IJIRSET
International Journal of Innovative Research in Science | Engineering | Technology

ISSN(O): 2319-8753
ISSN(P): 2347-6710



International Journal of Innovative Research in Science Engineering and Technology (IJIRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)



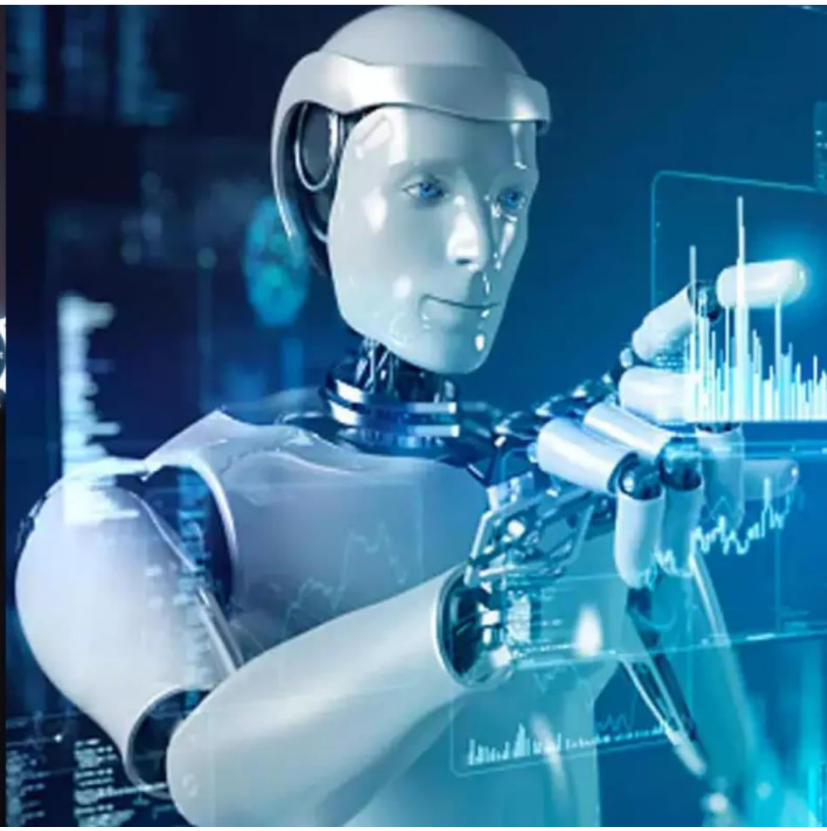

Impact Factor: 9.935

Volume 15, Issue 11, November 2026

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 9.935

|www.ijirset.com |A Monthly, Peer Reviewed & Refereed Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 15, Issue 11, November 2026

|DOI: 10.15680/IJIRSET.2026.1511034|

CRIP: An Explainable Churn Prediction and AI-Assisted Customer Retention Platform

J. Hari Krishna, P. Rithik, Mrs. A. V. Lakshmi Prasuna

Department of Information Technology, Mahatma Gandhi Institute of Technology, Hyderabad, Telangana, India

Assistant Professor, Department of Information Technology, Mahatma Gandhi Institute of Technology, Hyderabad, Telangana, India

ABSTRACT: Customer churn models can identify likely departures, but a numerical risk score alone does not tell a retention team what to do next. This paper examines the design and implementation of the Customer Retention Intelligence Portal (CRIP), an application that connects supervised churn prediction with customer-level explanations, business dashboards, and assisted outreach. A seven-feature Random Forest model is trained on the Telco Customer Churn dataset using a stratified holdout split. A FastAPI inference service supplies scores to an Express application backed by MongoDB, where customers, tickets and campaign information are managed. SHAP is used when the required library and operating conditions permit; the deployed code also includes fallback explanations. The application provides CSV import, a non-persistent What-If experiment, Gemini-assisted campaign drafting with a local fallback, and email delivery. A prior project evaluation reports 79.77% accuracy and 83.7% ROC-AUC on 1,409 held-out observations, although the repository README contains a different evaluation snapshot. The main contribution is a traceable operational workflow joining prediction, interpretation and retention action, rather than a new classification algorithm. We discuss the evaluation boundaries and the need for controlled campaign and generalization studies.

