

Predicting Voluntary Subscriber Churn through Engagement-Centric Machine Learning Models

Wael Breich
Sr Data Analyst
21320 Parthenia Street, # 109
Canoga Park, California 91304

ABSTRACT

Customer churn is a critical challenge in subscription-based industries due to its direct impact on revenue stability, customer lifetime value, and long-term growth. With the growing availability of behavioral data, churn prediction has become a prominent application for assessing both predictive performance and model interpretability in data mining and analytics. However, much of the existing literature relies on heterogeneous feature sets that combine engagement metrics with demographic, contractual, and pricing variables, which can limit interpretability and obscure the independent role of engagement behavior. This study investigates whether engagement-centric behavioral metrics alone are sufficient to predict voluntary customer churn. Using a real-world customer-level dataset, normalized engagement features were constructed to capture usage frequency, consumption intensity, engagement breadth, and service reliability while accounting for differences in exposure duration through tenure normalization. Logistic regression, random forest, and gradient boosting (XGBoost) models are benchmarked to evaluate predictive performance and to assess the importance of nonlinear and interaction effects. To enhance interpretability, SHapley Additive exPlanations (SHAP) are applied to the gradient boosting model to quantify the marginal contribution of individual engagement features. Results show that engagement-only features provide strong predictive performance, with nonlinear ensemble models substantially outperforming linear baselines. The SHAP analysis reveals that low engagement frequency and reduced interaction breadth are the dominant drivers of churn risk, exhibiting pronounced nonlinear effects at low activity levels. Overall, the findings demonstrate that engagement behavior offers a robust and interpretable foundation for explainable churn prediction.

General Terms

Machine Learning; Predictive Modeling; Classification.

Keywords

Customer churn prediction; Engagement features; XGBoost; Ensemble Learning; SHAP; Explainable AI.

1. INTRODUCTION

Customer churn poses a persistent challenge for subscription-based businesses, directly affecting revenue stability, customer lifetime value, and long-term growth. As competition across digital subscription services intensifies, organizations increasingly seek data-driven methods to identify at-risk customers and deploy targeted retention strategies[1]. Among the available predictors of churn, user engagement behavior is particularly valuable, as it reflects how customers interact with a service and can be influenced by product design, content strategy, promotional efforts, and customer experience interventions.

Existing churn prediction research has primarily focused on heterogeneous feature sets that combine demographic attributes, pricing plans, contractual terms, and service usage indicators[2]. Although such approaches often yield strong predictive performance, they limit interpretability and practical actionability by obscuring the specific contribution of engagement behavior. In practice, many demographic and contractual variables are static or difficult to influence, whereas engagement signals provide direct insight into evolving customer behavior.

This study investigated whether engagement metrics alone are sufficient to predict voluntary subscriber churn. Rather than emphasizing early-life churn exclusively, the analysis evaluates churn prediction across the subscriber lifecycle, while accounting for differences in exposure duration through tenure normalization. This study aimed to quantify the standalone predictive value of engagement signals by isolating engagement behavior from demographic, economic, and contractual confounders.

The contributions of this study are threefold. First, it demonstrates that engagement-centric features—capturing usage frequency, intensity, breadth, and reliability—provide strong predictive power for churn without reliance on demographic or pricing variables. Second, it benchmarks linear and nonlinear modeling approaches to assess the role of interaction and threshold effects in engagement-driven churn. Third, it applies SHAP-based interpretability techniques to identify engagement behaviors most strongly associated with churn risk. To the best of the author’s knowledge, few existing studies in applied subscription churn settings have isolated engagement-centric behavioral metrics as the sole predictors of churn while providing model-level explanations using SHAP.

The analysis was conducted on a real-world subscription dataset containing approximately 3,150 users with detailed behavioral activity and churn labels. Engagement measures are normalized by subscription tenure to isolate behavioral intensity from exposure duration. The results offer an applied and reproducible framework for understanding how engagement behavior drives churn in subscription-based services. By focusing exclusively on engagement behavior, the proposed framework emphasizes variables that are continuously observable, directly measurable, and actionable through product design, customer experience, and service level interventions.

