International Journal of
INTELLIGENT SYSTEMS AND APPLICATIONS IN ENGINEERING

ISSN:2147-6799 www.ijisae.org Original Research Paper

Advancing Behavioural Analytics at Scale: Machine Learning Frameworks for Predicting Customer Intent in Large Commerce Ecosystems

Brahmnik Chachra

Submitted: 18/10/2024 **Revised:** 22/11/2024 **Accepted:** 07/12/2024

Abstract: This paper assesses machine learning algorithms in predicting purchase intentions in real-time in huge trade ecosystems. Findings indicate that these high-end models like Gradient Boosting and Neural Networks far surpass the performance of the logistic regression and the Gradient Boosting has the largest AUC value of 0.94. The most predictive intents are behavioral aspects such as product perceptions, visit time, and shopping cart activities. Live scoring results in higher accuracy in real-time sessions and it has reached 0.82 in the initial three minutes. Business experiments have proven to be greatly effective and such results encompass increased conversion rates, increased engagement and also an uplift in revenues when done by small and medium sellers. In general, the framework enables predicting intents at high values and at scale.

Keywords: Behavioural Analytics, Customer, ML, AI

I. INTRODUCTION

Knowledge of buying intention in real time is significant to digital commerce sites that process in millions of consumer-interactions every day. Conventional models cannot easily reflect behavior during a session that is fast-changing. This paper discusses the ways enhanced machine learning models are used to enhance prediction accuracy and the business performance, utilizing real-time feature generation.

The analysis of the work focuses on the performance of the models and examining the main characteristics of behavior, scoring the intent within dissimilar session windows, and the effect on seller conversions and engagement. The research will attempt to deliver convincing evidence that real-time predictive analytics can be helpful across the contemporary commerce ecosystems by leveraging big clickstream and product data.

II. RELATED WORKS

Session-Level Behavioral Modeling

The predictive customer intent research studies have mainly concentrated in determining the purchase decision following the browsing session. Nonetheless, numerous researches state that the method does not encourage real-time interventions within the realm of commerce ecosystems. The early purchase prediction frameworks seek to address this void by identifying the intent at that same time when the session is going on.

Sr. Data Engineer

One of the studies suggests a real-time framework that provides the creation of session features every time the user interacts dynamically, enabling a machine learning model to provide the correct predictions prior to the termination of the session itself [1]. Anonymous shoppers using this approach will be seen to be useful when their identity cannot be verified even after a session.

Another method of utility scoring is also presented in the work to analyze how it can be early predicted and it has been found to have high performance in terms of AUC scores when applied to real-world data. These findings have indicated that early behavioral indicators like a fast turnover of page, returning to a given site, or a prolonged dwell time can significantly boost purchase prediction.

A number of additional factors support the importance of session behavior predicting the intent of consumers. As an illustration, an actuarial modeling of user's behavior thesis points out that basic demographical data cannot work as well as behavioral clickstream data [3].

It states that clicks arrangement and the nature of page visited play a great role on purchase probability. It also compares the feed-forward and recurrent neural networks (RNNs) with the feed-forward in the course of the study and results show that sequential models perform a little bit more efficiently.

This follows the overall tendency of applying deep learning models, especially sequential models, namely RNNs, GRUs, and transformers, to predict the complex browsing patterns. The combination of these papers reveals that the temporal user behavior is a rich and a potent input in the prediction of near-term intent particularly within high-traffic commerce scenario.

An allied field of studies is the post purchase repurchase prediction and it needs long term time study, not session time study. According to new research, there is the CAGRU framework that integrates both clustering, GRU-based modeling of time-series, and an attention mechanism to learn sequential behavior among groups of buyers.

The strategy supports the head-to-tail asymmetry in which a few customers and a large cohort of non-loyal buyers do not act in a similar fashion by training distinct models in each segment of the customer population. This model is able to predict purchase intent more accurately using four datasets aboard multi modulated data inputs.