KEYWORDS: Customer churn, Random Forest, explainable artificial intelligence, SHAP, customer retention, MongoDB, generative AI.

I. INTRODUCTION

Churn management poses two distinct questions: which customers are likely to leave, and which intervention is appropriate for each customer? The first is a prediction task; the second requires an account-level rationale, operational context, and a way to act on the finding. A system that addresses only classification leaves the business workflow incomplete. CRIP was developed to bring these steps into one administrator-facing interface.

The project combines a web dashboard, a Node.js/Express service, a MongoDB document store, and a Python inference service. In addition to risk scoring, administrators can inspect customer records, manage support cases, test hypothetical changes to selected features, and prepare retention messages. We investigate what the supplied implementation establishes, separate the reported empirical results from architectural capabilities, and identify weaknesses that remain before production or cross-industry use.

The specific contribution is an integrated implementation and an implementation-grounded analysis of its data path, model behavior, explanation fallback, and outreach mechanisms. It does not claim a novel learning algorithm, causal identification of churn drivers, or a demonstrated increase in customer retention.

II. RELATED WORK

A. Ensemble prediction on tabular customer records

Random Forest aggregates decisions across randomized trees to reduce the instability of individual trees [1]. Its ability to model nonlinear interactions without a fixed linear decision boundary makes it suitable for a compact tabular churn baseline. The CRIP implementation adopts a fixed-depth, 100-tree forest rather than introducing a new model family. The resulting probabilities support account ranking but should not be treated as calibrated intervention effects.

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 9.935

|www.ijirset.com |A Monthly, Peer Reviewed & Refereed Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 15, Issue 11, November 2026

|DOI: 10.15680/IJIRSET.2026.1511034|

B. Local explanations and decision support

SHAP assigns feature-level attributions to a particular model output and provides a principled framework for interpreting predictions [2]. In a retention workflow, these attributions can help an administrator distinguish a high-risk account associated with short tenure from one associated with contract or billing characteristics. Attribution describes how a fitted model uses features; it does not, by itself, establish that changing the feature will prevent churn. CRIP therefore presents explanatory evidence to support, rather than replace, human judgment.

C. From prediction to customer outreach

Prior churn studies often stop at performance reporting. Connecting predictions to customer records and messaging introduces new engineering questions: whether imported fields match the model inputs, whether an explanation is available, whether a generated message is grounded in the account, and whether an observed response is tracked. CRIP addresses these concerns through separate application routes and campaign records. The campaign module is evaluated here as an implemented workflow, not as evidence that language-model-generated messages produce superior retention.

III. SYSTEM ARCHITECTURE

CRIP uses a static HTML/CSS/JavaScript client with dashboard visualizations and a Node.js/Express application server. Express routes handle authentication, data ingestion and management, and the Grow/Campaign functions. MongoDB, accessed through Mongoose models, holds users, customers, campaigns, activity and tickets. A separate FastAPI process loads a serialized scikit-learn model to provide inference. This separation allows prediction code to evolve independently of the web application, although the supplied project does not establish distributed scaling or production service-level guarantees.