The remainder of this paper is organized as follows. Section 2 presents a synopsis of the existing body of literature on churn analysis. Section 3 discusses the methodology including the three classification techniques used in this study. A description of the data set, feature engineering, modeling approach, and experimental setup are discussed in this section. Section 4 presents the results and specifically interprets the output of the

XGBoost algorithm using SHAP. Section 5 is the discussion, Section 6 highlights some limitations of the research and issues for future research, and Section 7 presents the conclusions.

2. LITERATURE REVIEW

Customer churn prediction is a longstanding and active research area in data mining and analytics because of its direct impact on revenue stability, customer lifetime value, and sustainable growth in subscription-based industries such as media streaming, telecommunications, finance, and digital services. With the increasing availability of large-scale behavioral data, recent work has emphasized extracting actionable insights from customer engagement patterns rather than relying on static demographic descriptors[3-5]. As a result, churn prediction continues to serve as a representative application domain for evaluating both predictive performance and interpretability in applied machine learning research.

Recent studies predominantly frame the problem as a supervised classification task, with logistic regression commonly retained as a benchmark because of its transparency and ease of deployment in business environments[6]. However, empirical evidence consistently demonstrates that ensemble-based and boosting models outperform linear baselines when applied to large heterogeneous behavioral datasets[5, 7-8]. These performance gains are particularly pronounced in imbalanced churn settings, where minority churners must be accurately identified under noisy usage signals[9]. Consequently, tree based ensembles such as Gradient Boosting, XGBoost, and LightGBM have become dominant modeling choices in contemporary churn analytics[10-11].

In this literature, engagement-related behavioral variables consistently emerge as some of the most influential predictors of churn. Features capturing usage frequency, interaction regularity, service reliability, and customer support activity typically dominate feature importance rankings, whereas static demographic attributes provide limited incremental value once behavioral signals are incorporated[7, 11]. These findings align with broader data mining evidence that temporal and behavioral features encode early warning signals of disengagement more effectively than descriptive customer profiles do. Contemporary extensions demonstrate that unstructured behavioral data, such as customer communications or complaint text, can enhance churn prediction when integrated with structured engagement features[6]. Recent BDMA-published work corroborates this perspective, demonstrating that behavioral usage features play a vital role in evaluating churners across multiple machine learning techniques[5].

Beyond predictive accuracy, a growing body of work has emphasized the importance of decision-oriented and profit driven churn modeling. Rather than optimizing conventional performance metrics alone, these approaches explicitly incorporate intervention costs and expected retention values into the model training and evaluation[12-14]. This perspective is particularly relevant in churn contexts, where the operational objective is not merely classification accuracy, but also the effective allocation of retention resources. From a data mining standpoint, this study highlights the necessity of aligning analytical models with downstream decision processes.

Despite advances in predictive performance, the increasing complexity of ensemble and deep learning models introduces significant challenges related to transparency and interpretability. In applied churn settings, black-box predictions often fail to provide sufficient insight into the behavioral mechanisms driving churn, limiting stakeholder

trust and hindering actionable intervention design. This tension between accuracy and interpretability has been widely discussed in explainable artificial intelligence (XAI) literature[15-18]. For churn analytics in particular, explainability is not merely desirable but frequently a prerequisite for operational adoption.

Among the XAI techniques, SHapley Additive exPlanations (SHAP) have emerged as a widely adopted framework because of its solid theoretical foundation, additive consistency, and scalability for tree-based ensemble models[19-20]. SHAP enables both global and local interpretations of model predictions, making it especially suitable for large-scale churn datasets where heterogeneous engagement patterns interact in nonlinear ways. Recent churn studies leveraged SHAP to identify dominant behavioral drivers, uncover threshold effects, and analyze feature interactions that are otherwise obscured in black-box models[11, 21-22].

Empirical evidence increasingly suggests that engagement effects on churn are inherently nonlinear. Several studies document sharp increases in churn probability once the usage frequency or interaction regularity falls below the minimum thresholds, followed by diminishing marginal retention benefits at higher engagement levels[7, 21]. These findings reinforce the suitability of tree-based ensembles for churn modeling, as they naturally capture the threshold behavior and interaction effects among behavioral variables. From a data mining perspective, this underscores the importance of model classes capable of learning complex response surfaces from large behavioral feature spaces.