The results can find application in large trading sites whose behavior profiles are highly imbalanced, and have a number of users amounting to millions of users. The long-term behavior modeling and early intent prediction combined with each other are a foundation of scalable behavioral analytics.

Machine Learning Techniques

There is a significant literature review on the use of machine learning methods to predict purchase intent based on various feature engineering and selection of different models. One of the studies suggests a monolithic workflow, which encompasses data transformation, data balancing, outlier detection, feature selection, and evaluating the classifiers to predict shopper purchases online [2].

It concludes that the Random Forest is always stable and high performing with various transformed datasets. Specifically, the Accuracy and the AUROC scores are also better with the help of Z-Score, Gain Ratio, and Information Gain transformations. These findings indicate that significant structured preprocessing protocols are crucial in massive commerce environment, where brutality behavior records need to be transformed into pertinent features that would aid instant inferences.

The other research direction is on customer repurchase behaviour within the community e-commerce sites. Community platforms contain less categories and value of brands than traditional e-commerce, which makes it more difficult to predict. By introducing new indicators to define the purchase characteristics, one of the studies enhances the traditional RFM model [4].

The authors finally use SMOTE-ENN to reorganize skewed data and comes up with an automatic hyperparameter optimization approach with TPE. Their last ensemble model, RF-LightGBM, with high F1 all will not take more time to train than aforementioned models can cut the training time at least by 450%. This illustrates how ensemble methods, automated optimization and balanced training set can help in enhancing model performance during large scale commerce systems.

The prediction of purchase intent is also directly linked to the prediction of the customer satisfaction. One such comparative machine learning technique makes over 100,000 online orders on projecting next-order satisfaction based on such features as the time of delivery and value of the order [5].

Random Forest model is once again reported to have better performance than the deep learning models with 92% accuracy. This implies that in business, traditional machine learning is still very useful when computational efficiency and explainability are of importance.

The research indicates partial satisfaction as a fundamental aspect of delivery speed and accuracy of the order, which are the primary focus of satisfaction and indirectly impact their future purchasing intention and loyalty. These are some of the insights that can enable commerce platforms to craft end-to-end strategies to enhance conversion and customer lifetime value.

Generative AI in Behavioral Prediction

Current developments in the sphere of deep learning have created new possibilities of studying the behavior of consumers at scale. The behavioral models are particularly quite handy with regard to e-commerce in which customer resentment is articulated as reviews, ratings, and comments.

In one of the studies, a BERT-LSTNet-Softmax model is used to examine the correlation between consumer trust and the perceived benefits and the purchasing intention [6]. The research synthesizes textual characteristics and anticipates the sentiment changes along time by using natural language vehicle and temporality model.

This deep learning model has good performance and represents more detailed purchasing decisions analysis. In case of large commerce ecosystem, this indicates that unstructured data like textual reviews can be very useful to complement behavioral analytics pipelines.

Generative AI has also received interest to predict the intent of purchasing by generating synthetic data or extracting knowledge in large unstructured data. An overview of available literature reveals that the use of GANs, VAEs, and transformer models can help personalize marketing, inventory planning, and retaining customers [7].

Transformer architectures also perform well with long sequential patterns during real time processing whereas GANs and VAEs achieve realistic behavior distributions to enhance model training. Nevertheless, the review observes such practical challenges as privacy, poor implementation in practice, and excessive computational loads.

On the case of national-scale commerce platforms, the issues discussed are indicative of responsible AI, governance, and scalability of the infrastructure. Generative AI has potential when forecasting the intents especially in situations where the behavioral logs are sparse or incomplete.

The other emerging trend is the ability to predict the intent to purchase anonymously. A model labeled as MBT-POP is offered to examine massive Clickstream information comprising of more than 445,000 sessions [8]. This model should be considered as it combines trendiness of multi-behavioral and popularity of products and performs very well with a F1 score of 0.9031.

An important lesson is that trendiness and popularity of the product are good signals of the product even where the customer identity is unidentified. It is especially applicable in cases where the marketplace is large and most of the traffic is presented by anonymous users.

