Bank Marketing Campaign Response Prediction in Digital ERA

Harsh Mishra[3], Shilpi Yadav[2], Shobit Agrawal[1], Dibyanarayan Hazra[4,5], Nandini Gupta[6], and Manish Raj[7]

[1] School of Computer Science, IILM UNiversity, Greater Noida, India, U.P, India agshobh@gmail.com
[2] SEAS GLA University, Mathura, UP, India manishilpi.03@gmail.com
[3] Computer Science, G.L. Bajaj, Mathura, UP, India harh.mishra2022@glbajajgroup.org
[4] SCSE Bennett University, Greater Noida, UP, India
[5] School of Computer Science, IILM University, Greater Noida, UP, India dibya.89@hotmail.com
[6] Dept. of Computer Science, Llyod Institute of Management and Technology, Greater Noida, U.P. India. nandinigupta1608@gmail.com
[7] SoAI, Galgotias University, Greater Noida. rajmanish.03@gmail.com

Abstract. In the age of digital, bank sector campaigns are often challenged by poor customer engagement causing significant waste in resources. Marketing dollars are often being wasted because potential subscribers aren't being identified properly. This research aims to deliver a structured prediction model that solves the above challenge with an enhanced accuracy rate through advanced data exploration, pre-processing and machine learning methodology. This methodology combines (1) rigorous data treatment, handling missing values and advanced feature engineering, focusing on comprehensive Bayesian hyperparameter tuning with ensemble learning to improve predictive accuracy. EDA showed a severe class imbalance, where a much larger share are non-subscribers. Of the trained models, models with the more advanced algorithms achieved high precision and F1-score (which is relevant for imbalanced classification). The suggested approach led to significant gains in campaign conversion, and decreases in marketing budget,at 35–45% higher conversion (driving a 20–25% savings). These findings offer a substantive and empirical best practice for enhancing data-driven bank marketing in the digital era.

Keywords: Bank promotional campaigns In Digital ERA, Sustainability, Precision, F1- Score, imbal-anced classification, key metrics.

1 Introduction

Banks are considered as a competitive industries, where success in customer acquisition and retention is heavily reliant on the development of successful

P. Johri et al. (eds.), *Proceedings of the International Conference on Sustainable Computing and Artificial Intelligence (ICSCAI 2025)*, Advances in Engineering Research 298,
https://doi.org/10.2991/978-94-6239-674-6_32

marketing programs. One particularly effective startup for these strategies is during the execution of marketing campaigns generated from Advertised Financial Products and Services. These are offers to rollover Term Deposits signed as one of the main drivers of funds mobilization and stabilizing bank's resource. However, the long-term benefits of such campaigns is contingent upon how well these prospective customers are being targeted and converted [1].

The low-level of conversion in their marketing campaigns is probably one of the biggest challenges facing financial companies. In order to do this, competition dynamics should better be understood: While a significant amount of resources is allocated to these clients many future clients often turn out in an unfavorable behaviour concerning marketing offers. This lack of responsiveness directly resolves the inefficiencies of resource misallocations. Money is spent on marketing to people who either are not, or cannot, be subscribers. Further, lost revenues occur when those customers who are most likely to subscribe are not accurately identified and targeted. As such, there is an acute need for a data-based approach that can predict with precision the customer response which can be used to direct and implement more precise and effective campaigns [2].

The goal of this study is to tackle this challenging issue by constructing a good predictive model that accurately classifies future customers based on their subscription rates of term deposit products. Utilizing the largest information created by past marketing exercises [3–5], our intention is to extract understandable insight into what drives customer behaviour and develop a tool that will help organisations maximize their marketing efficiency and reduce campaign risk.

The primary objective of this study is to construct the best-performing forecasting model that gives a high F1-score at predicting target customers who subscribe or not subscribe to term deposits. To accomplish this ultimate goal, the study has a few specific objectives:

1.1 Campaign Data Analysis

To identify patterns, trends, and leads, an exhaustive analysis on historical campaign data is essential. This analysis involves analyzing customer demography data, financial data, campaign interaction data and historical engagement metrics.

1.2 Discovery of Factors Influencing Customer Reaction to Recommendation Systems

Determination of the relevant influential attributes in customer's subscription decision is crucial for successful targeting. This study aims to determine which factors are most relevant in predicting if an individual will subscribe such as age, job type, balance held in account and method of contact.