Event-driven behavioral variables (such as service failures, complaints, and negative support interactions) have also been shown to exert disproportionate influence on churn outcomes. Although these events occur relatively infrequently, their presence often signals acute dissatisfaction and significantly elevates churn risk, particularly when combined with declining engagement trajectories[6, 11]. These variables frequently act as churn accelerators rather than as continuous predictors, reinforcing the need for analytical approaches that account for interaction effects and conditional risk amplification.

Despite broad agreement on the centrality of engagement behavior, most existing churn studies incorporate engagement features along with tenure, pricing plans, contractual constraints, and billing attributes. While this improves the predictive accuracy, it complicates interpretation and limits the ability to isolate engagement-driven churn mechanisms. Moreover, many commonly used predictors (such as tenure or pricing) are difficult to directly influence through product or service design, thus reducing their strategic value for proactive retention. Recent surveys emphasize the need for more focused and interpretable churn modeling frameworks that balance predictive power with actionable insights[4].

This study contributes to the churn analytics literature in three ways. First, it explicitly isolates engagement-centric behavioral metrics and evaluates their predictive power independent of demographic, pricing, and contractual variables. Second, by benchmarking logistic regression against random forest and gradient boosting models, this study examines the role of nonlinearity and interaction effects in engagement-driven churn. Third, by integrating SHAP based interpretability, this study provides both global and local explanations that translate predictive patterns into actionable behavioral insights. In doing so, the study demonstrates how data mining methods can bridge predictive performance and decision support in real-world churn applications.

3. METHODOLOGY

3.1 Dataset and Preprocessing

The dataset consists of approximately 3,150 anonymized subscriber records collected from a telecommunications service provider and publicly available through the UCI Machine Learning Repository[23]. Each record corresponds to a single customer and includes a binary churn indicator reflecting the customer status at the end of a twelve-month observation period. The observed churn rate was approximately 16%, reflecting a realistic class imbalance commonly encountered in applied churn modeling settings.

The dataset contains no missing values in the engagement related features used in this study. Feature distributions were inspected for extreme values using summary statistics and boxplots. Given the use of tree-based models and normalized engagement measures, no additional outlier treatment was applied. All engagement-related features are aggregated over the first nine months of the observation window to capture customer behavior prior to churn. The churn label reflects whether a customer churned by the end of month 12, resulting in a three-month prediction horizon between the behavioral observation period and churn outcome. This temporal separation ensures proper ordering between predictors and outcomes and avoids look-ahead bias.

The feature set is composed exclusively of engagement centric behavioral metrics, which can be grouped into four categories: (i) engagement frequency (e.g., number of interactions per month), (ii) engagement intensity (e.g., total usage time per month), (iii) engagement breadth (e.g., diversity of interaction partners), and (iv) service reliability and friction signals (e.g., call failures and complaint events). Subscription length is recorded in months and used solely to normalize engagement measures to account for differences in exposure duration.

Consistent with common practices in churn modeling studies, all observed churn events in the dataset are treated as voluntary churn, as the data does not distinguish between voluntary and involuntary cancellations. No demographic, pricing, or contractual variables were included in the primary feature set to isolate the predictive contribution of engagement behavior. The dataset was fully anonymized and did not contain personally identifiable information.

3.2 Feature Engineering

To account for differences in exposure duration, all usage based engagement variables were normalized by subscription length, yielding monthly engagement rates. This transformation isolates behavioral intensity rather than cumulative opportunity, ensuring comparability across subscribers with different tenures. Engagement breadth was operationalized as the ratio of distinct interaction partners to total usage frequency, capturing the diversity of user interactions beyond raw activity levels. Service complaints were treated as a binary behavioral flag because of their sparse, event-driven nature and disproportionate impact on churn risk when present. The subscription length was excluded from the primary feature set to avoid conflating engagement effects with exposure duration and to preserve the interpretability of engagement-driven signals.

3.3 Modeling Approach

Three models were evaluated to assess churn predictability using engagement-centric features: logistic regression, random forest, and gradient boosting (XGBoost). Logistic regression served as an interpretable linear baseline, providing insight into the additive relationships between engagement metrics and

churn. Tree-based ensemble models were employed to capture nonlinear effects and interactions among engagement features that were not represented by linear models.

All models were trained using the same feature set and evaluated using a stratified train–test split to preserve the observed churn rate. Model performance was assessed using the Area Under the Receiver Operating Characteristic Curve (AUC) and the F1 score to balance ranking quality and minority-class classification performance under class imbalance.