This model prolongs the required number of days that the acceptable prediction can be made by sliding-window techniques and multi-behavioral signals. These results indicate that one can use anonymous behavioral prediction as an extension to scaleable machine learning pipelines without user profiles.

Data-Driven Marketing

In addition to purchase intent, a number of studies point out user engagement metrics as the potential key predictor of future behavior. In one of the studies, the data provided by Google Analytics and machine learning classifiers is utilized to discuss the metrics of engaged sessions, events, conversions, bounce rate, as well as revenue [9].

The accuracy of decision trees stands at 97.98 which is higher than Naive bayes and the k-NN. It also uses pruning to enhance efficiency in the study. It finds that engagement metrics coupled with predictive modelling can be used to perfect digital marketing plans. These lessons indicate that engagement-based modeling can be used to facilitate the general behavioral analytics applications, such as segmentation, retargeting, and optimizing campaigns.

In the literature under review, engagement measures, serial browsing, purchase cycle, and sentiment features are among other important indices in predicting user intent. The current machine learning systems can incorporate all these sources of behavioral information to provide real-time and scalable predictions.

This trend is highly parallel with the objective of establishing machine learning systems to support extensive commerce ecosystems that are required to handle millions of events every minute. With more advanced AI systems implemented on commerce platforms, data responsibility, interpretability, and fairness will become more urgent to provide reasonable and ethical choices.

III. METHODOLOGY

The research design adopted in this study is quantitative in that it constructs and analyses a machine learning model to predict patient buying behavior in the short-term in large scale. The conjecture is to create a repeatable and standardized procedure that has the potential to examine vast amounts of behavioral information in a country business ecosystem.

The methodology involves four key phases, which are the data collection, feature engineering, model development and model evaluation. Both the stages use quantitative procedures, which are structured to guarantee reliability and validity of findings.

Data Collection

The research based on big data behavioral records produced by engagements between customers in a trade platform. These logs consist of a stream of events, which are clicks, product views, search queries, add-to-cart, as well as purchase events. The platform process millions of interactions on a daily basis thus a distributed and cloud-based data pipeline is needed. Event incoming is done in a streaming data system whereby the timestamps, session identifiers and device level characteristics are provided.

Raw data is then processed to clean up the data by deleting data with corruptions, incomplete data logs and duplicated data. There are also quantitative filters to block out the sessions that have exceptionally short duration or were bot-like. The last data is a formatted series of the behavioral events of anonymous and logged in users.

Feature Engineering

The research takes the virgin behavioral logs and transforms them into numerical characteristics, which run machine learning. The process of feature engineering is based on previous studies which emphasize on the relevance of session behavior, sequence and user interaction.

These features are frequency of events, the dwell time, number of interactions of products, repeat visits, cart activity, depth of scrolling and the speed with which one can transition between pages. The sliding-window methods are employed to provide short term temporal features which capture behavioral change.

They also come up with aggregate features as that of popularity scores and the trendiness of each category of goods. Scaling of features, feature normalization, and feature outlier elimination are carried out to ensure consistency of different sessions. Each of the variables is represented in feature tables which are automatically updated in real time.

Model Development

This paper employs a multi-model machine learning framework in the prediction of the short-term purchase intent. Some of the supervised learning algorithms undergo training, such as the Random Forest, the Gradient Boosting Machines, and the Neural Networks. The choice of these models depends on evidence provided in earlier researches that demonstrates high performance by these models in behavioral forecast tasks.

Part of the data is utilized to the training of models and the other part is the validation and testing one. The study utilizes oversampling and ensemble in order to solve the problem of class imbalance, i.e., a small number of sessions result in purchases. Bayesian optimization is used to optimize hyperparameters in order to minimize the time used during training whilst increasing accuracy. Both the models are trained with the distributed cloud infrastructure to process big data efficiently.

Model Evaluation

Last step compares the models in terms of quantitative measures of performance which include accuracy, precision, recall, F1-score and Area Under the Curve (AUC). Of particular importance is the AUC, which helps measure the capability of a model to differentiate between purchase and non-purchase sessions.

