International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 09 Issue: 09 | Sept - 2025 SJIF Rating: 8.586 ISSN: 2582-3930

An Explainable AI Approach for Telecom Churn Prediction and Retention Strategy Design Using XGBoost and SHAP

[1]S. Pandikumar, [2]Mohammed Shoukath B H

[1]Dept. of MCA, Acharya Institute of Technology, Bengaluru-560107

[2]Dept. of MCA, Acharya Institute of Technology, Bengaluru-560107

Abstract: The problem of customer churn is a continuous challenge for profitability in the telecommunications sector, particularly given significant competition within the industry. Therefore, the ability to use advanced analytics to identify compressed subscribers in advance, and potentially retain them, is critical. This study describes an end-to-end explainable AI framework for predicting and intervening with telco churn. We used the IBM Telco Customer Churn dataset and simulated relevant behavioural features, so we could train our system via informative preprocessing (handling missing values and inappropriate encodings), feature engineering (where we devised aspects such as charge ratio, tenure, and service stability), and XGBoost modelling (where our modelling is optimized via stratified cross-validation and business-based target metrics) to monitor our predictions. Importantly, we provide SHAP-based explainability for every prediction, where our predictions contain both global and local explainability features to provide an actionable bridge between the power of algorithmic manipulation and business consideration for the customer. Our prediction offers the functionality of a business logic layer allowing churn probabilities to be mapped to risk formulations, loyalty tier definitions, and ways to intervene with the customer prior to their leaving as well as knowing later via a secure Flask-based web app with live analysis and audit of actions performed. Our experiment produced strong prediction values (AUC-ROC value of 0.95; retention of churn class recall up to 0.93), consistent with feature ranking providing interpretably relevant information, and identifiable space for practical business consideration. In conclusion, we have provided what we believe to be a scalable and transparent model for interpretable churn management with an aim of providing timely, informed decisions and accountability to operate the retention strategy in the telecommunications environment.

Index terms: Customer churn, Churn prediction, Telecom retention strategy, Explainable AI, XGBoost, SHAP (SHapley Additive exPlanations).

1. INTRODUCTION

1.1. Background and Motivation

The telecommunications industry is characterized by a high degree of competition, as customers are free to switch providers. Customer churn—when subscribers stop subscribing to their services—has many negative impacts on a company's revenue and profitability. Retaining existing customers is cheaper than adding new customers; research shows that the cost of acquiring a new customer can be six times more than retaining a customer. Predicting churn ultimately requires the capability of accurately predicting that customers will churn, allowing operators to make decisions to target subscribers that may churn with specific retention campaigns. Predictive analytics generally began with the application of classification algorithms in the telecommunications industry, including Logistic Regression, Support Vector Machine (SVM), Random Forests, and Gradient Boosted Trees [1].

1.2. Technology Developments in Churn Prediction

Telecom churn datasets are usually high dimensional, noisy, and unbalanced. Due to these complexities, accurate prediction of churn can be difficult. As research developed in this area, telecom researchers moved from using traditional statistical methods to utilizing modern machine learning solutions that can be applied to handle these complexities. Current approaches combine the use of data preprocessing techniques — including adaptive kk-means clustering — with supervised learning methods. The result is a supervised learning method that is able to address class imbalance and exemplify improvements in minority-class recall. Research has shown that the improvements in accuracy have led to improved dataset quality for more accurate and reliable prediction in real-life situations[2].

1.3. Research Base and Comparative Methods

Supervised learning will remain the foundation for customer churn prediction mainly because it takes advantage of labelled historical data pulled from CRM systems, billing records, and logs of customer usage. These are significant historical, structured datasets that provide information that allows models to learn mappings between customer behaviours and attributes and actual churns. From comparative works, the benefits of algorithms, for example, Decision Trees, Random Forest, SVM and Logistic Regression, come heavily as data specific tuning and feature engineering fit the model to features in the particular telecom data domain. Moreover, ensemble methods, mainly Random Forest, are the preferred methods in telecom churn analysis as they are more robust to complex nonlinear interactions and typically limit overfitting by aggregating results from many decision trees. Ultimately this leads to the more consistent and accurate churn predictions of customer churn and the significant use of ensemble models within the telecom churn prediction market[3].

