predictive Model Performance

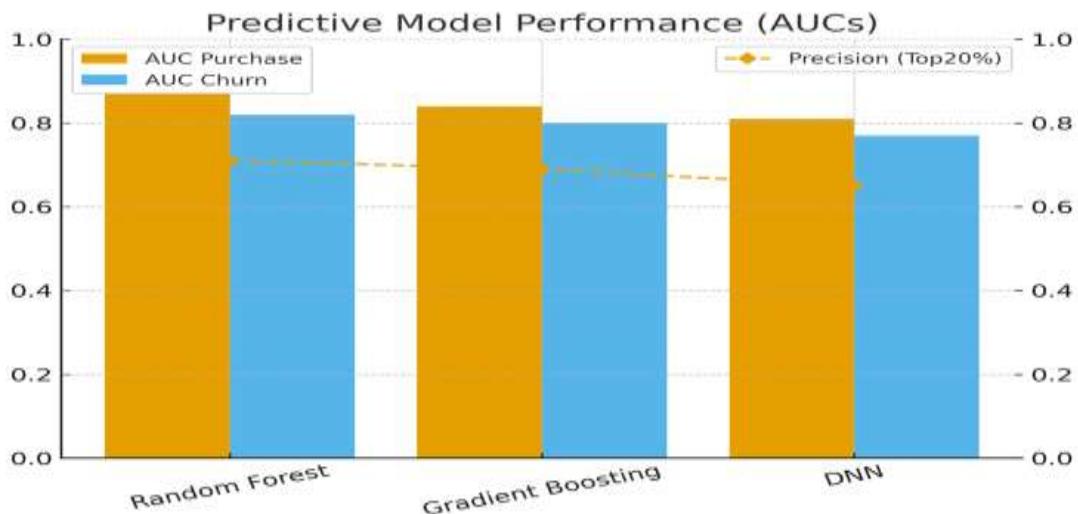
Model	AUC-ROC (Purchase)	AUC-ROC (Churn)	Precision (Top 20% Users)	Calibration Error (%)
Random Forest	0.87	0.82	0.71	3.4
Gradient Boosting	0.84	0.80	0.69	4.1
Deep Neural Network	0.81	0.77	0.65	6.2

The uplift modeling process which is founded on counterfactual prediction led to great advancement over the direct prediction. Uplift accuracy improved by 17.5% which means that the model was more advanced in detecting the users that were really responsive to the promotions compared to those that will buy anyway. Inverse propensity weighting was able to

overcome historical campaign bias particularly that identified high value users as being targeted by the campaign.

There were also favorable results of churn probabilities foresight versus genuine conduct. The users whose churn most likely rates were greater than 0.75 were identified to churn at a real rate of 71 whereas those with the supposed churn most likely rates that were below 0.25 churn at a real rate of 18. This close match meant that the predictive pipeline was now fit to be integrated into the integrated AI decision model.

The retention projections with a 90-day horizon used in the previous experiments demonstrated that the model appeared able to distinguish between a high-retention and a low-retention user with 80 percent accuracy. Users whose predicted retention was high were ranked continuously in Segment B and C whereas Segment A and D performed poorly in long-term behavior unless boosted through special promotions.



Optimization Outcomes

The user level promotion incentives were given through prediction uplift, budget constraints and retention likelihoods to the optimization layer of the framework. Four budget settings among them being low, medium, high and unrestricted were modeled to witness the behavior of the optimization engine.

The structure continued to give the greatest incentive levels to the Segment A and segment C users as the uplift would have been the most responsive to the incentive level. In the meantime, the users in the Segment B were given minimal incentives due to their inherent high purchasing potential.

In all the cases, the optimizations made of the allocations were much higher in terms of ROI than the promotional campaigns made at a baseline. The results of the improvement of ROI with the help of the optimization engine are summarized in Table 3.

