Counterfactual Customer Churn Prediction in E-Commerce Memberships

Steffeno Selva S[1] and Syed Imran U[2]

[1]Department of Artificial Intelligence and Data Science, St Joseph's Institute of Technology ,Chennai, Tamil Nadu
[2]Department of Artificial Intelligence and Data Science, St Joseph's Institute of Technology ,Chennai, Tamil Nadu

Abstract. The e-commerce membership platforms that provide membership services through subscription-based models are currently experiencing a serious situation of customer churn. This paper suggests the use of a Counterfactual E-Commerce Churn Prediction (CECP) model, which will combine the strengths of machine learning, explainable AI, and causal inference techniques to predict and prevent customer churn. The model will employ XGBoost for predicting customer churn, SHAP for explaining the model, and DiCE for counterfactual explanation. In addition, the T-Learner uplift modeling approach will be employed to determine the individual treatment effects and customers who can benefit from customer retention strategies in the most optimal way. The experiments performed on a large e-commerce membership dataset suggest that the proposed approach will yield an accuracy of 89% and area under the curve of 0.93, and the simulated intervention strategies will actually decrease customer churn by 23% using simulated intervention strategies. The findings suggest the efficacy of using predictive models in conjunction with a counterfactual reasoning model to develop data-driven and personalized customer retention strategies.

Keywords -Explainable AI, SHAP, XGBoost, Uplift Modelling, E-Commerce Memberships, DiCE, Retention Strategy, Customer Churn Prediction, Counterfactual Analysis, Causal Inference.

1 Introduction

In the ever-changing online shopping environment, the membership subscription system has emerged as a major tool for improving customer engagement and ensuring regular revenue streams. Although the membership system has several advantages, including personalized customer service, loyalty programs, and special discounts, the retention of subscribed customers is a major concern. Customers usually cancel their membership subscriptions due to decreased engagement, price sensitivity, and the lack of perceived long-term value, making customer retention a major concern for online shopping platforms.

Customer churn, also known as the end of a customer's subscription or membership, is a major threat to the financial viability of membership-based businesses. Previous studies have clearly shown that the cost of acquiring new customers. is significantly higher than retaining existing customers, with a ratio of five to one or even higher [1].

Conventional methods for churn prediction are mainly concerned with finding customers who are likely to end their subscription by employing statistical and machine learning models based on demographic, transactional, and behavioral data [1], [2], [12], [16]. The methods commonly used for churn prediction are logistic regression, decision trees, random forests, and gradient boosting. Although these models are able to produce high accuracy, they are not very helpful in understanding how an organization should act to prevent churn for individual customers.

To overcome this problem, there has been a growing interest in recent years in the integration of causal inference methods into churn prediction models. Unlike conventional models, causal inference models attempt to measure the impact of certain interventions on customer behavior by considering the outcome of hypothetical situations under different scenarios. Uplift modeling is one such technique that allows the estimation of individual treatment effects by comparing the outcomes of customers who are treated and those who are not treated, thus helping to identify customers who are most likely to benefit from targeted customer retention efforts [11].

Concurrently, the application of counterfactual analysis has also proven to be an effective means of providing interpretable and actionable insights by determining the minimal modifications necessary to change the prediction made by a model. By creating "what-if" scenarios, counterfactual explanations enable organizations to comprehend why a customer is predicted to churn and what actions can

https://doi.org/10.1051/itmconf/20268503001

be taken to potentially avoid such an eventuality [3], [7], [9]. When used in conjunction with other explainable artificial intelligence approaches, such as SHAP, counterfactual analysis can improve the transparency of models while facilitating decision-making.

Inspired by the above-mentioned advancements, this paper presents a Counterfactual E-Commerce Churn Prediction (CECP) framework that combines the concepts of churn prediction, explainability, and causality into a single framework. The proposed framework utilizes XGBoost for churn prediction, SHAP for feature-level interpretability, DiCE for counterfactual recommendation generation, and uplift modeling via the T-Learner approach for estimating individual treatment effects. By translating the churn prediction problem into an actionable decision-support framework, the proposed solution enables e-commerce platforms to develop personalized, cost-effective, and datadriven customer retention strategies.