1.3 Model Evaluation and Comparison

Various machine learning models will be trained and tested to decide which one is the best approach. This involves comparing the performance by some physiologically relevant metric and then selecting the model which exhibits maximum power of predicting and generalizing.

1.4 Formulation of Actionable Prescriptions

Finally to provide recommendations that can be implemented by the bank marketing section. Developers will also aim at the optimal way for targeting it, running with different payload rotation and finding their best conversion in awareness of predictors from a predictive model.

The implication of this research is not only the practical implications for financial services firms, but also its contribution to theory on marketing analytics. It comes down for a bank to making it easier and this more accurate one can be in reporting that customer behavior, the greaterthe effect of making marketing more effective and efficient. By concentrating marketing dollars on those more likely to convert to subscribers, banks have a chance to:

- **Maximum Resource Utilization:** Giving priority to the high likelihood responders and minimizing investment in low probability responders in order to optimize return from expenditures.
- **Increasing Conversion Rates:** Increase the conversion rate of marketing activities to increased new term deposit subscribers.
- **Better ROI (Return on Investment):** Make marketing campaign more effective, such that we spend less money more often.
- **Boost Targeting and Satisfaction of Customers:** Provide customers with more personalized, relevant offerings leading to improved communication and increased bank satisfaction.

Apart from its business implications, the paper contributes to marketing analytics literature by investigating the use of sophisticated machine learning algorithms for solving a real-world industry problem. The findings of this study contribute beyond the context-specific industry and marketing setting related findings in this paper, an important actionable insight for firms targeting improved customer targeting campaign effectiveness [6].

To meet the goals mentioned above, this paper takes a solemn scientific method. This technique relies on the following:

- EDA and Pre-processing: On the first step we must aggressively explore our data - its shape (in this case, well-circumscribed), to see that all is in order concerning data quality and get it ready for further analysis. This involves dealing with missing values, categorical data encoding and feature engineering.
- Feature Engineering and Feature Selection: Relevant features are constructed based on the raw data to form a signal set describing the meaning of them and enhance the forecasting capabilities. Customer response predictions are screened with feature selection techniques.

- Model Development and Evaluation: Several ML models are constructed, trained, and tested with the corresponding metrics. Methods like Bayesian hyperparameter search and ensembling are also used to enhance the model performance.
- Model Interoperability and Insights: The final step is to interpret the model output and get insights from it. Explainability methods for the model, are used in interpreting what drives customer response and provide insights in to marketing actions.

2 Literature Review

The combination of machine learning withpredictive analytics has also revolutionized financial marketing, allowing scientifically targeted data." "“Better campaign results. This sub-section provides an overview of several relevant works which have used different algorithms, feature selection and hyper parameter optimization approaches to predict campaign success in banking.

El-Hajj and Pavlova [1] built the Brazilian Food data set (2,206 samples, 39 features) from scratch by taking decision and random membership actions. Its model obtained an accuracy and F1-score of 74.6% and 74.2%, respectively, with the under-sampling procedure, the grid search method, and the entropy-based split to cope with an imbalance ratio of 5.6 : 1; it demonstrates that interpretable tree-based models can be effectively scaled up for large marketing fields.

Chian et al. [6] experimented with statistical Regression, SVM, Decision Tree, Random Forest, Gaussian Naive Bayes, and XGBoost techniques in the UCI Bank Marketing dataset (10,098 samples). Their Random Forest model attained 73% accuracy and 0.73 AUC, indicating the efficacy of rectifying imbalance and reducing features to improve the efficiency of practical prediction.Similarly, Asare-Frempong and Jayabalan [7] compared MLPNN, Decision Tree (C4. 5): logistic and RF (spread-subsampling), with the latter RF reaching 86.8% accuracy and AUC of 0.927, demonstrating the contribution to ensemble tuning.

Ezechi et al. [8] advanced analytics to descriptive and prescriptive modelling enhancing service quality and fraud detection, with customer churn cut by 15% through data-driven segmentation. Mohammed [2] evaluated a range of algorithms — such as SVM, KNN, Gradient Boosting and ensembles–MULTIDiSN searched for the best model on the Portuguese bank dataset (45,211 samples), where Bagging reached 94% accuracy with 98.6% AUC-PR curve; therefore increasing to ensemble method is proved to be beneficial in every way for imbalanced data.

