

Explainable AI (XAI) for Customer Retention: A Prescriptive Modelling Approach Using Tree-Based Ensembles and SHAP

Rutvik Varankar¹, Pranav Kamble², Ria Gawde³, Ashwini Kadam⁴

MCA 2nd Year, School of Humanities and Sciences, D.Y. Patil Deemed to be University

Abstract - Customer churn creates significant financial challenges for telecommunications companies because retaining existing customers is generally less costly than acquiring new ones. Although machine-learning models can predict churn, imbalanced data and limited interpretability often reduce their practical value. This study presents an Explainable AI (XAI) and prescriptive modelling framework using the IBM Telco Customer Churn dataset (N=7,043). Cost-sensitive Random Forest classification and a 0.40 probability threshold are used to improve churn detection, while SHAP explains the main contributing factors. Results on 1,409 test records achieve 65.50% churn recall. Tenure, month-to-month contracts, fiber-optic service, and monthly charges are important churn indicators. The framework also groups customers into Low, Medium, and High risk, identifying 413 customers for targeted retention: 220 High-risk and 193 Medium-risk customers.

INTRODUCTION

Artificial intelligence is increasingly used for practical business decisions, including customer churn prediction in telecommunications. Accuracy alone can be misleading when most customers do not churn, while business teams also need to understand why a customer is predicted to leave. This study addresses these issues by combining tree-based ensemble modelling, class-sensitive evaluation, SHAP explanations, and risk-based retention recommendations.

LITERATURE REVIEW

Oladapo et al. (2025) reviewed Explainable AI methods, including LIME, SHAP, counterfactual explanations, and model-agnostic techniques. Their review describes XAI use in critical areas such as cybersecurity, finance, manufacturing, and autonomous vehicles. It highlights the trade-off between model accuracy and explainability and recommends context- and user-focused explanations to support trustworthy decisions.

Ortigossa et al. (2024) reviewed XAI foundations, explanation techniques, challenges, and practical machine-learning applications. Their discussion notes the growing use of XAI across several sectors and the continuing difficulty of balancing interpretability with predictive performance. The authors emphasize reliable, context-aware explanations for trustworthy decisions.

Kalasampath et al. (2025) examined XAI applications across multiple domains. The review identifies trust, public acceptance, and regulatory requirements as important barriers, particularly where AI decisions affect people directly. It concludes that transparent explanations are needed to make AI systems easier to understand and trust.

Atakishiyev et al. (2023) surveyed XAI methods and future research for autonomous driving. Although AI supports real-time decisions in self-driving vehicles, its reasoning can remain unclear to users. The study highlights the need for transparent methods that improve understanding of vehicle decisions and address current limitations.

Bommer et al. (2024) proposed a quantitative framework for comparing local post-hoc XAI methods in climate science. They assessed feature-attribution methods using neural-network models and found that performance varied by explanation technique and architecture. The study recommends objective numerical evaluation rather than relying only on subjective judgement.

RESEARCH GAP

Despite extensive research on telecommunications churn, important gaps remain. Accuracy can conceal poor identification of churners in imbalanced datasets, so recall and cost-sensitive learning are important. In addition, churn predictions need clearer explanations and practical retention actions.

Existing churn studies often identify influential features without clearly explaining how those features affect individual predictions. SHAP can provide both overall feature importance and customer-level explanations, helping managers understand why a customer is considered at risk.

Many studies stop at predicting churn instead of translating probabilities into practical retention measures. A prescriptive framework can classify customers by risk and guide suitable actions, such as targeted discounts, contract offers, or customer support.