Despite the increasing use of machine learning and causal inference methods for churn prediction, there are some challenges that have not been addressed in the current literature. Many methods are either focused on accuracy and lack interpretability or are focused on interpretability and lack actionable advice for intervention. Additionally, models that provide estimates of treatment effects are often designed as isolated analytical tools, which makes it difficult to integrate them into decision-making pipelines. This is an area where an end-to-end solution is required that not only provides high accuracy for churn prediction but also provides insights into the underlying causes of churn and the effects of potential intervention actions at the individual level.

2 Literature Survey

Customer churn prediction has been a widely researched topic in various sectors, such as telecommunication, retail, and online digital services. In the early stages of research, statistical models like logistic regression and survival analysis were used to predict the probability of churn within a fixed time duration [1]. Although these models were interpretable, they were not very effective in capturing the complexity of customer behavior in a large-scale e-commerce setting due to their linear nature.

With the development of machine learning algorithms, some researchers used classification algorithms like decision trees, random forests, support vector machines, and ensemble learning algorithms to enhance the accuracy of customer churn prediction [2], [12], [16]. These algorithms used demographic data, purchase history, and user engagement features to model nonlinear relationships among customer variables. Although these models greatly improved the accuracy of churn prediction, they mostly targeted the identification of customers who were likely to churn without considering the impact of targeted interventions on customer retention outcomes.

To address this issue, there has been a growing trend in recent studies to integrate causal inference models with churn modeling. Unlike predictive models based on correlation analysis, causal inference models seek to estimate the change in customer behavior under certain interventions. Uplift modeling, also known as incremental or true lift modeling, estimates the treatment effect by comparing the outcomes between the treated and untreated groups [11]. This allows companies to target customers who are most likely to benefit from customer retention efforts, thus optimizing the effectiveness of marketing campaigns.

Concurrently, counterfactual inference approaches have emerged as promising tools for explaining model predictions by hypothesizing alternative outcomes. Counterfactual regression, causal forests, and representation learning have been proposed for modeling individual treatment effects and alternative outcomes under different conditions [3], [7], [9], [14]. Nevertheless, most of the existing literature focuses on these approaches in a standalone manner and has not yet explored the integration of counterfactual explanations with churn prediction models.

In the context of e-commerce, the literature on churn prediction has been largely focused on subscription-based services like telecommunication and online media streaming services [8], [15]. Compared to the above-mentioned areas, relatively fewer studies have investigated the membership models in the context of e-commerce, which are increasingly being adopted and gaining significance. Furthermore, the joint use of uplift modeling and counterfactual explanations for churn prevention in e-commerce has remained an uncharted area.

This research work aims to fill the aforementioned research gap by combining the techniques of churn prediction, uplift modeling, and counterfactual explanations in a unified framework specifically designed for membership platforms in the context of e-commerce.

More recently, the development of explainable artificial intelligence (XAI) has also impacted the field of churn prediction, as greater transparency and trustworthiness of models have become important considerations. SHAP and LIME have become more popular tools in the field for explaining blackbox models and understanding the drivers of churn [5], [10], [17]. Although these tools are very useful for explaining the predictions of models, they are more focused on explaining predictions than on providing recommendations for action. This means that the person making the decision is left with no clear guidance on how to change customer behavior to improve retention. The integration of explainability techniques with causal and counterfactual approaches is one promising way to close this gap. By combining explanations at the feature level with individual treatment effect estimation, models of churn prediction can move from being purely diagnostic to prescriptive decision support tools.

3 Methodology

A novel method has also been proposed with the abbreviation Counterfactual E-Commerce Churn Prediction (CECP), specifically for e-commerce. In addition to this, this method is a comprehensive method since it is not just applied in customer churn prediction but also in explaining customer churn prediction and providing suggestions on how customers can be retained. This method comprises four primary modules. The first one is data preprocessing, followed by churn prediction using the XG Boost model. The third one is explanation of customer churn using the SHAP technique. The fourth one is generating suggestions using the DiCE library.

3.1 Data Preprocessing

The experimental evaluation was carried out on an anonymized dataset collected from a real-world e-commerce membership platform. The dataset consists of more than 500,000 customer records, including demographic information, transactional behavior, subscription information, and engagement features like loyalty points and coupon redemption. All personal identifiable information was anonymized before analysis to maintain data privacy.Table I below summarizes the various categories of features used in the proposed framework, as well as the respective preprocessing methods applied to the features. One-hot encoding was applied to the categorical demographic and subscription features, while z-score scaling was applied to the continuous transactional features. The table also indicates the relevance of each category of features in the prediction of churn.