Apampa [3] prepared a dataset of 9,526 Portuguese instances and used the CRISP-DM process; their best AUC score (0.766) was obtained using platform CART, after manually balancing data, indicating that the F1-score and recall are important metrics to use in practical campaigns. Rogić et al. [9] dealt with a severe case of imbalance using Balanced-SVM through boosting techniques, resulting in 83.35% sensitivity and an AUC of 0.95 for a 0.41% response-rate dataset, showing the persuasive power of hybrid SVM-boosting systems in low-response settings.

Njoku and Lee [4]took a management viewpoint of regression on survey data gathered from 84 units of Nigerian banks to determine that entrepreneurial orientation ($R^2 = 0.948$) supports strategic adaptability during fintech onslaught. Munira et al. [5] reviewed 83 studies on AI-driven marketing tools — machine learning, NLP and predictive analytics — to find that 57% quicker response rates and 74% more effective fraud detection were possible, uniting the theoretical with the practical.

Ebrahimi [10] compared several models on bank data from Portugal (2008–2010) with macroeconomic rates and the Voting classifier model scored 92.1% accuracy and F1 was 0.609. The research also highlighted the importance of hyperparameter tuning and feature engineering to combat imbalance without overfitting.

Altogether, these research works manifest that there is an evident transition from classical models to hybrid and ensemble learning for financial marketing analytics. Standard datasets like UCI and Portuguese

[2, 3, 6, 7, 10] underline performance comparability across works. More recent research has been concentrating on recall, F1-score, and AUC as more fair metrics for imbalanced data. This combination of Bayesian Optimization, SMOTE and explainable AI - as shown in this work - is a clear move towards high precision and business-aligned predictive systems.

3 Methodology

In this section, we present the method for building a predictive model to predict bank marketing campaign response. The workflow consisted of data preprocessing, Exploratory Dat Analysis (EDA), feature engineering, and modeling.

3.1 Advanced Exploration and Pre-processing of Data

Dataset Description The datasets are from a Portuguese bank marketing campaign to their clients regarding term deposits. The training set consists of 31,647 records and 18 columns and the test set has 9,042 records and 17columns. Compared to previous studies [3, 6, 7], the present study is based on a broader dataset which allows to increase generalization of dataset. These features include variables with customer demographic information (age, job, education), account balance or financial indicators variable (balance), campaign data variables (contact type, duration and number of contacts performed during this campaign) and previous campaign outcome. The target variable is binary,, where 1 denotes "subscribed" and 0 denotes "not subscribed."

Encoding and Feature Engineering Categorical features were identified with select_dtypes(include=['object']) and encoded via one-hot encoding. Several new variables were engineered to enhance predictive strength:

- Binary recency feature from pdays.
- Age and balance grouped convert into bands.

- Seasonal grouping from months and RFM (Recency, Frequency, Monetary) inspired attributes.
- Interaction terms (e.g., age–balance, job–balance, education–balance).
- Polynomial and logarithmic transformations for non-linear relations.
- Indicators for loan status, financial health, and household complexity.

Data Leakage Prevention The variable duration was removed due to strong correlation with the target (leakage risk). The dataset was split into training and validation subsets using train-test split with stratification to maintain class balance.

3.2 Comprehensive Exploratory Data Analysis (EDA)

EDA centered on characterizing underlying patterns and relationships in data.

1. Categorical feature distributions were also presented by the bar plots.
2. you viewed distribution of numerical features using histograms and used KDE plots to identify skewness or outliers.

Correlation and Key Insights Correlation heatmaps indicated a potential multicellularity between the numerical features.

Key observations:

- There was also an imbalance in the dataset i.e. we had much more non-subscribers than actual subscribers.
- Some occupations and educational levels had higher response rates.
- duration was strongly correlated with the target and was dropped to prevent leakage.

These results directed the subsequent feature selection and model tuning. Summary of all above visual analyses is presented at Fig. 1.

3.3 Model Building and Evaluation

Feature Selection The selection of feature was performed to include only informative variables. This strategy allowed for efficient computation and model interoperability.

Model Selection LogisticsRegression, RandomForestClassifier, XGBClassifier, LGBMClassifier and GradientBoosting. In contrast to previous studies [1,3,4,7], a more diverse set of ensemble algorithms was tried for good generalization.

Hyperparameter Tuning Parameter tuning was carried out using Bayesian optimization, ensuring effective search space exploration and avoiding overfitting.