1.4. Advances in Feature Engineering and Model Enhancement

Recent developments highlight the importance of grounded feature engineering and hybrid classification approaches. We can incorporate attributes that are relevant to the domain such as tenure, contract type, billing practices, and usage figures for services to improve model precision and recall. We have reported examples with Decision Tree, Random Forest, K-Nearest Neighbours (KNN) where decisions to incorporate custom features derived from domain knowledge resulted in significant improvements in all of the customer segments on prediction capacity. Including interaction term attributes and time-related behavioural features allows models to estimate finer changes in customer usage behaviour that can normally signify potential churn. When applied in an ensemble or hybrid learning framework richer features are not only attached with improved predictive performance , but they strengthen generalizability across a set of situational circumstances. This allows churn prediction systems to be robust and responsive in the telecommunication environment, which is volatile and changing quite rapidly. [4]

1.5. Ensemble Learning and Suggested Explainable AI framework

Because of the heterogeneous behaviour of telecom subscribers, ensemble learning is a main strategy for robust churn predictions. Using base learners such as Gradient Boosting Machines, Random Forest, and neural networks, ensemble learning provides more stability in classifications and better minority churn detection. Adding clustering before classification allows the model to be run on pre-classification clusters, thus modelling the segments, which enhances the predictability and value of churn for decision-making in business[5].

Using these ensemble methods, the study has developed an Explainable AI (XAI)-driven churn prediction and retention framework that extends the current approaches to ensemble learning through model explainability. The framework developed will combine XGBoost learner with curated domain feature engineering as charge ratio, tenure value, and service stability, and will produce SHapley Additive exPlanations (SHAP) for explainability for the churn prediction modelling on a totalized global model level and with the individual customer level. Each end product of the churn prediction modelling will have attrition across features output to help business stakeholders better understand the churn score calculated for their customer base and individual customers. The churn score can be mapped to risk categories and loyalty tiers for specialized retention programs. The framework,

International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 09 Issue: 09 | Sept - 2025 SJIF Rating: 8.586 ISSN: 2582-3930

built into a flask based web application, has been implemented in either a real-time analytics, batch processing, or KPI tracking useful for business analytics and decision making for predictive power.

2. LITERATURE SURVEY

The field of customer churn prediction in telecommunications has seen a considerable volume of literature utilizing a variety of modelling techniques and data pre-processing techniques. Initially researchers employed deep learning techniques for the prospects of high accuracy with appropriate data pre-processing techniques. For example, a Multilayer Perceptron (MLP) deep learning model was augmented with preprocessing techniques such as Tukey's Fences and SMOTE combined with feature engineering and ultimately exhibited 92% accuracy with enhanced interpretability from removal of insignificant factors like gender[6].

After accuracy, there has been an emphasis on interpretability for use in real-world applications, interpretability is key for practical implementation of churn models. A recent study examined the applicability of SHAP (SHapley Additive exPlanations) on LightGBM models and found prominent features to account for in actionable retention strategies e.g., tenure and contract type. This paper serves to create a complementary resource in predictive performance by ensuring that model decisions are transparent and trustworthy[7].

Rather than purely reflect classification accuracy and explainability, survival analysis methods were engaged to frame the problem of churn within logistic regression and Kaplan–Meier and Cox Proportional Hazards models. This dual framework not only examines the likelihood of churn but also allows for characteristics describing churn timing, presenting opportunities for analysing customer behaviours and using retention methods applied to units of uncertainty[8].

In consideration of customer heterogeneity a hybrid approach combining K-means clustering for segmentation and Random Forest classifiers was suggested. Their ability to assess customer segments separately improved precision and accuracy as they provide an additional unsupervised learning layer to our original supervised prediction pipelines[9].

As a complement to the previous work on classification algorithms the authors used a cross-validated deception study to compare traditional models with machine learning on churn prediction. Gradient Boosting and XGBoost yielded the highest overall performance, leaving us with an adequate balance for accuracy and computational efficiency in a real-world telecom environment. Their findings reflect the trade-offs made between supervised machine learning models being more complex for predictive potential and interpretability [10].

As the authors included data balancing methods including Random Oversampling and SMOTE, along with creating features driven by domain expertise focusing on the preprocessing steps needed to obtain the best prediction possible. Preprocessing also enhances recall and area under the curve (AUC), which is one of the most important performance metrics for Gradient Boosting. Data engineering also needs to be acknowledged, regardless of the complexity of the model placed on the engineered data [11].

In a similar manner of emphasizing data quality, a data driven pipeline specifying data cleaning, correlation tests for multicollinearity has allowed more simplistic models such as a Random Forest and Logistic Regression to produce competent results. Ultimately stressing the role of data preparation as it relates to the management of predictive ability [12].