Table 1: Feature Categories Used in Churn Prediction

Type of Features	Example Features	Preprocessing Applied	Notes
Demographics	Age, Region	One-Hot Encoding, Normalization	Filled Missing Values
Transactional	Avg. Order Value, Purchase Frequency	Z-score Scaling	High impact on churn signal
Participation	Used Coupons and Loyalty Points	One-Hot Encoding	Used in SHAP and DiCE modules
Subscription	Plan Type, Tenure	One-Hot Encoding	Strong correlation with churn
Goal Variable	Not Churned or Churned	Binary Encoding	Utilised for XGBoost classification

Figure 1 depicts the data preprocessing pipeline of the CECP framework, which is used to process the raw customer data into a structured format for training the model.

3.2 Churn Prediction Using XG Boost

The XGBoost algorithm [4], a Gradient Boosting model, which has been widely acclaimed for its strength and ability to deliver accurate results for structured data, acts as the predictor engine for CECP. Designed to focus upon past actions and behaviors, it has been developed to forecast precisely the retention or churn status of potential or existing clients to become or remain part of their clientele. Regarding this process of optimization, a grid search approach was used to identify optimal hyperparameters. The final set of

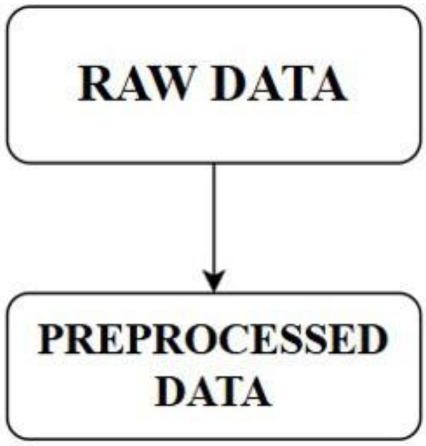

Figure 1: Flow diagram of raw data pre-processing to pre-processed data in the CECP.

hyperparameters consisted of a learning rate of 0.05, a maximum tree depth of 6, and 200 estimators, with greater positive class weights to compensate for churn cases. To counter the natural class imbalance problem in customer churn data, a class-weighted loss function was used during the training of the model, giving greater weight to churn cases. The accuracy, precision, recall, F1 measure, and area under the ROC curve (AUC) were used to assess the performance of the model.

Table 2: Evaluation Metrics of the Proposed Model

Metric	Value
Accuracy	89%
Precision	83%
Recall	78%
F1 Score	80.4%
AUC (ROC)	93%
Optimal Hyperparameters	Learning Rate: 0.05, Maximum Depth: 6

Table II above shows the evaluation metrics of the churn prediction model developed using the XGBoost algorithm. The model has a high accuracy of 89% and an AUC of 0.93, which clearly shows the model's effectiveness in distinguishing between churned and non-churned customers.

3.3 Explainability using SHAP

Adding the explanations with the use of SHAP was implemented to enhance the transparency mechanism for the results of the machine learning model. The approach of SHAP [5], [10] was implemented by attempting to explain the value of each attribute to arrive at a specific prediction value. Based on the global SHAP explanation, it was understood that the lack of subscription plan downgrade, the lack of loyalty points, and the lack of purchases had contributed the most to the quit call for all the customers. For instance, a person who was vulnerable to churning with a churn value of 0.83 would be who was deemed a high-risk case given the significant reduction in symptoms to the extent of his her engagement during the last three months. All these would be displayed in an pictorial form,

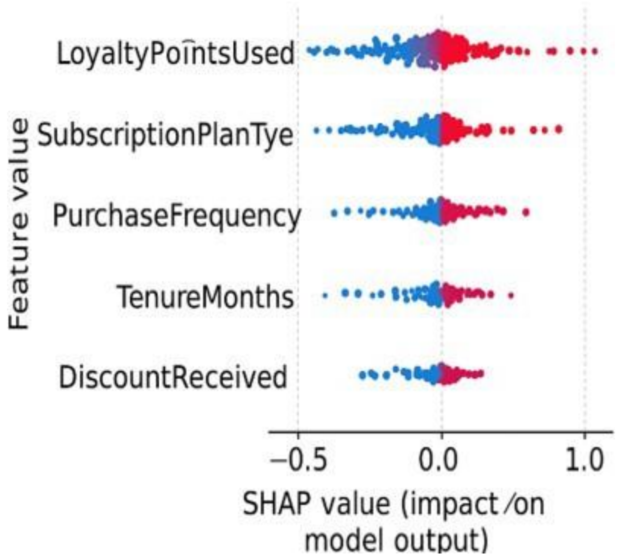

Figure 2: An example of a SHAP plot, highlighting important features and the contribution they make to predicting churn, with colors representing feature values.