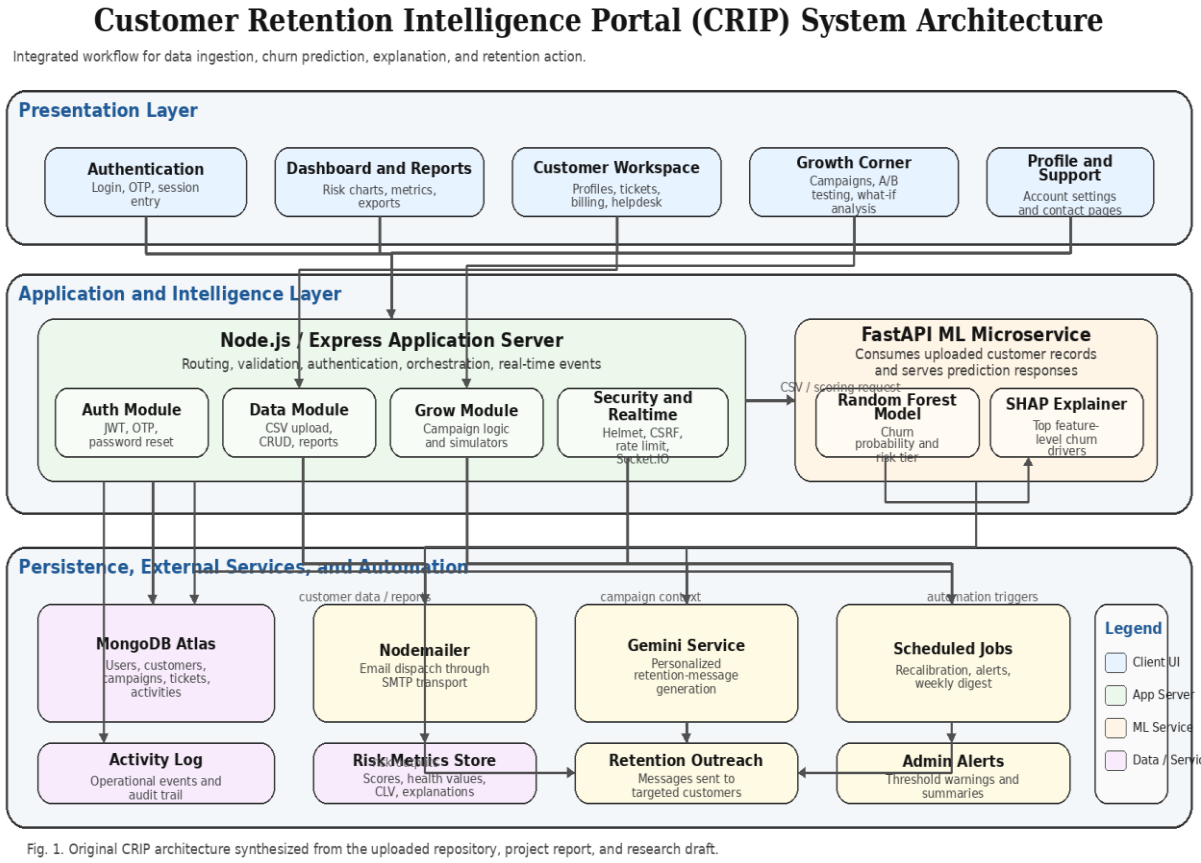

Fig. 1. Original CRIP architecture synthesized from the uploaded repository, project report, and research draft.

Fig.1. Original system architecture of CRIP

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 9.935

|www.ijirset.com |A Monthly, Peer Reviewed & Refereed Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 15, Issue 11, November 2026

|DOI: 10.15680/IJIRSET.2026.1511034|

Component	Implementation	Function
Presentation	HTML, CSS, vanilla JS, Chart.js	Dashboard, customer and report pages
Application API	Node.js, Express, Mongoose	Authentication, CRUD, CSV import and orchestration
Inference	Python, FastAPI, scikit-learn	Churn probabilities and risk labels
Explanation	SHAP with fallback	Feature-level reasons where available
Persistence	MongoDB	Customer, campaign and support records
Outreach	Gemini SDK, Nodemailer	Assisted campaign drafts and mail delivery
Operations	node-cron	Scheduled application tasks

TABLE I. MAIN COMPONENTS CONFIRMED IN THE PROJECT SOURCE

The principal flow is: administrator authentication → customer import or selection → validation and feature mapping → inference request → risk assessment and explanation → stored/customer-facing dashboard result → optional message draft and dispatch. The What-If endpoint constructs a temporary prediction input and returns a hypothetical score without persisting the simulated changes to customer records. Campaign generation can return a local draft when a configured Gemini service is unavailable; consequently, not every displayed draft is necessarily model-generated.