OBJECTIVES OF RESEARCH

- 1) To develop a tree-based ensemble model for the initial identification of customers at risk of churn.
- 2) To improve churn detection by utilizing class-weighted learning and recall-focused model optimization to identify a bigger proportion of potential churners.
- 3) To apply SHAP-based Explainable AI to identify and interpret the fundamental aspects affecting customer churn at both overall and individual customer levels.
- 4) To categorize customers into Low, Medium, and High churn-risk types established on their predicted probability of exiting.
- 5) To develop a prescriptive retention framework that converts churn predictions and their fundamental reasons into targeted, cost-effective retention actions similar as discounts, offers, or customer support.

Hypothesis

H₀: Customer tenure and another service-related aspects do not play a substantial role in determining customer churn.

H₁: Customer tenure has a role in customer churn, in that customers who have a shorter tenure have a greater probability of churning, and alternative aspects similar as total charges, contract type, internet service and monthly charges are furthermore aspects contributing to churn.

RESEARCH METHODOLOGY

1. Research Design

This quantitative, experimental study follows a deductive approach. It tests cost-sensitive learning, model evaluation, explainability, and risk-based customer segmentation using recorded telecommunications data.

1.1 Research Strategy & Methodological Choice

- **Methodological Choice:** Mono-method quantitative study established on empirical subordinate data.
- **Research Type:** Experimental and evidence-based study.
- **Time Horizon:** Cross-section, studying one snapshot of recorded telecommunications data of customers.

1.2 Population, Sampling, & Data Source

- **Data Source:** Secondary telecommunications data set (N=7043).
- Stratified sampling is used for an 80/20 train-test split (5,634 training records and 1,409 test records), maintaining similar class distributions in both sets.

1.3 Operationalization of Variables & Preprocessing

- The dependent variable is customer churn status, represented as a binary target: churned or not churned.
- Predictors include customer demographics, subscribed services, contract and payment details, tenure, monthly charges, and total charges.
- Independent variables include continuous features such as Tenure Months, Monthly Charges, and Total Charges, alongside categorical features such as Contract Type, Internet Service, Payment Method, Online Security, and Tech Support.
- To reduce data leakage, columns containing information unavailable at prediction time, including Churn Reason, Churn Score, and CLTV, were removed. CustomerID was also excluded as an identifier.

1.4 Computational & Analytical Procedures

1. **Algorithmic Model Training:** Development of tree-based community classifiers (Random Forest, XGBoost). Cost-sensitive learning (class_weight='balanced') is merged into the loss function in order to punish false-negative errors in minority-class.
2. **Decision Boundary Calibration:** Empirical tuning of the threshold (0.40 of probability) in order to give priority to Recall over standard threshold values (0.50).
3. **Model Transparency Framework (XAI):** Introduction of the Game-Theoretic SHapley Additive clarifications (structure) to approximate localized Shapley values (ϕ) for analysis of directionality and assessment of significance of features.
4. **Prescriptive Analytics System:** Mapping of expected and foreseen posterior probabilities $P(Y=1|X)$ into distinct levels of working risk via the application of a rule-based algorithm through risk tiers.

Risk Tier = {Low Risk (Safe), $P(Y = 1|X) < 0.40$ Medium Risk (Targeted Discount), $0.40 \leq P(Y = 1|X) < 0.60$ High Risk (Immediate Intervention) $P(Y = 1|X) \geq 0.60$

2.1 Data Preprocessing

The IBM Telco Customer Churn dataset was prepared by encoding categorical features and converting financial fields, including Total Charges and Monthly Charges, into numeric values. Leakage-related columns were removed before model training.


2.2 Addressing Class Imbalance

Initial models achieved reasonable accuracy but missed many churners, with the first Random Forest model recording 52.25% churn recall. Class weights were therefore set to 'balanced' to give greater importance to the minority churn class.

2.3 Model Training and Evaluation

The data were split 80/20, and Random Forest and XGBoost were selected for their ability to model non-linear relationships. Evaluation emphasized churn-class metrics rather than potentially misleading overall averages. The classification threshold was tuned to 0.40 to improve recall.