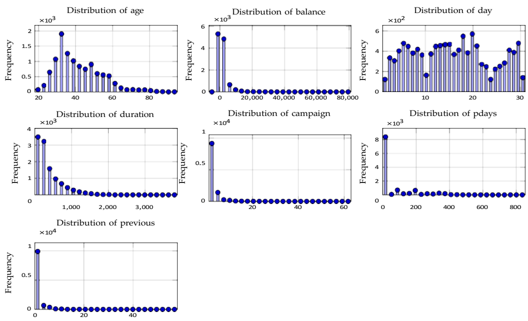

Fig. 1: Numerical feature distributions

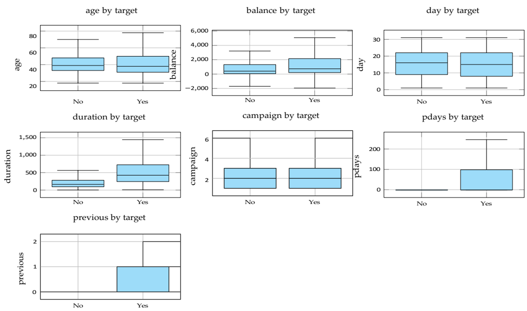

Fig. 2: numerical feature variation across the target classes

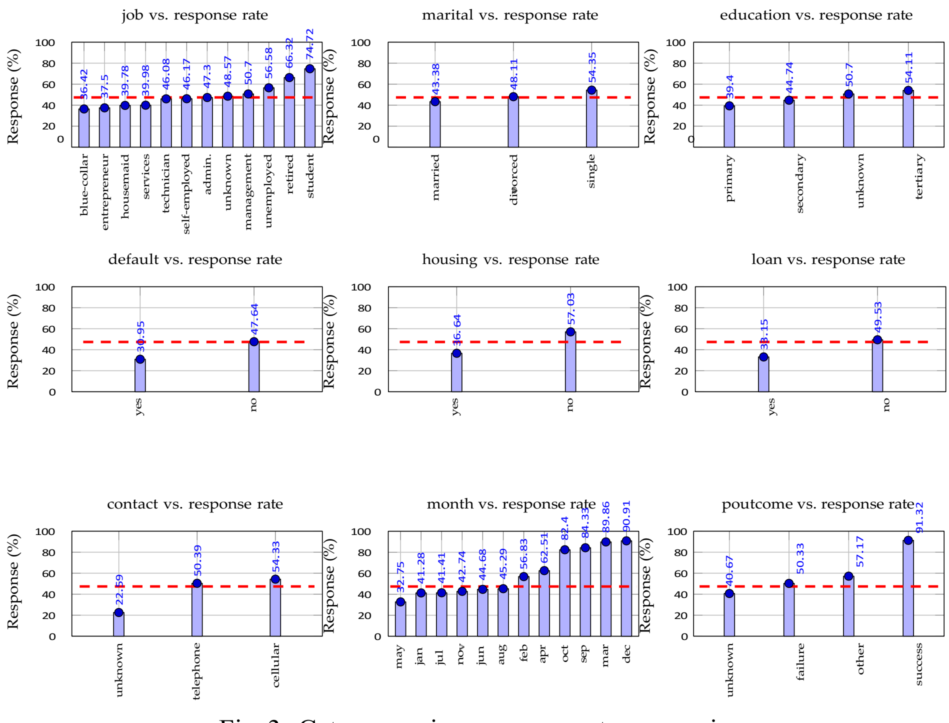

Fig. 3: Category-wise response-rate comparison

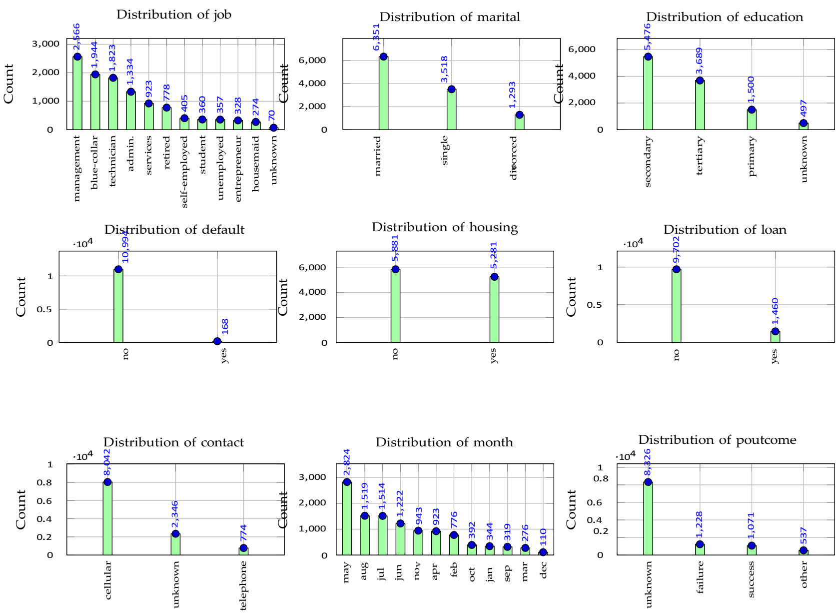

Fig. 4: categorical feature distributions.