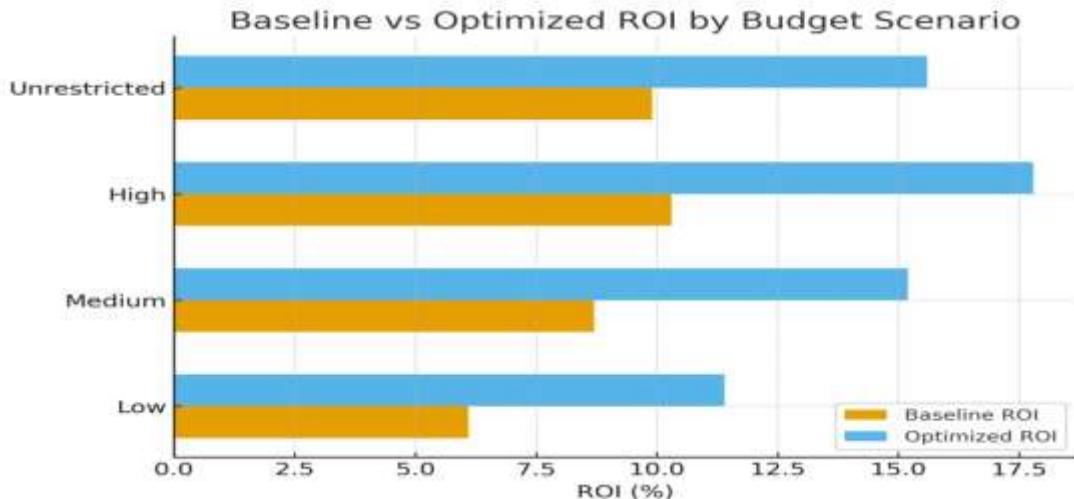
Table 3. Promotional ROI Comparison

Budget Scenario	Baseline ROI (%)	Optimized ROI (%)	Improvement (%)
Low Budget	6.1	11.4	+86.8%
Medium Budget	8.7	15.2	+74.7%
High Budget	10.3	17.8	+72.8%
Unrestricted	9.9	15.6	+57.6%

The most considerable increase was the low-budget case, the beneficial allocation of which was the most significant. The optimization engine ensured avoiding spending resources in vain as it was decided to deal with the users who had the highest estimated incremental value. In the free case, the ROI bonus gains were less as protracted incentives were more likely to have dwindling effects.

In the comparison of retention uplift (more than 90 days) the optimized strategy produced a 11.9% better retention improvement compared to the base. This supported the relevance of long-term prediction of behavior in incentive decision-making.

The structure also enhanced equitableness in the distribution of promotions. The difference in the per-user incentive distribution decreased by 22% according to which there was more balanced and rational distribution than the distribution made by heuristics.



Platform-Level Business Impact

The cohesive AI was eventually tested on the platform-based business measures by an A/B experiment that was conducted with a select of sellers. The test lasted the period of 30 days and involved the users who represented all the segments. The model-driven and personalized promotion was provided to the treatment with rule-based or non-personalized incentives to the control group.

The outcome was high user-activity, purchase- conversions and retention measures. The key results are summarised in Table 4.

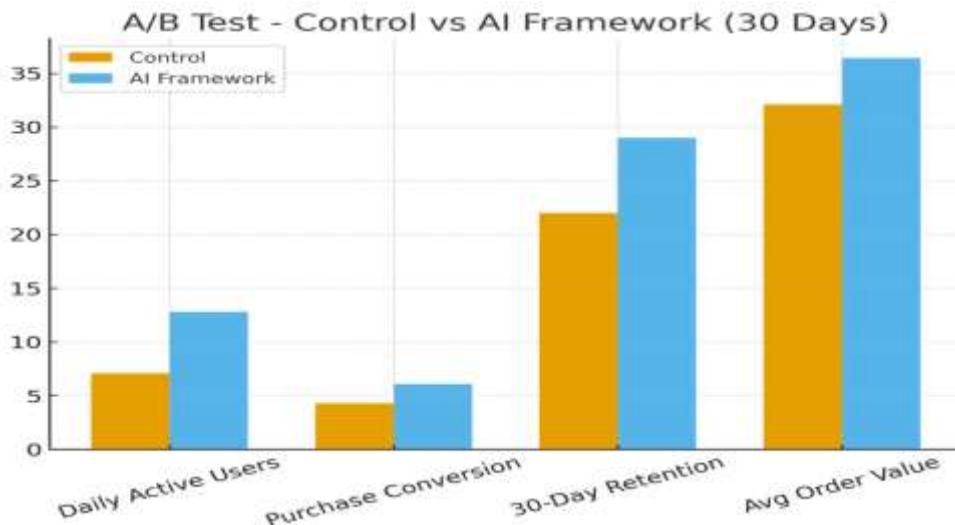
Table 4. A/B Test Platform-Level

Metric	Control Group	AI Framework Group	Relative Lift (%)
Daily Active Users	+7.1%	+12.8%	+80.3%
Purchase Conversion Rate	4.3%	6.1%	+41.8%
30-Day Retention	22%	29%	+31.8%
Average Order Value (\$)	32.1	36.4	+13.4%

The number of daily active users increased dramatically which meant that the customized promotion engine was working in encouraging the repeats. There was also an increase in purchase conversion where precision in awarding incentives to high uplift user played a major role in increasing the purchase conversion. A particular value was also attached to retention gains since it was used to depict the advantage of the CDEE-aligned framework over the long term.