More evidence of the value of feature engineering was also obtained through the application of more advanced Relief-F feature selection techniques in conjunction with CART, and ANN classifiers. While the ANN model did have exceedingly strong dimensionality reduction and therefore, higher reported accuracy and precision than the CART model, this only further supports the assertion that feature selection imbedded in neural methods can yield great additional perspectives[13].

An optimized weighted ensemble of the K Nearest Neighbours, XGBoost, and Random Forest models outperformed any of the single models and all standard ensemble models with their algorithmically optimized ensemble weights.

Although this method did not provide gradient based optimization, it did yield consistent performance improvements on a comparatively heavy preprocessing, class imbalance, and explainer framework[14].

The XGBoost framework in conjunction with the Genetic Algorithm (GA-XGBoost) - sampling based hyperparameter optimization approach with ADASYN for handling class imbalance demonstrated high accuracy and AUC as well as increasing the interpretability of the model. It was the clarity of the churn predictors and their relationship together with SHAP that made the result of predictive accuracy and observability resonate the strongest[15].

From a business value perspective, we created a model that included dimensions of customer value and service attributes, investigated churn and its association and dependency with ARPU, complaints and bundled contracts, acted to define retention and win-back strategies for high-value segments[16].

Practical model evaluations on a Ukrainian telecommunications data collected for consumers produced C5.0 and Neural Net Ensemble models with perfect accuracy and AUC. We found predictive features for churn like average call duration, and a robust share of calls to other operators, which enables scalable or operational model deployment, but required regular model updating[17].

Lastly, we leveraged a big data infrastructure, and in-depth Social Network Analysis (SNA) feature use to significantly raise performance of a churn prediction model using XGBoost. This model moved to production deployment and has allowed us to proactively act on retention strategies that delivered longer customer capture (churn) times, and increased revenue [18].

In summary, the literature review has revealed that, to date, almost all studies have ultimately only covered disconnected parts as it relates to churn prediction – accuracy, interpretability, or preprocessing. Very few studies have developed an end-to-end deployed churn prediction system that brings the aspects together into the strengths of all three areas. My study is an example of a study that does this and implemented Agile feature engineering with domain knowledge and sophisticated ensemble learning using XGBoost for complex, imbalanced telecommunication data. My study also utilized SHapley Additive exPlanations (SHAP) for a transparent and interpretable churn prediction system and mapped the churn score outcomes to actionable, meaningful business categories. The system is also delivered through an interactive Flask application that connects technical innovations to business actions and creates a platform for actionable real-time strategic interventions to help retain customers.

3. METHODOLOGY

The proposed solution is an end-to-end, explainable AI framework to improve the predictive capabilities of customer churn within the telecommunications industry and enhance retention. It encompasses enriched customer datasets—such as your team's behavioural measurements (e.g., payment history, customer support, and service changes)—and sophisticated feature engineering, to uncover the underlying risk patterns. A combination of feature engineering and enriched datasets comprises the modelling backbone, which uses an XGBoost classifier trained on normalized, encoded datasets, with a measure of model transparency by leveraging SHAP-based explainability for global and individual risk factors. The predicted churn probabilities were then mapped into actionable components for the business, by ultimately leading to retention scoring, retention offers, and customer tiering using an automated business logic layer. Lastly, the entire workflow is deployed in a secure, interactive web application that allows business users to input customer details, create-geco atm (automatically generated, and least take) policies with real-time churn risk probabilities and interpretability visualizations, and keep a check on performance and retention through dashboard analytics with persisted data storage - while everything remains safe.

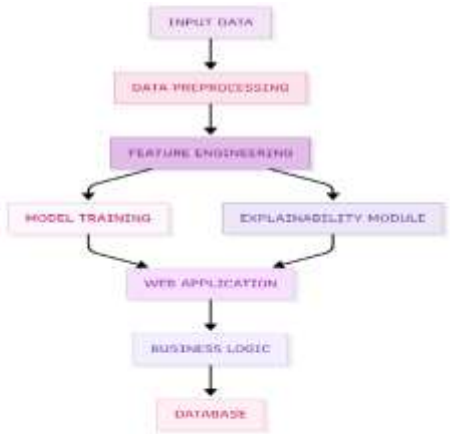

Fig 1. System Architecture and Design

3.1. Data Collection and Enrichment

The primary source of data for this work is the IBM Telco Customer Churn dataset, whose suitability as a good representation of real-world telecom customer data has caused it to be widely adopted in the professional literature of predictive analytics. The dataset describes demographics for over 7,000 customers, and it includes demographic characteristics (e.g. gender, age, and senior citizen), subscription types (e.g. contract duration, internet and phone), monthly and total charges, tenure, and the historic churn labels. These characteristics support foundational churn modelling based on many types of customer attributes and service consumption.

