Explainable Optimized Machine Learning for Predicting Customer Churn

Konstantinos Zervoudakis[1[0000-0002-2146-1493]] and Stelios Tsafarakis[2[0000-0001-5535-4787]]

School of Production Engineering and Management, Technical University of Crete, Chania, Greece
[1]kzervoudakis@tuc.gr
[2]tsafarakis@tuc.gr

Late Breaking Abstract

Customer churn refers to the case where customers stop using the products or services of a company. High churn rates can lead to great losses in revenue and market share. Since acquiring new customers is more expensive than retaining existing ones, predicting customer churn has become a primary focus for many corporations. Machine Learning (ML) offers powerful tools to analyze large databases and identify patterns linked to customer churn (Bansal et al., 2022). By building predictive models, companies can identify at risk customers and take targeted actions to improve retention. The aim of this work is to develop an explainable, optimized ML approach to accurately predict customer churn, aiming to support decision-making processes.

The main contribution of this work is the combination of ML with optimization techniques, specifically metaheuristics from the field of Computational Intelligence, for hyperparameter tuning and feature selection. The optimized model is then analyzed using SHapley Additive exPlanations (SHAP) to explain how different features affect the predictions (Kim & Kim, 2022). This approach helps researchers to build more accurate models for predicting customer churn, while also providing clear insights into the factors that cause churn. By combining ML, optimization and explainability, a practical tool that companies can use for both predicting customer churn and understanding their behavior is proposed.

We employ several ML classifiers to predict customer churn, including Logistic Regression (LR), Decision Trees (DT), Random Forest (RF), Gradient Boosting (GB), Support Vector Machine (SVM) and Neural Networks (NNs). The models were trained on a dataset retrieved from Kaggle, which contains detailed information about 7043 telecom customers in the United States, including their demographic details, subscription plans, and churn status. Cross-validation is also used to evaluate the performance of each model.

To improve models' performance, Multi-Objective (MO) optimization algorithms were used, including Particle Swarm Optimization (PSO), Non-dominated Sorting Genetic Algorithm II (NSGA II) and the Mayfly Algorithm. These algorithms simultaneously tune hyperparameters and select import features, aiming to maximize both the accuracy of the model and the F1 score. The performance of the MO optimizers was compared using widely accepted multi-objective performance metrics.

To investigate the impact of different features on predictions, SHAP analysis was used. SHAP values show the contribution of each feature to the model's output for individual predictions. SHAP is applied to the best-performing optimized ML model to identify the most important factors in predicting churn, making the model's decisions more transparent and interpretable for business applications.

The results show that all MO optimizers improved the performance of all ML models, with MO Mayfly Algorithm (MA) exhibiting superior dominance and faster convergence based on MO metrics. The optimized RF classifier emerged as the best model, achieving the highest accuracy and F1-score. Specifically, MO MA detected a set of hyperparameters and features that resulted in an accuracy of 0.82 and an F1-score of 0.83, which represents a significant improvement over the default RF, which had both metrics below 0.80.

The SHAP analysis provided clear insights into the factors that most influenced churn predictions. The beeswarm plot showed that contract type, tenure months dependents, internet service, and monthly charges were the strongest drivers of the model's output. Dependence plots revealed more details. For example, customers with shorter tenure and month-to-month contracts were more likely to churn, while longer contracts reduced churn risk. Higher monthly charges combined with Fiber optic internet service were also associated with increased churn probability. Waterfall plots were also generated to illustrate how individual feature contributions sum to the final churn prediction for each customer. These plots break down the prediction step by step, starting from the average model output and incrementally adding or subtracting the impact of each feature. This makes it easier to understand exactly why the model predicted a high or low risk of churn for a specific customer. Such detailed explanations can support case-by-case decisions, such as targeting at-risk customers with personalized offers.

This study demonstrates that combining machine learning with multi-objective optimization and SHAP explainability results in more accurate and transparent churn prediction models. By adopting this comprehensive approach, companies can more effectively predict which customers are likely to leave, while also understanding the key reasons behind these risks. This, in turn, enables the design of targeted strategies to improve customer retention and minimize losses.

References

Bansal, M., Goyal, A., & Choudhary, A. (2022). A comparative analysis of K-Nearest Neighbor, Genetic, Support Vector Machine, Decision Tree, and Long Short Term Memory algorithms in machine learning. *Decision Analytics Journal*, *3*, 100071. https://doi.org/10.1016/J.DAJOUR.2022.100071

Kim, Y., & Kim, Y. (2022). Explainable heat-related mortality with random forest and SHapley Additive exPlanations (SHAP) models. *Sustainable Cities and Society*, *79*, 103677. https://doi.org/10.1016/J.SCS.2022.103677