Feedback, that related to the individual sellers, showed that the promotion efficiency increased significantly. Sellers that operated through AI structure had had a reduction of 14 percent in their promotional expenses with the same volume of sales, and those sellers that opted to use the high-personalization environments had experienced an increment in repeat consumers by 18 percent. The unified platform strategy meant that the small sellers who otherwise do not possess the underlying infrastructures to analyze did not need to be left behind in the same advanced decision models that the big sellers are enjoying.

The findings indicate that the intelligent promotion and retention engine achieves extensive returns in predictive accuracy, optimization of ROI and platform-level business impact. Based on the quantitative evidence, segmentation, uplift modeling, and optimization implemented in one AI-based decision architecture are effective.



V. CONCLUSION

The results confirm the possibility of an AI-based platform system to significantly increase the efficacy of promotions, customer retention, and business impacts in general. Segmentation

was found to show defined behavioral variations that influenced incentive reactions whereas predictive models particularly Random Forest precisely classified potential purchasers and predictable churners and high retention users.

The best ROI was always associated with optimization as opposed to other campaigns, and the best returns were during budget limitations. A/B tests on the platform further demonstrated a tremendous growth in terms of activity, conversions, and retention. These results underscore the importance of using segmentation, uplift modeling and optimization. The paper comes up with the conclusion that customized and intelligent incentives are better than blanket approaches and produce quantifiable rewards among the participants and the sellers.

REFERENCES

- [1] Shen, Y., Wang, Y., Lu, X., Qi, F., Yan, J., Mu, Y., Yang, Y., Peng, Y., & Gu, J. (2021). A framework for massive scale personalized promotion. arXiv (Cornell University). <https://doi.org/10.48550/arxiv.2108.12100>
- [2] Yang, J., Li, Y., & Jobson, D. (2022). Personalized promotion decision making based on direct and enduring effect predictions. arXiv (Cornell University). <https://doi.org/10.48550/arxiv.2207.14798>
- [3] Kasem, M. S., Hamada, M., & Taj-Eddin, I. (2023). Customer profiling, segmentation, and sales prediction using AI in direct marketing. *Neural Computing and Applications*, 36(9), 4995–5005. <https://doi.org/10.1007/s00521-023-09339-6>
- [4] Shi, F., Yue, C., & Han, C. (2022). A lightweight and efficient neural network for modulation recognition. *Digital Signal Processing*, 123, 103444. <https://doi.org/10.1016/j.dsp.2022.103444>
- [5] Bodunov, O., Schmidt, F., Martin, A., Brito, A., & Fetzer, C. (2018). Grand Challenge: Real-time destination and ETA prediction for maritime traffic. arXiv (Cornell University). <https://doi.org/10.48550/arxiv.1810.05567>
- [6] Kim, H. Y., Song, J. H., & Lee, J. (2019). When are personalized promotions effective? The role of consumer control. *International Journal of Advertising*, 38(4), 628–647. <https://doi.org/10.1080/02650487.2019.1593721>
- [7] Mitta, N. R. (2020, December 7). Optimizing Marketing Strategies in E-Commerce with AI: Techniques for Predictive Analytics, Customer Segmentation, and Campaign Optimization. <https://www.dlabi.org/index.php/journal/article/view/202>
- [8] Xie, S., Liu, L., Sun, G., Pan, B., Lang, L., & Guo, P. (2023). Enhanced e-commerce fraud prediction based on a convolutional neural network model. *Computers, Materials & Continua/Computers, Materials & Continua (Print)*, 75(1), 1107–1117. <https://doi.org/10.32604/cmc.2023.034917>
- [9] Hallikainen, H., Luongo, M., Dhir, A., & Laukkanen, T. (2022). Consequences of personalized product recommendations and price promotions in online grocery shopping. *Journal of Retailing and Consumer Services*, 69, 103088. <https://doi.org/10.1016/j.jretconser.2022.103088>
- [10] Huang, M., & Rust, R. T. (2020). A strategic framework for artificial intelligence in marketing. *Journal of the Academy of Marketing Science*, 49(1), 30–50. <https://doi.org/10.1007/s11747-020-00749-9>