To better mimic the realities of operational practice in today's telecommunications environment, the base dataset was further enriched with synthetic behavioural characteristics that have the potential to simulate such ongoing customer behaviours and signals of potential risk. Specifically, new variables that provide for the number of late bill payments per customer, the total number of customer support requests and formal complaints and how recently they were made, and the rate of service plan changes or subscribing to an add-on service. These behavioral characteristics were synthesized using probabilistic sampling, with the specific industry incidence rates during the creation of the variables ensuring that the behaviour characteristics were relevant to the statistical structure and organization of the data, while ensuring variability across all customers in the base dataset.

By integrating these active behavioural layers with the preceding static dimensions, now we have a more full-bodied and actionable stamp of each customer. The full data frame allows for the downstream predictive model to recognize underlying indications of disengagement and intent of churn (for instance, revealing behaviours that would not have been picked up in a transactional only dataset). Overall, this approach is more representative in driving better generalization, and contextualizing customer experience beyond snapshot predictions, but to illustrate an ongoing experience with the customer.

3.2. Data Preprocessing

The enriched telecom customer dataset underwent extensive preprocessing as a pathway to model training, ensuring both data quality and the robustness of analytics. The first steps of the process involved assessing and resolving missing values, in addition to removing any duplicate records, to provide a clear, consistent foundation for subsequent operations. The numeric features representing tenure, monthly charges, total charges and charge ratio, as well as all simulated behaviours, were power transformed to normalize the features and reduce skewness in feature distributions, creating more symmetric distributions and lending to more stability and accuracy of prediction by our model.

The categorical variables: contract type, payment method, internet service, gender, and other customer characteristics were one hot encoded. One hot encoding converts each possible categorical value into separate binary columns for processing, allowing the machine learning model to identify categorical information without arbitrary ordinal assumptions or bias.

While the model addresses the issue of class imbalance directly during model training, via class weights built into the XGBoost algorithm; preprocessing will ensure both the feature sets are completely and accurately established, to prepare for processing. The above operations collectively convert raw data and behavioural datasets into a noise reduced, normalized, and model ready dataset, thereby enhancing the accuracy and usability of feature information available.

3.3. Feature Engineering

To enhance the predictive power and interpretability of the churn prediction model, a consistent feature engineering approach was taken across both the original attributes of the consumers, and the newly simulated behavioural facets. Raw data were transformed into business metrics that were meant to reveal hidden customer risk patterns and provide insights for action. Important engineered features are the charge ratio which provided some contextual understanding on consumer value based on average spend over subscription tenure so that high- and low-value subscribers could be differentiated, the tenure value which is a straightforward interpretation of loyalty and engagement metrics, since higher tenure values, generally indicates less risk of churn, and service stability which was defined as how frequently service plans were altered or service was disrupted, with higher service instability potentially inferring a negative experience or dissatisfaction.

Moreover, variables were created to capture high-risk managed trading contract and managed trading payment flags which allow identification of customers that use either types of contracts or payment methods that historically have worse churn risk (, e.g. month-to-month contracts; and manual billing). Autopay status was included because customers that use autopay, on average, have a better retention history. Other behavioural variables like: total late payments; customer support interactions; and modifications to service were included to provide more of a dynamic view of ongoing engagement or happiness. The choice and development of the variables considered our telecom domain understanding and iterative modelling, including checks for multicollinearity to verify each variable had a unique contribution and utility. By also including very well designed potential variables allowed the model to become more aware of arbitrary and nuanced signals of churn risk; and will likely yield more accurate predictions and an enforceable explanation or foundation to support data driven retention strategy design.

3.4. Model training with XGBoost

XGBoost (Extreme Gradient Boosting) is the predictive engine that is involved in our framework and utilizes a robust implementation of gradient boosting for structured/tabular data. XGBoost builds an ensemble of decision trees sequentially while training each tree to the learn the residuals of previous tree ensemble prediction, thus creating a better predictive model. This combination of sequentially learning and predictions make the algorithm more computationally efficient by making the ensemble learn structure, while also reducing the risk of overfitting with both parallelized training and shrinkage via two regularization term.