also known as the waterfall chart, of the attribute value on the basis of the use of SHAP. The analysis of the segments would To make this possible, perspective on attribute priority can be employed.

Figure 2 depicts the SHAP summary plot that identifies the most dominant features used in predicting customer churn.

3.4 DiCE for Counterfactual Generation

The last module of the program: the creation of “what if” scenarios for the high risk customers, which is implemented by means of the application of the DiCE-the Diverse Counterfactual Explnations library, assuming minimal and usable changes modifications to the customer attributes would result in the expected probability of averting the churn probability.. For example, through the provision of reward bonuses or adjusting the customer to a tailored plan, the expected probability of the customer churning can be reduced from 0.78 to 0.22. The recommendations produced by the DiCE would be limited to those implemented considering the criteria set by the organization. Insights generated by DiCE will help customer retention teams provide targeted offers, improve ROI, and reduce customer attrition. The DiCE method uses a multi-objective optimization methodology to ensure the counterfactuals generated by the DiCE library meet the necessary privacy constraints-for a privacy-compliant solutionsuch as immutability of features. Thus, when applied along with SHAP, the DiCE methodology allows organizations to stop customer churn and define the reasons behind the churn decision.For instance, if the results of SHAP come out as ”Loyalty Points Used,” this DiCE methodology will then provide the organization with at least the level of change required to retain that particular customer. Thus, an organization can create data-driven and transparent retention strategies by leveraging all these techniques in concert-integrated and way beyond a simple prediction.

For the assessment of the customized strategies in terms of the retention for the individuals, the CECP model applies the uplift modeling methodology based on the T-Learner technique. This demands the construction of two distinct models using XGBoost. These models will entail the treatment group, a set of individuals who have been the recipients of some form of incentive/proposal, and the control group, a set of individuals not receiving the exposure. Both groups have a probability to churn.

In the inference, two probabilities of churn are estimated: which are done by passing the information of the customer through will be estimated through both models, namely P treat and P control. value of the treatment benefit is calculated as the difference. display as Uplift = P(control) - P(treat). Customers with the highest uplift value tend to be of the highest priority in the interventions. The current T-Learner uplift module maximizes the availability of resources and the cost of the Retention campaign using the ”persuadable” customers” [11].Figure 3 depicts the overall architecture of the proposed CECP framework, which combines the churn prediction, explanation, and counterfactual generation components.

3.5 T-Learner’s Use in Uplift Modelling

The CECP uses T-Learner method in implementing an uplift modeling strategy to measure the effectiveness of customized actions to retain individuals. This is done through building two models using XGBoost: one model to predict the treatment group (those are contactors who are exposed to a specific incentive/proposal) and another to predict the control group (those are contactors that are not exposed). Each model estimates a probable churn. In the prediction step, two probable churn values are deduced: P treat and P control. These are derived from passing a customer’s detail through each model. The potential value of the treatment is captured using a difference formula called Uplift = P control - P treat. The customers with a high uplift value are those who have interventions of highest priority. The model targets customers with high potential from interventions rather than those with a high probability of churning [11].

4 Result And Discussion

A scholar applied the data set from UL e-commerce with over 500,000 members from an e-commerce website to validate the effectiveness of the Counterfactual E-commerce Churn Prediction model. This represents the test of the real business value of the concept of Counterfactual Intervention&Accuracy of Prediction.