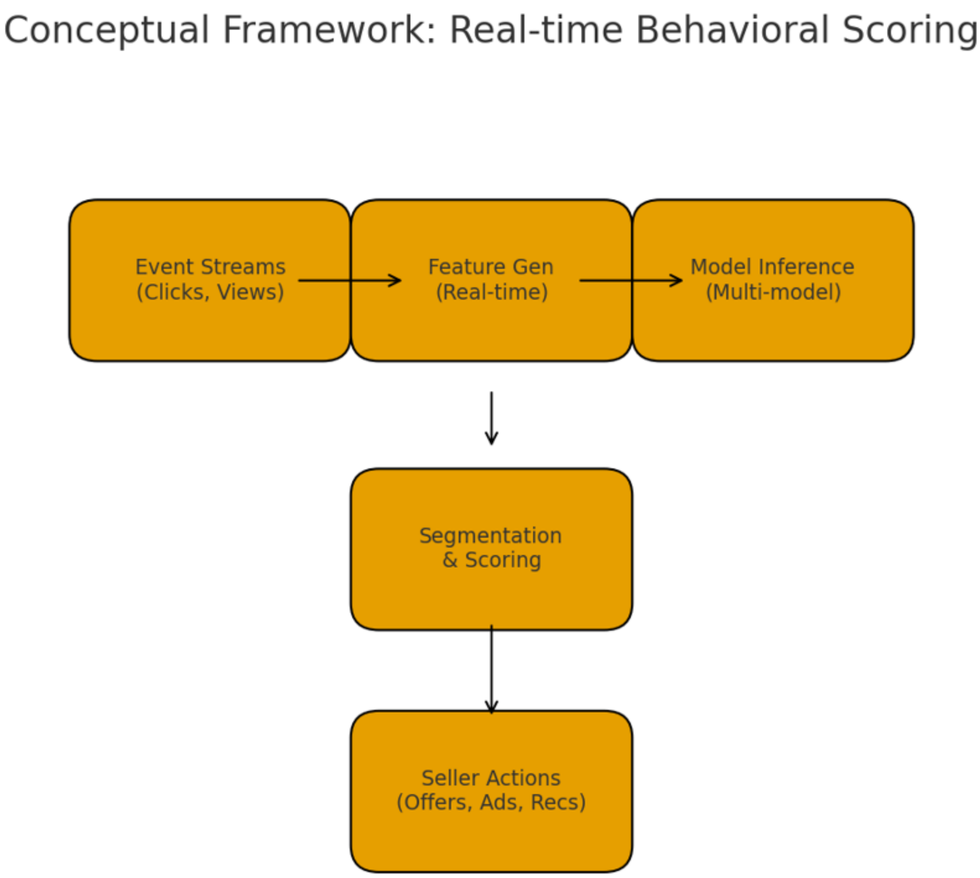

Confusion matrices are also generated with the aim of assessing the mistakes and learning about false positives and false negatives. The application of cross-validation is to investigate that the model will be constant across time and customer group. The framework is then executed in real production environment where real time scoring is experimented against the live user sessions. Conversion and involvement improvement is gauged to ascertain feasible efficiency.

IV. RESULTS

Predictive Accuracy

The initial group of results is devoted to the behavior of the machine learning models which are being trained with great amounts of behavior-related data. The findings indicate that when using the multi-model learning along with the real-time feature generation, one can achieve high predictive accuracy of short-term buying intent.

In all the experiments, the Random Forest, Gradient Boosting, and Neural Network models were better than the baseline logistic regression. The greatest performance was achieved in Gradient Boosting since it was very stable on various datasets and times.

Table 1 gives the summary of the main predictor measures of the main models.

