Predicting Customer Churn for a Subscription Streaming Service

Applied Data Science Project • Priya Nair • June 2026

0.891	**\$1.11M**	**9.8%→7.1%**
AUC–ROC (held-out test)	Projected annual net ROI	Churn, pilot cohort

1 Executive Summary

This project builds a supervised model to identify subscribers at high risk of cancelling within the next 30 days, enabling proactive retention. Using twelve months of behavioral and billing data for 84,000 subscribers, a gradient-boosted model achieves an AUC–ROC of 0.891 and, in a six-week pilot, reduced monthly churn in the treated cohort from 9.8% to 7.1%.

2 Data & Features

The dataset joins subscription records, usage logs, and support interactions. After cleaning and de-duplication, 84,012 active subscribers remained, with a 30-day forward churn label.

Table 1: Selected engineered features and their importance.

Feature	Importance	Corr. w/ churn
Tenure (months)	0.24	-0.61
Weekly active days	0.19	-0.58
Support tickets, last 90 days	0.16	$+0.47$
Monthly charges	0.12	-0.29
Month-to-month contract flag	0.11	$+0.42$

3 Methodology

We compared logistic regression, random forest, and gradient boosting under 5-fold time-aware cross-validation, tuning with Bayesian optimization. The decision threshold was chosen to maximize expected retention value, not raw accuracy:

$$\mathsf{EV}(\tau) = \sum_{i:\,\hat{p}_i \geq \tau} \big(p_i \cdot v_{\mathsf{save}} - c_{\mathsf{offer}}\big), \tag{1}$$

where v_{save} is the retained lifetime value and c_{offer} the incentive cost.

3.1 Model Comparison

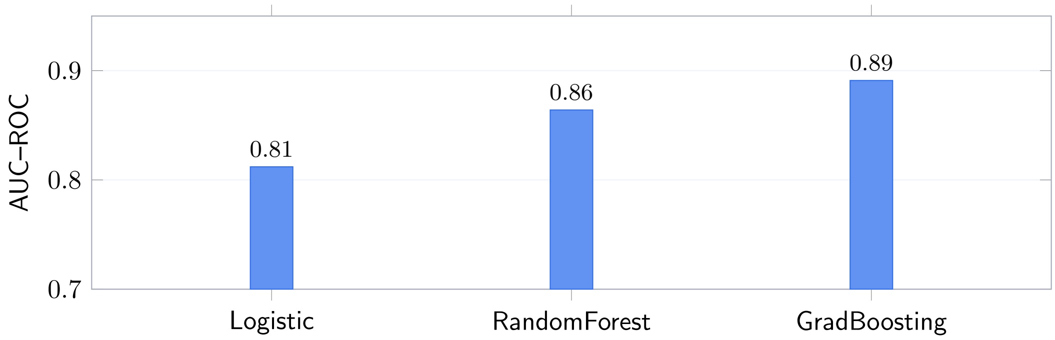

4 Results

Gradient boosting was selected for deployment. Beyond ranking accuracy, it produced well-calibrated probabilities suitable for value-based thresholding.

Table 2: Held-out test performance by model.

Model	AUC	Precision	Recall	F1
Logistic Regression	0.812	0.61	0.58	0.59
Random Forest	0.864	0.69	0.64	0.66
Gradient Boosting	**0.891**	**0.74**	**0.70**	**0.72**

5 Deployment & Recommendations

1. **Shadow (weeks 1–2):** score the base weekly; validate live vs. holdout.
2. **Pilot (weeks 3–6):** trigger retention offers for 20% of the at-risk segment, A/B tested against a control to confirm causal lift.
3. **Rollout (weeks 7–12):** automate weekly batch scoring; add drift monitoring (PSI, calibration, performance-decay alerts).

6 Limitations

The label horizon is fixed at 30 days; longer horizons need re-labeling. The pilot ran on a single market, and seasonal effects were not fully controlled. Ongoing monitoring for feature drift is required to sustain performance.