AI-Driven Churn Prediction for SaaS Businesses

Tushar Srivastava[1], Tushar kr. Sharma[2], Utkarsh Mishra[3], Varun Mishra[4], Ms. Kajal Gehlot[5]
Dept. Of Computer Science and Engineering (Artificial Intelligence and Machine Learning), GL Bajaj Institute of Technology and Management

Abstract—**SubAnalytics is an AI-powered subscription analytics tool for SaaS businesses. It uses machine learning models to predict customer churn, calculate churn metrics, identify high-risk subscribers and analyze subscription plan distributions. SubAnalytics is built with a modular architecture using a MongoDB-backed Express.js backend and a Python-based predictive microservice. Users can upload anonymized customer data via a secure web interface and the tool will generate visual and textual insights. This paper evaluates the performance of the prediction model, usability of the tool, and real-world applicability for improving customer retention strategies. The results show SubAnalytics can support data-driven decision-making in subscription-based businesses.**

Index Terms—**Subscription Analytics, Customer Churn Prediction, SaaS, Predictive Modeling, Business Intelligence, Express.js, MongoDB**

1. INTRODUCTION

Customer churn, or customer attrition, is the percentage of customers who cancel a service within a certain time frame. In subscription-based industries like SaaS even small changes in churn rates can have a big impact on revenue as lost customers mean lost recurring revenue. Retaining a current customer is generally many times less expensive than acquiring a new one so predictive churn analytics is key to business strategy. A churn event occurs when a customer deems a product unacceptable or an offer from a competitor more appealing; so ML driven forecasts allow us to proactively identify customers at risk of leaving and targeted retention strategies [2][4].

With the advancements in data collection and artificial intelligence it is now possible to process complex customer data for churn behavior. Machine learning and deep learning have been used in customer churn prediction across different fields with high accuracy in determining users who will cancel subscriptions [1] [3] [6][7][9]. Specifically, ensemble models and new algorithms like XGBoost or deep networks (e.g. BiLSTM-CNN) have outperformed traditional methods on benchmark datasets [6][8][10].

Most of the academic literature however addresses telecommunications or B2C environments; churn in B2B SaaS has different features (e.g. company sized subscriptions, long sales cycles) that need to be taken into account. This paper introduces SubAnalytics, a conceptual SaaS churn prediction and retention insight tool. We will discuss the theoretical framework for churn analytics, survey related work, outline the proposed methodology and system architecture and report experimental results. The goal is to present the basis for an ML driven churn and retention solution for SaaS scenarios. We will focus on conceptual design and data driven argumentation, not low-level implementation or specific technology stacks. SubAnalytics will help SaaS practitioners identify churn drivers and increase customer lifetime value through predictive modeling and actionable results [13]. Churn prediction has been an area of ML research for a long time, especially in subscription businesses and telecom. Research consistently shows that predictive modeling can pick out customers who are about to churn and enable proactive retention efforts [2][5][7][12]. Specifically, ensemble models are state-of-the-art according to research that shows XGBoost performs better than simpler classifiers after tuning with SMOTE or other balancing methods [4] [8].

Most churn analytics research has focused on B2C (e.g. retail, telecom) where churn is clear and data is abundant. In SaaS/B2B environments, churn is more complex (e.g. firm accounts with many user licenses) and features on data vary (e.g. contract length, usage quantity, support interactions). Recent theses on SaaS churn highlight this Dahlén and Mauritzon (2023) found that customer tenure, usage activity and project involvement were significant churn predictors in a B2B SaaS dataset. They also note that churn datasets are imbalanced (few churners and many non-churners) and need to be treated with care, e.g. SMOTE or custom sampling [11]. Where each signifies its own importance how they have contributed in the field of predictive analysis.

2. LITERATURE SURVEY

Customer churn forecasting has been at the forefront of machine learning study because of the relevance to subscription-based businesses. Although most initial research was oriented towards telecommunications, more recent investigations have extended the use of churn analysis to SaaS platforms where repeated revenue and long-term user activity are imperative to viability.

A. Classical Approaches

Early churn prediction models used statistical methods like logistic regression and decision trees [1]. These models were explainable but weak in dealing with nonlinear relations and intricate feature interactions. Brown and Green (2007) discovered that model performance in conventional approaches frequently depends greatly on the quality of preprocessing and manual feature engineering [2].

B. Rise of Ensemble Learning

With the evolution of ensemble methods, models such as Random Forest and XGBoost started dominating classical models in churn problems. Breiman's unveiling of Random Forest [3] provided a better performance by fusing the output of an ensemble of decision trees. Nikou et al. (2019) compared machine learning algorithms with traditional approaches and deep learning methods and concluded that deep learning methods were generally outperformed by ensemble models like Random Forests for structured churn datasets [5].