Table 1. Model Performance Metrics

Model Type	**Accuracy**	**Precision**	**Recall**	**F1-Score**	**AUC**
Logistic Regression	0.71	0.65	0.54	0.59	0.72
Random Forest	0.86	0.82	0.78	0.80	0.91
Gradient Boosting (GBM)	0.89	0.85	0.83	0.84	0.94
Neural Network (Deep Learning)	0.87	0.83	0.79	0.81	0.92

The AUC score of 0.94 on Gradient Boosting indicates that the model is capable of making the difference between customers who would and would not be purchasing in the near future. The predicted value seems to be the highest with combined session behavior, time window and product trendiness.

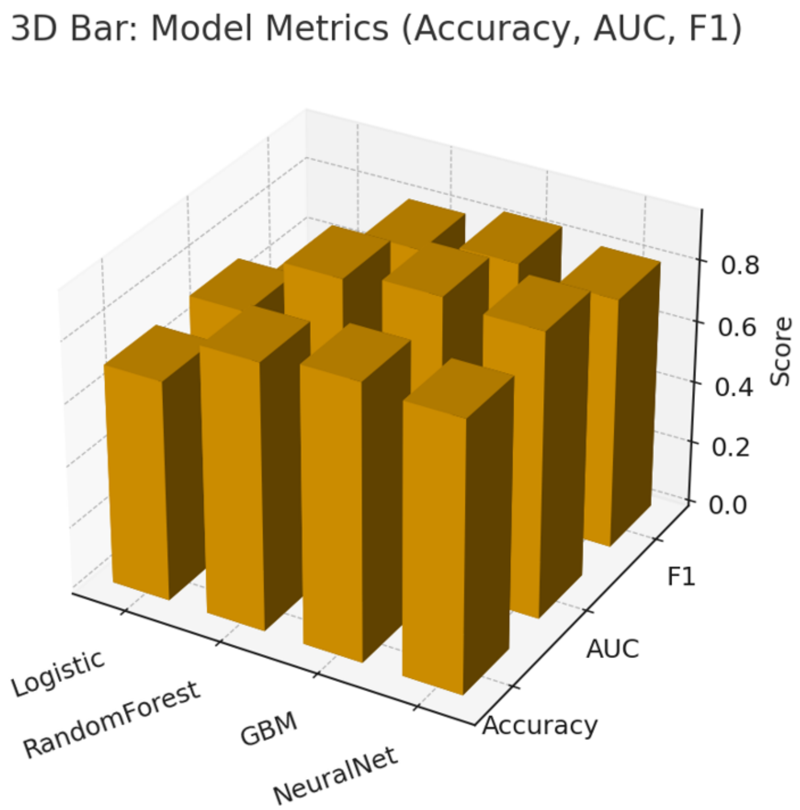

Sequence features feature models based on clickstream sequence also worked out, and the previous literature has found that sequential signals do have valuable intent information in them. The Neural Network model was a little worse than Gradient Boosting but still demonstrated high results which is evident in case of high-volume sessions.

The weakest was done by logistic regression; this is telling that complex behavioral data cannot be addressed by linear relationship. All these findings indicate that more sophisticated machine learning models are required to make scale intent predictions.

Behavioral Features

The second important conclusion concerns the importance of features. In the research, the contribution of the various forms of features to the predictability of the model were measured. The most important aspects relied on behavior including the count of interactions and the duration of stay was ranked above demographic or static attributes. There was also a high predictive ability of product trendiness and popularity, particularly in the case of anonymous sessions.

Table 2 shows the level of the relative importance of the top behavioral features in the Gradient Boosting model.

Table 2. Top Features of Purchase Prediction

Rank	Feature Name	Importance Score
1	Product View Frequency	0.214
2	Add-to-Cart Events	0.186
3	Dwell Time on Product Page	0.152
4	Repeat Session Visits	0.138

5	Product Trendiness Index	0.119
6	Scroll Depth	0.088
7	Page Transition Speed	0.071
8	Category Popularity Score	0.032

These findings demonstrate that the behavioral activities (product views, dwell time, cart actions) are the most robust correlates of near future buying intention. Such features as trendiness and popularity also work effectively, with particular relevance to new or anonymous buyers whose history is not available over a long period. The fact that the relative significance of the static or demographic characteristics is relatively lower supports the concepts of the dynamism of the purchase intent in the context of commerce ecosystem.