The output for a specific observation is the summation of output from many trees, mathematically; this can be represented as:

$$\hat{y}_i = \sum_{k=1}^{K} f_k(x_i)$$

where $\hat{y}_i$ is the predicted outcome for the observation i , K is the number of trees and $f_k(x_i)$is the prediction of k-th tree[19].

The below figure provides an illustration of how XGBoost works including the input data pre-processing, sequentially boosting weak learners (increasing predictive performance), and the final aggregation of outputs that combine the predicted outputs to provide more robust churn predictions.

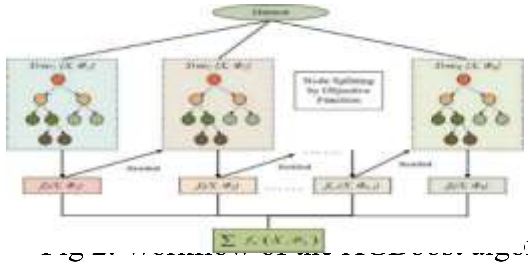

Fig 2. Workflow of the XGBoost algorithm showing sequential boosting and aggregation[20].

3.5. Model Explainability with SHAP

SHAP (SHapley Additive exPlanations) was used to support model interpretability, which was built into the XGBoost inference process. After every churn prediction, SHAP values are computed based on the TreeExplainer module, which takes advantage of the inner workings of the trained XGBoost model to compute feature attribution values. When applied to a single input instance, the SHAP algorithm will return a vector where each entry corresponds to the marginal contribution of each of the features to the model's churn probability prediction.

In practice, local SHAP values are generated for each individual customer prediction and the top positive and negative feature contributions that achieved that prediction can be identified. Global feature importance is calculated on the sample population, in summary statistics, in order to reflect, on average, which of the variables potentially contributes to higher or lower churn risk.

The system displays these SHAP values as bar plots (using Matplotlib or Plotly libraries) and the display is a part of the web application results page. In this way, users can readily refresh and view not just the predictive score, but also have available to them a clear, ranked explanation for the drivers of each prediction. The explainability workflow is fully automated, as after inference is conducted, SHAP computations and visualizations are automatically executed in real time, meaning any latency to the end user is negligible. Incorporating this explainability supports a complete view of transparent, traceable churn predictions and recommendations to act on.

3.6. Integrating Business Logic

To make operational use of model outputs, the churn probabilities predicted by XGBoost are mapped to discrete risk levels - High, Medium or Low - using defined probability thresholds predetermined for business relevance. Each risk level is also linked programmatically to context-dependent retention offers - in other words, high risk customers can be identified within the high-risk segment for urgent intervention, specific incentives or activity tailor. Conversely, low risk customers will be monitored for continued satisfaction.

In addition to risk, the system maps risk into a networked loyalty scoring system whereby, each customer's churn score, tenure value and behavioural measures are aggregated to yield a loyalty index. Customers are ranked on the loyalty index (e.g. Platinum, Gold, Silver, Bronze) by value as a means of further segmenting retention programs.

This layer will also compute other key performance indicators such as customer lifetime value estimate, recent engagement score, described in order to assist in prioritizing customer contact and operational workload for retention workers.

All business logic mapping and decision support outputs are presented back to the end user within the web app interface, along with the core churn predictions and SHAP interpretability results, ensuring that every prediction is also qualified by a factual data driven business recommendation.

3.7. Deployment and System Integration

In this phase, we operationalized the entire churn prediction pipeline, from data preprocessing, feature engineering, XGBoost inference, SHAP explainability, and business logic into a secure and interactive web application built using Flask. The XGBoost model, the preprocessing workflow, and SHAP explanation workflow were all put into containers and deployed to the backend so that the predictions and explanations are reproducible and consistent.

As users, administrators log in to this system using an authenticated web interface to either input customer data manually or upload files for batch processing. Once any user workflow was initiated, the application took a user through executing predictions of real-time inference. After the dataset was submitted, the application predicted churn probability, risk categorization, and provided SHAP visualization plots, and provided recommended retention strategies back to the user interface. The application automatically persisted the results in an SQL database, along with the key business metrics, which enabled business analytics dashboard capabilities, and preserved audit functionality for the application. Finally, the shoebox deployment solution provided the ability to capabilities for long-term monitoring of the model for drift, as well as the ability to scale the solution to ingest and consume live streams of customer information, if required and desired. The operationalized infrastructure, thus bridged the gap between analytics and business action in a transparent, actionable, and scalable approach to telecom churn for Organization enterprise users.