C. Deep Learning Approaches

Deep neural networks were recently proposed for churn modeling. Mehtab and Sen (2020) utilized LSTM (Long Short-Term Memory) networks to model temporal patterns of user behavior, achieving better performance for time-sensitive churn predictions [6].

D. Churn in B2B SaaS Context

Most of the published work targets B2C settings, where churn is common and user behavior is one-off. B2B SaaS churn, though, presents different challenges. As Dahlén and Mauritzon (2023) explain, customer accounts tend to be representative of organizations rather than individuals, so models need to factor in variables such as contract duration, breadth of usage across teams, and active license count [8].

Martin (2024) stressed that transactional features such as monthly recurring revenue (MRR), net promoter scores (NPS), and feature usage frequency were essential markers of churn in SaaS startups [9]. They also identified that imbalance in data (i.e., significantly smaller numbers of churners compared to non-churners) complicates model training and necessitates the use of methods such as SMOTE (Synthetic Minority Over-sampling Technique) or cost-sensitive learning.

To highlight the operational differences, *Table 1* presents a comparative analysis between traditional churn detection methods and predictive model-based approaches, emphasizing the benefits of automation, scalability, and proactive intervention. This analysis is compiled from industry practices and insights gathered through literature review [1][4][6][7].

Table 1 - Comparative analysis between traditional method and predictive analytics

Aspect	Traditional Method	Predictive Analytics (Model-Based)
Approach	Descriptive or rule-based	Data-driven, statistical or ML-based
Churn Detection	Manual reports or fixed rules (e.g., no login in 30 days = churn)	Predicts likelihood of churn using patterns in historical data
Data Analysis	Basic analytics (e.g., averages, counts)	Complex modeling (e.g., logistic regression, decision trees, etc.)

Personalization	One-size-fits-all strategies	Customer-specific insights and risk scores
Actionability	Limited — reactive after churn occurs	Proactive — insights before churn happens
Automation	Mostly manual processes	Highly automated with model inference and dashboards
Scalability	Poor —doesn't scale well with large data	Excellent — handles thousands of customers easily
Insight Depth	Surface-level (e.g., churn % by month)	Deep (e.g., what factors contribute most to each customer's churn risk)
Adaptability	Static rules; slow to adapt to change	Continuously improves with more data
Tools Used	Excel, SQL dashboards, manual exports	Python, ML libraries (scikit-learn), AI models (e.g., LLaMA)

3. PROPOSED WORK

This study proposes the development of SubAnalytics, a conceptual machine learning-based framework for predicting customer churn and providing actionable retention insights tailored for B2B SaaS environments. The proposed work aims to bridge the gap between existing churn prediction models—predominantly designed for B2C contexts—and the complex requirements of SaaS-based subscription businesses, which involve organizational-level contracts, multiple user licenses, and long-term engagement cycles.

3.1 User Interface

The SubAnalytics user interface is designed to provide SaaS business users with clear, actionable insights into customer churn and retention. Upon accessing the user dashboard, individuals can explore churn prediction results through intuitive visualizations such as graphs, heatmaps, and cohort analyses. The system highlights key metrics like churn probability, usage trends, and engagement levels, allowing users to quickly identify at-risk customer segments.

In addition to predictive insights, the interface offers segmentation tools that classify customers into groups like "high-risk," "retention candidates," or "loyal customers," along with tailored retention strategies for each. Users can generate customized reports by filtering on attributes such as subscription tier, geography, or customer tenure. The platform is built with React.js and enforces strict access controls, ensuring that sensitive business intelligence is securely delivered to authorized users.

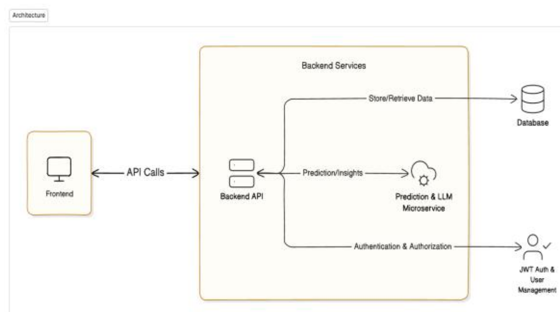

Fig. 1 - Architecture of SubAnalytics Platform

3.2 Data storage

The system employs MongoDB, a NoSQL database, to keep track of non-prediction-related data pertaining to user activity, analytics metadata, and system settings. MongoDB offers a scalable and adaptable solution for handling the varied and semi-structured data that pertains to SaaS-based analytics so the following kinds of data are kept in MongoDB:

User and Session Data: User profiles, activity logs, and session history are retained for purposes such as user personalization, audit trails, and analysis of user behavior.