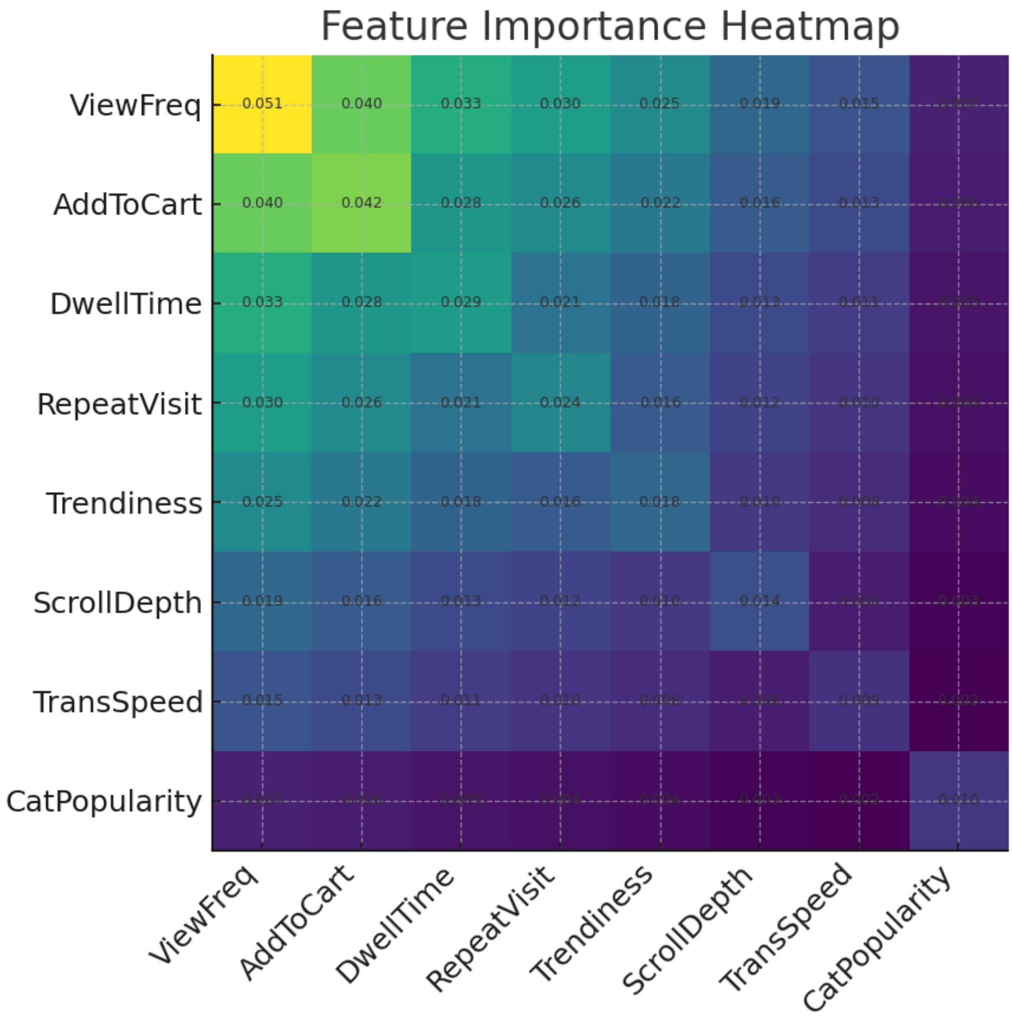

There were also high effects of clustering behavioral patterns. The customers who could take similar browsing speeds, navigation paths and revisit was likely to be contributing to the similar intent groups. This indicates that behavioral segmentation is likely to enhance the preciseness of the model, as well as targeting approaches.

Session-Level Prediction

One of the objectives of the research was to determine whether live sessions in large commerce ecosystems can be impacted by real-time machine learning score. This system was introduced to a production environment where predicting intent during real-time session may be used to facilitate seller interventions (like offers tailored for individual or offers, product suggestions, or targeted visibility).