4. RESULT ANALYSIS

4.1. Model Development Workflow

The proposed system includes extensive data pre-processing and simulated behaviour characteristics enrichment, data transformation, model training using an XGBoost classifier. We applied a power transformation on the input data to circumvent skewness where applicable and did one-hot encoding on categorical variables of high cardinality for stability in our data and for the model. The pipeline included domain-driven cost-sensitive business rules interpreting risk based on company retention objectives with a consultative loss matrix. Behavioural variables like late payments, support contact, service downtime, plan-change behaviour were simulated to represent nuanced risk indicators in addition to the primary demographic and transactional data. Meaningful feature engineering introduced new performance matrices like charge ratio, tenure value, and service stability to create further predictive representation in the dataset and contextualize the outcome of the supervised learning algorithm. Due to the modular nature of the development, SHAP explainability, real-time inference, and an automated dashboards and reports were easily integrated producing a transparent churn management plan and a product ready to be put live.

4.2. Model Performance Evaluation

Model evaluation was done using a stratified 5-fold cross-validation protocol which ensured the folds would have the same approximate group distributions of churners and non-churners. This resulted in a mean area under the receiver operating characteristic curve (AUC-ROC) value of 0.84 ± 0.00 indicating the model was able to differentiate between customers that are likely and unlikely to churn with good accuracy.

4.2.1. Default Threshold Performance (0.5)

At the default classification threshold of 0.5—and where customers with predicted churn probabilities greater than 0.5 will be labelled as churners—the XGBoost model achieved an accuracy of 0.86 and strong AUC-ROC score of 0.95. Performance evaluation metrics for each class can be seen in Table I. The model had strong recall (0.93) performance for the churn class (which is more important than precision as we want to minimize missed churners) and had moderately precision (0.66). The other class of no-churn had very high precision (0.97) but lower recall (0.83), which demonstrated the model learned nominally well overall.

Table 1: Results of classification at threshold 0.5

Class	Precision	Recall	F1-score	Support
No Churn (0)	0.97	0.83	0.89	5,163
Churn (1)	0.66	0.93	0.77	1,869

Macro avg (F1): 0.83 Weighted avg (F1): 0.86

4.2.2. Threshold Optimization Performance (0.584)

In order to further improve the model performance to meet business needs, the classification threshold was optimized based on the maximum F1-score to create a threshold of 0.584. The new threshold increases the accuracy to 0.88 where recall (0.87) and precision (0.73) for churn customers improved as well to create a better trade-off between identifying and contacting at-risk customers while not contacting non-risk customers unnecessarily.

Table 2: Classification results at the optimized threshold 0.584

Class	**Precision**	**Recall**	**F1-score**	**Support**
No Churn (0)	0.95	0.88	0.91	5,163
Churn (1)	0.73	0.87	0.79	1,869

Macro avg (F1): 0.85 Weighted avg (F1): 0.88

4.2.3. Business Impact Analysis

We further assessed the relevance of these classification outcomes through a business impact analysis. In this analysis, each false positive (wasted retention effort) incurred a loss of $100; and each false negative (missed churner) resulted in a loss of $500.

In many cases where optimizing the threshold increases true positives (providing a benefit) for the churners, enhances recall (remedying another cost) now means a larger intervention cost because of the increased false negatives. This numeric analysis presents a tangible illustration of the important trade-off faced by practitioners between costs incurred by the operational side of things, and the business need to proactively engage with customers most likely to churn.

4.3. Explainability and Important Features

Feature importance analysis showed contract-related variables and engineered behavioral features or attributes contribute strongly (as expected) to churn risk prediction. Here are the top ten features ordered by model importance:

Rank	Feature	Importance
1	Contract_Two year	0.190710
2	HighRisk_Flag	0.118069
3	Service_Stability	0.080062
4	Contract_One year	0.052120
5	InternetService_Fiber optic	0.047063
6	InternetService_No	0.045516
7	StreamingMovies_Yes	0.027017
8	OnlineSecurity_No internet service	0.024549
9	PhoneService_Yes	0.020896
10	MultipleLines_No phone service	0.020853

This transparent feature ranking is consistent with intuition around running a telecommunications business, since customers with longer term contracts, or with stable usage service attributes, are less likely to churn and negative behavioural signals, such as having a high risk payment profile, increases the probability to churn.

4.4. Customer Segmentation and Insights for Retention

The deployed solution provides the ability to operationalize and execute the model output by placing customers in assigned groups (or tiers) related to churn risk levels, such as low, medium, and high risk, based on their absolute proportionate output for individual churn risk probabilities. Each risk tier indicates the level of urgency associated with retention activity, which varies in type, depending on the risk for the customer. High-risk customers may elicit immediate personalized offers, while low-risk groups may be monitored to ensure they remain satisfied.