To ensure robustness of the reported results, model performance was additionally evaluated using five-fold stratified cross-validation on the training data. Reported performance metrics reflect the mean values across folds. This procedure mitigates potential variance arising from a single train–test split while preserving the original class distribution.

The logistic regression model was estimated using L2 regularization with the regularization strength selected using default library settings. Random forest and gradient boosting models were trained using commonly adopted hyperparameters to balance predictive performance and interpretability. For gradient boosting, tree depth, learning rate, and number of estimators were tuned using a limited grid search on the training set only, with final parameters selected based on cross-validated AUC. All modeling was conducted using Python’s scikit-learn and XGBoost libraries.

3.4 Model Interpretability

To interpret the predictions of the nonlinear models, SHapley Additive exPlanations (SHAP) were applied to the gradient boosting model. SHAP values quantify the marginal contribution of each engagement feature to individual churn predictions, providing both global feature importance and directional and instance-level effects. This approach enables a transparent interpretation of complex engagement-driven churn dynamics, allowing the identification of dominant behavioral drivers, nonlinear thresholds, and interaction patterns underlying model predictions.

3.5 Evaluation metrics and setup

The model performance was evaluated using a stratified train test split, with 80% of the data used for training and 20% used for testing. Stratification was applied to preserve the observed churn rates in both subsets. All models were trained using the same engagement-centric feature set to ensure comparability across approaches.

Predictive performance was assessed using the Area Under the Receiver Operating Characteristic Curve (AUC) and the F1 score (which incorporates precision and recall). AUC was selected to measure the overall ranking performance under class imbalance, while the F1 score was used to evaluate the balance between precision and recall for the churn class. Accuracy was not emphasized because it can be misleading in imbalanced classification settings, where a model may achieve high accuracy by predominantly predicting the majority (non-churn) class while failing to correctly identify churners.

No resampling or class reweighting techniques were applied because the observed class distribution reflected realistic subscription churn dynamics and preserved the practical relevance of model performance.

Figure 1 presents the descriptive distributions of the selected engagement features stratified by churn status. While these visualizations are not intended to establish causal relationships, they provide preliminary evidence of systematic differences in

engagement behavior between churned and retained subscribers and motivate the engagement-centric modeling approach adopted in this study.

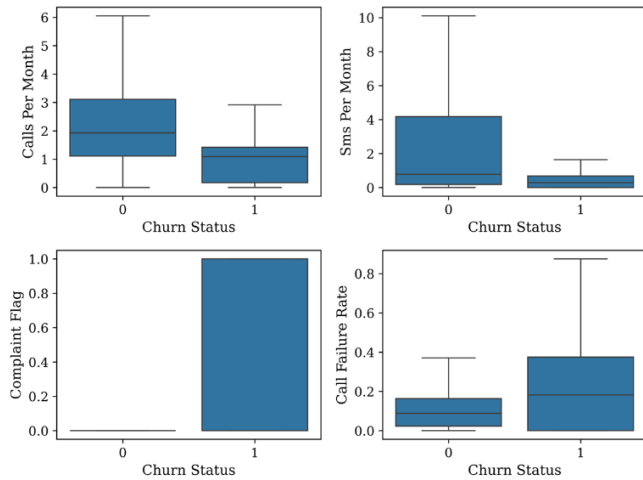


Figure 1: Distributions of selected engagement-centric features stratified by churn status

4. RESULTS

4.1 Predictive Performance

Table 1 summarizes the predictive performance of the evaluated models using engagement-centric features. All three models demonstrate strong discriminative ability, confirming that engagement behavior provides a substantial signal for churn prediction.

Table 1. Predictive performance of machine learning models using engagement-centric features

Model	AUC	F1 Score	Precision	Recall
Log regression	0.90	0.58	0.45	0.82
Random forest	0.96	0.58	0.98	0.41
XGBoost	0.97	0.70	0.92	0.57

The logistic regression baseline achieves an AUC of 0.90 and an F1 score of 0.58. The model exhibits high recall (0.82) but relatively low precision (0.45), indicating that it identifies a large proportion of churners at the expense of a high number of false positives. This behavior is consistent with linear models that capture general trends but are unable to model nonlinear engagement dynamics.