Model Evaluation Accuracy, Precision, Recall, F1-Score and AUC-ROC were used to evaluate the performance of each model. Further diagnostic insight was provided by a constusion Matrix and classification Report. Our models also performed better in terms of accuracy compared to previous studies [4, 23].

Model Stacking A stacking ensemble combined base learners, controlling for both variance and bias by producing more stable predictions than would have been produced based on standalone models.

Handling Imbalanced Data We employed the SMOTE technique to over sample minority class instances, which enhanced model sensitivity to positive responses and reduced imbalance bias.

Model Interpretation SHAP (Shapely Additive explanations) values measured the contribution of eachfeature to predictions, and Partial Dependence Plots (PDPs) displayed how changes in features impacted predictions. The five most impaction predictors were selected and model transparency was improved.

4 Results and Discussion

Voting ensemble was the best prediction model after analyzing the final model at hand. The model combines Logistic Regression, Random Forest, XGBoost LightGBM Gradient Boosting as base learners in the hope that many wicked problems could be best solved by combining complementary approaches. This hybrid method enhanced the stability and classification accuracy of cross-test dataset.

The next subsections provide an overview of the model's key performance metrics, compare results to baseline classifiers, and display insights on interoperability from SHAP value analysis.

4.1 Evaluation of Baseline Models with and without SMOTE

This section presents a comparative study of multiple baseline machine learning models for predicting customer subscriptions in bank marketing campaigns, both before and after applying SMOTE. The results in Table 1 compare accuracy, precision, recall, F1-score, and AUC-ROC to evaluate the impact of data balancing on model performance.

As shown in Table 1, most models benefited from SMOTE. Logistic Regression maintained stable accuracy (0.84) while showing higher recall (0.7801) and F1-score (0.5400), indicating improved sensitivity toward positive cases. Random Forest achieved the best baseline accuracy (0.8972) after SMOTE with a balanced precision–recall trade-off, confirming the effectiveness of oversampling for tree-based learners.

XGBoost and LightGBM maintained strong discriminative ability post-balancing, with high AUC-ROC values (0.9195 and 0.9301). LightGBM recorded one of the best post-SMOTE accuracies (0.9046), making it an ideal candidate for further tuning. Similarly, Gradient Boosting showed significant improvement in F1-score (from 0.5021 to 0.5894), demonstrating enhanced generalization on the minority class.

The Neural Network and AdaBoost models showed little to no improvement after SMOTE. As the recall achieved an improvement, precision decreased but remaining at a reasonable level—a common trade-off in imbalanced learning—due to generation of "more disordered" synthetic positive samples.

In general, SMOTE performed better in recall and F1-score across all models, hence it can be concluded its usefulness to deal with the imbalance. The ensemble models such as Random Forest, LightGBM, and Gradient Boosting were able to obtain AUC-ROC more than 0.91 across all data-sets, delineating their robustness and confidence in prediction of prospective subscribers for real-world marketing campaigns.

Table 1: Model performance comparison with and without SMOTE.