2.4 Working



3. EXPLAINABLE AI (XAI) INTEGRATION

The framework uses SHAP (SHapley Additive Explanations) to show how individual features contribute to churn predictions. SHAP summaries help identify important customer and service factors, support risk classification, and guide targeted retention actions.

Data Analysis



Data Interpretation

- Features on the y-axis are ordered by importance. Tenure Months is the strongest predictor in this model.
- SHAP values on the x-axis show each feature's effect on the prediction: positive values push the result toward churn, while negative values push it toward staying.
- Colour represents feature values: red indicates high values (or Yes/1), while blue indicates low values (or No/0). The SHAP summary shows non-linear relationships between customer attributes and churn. Short tenure increases churn risk, while longer tenure is associated with retention. Two-year contracts strongly reduce risk. Fiber-optic service and high monthly charges are associated with greater churn risk; dependents and several support-related services are associated with lower risk.

Statistical Hypothesis Testing

```
Optimization terminated successfully.
Current function value: 0.428953
Iterations: 8

Logit Regression Results
=====
Dep. Variable: Churn Value No. Observations: 7043
Model: Logit DF Residuals: 7035
Method: MLE DF Model: 7
Date: Wed, 30 Sep 2026 Pseudo R-squ.: 0.2586
Time: 17:07:26 Log-Likelihood: -3021.1
Converged: True LL-Null: -4075.1
Covariance Type: nonrobust LLR p-value: 0.000
=====
coef std err z P>|z| [0.025 0.975]
-----
const 0.0212 0.193 0.110 0.913 -0.357 0.399
Tenure Months -0.0579 0.006 -9.464 0.000 -0.070 -0.046
Total Charges 0.0003 7e-05 4.553 0.000 0.000 0.000
Monthly Charges -0.0036 0.003 -1.056 0.291 -0.010 0.003
Contract_One_year -0.0690 0.104 -0.376 0.000 -1.072 -0.666
Contract_Two_year -1.7816 0.171 -10.395 0.000 -2.118 -1.446
Internet_Service_Fiber_optic 1.1059 0.128 8.632 0.000 0.855 1.357
Internet_Service_No -1.0925 0.155 -7.031 0.000 -1.397 -0.788
=====
```

```
===== HYPOTHESIS TEST RESULTS =====
Feature Coefficient P-Value Odds Ratio CI Lower CI Upper Significance
Tenure Months -0.0579 0.000000 0.9437 0.9324 0.9551 Significant (Reject H0)
Total Charges 0.0003 0.000005 1.0003 1.0002 1.0005 Significant (Reject H0)
Monthly Charges -0.0036 0.290758 0.9964 0.9897 1.0031 Not Significant (Fail to Reject H0)
Contract_One_year -0.0690 0.000000 0.4194 0.3422 0.5140 Significant (Reject H0)
Contract_Two_year -1.7816 0.000000 0.1684 0.1203 0.2356 Significant (Reject H0)
Internet_Service_Fiber_optic 1.1059 0.000000 3.0221 2.3510 3.8847 Significant (Reject H0)
Internet_Service_No -1.0925 0.000000 0.3354 0.2473 0.4548 Significant (Reject H0)
=====
```

```
===== TENURE HYPOTHESIS TEST =====
Tenure coefficient : -0.05794595222650136
Tenure p-value : 2.968253278786819e-21
Tenure odds ratio : 0.943700950997356
Result: Tenure Months has a statistically significant association with churn.
Direction: Higher tenure is associated with lower odds of churn.
This supports the hypothesis that shorter-tenure customers have higher churn risk.
```

```
===== OVERALL HYPOTHESIS TEST =====
Likelihood Ratio Statistic : 2107.9137428878075
Degrees of Freedom : 7
Overall p-value : 0.0
Decision: Reject H0
Conclusion: The predictor variables jointly provide statistically significant information about customer churn.
```

Interpretation

A regression-based statistical test examined the relationship between tenure and churn.