Tree-based models substantially outperform the linear baseline. The random forest model improves AUC to 0.96 but maintains a similar F1 score (0.58), reflecting an imbalance between precision and recall. Specifically, the model achieves very high precision (0.98) but low recall (0.41), indicating a conservative classification strategy in which churn is predicted only when the model is highly confident. While this minimizes false positives, it results in a large number of missed churners.

The gradient boosting (XGBoost) model achieves the strongest overall performance, with an AUC of 0.97 and an F1 score of 0.70. Unlike the other models, XGBoost achieves a more balanced trade-off between precision (0.92) and recall (0.57), indicating improved ability to identify churners while maintaining a relatively low false positive rate. These results

suggest that churn risk is governed by nonlinear and interacting engagement patterns that are not captured by purely additive models.

Given the class imbalance in the dataset (approximately 16% churn rate), precision and recall provide a more informative evaluation than accuracy. The observed differences across models highlight the importance of balancing false positives and false negatives in churn prediction, particularly in practical retention scenarios.

To ensure robustness, model performance was additionally evaluated using five-fold stratified cross-validation. The results were consistent with the hold-out test set, with standard deviations below 0.02 for AUC across all models, indicating stable performance. These findings suggest that the observed performance improvements are not driven by a specific train-test split but reflect stable and generalizable engagement patterns across data partitions.

4.2 Confusion Matrix Analysis

To further examine classification behavior, Figure 2 presents the confusion matrix for the XGBoost model. The model correctly classifies a large number of non-churners (true negatives = 526) while maintaining a very low number of false positives (5), reflecting its strong precision. At the same time, it correctly identifies a substantial portion of churners (true positives = 56), although some churners remain undetected (false negatives = 43), indicating a moderate recall.

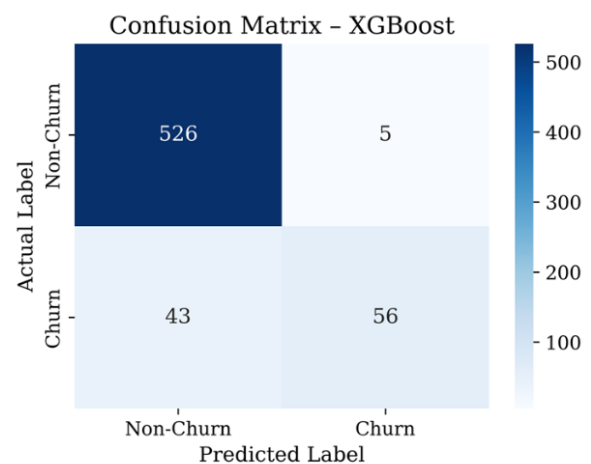


Figure 2: Confusion Matrix for XGBoost Model

This distribution highlights an important practical trade-off. The low false positive rate ensures that retention efforts are not unnecessarily applied to customers who would not churn, while the remaining false negatives suggest opportunities for improving churn coverage through threshold tuning or cost sensitive optimization. Overall, the confusion matrix confirms that XGBoost achieves a balanced classification strategy, aligning with its superior F1 score and making it well-suited for real-world deployment.

4.3 Global Feature Importance

Figure 3 presents the global feature importance derived from SHAP values for the XGBoost model. Engagement frequency, measured by calls per month, emerges as the dominant predictor of churn, followed by SMS activity per month. These frequency-based measures contribute more to model predictions than aggregate usage volume, highlighting the importance of consistent interaction over total consumption.

The complaint flag ranks as one of the most influential features despite its low prevalence, indicating that complaint events exert a disproportionately large impact on churn risk. Usage intensity also contributes meaningfully, while call failure rate and engagement breadth play secondary but non negligible roles.

Overall, the feature importance results suggest that churn is primarily driven by behavioral engagement patterns rather than isolated service quality indicators.