IV. PROPOSED METHOD

A. Data preparation and model fitting

The training code uses seven input variables: tenure, MonthlyCharges, TotalCharges, Contract, InternetService, PaymentMethod and SeniorCitizen. TotalCharges is converted to a numeric type with invalid values replaced by zero; remaining missing values are filled with zero. Object-valued columns are label encoded. Churn is mapped to a binary target and separated using a stratified 80/20 split with random_state = 42. The classifier uses 100 trees, max_depth = 10, random_state = 42, and parallel execution. Model performance is calculated on the holdout set before the classifier is fitted again on all observations for deployment.

For a customer feature vector x, the estimated churn risk is $p(x) = (1/T) \Sigma I[h_t(x) = 1]$, where T is the number of trees and h_t is a tree vote. The system labels $p < 0.30$ as Safe, $0.30 \leq p < 0.70$ as Warning, and $p \geq 0.70$ as High-Risk. These tiers are operational thresholds, not independently optimized or clinically meaningful probability categories.

B. Prediction explanation

The prediction module attempts to construct a SHAP TreeExplainer and extracts the leading contributions for eligible non-Safe cases. In the supplied source, SHAP processing is conditional on the package being available and on a batch size not exceeding 5,000. Exceptions also lead to fallback behavior. This distinction matters: a reason returned by the application may be a SHAP attribution or a heuristic explanation rather than a mathematically equivalent SHAP result. Explanations should therefore be labeled with their provenance in future versions.

C. Retention campaign workflow

An administrator can request a draft associated with customer characteristics and detected risk. The Grow route configures Google Gemini when API credentials are available, using a time-limited request; on failure or missing configuration, it can supply local fallback text. Campaigns and subsequent outcomes can be recorded, and the software includes variant-oriented fields for A/B comparisons. A/B support is an instrumentation capability; a statistically established causal lift requires randomized assignment, adequate sample size and a prespecified observation window.

D. Risk simulation and persistence

The What-If route accepts a proposed customer configuration and calls the scoring service using a temporary input. The result is a counterfactual model prediction: it says how the fitted model responds to altered input fields, not how an

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 9.935

|www.ijirset.com |A Monthly, Peer Reviewed & Refereed Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 15, Issue 11, November 2026

|DOI: 10.15680/IJIRSET.2026.1511034|

actual customer would respond to a business intervention. The distinction is necessary because changing contract type or monthly charges in software does not establish real-world causation.

V. EXPERIMENTAL EVALUATION

A. Dataset and evaluation protocol

The study materials describe a Telco Customer Churn benchmark with 7,043 rows and an 80/20 stratified split (5,634 training and 1,409 test records). The model evaluation is performed before production refitting. Table II transcribes the held-out results reported in the supplied Nova chapter; they are not presented as newly reproduced measurements. The repository README records another snapshot (79.56% accuracy and 84.12% ROC-AUC), and neither file identifies a definitive shared run artifact for reconciling the two. A future release should attach the dataset checksum, dependency versions and output metrics to each run.

Evaluation measure	Reported holdout value
Accuracy	79.77%
Precision (churn class)	65.84%
Recall (churn class)	49.47%
F1-score (churn class)	56.49%
ROC-AUC	83.7%
Confusion matrix	TN 939; FP 96; FN 189; TP 185

TABLE.II. RESULTS REPORTED IN THE NOVA PROJECT DRAFT (NOT INDEPENDENTLY REPRODUCED)

The reported confusion-matrix counts total 1,409 examples. Recall below 50% is the central operational limitation: at the evaluated decision threshold, approximately one in two actual churners is missed. Accuracy alone therefore gives an incomplete account of usefulness. Precision–recall analysis, cost-weighted thresholds and calibration testing are needed when the goal is allocating retention resources.