The p-value for Tenure Months was 2.968×10^{-21} , below the 0.05 significance level. Therefore, the null hypothesis was rejected and the alternative hypothesis supported: customers with shorter tenure are more likely to churn. Random Forest was used for prediction, while SHAP explained the effects of individual features.

Major Findings

- **Tenure Months:** Low tenure is associated with positive SHAP values and higher predicted churn risk, while longer tenure generally reduces risk.
- **Total Charges:** Lower total charges are generally associated with higher churn risk, while higher total charges tend to be associated with lower risk.
- **Internet Service—Fiber optic:** Fiber-optic service is associated with increased predicted churn risk in the model.

- **Contract—Two year:** A two-year contract is associated with lower predicted churn risk.
- **Dependents—Yes:** Having dependents generally contributes to a lower predicted churn probability.
- **Monthly Charges:** Higher monthly charges are associated with increased predicted churn risk.
- **Payment Method—Electronic check:** Electronic-check payment is associated with higher predicted churn risk.
- **Contract—One year:** A one-year contract is associated with lower churn risk than the baseline contract category.
- **Online Security—Yes:** Having online security generally reduces predicted churn risk.
- **Paperless Billing—Yes:** Paperless billing is associated with increased predicted churn risk.
- **Tech Support—Yes:** Having tech support generally reduces churn risk, though its effect is smaller than tenure, fiber-optic service, and contract type.
- **Multiple Lines—Yes:** This feature has a relatively small and mixed contribution to churn predictions.
- **Online Security—No internet service:** This category is associated with a lower predicted churn risk in the model.
- **Device Protection—No internet service:** This feature has little influence on the model's churn predictions.
- **Streaming TV—Yes:** This feature has a small contribution toward increased churn risk.
- **Gender—Male:** Gender has little influence on churn predictions in this model.
- **Tech Support—No internet service:** This category has little influence on predicted churn risk.
- **Partner—Yes:** Having a partner has a modest, mixed influence on churn predictions.
- **Online Backup—Yes:** Having online backup is generally associated with a lower predicted churn risk.
- **Senior Citizen—Yes:** This feature has a limited effect, with SHAP values mostly near zero.

4. RESULTS & DISCUSSION

4.1 Predictive Performance


Using balanced class weights and a 0.40 decision threshold improved the model's ability to identify churners. On the test set, the model achieved 79.49% accuracy and 65.50% churn recall, correctly identifying 262 actual churners.

4.2 Prescriptive Analytics & Risk Tier Distribution

The model's predicted probabilities were converted into three practical risk groups for the 1,409 test customers. A total of 996 customers (70.7%) were classified as Low Risk, 193 (13.7%) as Medium Risk, and 220 (15.6%) as High Risk. This segmentation helps focus retention spending on the 413 customers most likely to leave.

- **Probability ≥ 0.60 :** High Risk (Immediate Intervention)
- **Probability ≥ 0.40 :** Medium Risk (Targeted Discount)
- **Probability < 0.40 :** Low Risk (Safe)

The model identified 413 at-risk customers: 220 High Risk and 193 Medium Risk. Targeted outreach and suitable incentives for these groups can make retention efforts more efficient.



Key Data Breakdown

- **Low Risk (Safe) — 996 Customers (70.7%)**
 - **Criteria:** Churn probability < 0.40 .
 - **Business Action:** **Do nothing.** These customers are stable. Leaving them alone avoids unnecessary discount spending or unwanted spam.
- **Medium Risk (Targeted Discount) — 193 Customers (13.7%)**
 - **Criteria:** Churn probability between 0.40 and 0.59.
 - **Business Action:** **Nudge with minor perks.** These customers are on the fence. Re-engaging them with small incentives (e.g., a modest bill discount or contract extension deal) can convert them back into loyal accounts.
- **High Risk (Immediate Intervention) — 220 Customers (15.6%)**
 - **Criteria:** Churn probability ≥ 0.60 .