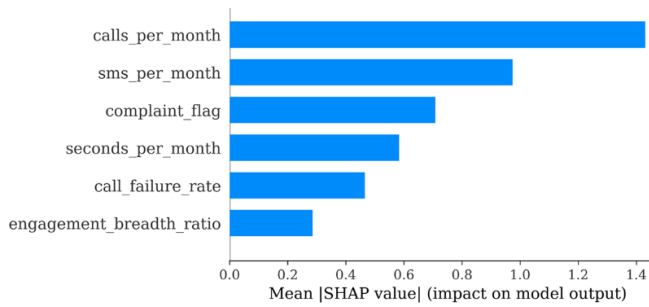


Figure 3: Global Feature Importance

4.4 Directional Effects/Behavioral Patterns

Figure 4 illustrates the directional effects of engagement features on churn predictions using SHAP summary plots. Lower values of calls per month and SMS per month consistently push predictions toward churn, whereas higher values shift predictions toward retention. These effects are nonlinear, with steep increases in churn risk at low engagement levels and diminishing marginal returns beyond moderate usage thresholds.

The complaint flag exhibits a clear asymmetric effect: the presence of a complaint strongly increases churn probability, whereas the absence of complaints contributes minimally to retention. This pattern reflects the event-driven nature of complaints and their role as accelerators of disengagement rather than continuous predictors.

Higher call failure rates are associated with increased churn risk, particularly when combined with low engagement frequency. In contrast, engagement breadth, measured as the ratio of distinct interaction partners to total usage, shows a protective effect: broader engagement is associated with reduced churn probability, although its impact is smaller relative to frequency-based measures. These nonlinear patterns indicate that modest declines in engagement among low-activity users substantially elevate churn risk, suggesting that early, targeted interventions at low engagement levels may yield disproportionately large retention benefits.




Figure 4: Directional Effects of Engagement Features

4.5 Summary of Findings

Taken together, the results demonstrate that engagement centric features alone provide strong predictive power for churn across the subscriber lifecycle. Nonlinear ensemble models significantly outperform linear baselines, indicating that churn risk is driven by interacting patterns of engagement frequency, intensity, and breadth.

Importantly, the results reveal distinct trade-offs across models: logistic regression prioritizes recall, random forest prioritizes precision, and XGBoost achieves the most balanced performance. This balance is critical in practical applications, where both accurate identification of churners and efficient allocation of retention resources are required.

The interpretability analysis confirms that low engagement frequency is the primary driver of churn, with complaint events and service reliability acting as important secondary signals. These findings support the use of engagement-driven modeling frameworks and highlight the value of interpretable machine learning in translating predictive performance into actionable insights.

5. DISCUSSION

This study demonstrates that engagement-centric behavioral metrics provide strong predictive power for subscriber churn across the customer lifecycle. Using normalized engagement features alone, all evaluated models achieved high discriminatory performance, with nonlinear ensemble methods substantially outperforming a linear baseline. These results suggest that how subscribers interact with a service (how often, how intensely, and how broadly) captures meaningful information about churn risk independent of demographic or pricing attributes.

The performance gains observed in the tree-based models indicate that churn is not governed by simple additive relationships between the engagement metrics. Instead, churn risk emerges from nonlinear and interacting engagement patterns, particularly at low activity levels. The SHAP-based interpretability analysis reveals that usage frequency is the dominant driver of churn predictions, followed by secondary engagement channels and consumption intensity. This finding highlights the importance of consistent interaction over sporadic or high-volume usage, suggesting that regular engagement may play a greater role in retention than aggregate consumption alone.

The analysis further shows that complaint events, while infrequent, have a disproportionate impact on churn risk. The asymmetric effect observed in the SHAP plots indicates that complaints function as accelerators of churn rather than as gradual predictors, reinforcing their role as critical intervention signals. Similarly, service reliability, as captured by call failure rates, meaningfully contributes to churn predictions when combined with low engagement levels, suggesting that quality issues are particularly damaging among already disengaged users.

Engagement breadth, measured through interaction diversity, exhibited a smaller but consistent protective effect. Subscribers who interact across a broader set of contacts or usage contexts are less likely to churn, supporting the notion that deeper behavioral embeddedness increases switching costs and perceived value. Although breadth is less influential than frequency-based measures, its contribution underscores the multidimensional nature of engagement-driven churn dynamics.

From a practical standpoint, these findings suggest that retention strategies should prioritize the early detection of declining engagement frequency and narrowing usage patterns, rather than relying solely on tenure or customer value metrics. The nonlinear effects identified through the SHAP analysis imply that modest increases in engagement among low-activity users may yield outsized retention benefits, whereas additional engagement among already active users may have diminishing returns. This has direct implications for targeted interventions such as personalized onboarding, engagement nudges, and proactive service outreach.