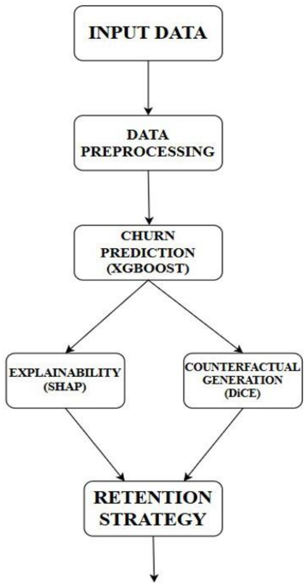

Figure 3: In this Workflow of the proposed CECP framework comprising data preprocessing tasks, prediction of churn using the XGBoost model, explainable analysis of the result using SHAP, and generation of counterfactual explanations through DiCE.

4.1 Quantitative Evaluation

When evaluating three different groups of metrics (performance), the overall best-performing predictive model (XGB cleaned - an XGBoost based predictive model) based on a standardised classification metric (accuracy, precision and recall) shown in Figure 4 outperformed all of its competitors with an accuracy score of 0.89 which indicates this predictive model was able to accurately predict customer churn. The precision score for this predictive model was 0.83 and recall score for was 0.78 indicating that the combination of these two scores provides this predictive model with a balance of False Positives (Those customers that will not churn) and False Negatives (Those customers that will churn).

4.2 Real-Time Counterfactual Application

For the purpose of testing how well the given CECP framework could function with actual time constraints, we moved forward with running member data on a weekly basis through batch processing. If any member has a probability of churn exceeding 0.60, we considered such members as high-risk for member churn. For example, for a member who initially has a probability for member churn at 0.78, we can expect them to reduce such a member's probability for member churn to 0.22 with a discount or loyalty reward strategy through counterfactual support.

4.3 Business Impact Simulation

A simulated retention task was performed to assess the potential business value of the proposed CECP approach.

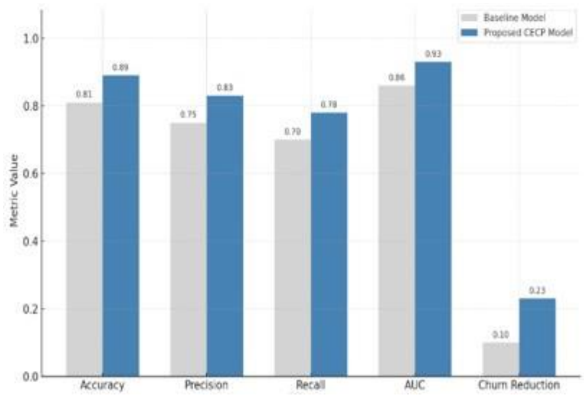

Figure 4: Comparison of the baseline model and CECP model based on accuracy, precision, recall, AUC, and churn reduction.

High-priority customers identified according to uplift values were targeted with intervention strategies informed by counterfactual suggestions. Relative to a control group that did not receive targeted interventions, the treatment group showed a 23% lower probability of churn. Note that this finding is based on simulated intervention analysis and provides an illustrative estimate of the value of counterfactual-informed retention strategies.Figure 4 compares the performance of the baseline model with the proposed CECP framework on the key performance metrics of accuracy, AUC, and churn reduction.

4.4 Visualization

This is shown in this graph:This bar chart compares very high AUC, Very High Accuracy, and is also shown to have a positive effect of the framework for business use regarding fewer customers opting to leave. This SHAP summary plot explains and interprets feature contributions towards the churn prediction value [5]. Each dot corresponds to customer instances, and these dots are colored according to feature values, for instance, high and low points utilized in rewards. Characteristics with top contributions towards influencing churning predictions include Delivery Frequency, Subscription Plan Type, and Loyalty Points Used.Figure 5 summarizes the performance metrics of the proposed CECP model on the churn prediction task.

5 Conclusion

This paper introduced a Counterfactual E-Commerce Churn Prediction (CECP) framework intended to tackle the churn issue in subscription-based e-commerce sites by combining predictive modeling, explainable AI, and counterfactual reasoning methods. Unlike conventional churn prediction models that concentrate on churn probability prediction alone, the proposed framework puts more emphasis on the provision of actionable insights regarding the most effective and personalized strategies for retaining customers.