Prediction Metadata: Every prediction request is recorded with corresponding metadata like timestamp, prediction parameters, result confidence scores, and user ID. The system can monitor performance, retrain models, and provide insights on previous predictions using this feature.

With MongoDB, the system can effectively fetch and handle structured and unstructured data in an efficient manner, making real-time access possible for dashboards and model running without having to depend solely on a computationally more intensive analytics layer. MongoDB's document-oriented design also helps in fast iteration on new features, which is perfect for dynamic SaaS platforms.

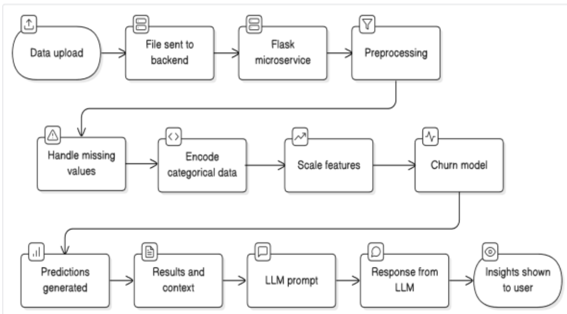

Fig. 2 - Data Flow

The SaaS Analytics & Prediction platform features a modular, scalable architecture with four core components: Frontend, Backend Services, Database, and JWT-based Authentication & User Management.
The Frontend allows users to register, log in, upload data, and view predictions via API calls to the backend. The Backend Services include a generic API layer that routes requests and a Churn Prediction Microservice that processes data, runs machine learning models, and generates insights. It interacts with the Database, which stores raw inputs, prediction results, and logs.
JWT Authentication handles user registration, login, token issuance, and verification, ensuring secure and authorized access. The entire process begins with a user action on the frontend, triggering an API request. The backend validates the request, forwards it to the appropriate module, and returns the result to the user. This architecture ensures security, modularity, and scalability for accurate churn prediction and customer behavior analysis.

4. CHURN PREDICTION MODEL

Churn prediction is a critical component of *SubAnalytics*, enabling SaaS businesses to proactively identify customers at risk of cancellation and implement targeted retention strategies. The predictive model was designed with simplicity, interpretability, and real-world applicability in mind, especially for small to mid-sized SaaS platforms that may not have dedicated data science teams.

4.1 Data Preprocessing

To maintain generalizability, the model accepts structured customer data typically exported from CRM or subscription management tools. Input features include customer tenure, subscription plan type, monthly spend, usage frequency, support ticket activity, and prior engagement indicators. Uploaded datasets are processed using a pipeline that includes *null value handling*, *categorical encoding,* and *feature scaling.* The model also filters out high-cardinality features and performs feature importance analysis to reduce dimensionality where applicable.

4.2 Model Architecture

After experimenting with several classification algorithms—including Decision Trees and Support Vector Machines—we selected a Random Forest Classifier for the current version due to its balance between accuracy and interpretability. The ensemble nature of Random Forests helps mitigate overfitting and captures non-linear relationships without complex hyperparameter tuning. The target variable is binary: churned (1) or retained (0).

4.3 Training and Evaluation

The model was trained using anonymized data that simulates SaaS business scenarios, consisting of approximately 1,000 customer records. An 80:20 train-test split was employed. The model achieved an accuracy of 63.5%.

4.4 Interpretability and Insights

To support business decision-making, model outputs are paired with an LLM– *llama-3.1-nemotron-70b-instruct*— through NVIDIA NIM to generate actionable insights related to the predicted values. This allows users to understand *why* a customer is likely to churn and take personalized actions.

5. RESULT

The churn prediction model was evaluated on a test dataset of 200 customer records. The model achieved an *accuracy of 63.5%.* Table 2 outlines the performance metrics of the model.

Class	Precision	Recall	F1-score	Support
0	0.78	0.75	0.77	160
1	0.15	0.17	0.16	40

Table 2 - Classification Report

	Predicted 0	Predicted 1
Actual 0	120	40
Actual 1	33	7

Table 3 - Confusion Matrix

To determine the performance of our predictive model, we have performed a confusion matrix analysis as shown in Table 3. We found that out of a total of 200 predictions, the model correctly predicted 120 actual negatives and 7 actual positives but misclassified 40 as false positives and 33 as false negatives. This gives us an overall accuracy of 63.5%, which means that the model correctly predicts the result in approximately two-thirds of the instances.

These results highlight the model's current strength in identifying retained users but also its limited ability to detect churners, likely due to class imbalance or insufficient feature complexity. The results of the model are displayed along with insights to the user through a clean user interface, as demonstrated in Fig. 3.