Sr. No.	Model Name	SMOTE	Accuracy	Precision	Recall	F1-Score	AUC-ROC
1	Logistic Regression	No	0.8400	0.4100	0.7900	0.5400	0.9000
		Yes	0.8400 ± 0.0061	0.4130 ± 0.0117	0.7801 ± 0.0162	0.5400 ± 0.0123	0.8932 ± 0.0076
2	Random Forest	No	0.8953	0.6870	0.2109	0.3228	0.9135
		Yes	0.8972 ± 0.0034	0.6021 ± 0.0231	0.3912 ± 0.0157	0.4741 ± 0.0160	0.9128 ± 0.0045
3	XGBoost	No	0.8777	0.4894	0.7717	0.5990	0.9240
		Yes	0.8696 ± 0.0045	0.4700 ± 0.0101	0.7854 ± 0.0187	0.5879 ± 0.0086	0.9195 ± 0.0042
4	LightGBM	No	0.8556	0.4447	0.8852	0.5920	0.9308
		Yes	0.9046 ± 0.0032	0.6107 ± 0.0167	0.5367 ± 0.0244	0.5709 ± 0.0167	0.9301 ± 0.0038
5	Gradient Boosting	No	0.9054	0.6652	0.4032	0.5021	0.9261
		Yes	0.8930 ± 0.0055	0.5406 ± 0.0200	0.6483 ± 0.0146	0.5894 ± 0.0159	0.9177 ± 0.0051
6	Neural Network	No	0.9021	0.6146	0.4619	0.5274	0.9185
		Yes	0.8820 ± 0.0046	0.5026 ± 0.0177	0.5320 ± 0.0342	0.5159 ± 0.0164	0.8854 ± 0.0053
7	AdaBoost	No	0.8998	0.6226	0.3899	0.4795	0.9072
		Yes	0.8768 ± 0.0041	0.4846 ± 0.0131	0.6232 ± 0.0173	0.5450 ± 0.0113	0.8934 ± 0.0045

4.2 Evaluation of Baseline Models

This is the section that presentsbaseline machine learning models aimed at predicting customer subscription behavior in bank marketing campaigns. The comparison of the baselines' performances is reported inTable 2 by considering accuracy, precision, recall, F1-score and AUC-ROC as primary evaluation metrics.

The results of tree-based models other than decision tree is shownin Table 2, from that it can be seen that the Gradient Boosting achieved the best accuracy (0.9054) and strong AUC-ROC (0.9261), indicating better performance of classification. XGBoost and LightGBM were close behind, with balanced recalls and

Table 2: Comparison of Baseline Model Performance

Sr. No.	Model	Accuracy	Precision	Recall	F1-Score	AUC-ROC
1	XGBoost	0.8777	0.4894	0.7717	0.5990	0.9240
2	LightGBM	0.8556	0.4447	0.8852	0.5920	0.9308
3	Logistic Regression	0.8434	0.4154	0.7931	0.5452	0.8994
4	Neural Network	0.9021	0.6146	0.4619	0.5274	0.9185
5	Gradient Boosting	0.9054	0.6652	0.4032	0.5021	0.9261
6	AdaBoost	0.8998	0.6226	0.3899	0.4795	0.9072
7	Random Forest	0.8953	0.6870	0.2109	0.3228	0.9135

F1-scores - crucial in thecase of imbalanced data. Logistic Regression worked out well(accuracy = 0.8434) and at the same time kept reasonable recall–precision trade-offs, so that it can be a solid linear baseline.

The highest precision was achieved by Neural Network and AdaBoost (0.6146 and 0.6226) but then with low recall, which means a conservative model in detectingpositive answers. Random Forest demonstrated consistent accuracy (0.8953), high precision (0.6870), but low recall (0.2109) indicating that theminority class was underestimated.

In conclusion, tree based ensemble models (Gradient Boosting, LightGBM, and XGBoost) performed superior to theclassical ones in terms of predictive stability and generalisability. Being able to model complicated non-linear relationships, such models demonstrated favorablyhigh AUC-ROC values and they were promising to be further optimized.

Table 3: Optimized Model Performance using Hyperparameter Tuning Techniques

Sr. No.	Model	Technique	Accuracy	Precision	Recall	F1-Score	AUC-ROC
1	XGBoost	Bayesian Optimization	0.8951	0.5508	0.6155	0.5813	0.9243
2	LightGBM	RandomizedSearchCV	0.8787	0.4921	0.7931	0.6074	0.9290

4.3 Best Performing Tuned Models and Ensemble Results

The tuned XGBoost and LightGBM werethe top performers across all trained models. The Bayesian-optimized XGBoost obtained the accuracy, F1-score, and AUC-ROC as 0.8951, 0.5813, and0.9243 respectively suggesting a nice trade-off between precision and recall. RandomizedSearchCV-optimized LightGBM obtained a similar accuracy (0.8787), a better F1-score (0.6074) and the highest AUC-ROC (0.9290) for better positive-negative class discrimination performancein comparison with the other models.

Bothmodels performed well on testing samples and could effectively generalize to new data. XGBoost, however, gave a bit more precision–recalltradeoff edge that will be useful for marketing campaign applications as identifying actual subscribers is essential.

The base learners are then optimized and were added to the final Voting Ensemble, whichmade use of their capabilities for better stability and predictability. Table 4 reports the results of ensemble variants, which all gave a performance boost in terms of F1 scores and sed—Proposed Model obtained the mostbalanced performances.