- **Business Action:** **Priority outreach.** These customers are actively heading out the door. The customer success or marketing team must intervene immediately with aggressive retention strategies (e.g., direct calls, substantial contract upgrades).

RECOMMENDATION

1. **Focused Retention Budget:** Rather than spending funds on marketing to all 1,409 members of the cohort, an organization aims the budget at simply 413 at-risk customers (Medium +High Risk).
2. **Tiered Resource Allocation:** Organizations should provide expensive retention provides solely to the 220 High Risk customers and low cost rewards to the 193 Medium Risk ones.
3. **Early Customer Intervention:** The model enables businesses to identify customers probable to leave before it occurs, therefore permitting initiative-taking retention measures instead than reactive ones.
4. **Personalized Retention Strategies:** Thanks to shape clarifications, firms learn what influences the predictions for each separate churn case. Therefore, interventions can be customized to customer-specific risk aspects as opposed to implementing the matching strategies to each and every customer.
5. **Reduction of Unnecessary Incentives:** While low-risk clients can be tracked, it does not mean that they need to be supplied with expensive retention offers. This helps minimize superfluous marketing costs, avoiding incentivization of customers whose predicted churn probability is relatively low.

5. FUTURE SCOPE

Future work could use longitudinal customer data to capture changing behaviour and identify churn earlier. Risk thresholds could be adjusted using retention costs and customer lifetime value. Further evaluation could compare XGBoost, AdaBoost, LightGBM, and CatBoost, combine SHAP with counterfactual explanations, and test the framework on other telecom datasets. A web dashboard could also support practical use and monitoring.

6. LIMITATIONS

Static cross-sectional dataset: The IBM Telco dataset represents one time period and cannot fully capture changing customer behaviour, time-based interactions, or economic changes.

Heuristic risk thresholds: The fixed 0.40 and 0.60 thresholds do not dynamically account for retention

costs, customer lifetime value, or campaign-specific budgets.

7. CONCLUSION

This study presents an end-to-end churn analysis pipeline that prioritizes recall for the minority churn class, explains predictions with SHAP, and converts probabilities into actionable risk tiers. The model achieved 65.50% churn recall and identified 413 at-risk customers for targeted retention. Future work should validate the approach on changing and diverse datasets and support deployment through an interactive dashboard.

8. REFERENCES

- 1) Oladapo, W. O. et al (2025). Explainable AI (XAI) Methods: Interpretability, Trust, and Applications in Critical Systems: A Systematic Literature Review. *International Journal of Computers*, 10, 303-311.
- 2) https://www.researchgate.net/publication/397016573_Explainable_AI_XAI_Methods_Interpretability_Trust_and_Applications_in_Critical_Systems_A_Systematic_Literature_Review
- 3) Khushi Kalasampath et al (2025). A Literature Review on Applications of Explainable Artificial Intelligence (XAI). *IEEE Access*. Volume 13, page No. 41111-41140. https://www.researchgate.net/publication/389457593_A_Literature_Review_on_Applications_of_Explainable_Artificial_Intelligence_XAI
- 4) EXplainable Artificial Intelligence (XAI); From Theory to Methods and Applications. *IEEE Access*, vol. 12, pp. 80799-80846. <https://ieeexplore.ieee.org/document/10549884> Shahin
- 5) Atakishiyev et al (2023). Explainable Artificial Intelligence for Autonomous Driving: A Comprehensive Overview and Field Guide for Future Research Directions. *IEEE Access*, 12, 101603-101625. <https://ieeexplore.ieee.org/document/10604830>
- 6) Philine Lou Bommer et al (2024). Finding the Right XAI Method; A Guide for the Evaluation and Ranking of Explainable AI Methods in Climate Science. *Artificial Intelligence for the Earth Systems*, Volume 3, Issue 3. <https://journals.ametsoc.org/view/journals/aies/3/3/AIES-D-23-0074.1.xml>