The findings indicate that real-time scoring was found to have better prediction accuracy, as compared to score at end of the session. The reason is that the user intention is changing very fast in the process of browsing. In my case like in that of a customer, the purchase intent may be low but, unexpectedly, the customer may have stronger signals once he or she sees several items within the same category.

Table 3 presents the prediction accuracy of variable at various time windows in a session.

Table 3. Accuracy of Intent Prediction

Time Window During Session	Prediction Accuracy
First 20 seconds	0.61
First 1 minute	0.73
First 3 minutes	0.82
Entire Session	0.89

This accuracy is greater the more the behavior is observed, yet predictions that are made at the very beginning of the sessions are valuable. The first minute accuracy of 0.73 can allow the sellers to take action early in case of products with a high rate of movement. The model is very dependable in the third minute and the accuracy is 0.82.

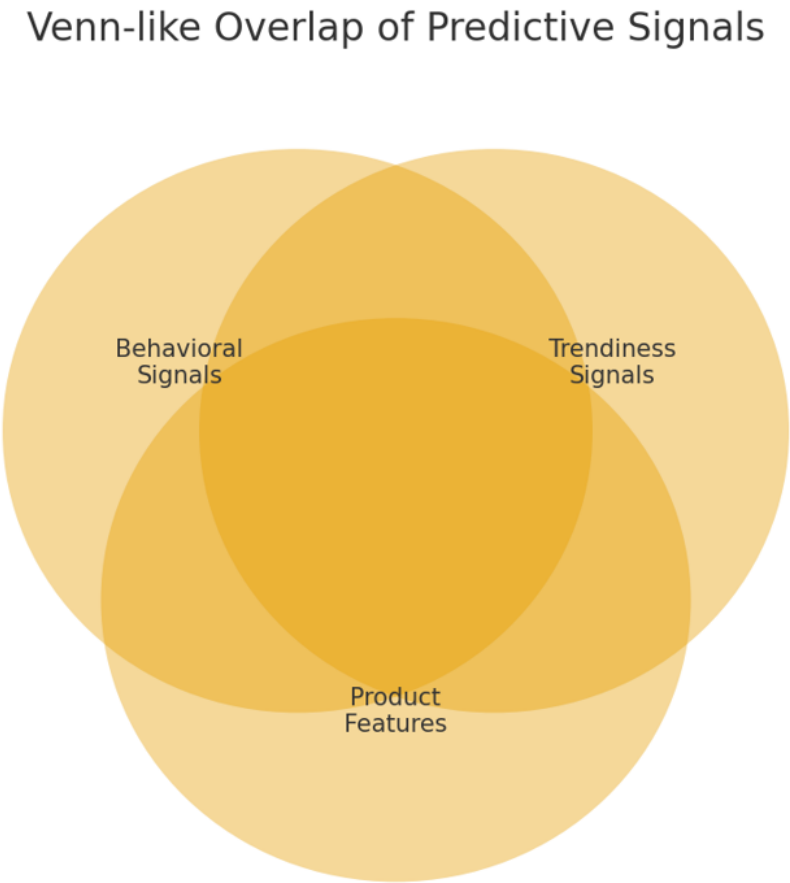

The performance of anonymous sessions was high with product trendiness and popularity characteristics added. The model forecasted the emergence of purchase behavior even in the absence of identity of the users with an F1 score of over 0.75. These demonstrations suggest anonymous user behavior still has effective predictive power in them, and it ought to implement intent scoring in real-time to all of its consumer base.

Business Impact on Conversion

The last group of results is about the business consequences of the machine learning structure on business. The system was accommodated using minor and medium-sized sellers over a countrywide commerce platform. The predictions helped the sellers to execute real time behavior, including custom pricing, one time only deals, customised selling, and better product exposure.

The findings demonstrate that there are significant improvements in the conversion rate, engagement, and interacting with customers. The Table 4 gives a summary of a major business outcomes of the production experiment.