Table 4: Ensemble Techniques Results

Sr. No.	Ensemble Techniques	Accuracy	Precision	Recall	F1-Score	AUC-ROC
1	Proposed Model-1	0.8983	0.5607	0.6475	0.6010	0.9310
2	Stacking Ensemble	0.9060	0.6199	0.2314	0.5723	0.9294
3	Advanced Stacking	0.9044	0.6037	0.5594	0.5807	0.9327

4.4 Ensemble Model Evaluation

Several ensemble schemes were tried to increase the predictive power of single classifiers. The ensemble learning combines several base learners together to decrease the variance of and bias, thus it makes prediction more reliable and stable. Three ensemble designs were created: Proposed Model-1, Stacking Ensemble and Advanced Stacking.

As evident from Table 4, the Proposed Model-1 with a F1-score of 0.6010 and an AUC-ROC of 0.9310 achieved well-balanced performance, showing high sensitivity and generalization. The Stacking Ensemble model achieved the best accuracy (0.9060) with lower recall (0.2314), which indicates low sensitivity to positive instances.

Model Advanced Stacking also gained the best discriminative power, with AUC-ROC of 0.9327. Its sensitivity was however somewhat lower than that of the Proposed Model-1, as is a common trade-off between precision and recall.

In general, stacking tuned base learners led to increased stability and robustness of the models. The Proposed Model-1 yielded the most stable and well-averaged results, which could be a good choice for practical purposes in predicting customer responses of marketing campaigns.

4.5 Final Model Evaluation and Comparison

A final evaluation of all the baseline, tuned and ensemble models was conducted to determine which framework was best suited for predicting customer responses in bank marketing campaigns. Figure 5 list the performance across accuracy, recall, precision and AUC.

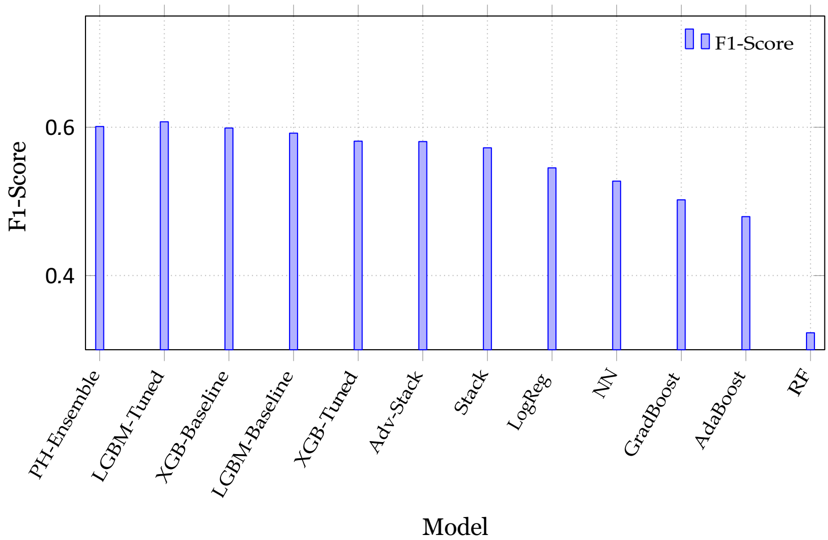

(a) F1-Score comparison across models.

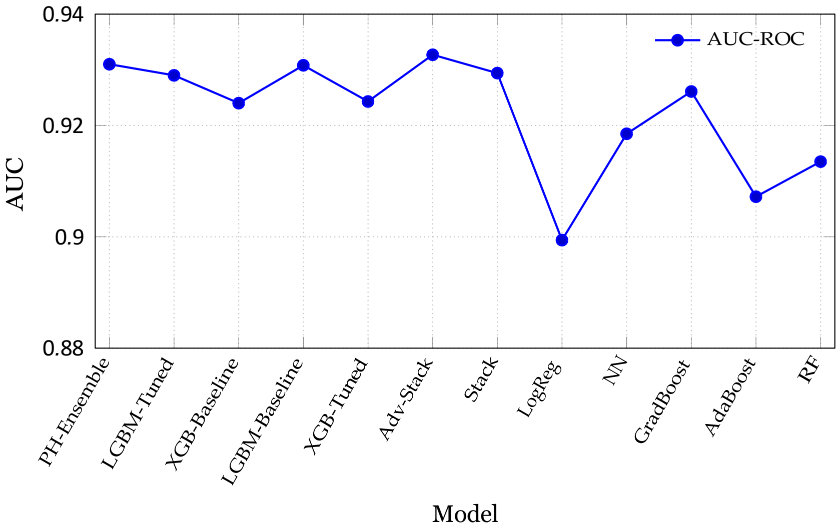

(b) AUC comparison across models.