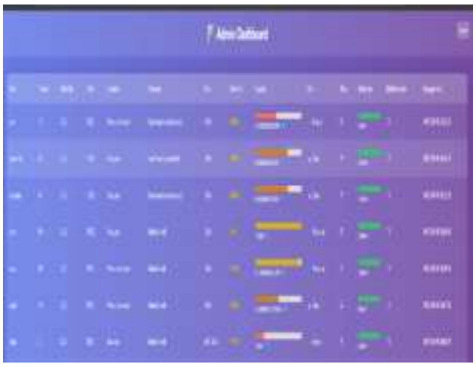

Fig 3. Admin dashboard with model prediction of churn

As depicted in Figure 4, the admin dashboard displays a summary of churn probability, loyalty scores, customer tier group (e.g., Bronze, Silver, Platinum), and retention status by individual users in the system. The admin dashboard displays other useful fields (tenure, payment method, lifetime value, etc.) for their business context; it allows the team members to prioritize their retention activity accurately. The dashboard allows managers to view, compare and strategize across customer groups with insight, while making data-driven decisions to improve customer retention, minimize customer churn, and obtain the most retention ROI (return on investment).

4.5. Real-Time Explainability and Action at the Customer Level

Beyond group-level segmentation, the implementation of the deployed platform offers detailed real-time explainability for individuals customers. Once a retention specialist or business analyst selects an individual customer record, the system provides the retention specialist with a detailed intelligence summary about that individual customer and the models' churn prediction, risk designation (low, medium, high), risk and retention scores, estimated lifetime value, and engagement score determined by the current data inputs.

Fig 4. Individual Customer Intelligence Analysis

Most relevant to note, the interface also features a SHAP based, feature impact plot which displays the top features influencing that customers, churn prediction. For instance, and as shown in Figure 5, "Service_Stability" and the contract period were the main drivers which accounted for a customer's churn risk. By presenting risk drivers in this way, the platform enables business users to understand and justify model decisions, transparently convey risk factors, and develop individualized retention plans with knowledge and certainty.

This instance-level explainability creates an ecosystem that builds trust in automated predictions and enables equitable, strategic action, as well as helping comply with and facilitate regulations around transparency.

5. CONCLUSION AND FUTURE WORK

This project has created a scalable and interpretable AI model to predict churn by telecom customers that achieves maximum predictive accuracy and trustworthiness of predictions. This framework contains systematic applications of feature engineering, power-transformed normalization, one-hot encoding, and XGBoost to measure customer churn risk across customer segments. The SHAP-based explainability provides transparent explanations of predictions that overcome trust and confidence limitations typical of black-box models, and eliminates friction points for industry adopters. Running as a secure, interactive web application provides the advantage of interfacing with live analytics in real time, generating batch scores, and translating risk scores with actionable strategies for retention. We will complete business impact assessments on adoption of our AI structured to offer a consistent framework for reasonable decisions on trade-offs between recall, precision, and cost of intervention. Persistent logging, an easily used dashboard and individual explanations offer the potential for final learning and appropriate regulation of risk-and-return for retention strategy. In the future, the platform can be extended to allow for real-time streaming data, continual auto-learning regarding automated retraining, and higher level explainability methods including counterfactuals and natural language explanations. The relevance through A/B testing within the actual telecom space would validate business impact as well as set the calibration thresholds for proactive outreach-work. There are further opportunities for unsupervised clustering for personalized retention, horizontally scalable cloud-based APIs to increase reach, and fairness audits to combat bias. Taken together, these advancements will position the framework to be highly competitive, if not the best in industry, to create sustainable customer loyalty and market differentiation.

REFERENCES

[1] Gaur, A., Dubey, R. (2018). Predicting Customer Churn Prediction In Telecom Sector Using Various Machine Learning Techniques. 2018 International Conference on Advanced Computation and Telecommunication (ICACAT).

[2] Aravinda, N., L., Reddy, R., A., V, B., D., Nayak, J., P., R., & Husseen, A. (2024). Customer Churn Prediction in the Telecom Sector with Machine Learning and Adaptive k-Means Cluster using Imbalance Data. 2024 International Conference on Distributed Computing and Optimization Techniques (ICDCOT).