More broadly, this study contributes to the literature on churn prediction by isolating engagement behavior as a standalone predictive signal and demonstrating its effectiveness across model classes. By combining strong predictive performance with interpretable machine-learning techniques, the proposed framework bridges the gap between churn prediction accuracy and actionable behavioral insight. While prior studies often integrate engagement features alongside demographic and contractual variables, the results show that engagement alone captures a substantial portion of churn risk, reinforcing its central role in subscription analytics.

6. LIMITATIONS AND FUTURE WORK

This study has several limitations that should be considered when interpreting the results. First, the analysis relies on observational data, which limits the ability to draw causal conclusions about the relationship between engagement behavior and churn.

Although the models identify strong predictive associations, they do not establish whether changes in engagement directly cause churn or merely precede it. Second, this study was based on a single dataset drawn from a specific subscription context. Although the engagement-driven framework is designed to be broadly applicable, the churn dynamics may differ across industries, platforms, and markets. Future work should validate the proposed approach using datasets from other subscription domains, such as media streaming, telecommunications, or SaaS platforms, to assess generalizability. Third, engagement behavior was captured using a finite set of proxy metrics. While these features reflect the key dimensions of usage frequency, intensity, breadth, and reliability, they may not fully capture the qualitative aspects of user experience such as content relevance, satisfaction, or perceived value. Incorporating richer behavioral signals or temporal engagement trajectories may further improve predictive performance.

Another limitation concerns data availability: publicly accessible churn datasets with sufficiently granular early-stage engagement data are limited, constraining the ability to model churn exclusively within the initial subscription period.

Future research could extend this work by explicitly modeling early-stage churn using datasets with higher-resolution engagement data during the initial subscription or trial period. Such data would enable fine-grained analysis of onboarding behaviors, early habit formation, and rapid disengagement signals. From a practical perspective, this would support more proactive retention strategies by allowing teams to identify at risk users during trial or early lifecycle stages, when targeted interventions, such as guided onboarding, personalized content exposure, or early engagement nudges, are most likely to influence long-term retention outcomes. Future research could also incorporate demographic and pricing variables as controls to assess how engagement-driven churn signals interact with the subscriber characteristics. Additionally, causal inference methods or controlled experiments could be employed to

evaluate the effectiveness of engagement-based interventions, as suggested by the predictive and interpretability results presented here.

7. CONCLUSION

This study demonstrates that engagement-centric behavioral metrics alone provide strong predictive power for subscriber churn across the customer's lifecycle. By isolating engagement behavior from demographic, pricing, and contractual factors, the analysis showed that patterns of usage frequency, consumption intensity, engagement breadth, and service reliability capture meaningful signals of churn risk.

Comparative modeling results indicate that nonlinear ensemble methods substantially outperform linear baselines, highlighting the importance of interaction and threshold effects in engagement-driven churn dynamics. The SHAP based interpretability analysis further reveals that low engagement frequency is the dominant contributor to churn risk, with complaint events and service reliability acting as critical secondary accelerators.

Together, these findings support the use of engagement focused modeling frameworks for churn prediction and emphasize the value of interpretable machine learning in translating predictive accuracy into actionable insights. By demonstrating that engagement behavior alone can effectively identify at-risk subscribers, this study provides a practical and reproducible approach for subscription-based platforms seeking to design targeted, behavior-driven retention strategies.