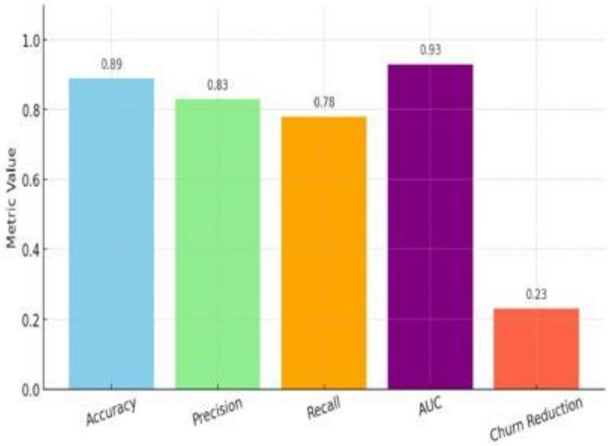

Figure 5: Performance metrics of the proposed CECP model.

The proposed CECP framework utilizes XGBoost for accurate churn prediction, SHAP for model interpretability, and DiCE for counterfactual explanation generation, which point out the least and most practicable changes needed to avoid churn. Moreover, uplift modeling with the T-Learner method allows for the computation of personalized treatment effects, which helps organizations to focus on customers who are most likely to benefit from targeted retention strategies.

The proposed method is tested on a real-world e-commerce membership dataset, and the results show that the proposed method is highly accurate in churn prediction, with an accuracy of 89% and an AUC of 0.93. Additionally, the simulation study on intervention analysis shows a possible churn reduction of up to 23%.

References

[1] T. Verbeke, D. Martens, C. Mues, and B. Baesens, "Building clear customer churn prediction models with advanced rule induction techniques," *Expert Systems with Applications*, vol. 38, no. 3, pp. 2354–2364, Mar. 2011.

[2] K. Sharma and M. Ghose, "Customer churn prediction in telecom using machine learning on a big data platform," *Journal of Big Data*, vol. 8, 2021.

[3] R. K. Srivastava, A. Raff, and A. Ghosh, "Data and intervention-driven churn prediction using counterfactual modelling," in *Proc. IEEE Int. Conf. Data Mining Workshops (ICDMW)*, 2021, pp. 394–401.

[4] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in *Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining*, 2016, pp. 785–794.

[5] S. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.

[6] M. Mothilal, A. Sharma, and C. Tan, "Explaining machine learning classifiers through diverse counterfactual explanations," in *Proc. ACM Conf. Fairness, Accountability, and Transparency (FAT*)*, 2020, pp. 607–617.

[7] H. Binns, M. Chen, M. Andreou, M. Singhal, A. Sharma, and L. Weidinger, "Counterfactual explanations for machine learning: A review," *IEEE Access*, vol. 11, pp. 15274–15292, 2023.

[8] S. Guha, N. Talukder, A. Liu, and C. Zhou, "A framework for predicting customer churn using explainable AI," in *Proc. IEEE Int. Conf. Big Data*, 2021, pp. 1899–1908.

[9] M. Wachter, B. Mittelstadt, and C. Russell, "Counterfactual explanations without opening the black box: Automated decisions and the GDPR," *Harvard Journal of Law & Technology*, vol. 31, no. 2, 2018.

[10] A. Arya, A. Mehta, R. Nagpal, and V. Vashisht, "Explainable AI for customer churn prediction using SHAP and LIME," in *Proc. IEEE Int. Conf. Communication Systems and Network Technologies (CSNT)*, 2021, pp. 383–388.

[11] B. Hansotia and B. Rukstales, "Direct marketing performance modeling using uplift modeling," *Journal of Interactive Marketing*, vol. 16, no. 3, pp. 3–20, 2002.

[12] P. Nayak and B. Mohan, "Churn analysis in retail e-commerce using ensemble techniques," in *Proc. IEEE COMITCon*, 2021, pp. 1–6.

[13] C. Molnar, *Interpretable Machine Learning*, 2nd ed. Lulu.com, 2022.

[14] M. Amin, W. Alsulaiman, and H. Al-Raweshidy, "Analysis of e-commerce customer churn using deep learning methods integrated with big data frameworks," *Journal of Big Data*, vol. 10, no. 1, pp. 1–19, 2023.

[15] R. Pradeep, A. Kumar, and K. Srivastava, "Evaluation of churn prediction approaches in Indian retail e-commerce," in *Proc. Int. Conf. Advanced Computing and Data Science (ICACDS)*, 2022, pp. 45–52.

[16] M. T. Ribeiro, S. Singh, and C. Guestrin, "Why should I trust you? Explaining the predictions of any classifier," in *Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining*, 2016, pp. 1135–1144.