B. Operational functionality

The project documentation includes functional tests for login and OTP, accepted and rejected CSV imports, customer operations, risk output, campaign generation and What-If analysis. The Nova chapter states that eleven documented scenarios met their expected outcomes, and describes a 29-record dashboard pilot with 9 Safe, 7 Warning and 13 High-Risk accounts. These are application-level observations; they do not provide an independent measurement of churn reduction, email deliverability, model fairness, or reliability under sustained load.

Evaluation level	Available evidence	Interpretation
Offline predictive evaluation	Held-out Telco metrics in project draft	One reported benchmark run
Functional acceptance	Eleven scenarios in project documentation/draft	Exercised workflow behavior
Pilot dashboard	29 imported customer records	Risk segmentation display
Campaign effectiveness	No controlled outcome dataset supplied	Retention lift not established
Industry transfer	Additional datasets in ZIP	Generalization not established by tests

TABLE.III. EVIDENCE BOUNDARIES

VI. DISCUSSION

The implementation offers three practical strengths. First, a separate prediction service keeps feature processing and model loading out of the presentation code. Second, account records, support interactions and campaign activity are connected in one workspace, reducing repeated manual transfers between analysis and action. Third, the ability to

International Journal of Innovative Research of Science, Engineering and Technology (IJIRSET)

Impact Factor 9.935

|www.ijirset.com |A Monthly, Peer Reviewed & Refereed Journal| e-ISSN: 2319-8753| p-ISSN: 2347-6710|

Volume 15, Issue 11, November 2026

|DOI: 10.15680/IJIRSET.2026.1511034|

inspect model-driven reasons and run What-If inputs gives administrators more context than a simple binary churn label.

The design also has important boundaries. Label encoding can impose artificial category orderings, so comparisons with one-hot encoding would be valuable. The training feature set is tied to Telco field names even though MongoDB can store flexible documents; database flexibility does not imply the predictive model can score arbitrary industry schemas. SHAP can be unavailable or skipped, so explanation completeness must be measured. An external language model may generate unsuitable claims or offers and should be subject to user review and privacy controls. Finally, no campaign randomization or observed incremental retention effects were demonstrated by the supplied evidence.

VII. CONCLUSION AND FUTURE WORK

CRIP demonstrates a working design for joining churn classification with account-level interpretation, customer management, hypothetical risk exploration and assisted outreach. Its strongest supported result is an integrated application workflow with an offline Random Forest benchmark, rather than proven improvement in real-world retention. The reported churn recall indicates that threshold choice and class-imbalance handling are high-priority research questions. Next steps include reproducible multi-model evaluation, stratified cross-validation, probability calibration, explicit explanation provenance, model monitoring across datasets, and a prospectively randomized A/B study of retention outcomes. These steps would permit a more defensible transition from a functional project to a validated decision-support system.

REFERENCES

[1] L. Breiman, “Random forests,” Machine Learning, vol. 45, pp. 5–32, 2001. doi: 10.1023/A:1010933404324.
[2] S. M. Lundberg and S.-I. Lee, “A unified approach to interpreting model predictions,” Advances in Neural Information Processing Systems, vol. 30, 2017.
[3] F. Pedregosa et al., “Scikit-learn: Machine learning in Python,” Journal of Machine Learning Research, vol. 12, pp. 2825–2830, 2011.
[4] CRIP project team, “Customer Retention Intelligence Portal,” Industry Oriented Mini Project Report, Mahatma Gandhi Institute of Technology, 2025–2026. Unpublished project documentation.
[5] CRIP project team, “CRIP: A Machine Learning and Explainable AI Framework for Proactive Customer Churn Prediction and Automated Retention Campaigns,” Nova Science chapter draft, unpublished.
[6] CRIP project repository, “CRIP — Customer Retention Intelligence Portal,” uploaded source archive, version dated May 2026. Unpublished software artifact.



INTERNATIONAL STANDARD SERIAL NUMBER INDIA

INNO SPACE SJIF Scientific Journal Impact Factor

doi crossref

निस्केयर NISCAIR

INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN SCIENCE | ENGINEERING | TECHNOLOGY

9940 572 462 6381 907 438 ijirset@gmail.com

www.ijirset.com

Scan to save the contact details