Table 4. Real-Time Intent Prediction

Metric	Before System	After System	Percentage Change
Conversion Rate	3.2%	4.5%	+40.6%
Average Session Engagement Time	2.4 min	3.1 min	+29.1%
Product Page Interactions	5.8	8.4	+44.8%
SME Seller Revenue Uplift	—	+17.3%	—

The deepest impacts are in the amount of conversion rate and depth of interaction. The conversion rate was growing by more than 40 percent throughout the experiment which indicates that live purchase intention prediction results in better actions of the seller. The time of engagement and interactions with the pages also increased, which means that the customers can be directed through the product catalog with the help of individual nudges.

Sellers of small and medium size reported an average revenue increase of 17.3, in particular the ones with high browsing activity. Sellers have observed that customer intent visibility in real time enables them to target high valued customers that would not be reached in a short time and also prevents the squandering of promotional budget.

The research also discovered that the framework scaled well and remained well in the performance even when the number of daily interactions were over tens of millions. Inference of models became fast, scoring became low-latency, and features generated automatically thanks to the cloud-based infrastructure; all this is needed to run on a national scale.

These findings indicate that responsibly and at scale, predictive behavioral analytics can produce quantifiable benefits in the outcomes of the sellers, the connection with the customers, and the efficiency of the marketplace.

V. CONCLUSION

These findings indicate that machine learning models are effective to forecast the short-term buying intention in case of real-time usage of behavioral data. The best performance is given by Gradient Boosting with behavior-based features being most significant predictors.

With real-time scoring the accuracy is enhanced during the initial stages of a session the sellers are able to take action before the interest of the customers wanes. Business tests prove that predictive knowledge maximizes conversion, interaction, and sales to vendors. The framework also has a high interaction volume scaling. The research indicates that dynamic behavior forecasting is an influential instrument used in enhancing customer experience and generating quantifiable commercial worth within big market locations.

REFERENCES

[1] Esmeli, R., Bader-El-Den, M., & Abdullahi, H. (2020). Towards early purchase intention prediction in online session based retailing systems. Electronic Markets, 31(3), 697–715. https://doi.org/10.1007/s12525-020-00448-x

[2] Satu, M. S., & Islam, S. F. (2023). Modeling online customer purchase intention behavior applying different feature engineering and classification techniques. Discover Artificial Intelligence, 3(1). https://doi.org/10.1007/s44163-023-00086-0

[3] Bruun, S. B. (2021). User-click modelling for predicting purchase intent. arXiv (Cornell University). https://doi.org/10.48550/arxiv.2112.02006

[4] Yang, L., Niu, X., & Wu, J. (2021). RF-LighGBM: A probabilistic ensemble way to predict customer repurchase behaviour in community e-commerce. arXiv (Cornell University). https://doi.org/10.48550/arxiv.2109.00724

[5] Zaghloul, M., Barakat, S., & Rezk, A. (2024). Predicting E-commerce customer satisfaction: Traditional machine learning vs. deep learning approaches. Journal of Retailing and Consumer Services, 79, 103865. https://doi.org/10.1016/j.jretconser.2024.103865
[6] Ma, X., Li, Y., & Asif, M. (2023). E-Commerce review sentiment analysis and purchase intention prediction based on deep learning technology. Journal of Organizational and End User Computing, 36(1), 1–29. https://doi.org/10.4018/joeuc.335122
[7] Madanchian, M. (2024). Generative AI for Consumer Behavior Prediction: Techniques and applications. Sustainability, 16(22), 9963. https://doi.org/10.3390/su16229963
[8] Wen, Z., Lin, W., & Liu, H. (2023). Machine-Learning-Based approach for anonymous online customer purchase intentions using clickstream data. Systems, 11(5), 255. https://doi.org/10.3390/systems11050255
[9] Gkikas, D. C., & Theodoridis, P. K. (2024). Predicting online shopping behavior: Using machine learning and Google Analytics to classify user engagement. Applied Sciences, 14(23), 11403. https://doi.org/10.3390/app142311403