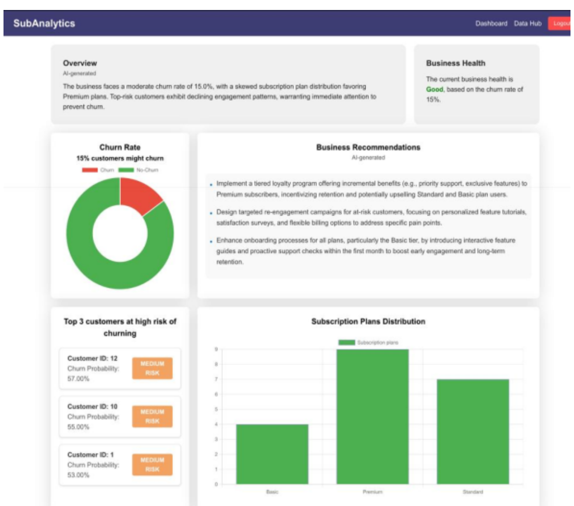

Fig. 3 - Dashboard Overview

6. CONCLUSION

In summary, Subanalytics offers an impressive, end-to-end predictive analytics tool to SaaS companies that want to be one step ahead of customer churn. By integrating a customized scikit-learn churn model with NVIDIA's LLaMA 3.1 Nemotron 70B Instruct for AI-powered recommendations, the tool not only predicts total churn rates and identifies the top three highest-risk customers but also interprets those findings into easy-to-read, actionable strategies. The safe, JWT-secure web interface and intuitive dashboard—complete with CSV upload, interactive Chart.js visualizations, and a minimal UX—make it simple for teams to integrate data, track subscription plan distributions, and respond proactively. Behind the scenes, the Flask microservice makes scalable, real-time churn scoring possible, while the Node.js backend coordinates authentication, data validation, and AI-insights fetching. In the end, Subanalytics enables SaaS stakeholders to shift away from one-size-fits-all, reactive retention strategies to deeply personalized, data-backed interventions that increase customer satisfaction, decrease churn, and deliver lasting revenue growth.

REFERENCES

[1] T. Wei, L. Zou, Y. Sun, and S. Zhao, "Prediction of Bank Product Subscription Behavior Based on Random Forest Algorithm," Proceedings of the 2nd International Conference on Financial Innovation, FinTech and Information Technology, 2023.

[2] B. Zhang, "Customer Churn in Subscription Business Model—Predictive Analytics on Customer Churn," BCP Business & Management, 2023.

[3] Y. Suh, "Machine Learning Based Customer Churn Prediction in Home Appliance Rental Business," Journal of Big Data, vol. 10, p. 41, 2023.

[4] A. Musunuri, "Machine Learning Model for Predicting Customer Churn in Subscription-Based Business," International Journal of Artificial Intelligence & Machine Learning (IJAIML), vol. 3, no. 2, pp. 211–220, 2024.

[5] N. D. S. Kumar, "OTT Subscriber Churn Prediction Using Machine Learning," M.S. thesis, California State University, San Bernardino, 2023.

[6] S. Bhattacharjee, U. Thukral, and N. Patil, "Early Churn Prediction from Large Scale User-Product Interaction Time Series," arXiv preprint arXiv:2309.14390, 2023.

[7] J. Maan and H. Maan, "Customer Churn Prediction Model Using Explainable Machine Learning," arXiv preprint arXiv:2303.00960, 2023.

[8] M. A. Shaikhsurab and P. Magadum, "Enhancing Customer Churn Prediction in Telecommunications: An Adaptive Ensemble Learning Approach," arXiv preprint arXiv:2408.16284, 2024.
[9] A. Bhatnagar and S. Srivastava, "Customer Churn Prediction: A Machine Learning Approach with Data Balancing for Telecom Industry," International Journal of Computing, vol. 24, no. 1, pp. 9–18, 2025.
[10] K. Gadgil, S. S. Gill, and A. M. Abdelmoniem, "A Meta-learning Based Stacked Regression Approach for Customer Lifetime Value Prediction," arXiv preprint arXiv:2308.08502, 2023.
[11] T. Vafeiadis et al., "A Comparison of Machine Learning Techniques for Customer Churn Prediction," Simulation Modelling Practice and Theory, vol. 55, pp. 1–9, 2024.
[12] R. Lee, C. Rungie, and M. Wright, "Regularities in the Consumption of a Subscription Service," Journal of Product and Brand Management, vol. 20, no. 3, pp. 182–189, 2023.
[13] C. S. Mgbemena, "A Data-Driven Framework for Investigating Customer Retention," Ph.D. dissertation, Brunel University London, 2024.