8. REFERENCES

- [1] E. Ascarza, "Retention futility: Targeting high-risk customers might be ineffective," *J. Marketing Res.*, vol. 55, no. 1, pp. 80–98, 2018.
- [2] M. Bogaert and L. Delaere, "A systematic review of customer churn prediction using machine learning techniques," *Mathematics*, vol. 11, no. 5, p. 1137, 2023.
- [3] P. Wachwanakijkul, S. Junsiritrakhoon, N. Kantanantha, G. Narayanamurthy, and P. Jarumaneeroj, "Customer churn prediction: A comprehensive survey of machine learning approaches," *Big Data Cogn. Comput.*, vol. 9, no. 3, 2025.
- [4] M. Imani, M. Joudaki, A. Beikmohammadi, and H. R. Arabnia, "A survey of customer churn analysis in data driven environments," *Big Data Cogn. Comput.*, vol. 7, no. 3, p. 105, 2025.
- [5] S. Kaisar, M. Rashid, A. Chowdhury, S. S. Shafin, J. Kamruzzaman, and A. Diro, "Enhancing telemarketing success using ensemble-based learning," *Big Data Mining and Analytics*, vol. 7, no. 1, pp. 1–14, 2024.
- [6] A. De Caigny, K. Coussement, K. W. De Bock, and S. Lessmann, "Incorporating textual information in customer churn prediction models based on a convolutional neural network," *Decision Support Syst.*, vol. 129, p. 113189, 2020.
- [7] T. Xu, Y. Ma, and K. Kangchul, "Customer churn prediction using ensemble learning and feature selection," *Appl. Sci.*, vol. 11, no. 11, p. 4742, 2021.
- [8] S. M. Shrestha and A. Shakya, "Customer churn prediction using machine learning techniques," *Procedia Comput. Sci.*, vol. 198, pp. 86–93, 2022.
- [9] A. K. Ahmad, A. Jafar, and K. Aljoumaa, "Customer

- churn prediction in telecom using machine learning in big data platform,” *J. Big Data*, vol. 6, no. 1, 2019.
- [10] S. M. Mirabdolbaghi and B. Amiri, “Model optimization analysis of customer churn prediction using machine learning algorithms with focus on feature reduction,” *Expert Syst.*, vol. 39, no. 6, 2022.
- [11] V. Chang, K. Hall, Q. A. Xu, F. O. Amao, M. A. Ganatra, and V. Benson, “Explainable AI for customer churn prediction in digital services,” *Information*, vol. 17, no. 6, p. 231, 2024.
- [12] E. Stripling, S. Vanden Broucke, K. Antonio, B. Baesens, and M. Snoeck, “Profit-driven decision trees for churn prediction,” *Decision Support Syst.*, vol. 111, pp. 28–41, 2018.
- [13] S. Höppner, E. Stripling, B. Baesens, S. Vanden Broucke, and T. Verdonck, “Profit-driven churn prediction with customer lifetime value,” *Eur. J. Oper. Res.*, vol. 284, no. 3, pp. 920–934, 2020.
- [14] S. Maldonado, J. López, and C. Vairetti, “Profit-based feature selection for churn prediction,” *Eur. J. Oper. Res.*, vol. 283, no. 3, pp. 1086–1098, 2020.
- [15] R. Guidotti, A. Monreale, S. Ruggieri, F. Turini, F. Giannotti, and D. Pedreschi, “A survey of methods for explaining black box models,” *ACM Comput. Surv.*, vol. 51, no. 5, pp. 1–42, 2018.
- [16] A. B. Arrieta et al., “Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges,” *Information Fusion*, vol. 58, pp. 82–115, 2020.
- [17] P. Linardatos, V. Papastefanopoulos, and S. Kotsiantis, “Explainable AI: A review of machine learning interpretability methods,” *Entropy*, vol. 23, no. 1, p. 18, 2021.
- [18] P. P. Angelov and E. A. Soares, “Towards explainable deep neural networks (xDNN),” *Wiley Interdiscip. Rev.: Data Mining Knowl. Discov.*, vol. 10, no. 1, e1346, 2020.
- [19] S. M. Lundberg and S.-I. Lee, “A unified approach to interpreting model predictions,” in *Proc. 31st Int. Conf. Neural Information Processing Systems (NeurIPS)*, 2017, pp. 4765–4774.
- [20] S. M. Lundberg et al., “Consistent individualized feature attribution for tree ensembles,” *arXiv preprint arXiv:1802.03888*, 2018.
- [21] S. C. Tékouabou, Ş. C. Gherghina, H. Touluni, P. Neves Mata, and J. M. Martins, “Explainable machine learning models for customer churn prediction,” *Mathematics*, vol. 10, no. 14, p. 2379, 2022.
- [22] S. S. Poudel, S. Pokharel, and M. Timilsina, “Explainable machine learning for customer churn prediction,” *Machine Learning with Applications*, vol. 14, 2024.
- [23] UCI Machine Learning Repository, “Iranian churn dataset,” 2020. [Online]. Available: <https://doi.org/10.24432/C5JW3Z>