[3] Pandithurai, O., Ahmed, H., H., S, H., N., Sriman, B., & R, S. (2023). Telecom Customer Churn Prediction Using Supervised Machine Learning Techniques. 2023 International Conference on Advances in Computing, Communication and Applied Informatics (ACCAI).

[4] Verma, G. (2024). Telecom Customer Churn Prediction using enhanced Machine Learning

Classification Techniques. 2024 First International Conference on Innovations in Communications, Electrical and Computer Engineering (ICICEC).

[5] Liu, Y., Fan, J., Zhang, J. *et al.* Research on telecom customer churn prediction based on ensemble learning. *J Intell Inf Syst* **60**, 759–775 (2023).

[6] Rabih, R., Sun, W., Ayoubi, M., & Jamal, K. (2024). Highly Accurate Customer Churn Prediction in the Telecommunications Industry Using MLP. *International Journal of Integrated Science and Technology*, 2(10), 931–946. https://doi.org/10.59890/ijist.v2i10.2564

[7] Noviandy, T. R., Idroes, G. M., Hardi, I., Afjal, M., & Ray, S. (2024). A Model-Agnostic Interpretability Approach to Predicting Customer Churn in the Telecommunications Industry. *Infolitika Journal of Data Science*, 2(1), 34–44. https://doi.org/10.60084/ijds.v2i1.199

[8] Denayer, J. (2024). Customer Churn Analysis in Telecom Industry using Kaplan–Meier and Cox Proportional Hazards Model. *Indian Scientific Journal Of Research In Engineering And Management*, 08(06), 1–5. https://doi.org/10.55041/ijsrem36121

[9] Ouf, S., Mahmoud, K. T., & Abdel-Fattah, M. A. (2024). A Proposed Hybrid Framework to Improve the Accuracy of Customer Churn Prediction in Telecom Industry. *Journal of Big Data*, 11, 1–27. https://doi.org/10.1186/s40537-024-00922-9

[10] Wu, S. (2024). Customer Churn Prediction in Telecom Based on Machine Learning. *Highlights in Science Engineering and Technology*, 94, 113–118. https://doi.org/10.54097/snc09915

[11] Sana, J. K., Abedin, M. Z., Rahman, M. S., & Rahman, Md. S. (2022). A Novel Customer Churn Prediction Model for the Telecommunication Industry Using Data Transformation Methods and Feature Selection. *PLOS ONE*, 17(12), e0278095. https://doi.org/10.1371/journal.pone.0278095

[12] Zhang, T., Moro, S., & Ramos, R. F. (2022). A Data-Driven Approach to Improve Customer Churn Prediction Based on Telecom Customer Segmentation. *Future Internet*, 14(3), 94. https://doi.org/10.3390/fi14030094

[13] Abdulsalam, S. O., Arowolo, M. O., Saheed, Y. K., & Afolayan, J. O. (2022). Customer Churn Prediction in Telecommunication Industry Using Classification and Regression Trees and Artificial Neural Network Algorithms. *Indonesian Journal of Electrical Engineering and Informatics*, 10(2). https://doi.org/10.52549/ijeei.v10i2.2985

[14] Khoh, W. H., Pang, Y. H., & Ooi, S. Y. (2023). Predictive Churn Modeling for Sustainable Business in the Telecommunication Industry: Optimized Weighted Ensemble Machine Learning. *Sustainability*, 15(11), 8631. https://doi.org/10.3390/su15118631

[15] Peng, K., & Peng, Y. (2022). Research on Telecom Customer Churn Prediction Based on GA-XGBoost and SHAP. *Journal of Computer and Communications*, 10(11), 107–120. https://doi.org/10.4236/jcc.2022.1011008

[16] Zhao, M., Zeng, Q., Chang, M., Tong, Q., Su, J., & Su, J. (2021). A Prediction Model of Customer Churn Considering Customer Value: An Empirical Research of Telecom Industry in China. *Discrete Dynamics in Nature and Society*, 2021, 1–12. https://doi.org/10.1155/2021/7160527

[17] Fareniuk, Y., Zatonatska, T., Dluhopolskyi, O., & Kovalenko, O. (2022). Customer Churn Prediction Model: A Case of the Telecommunication Market. *Economics (Bijeljina)*, 10(2), 109–130. https://doi.org/10.2478/eoik-2022-0021

[18] Ahmad, A. K., Jafar, A., & Aljoumaa, K. (2019). Customer Churn Prediction in Telecom Using Machine Learning and Social Network Analysis in Big Data Platform. *arXiv: Computers and Society*. https://doi.org/10.1186/S40537-019-0191-6