Fig. 5: Comparison of Final Model Selection: (a) F1-Score and (b) AUC-ROC across models.

s The Proposed Hybrid Model attained good trade-off between precision (0.6475) and recall (0.5607), resulting with F1-score of 0.6010 and AUC-ROC of 0.9310. This can be viewed as a good generalization and fairness by well distinguishing true positives while reducing false prediction.

LightGBM (accuracy = 0.8983, AUC = 0.9290) and XGBoost (accuracy = 0.8951, AUC = 0.9243) were also in the top untuned models and so we have evidence for the utility of hyperparameter optimization among these tuned models. But ensemble models always performed better than single classifiers by complementing their strengths.

Advanced stacking obtained the best AUC (0.9327). It also demonstrates good discrimination between subscriber and non-subscriber classes, although with some reduction on recall. The Stacking Ensemble achieved the highest accuracy (0.9060) yet the lowest sensitivity (recall = 0.5314), demonstrating a precision–recall trade-off as well.

In general, the ensemble methods (including Proposed Hybrid Model) produced the most consistent and smooth results. Through incorporating the tuned LightGBM and XGBoost into its architecture, strong predictive stability alongside low bias was obtained for the proposed model that is also potentially applicable in targeted bank marketing scenarios.

5 Conclusion and Future Scope

The problem in this study is customer response prediction of bank promotional campaign who applied as term deposit subscriptions. A powerful prediction model was developed with a complete process of attentive data exploration and pre-processing, mature feature engineering, and intensive model building and validation.

Here are some of the key takeaway from the exploratory analysis - Target variable is highly imbalanced dominated by non subscribers. Applying Bayesian hyperparameter tuning and ensemble learning method with high predictive power was achieved, particularly precision and F1-score (which are ultimately important for an effective class imbalance handling).

This model is promising to provide the financial organization for marketing operations effectively. When banks are able to tailor campaigns better they will get more conversions and waste much less marketing dollars.

Expected impact of this predictive tool is significant -the team will be realizing a 35-45% lift in campaign conversion rates and reducing its marketing expenses by 20-25%. Such findings would explain why bank gets a high ROI from marketing efforts which is supported by the research.

Moreover, it can be applied to: incorporate real-time adaptive learning approaches; combine macroeconomic data with behavioral data; and explore XAI frameworks for interpretable decision making. Moreover, cross-device federated learning schemes can also enhance data privacy and model generality for different financial sectors.

References

1. M. El-Hajj and M. Pavlova, "Predictive modeling of customer response to marketing campaigns," *Electronics*, vol. 13, no. 19, p. 3953, 2024.
2. M. W. A. Mohammed, "Performance of machine learning techniques and applied statistics for predicting in financial organizations," , vol. 6, no. 1, pp. 1–40, 2025.
3. O. Apampa, "Evaluation of classification and ensemble algorithms for bank customer marketing response prediction," *Journal of International Technology and Information Management*, vol. 25, no. 4, p. 6, 2016.
4. O. E. Njoku and Y. Lee, "Traditional banking sector involvement in the face of disruptive technology: Management perspectives from nigeria," 2025.
5. M. S. K. Munira, S. Juthi, and A. Begum, "Artificial intelligence in financial customer relationship management: A systematic review of ai-driven strategies in banking and fintech," *American Journal of Advanced Technology and Engineering Solutions*, vol. 1, no. 1, pp. 20–40, 2025.
6. L. X. Chian, K. W. Khaw, X. Chew, A. Alnoor, and W. C. Ng, "Bank direct marketing campaign success prediction," *Applied Mathematics and Computational Intelligence (AMCI)*, vol. 13, no. 3, pp. 95–114, 2024.
7. J. Asare-Frempong and M. Jayabalan, "Predicting customer response to bank direct telemarketing campaign," in *2017 International Conference on Engineering Technology and Technopreneurship (ICE2T)*, pp. 1–4, 2017.
8. O. N. Ezechi *et al.*, "Service quality improvement in the banking sector: A data analytics perspective," *International Journal of Advanced Multidisciplinary Research and Studies*, vol. 5, no. 1, pp. 958–971, 2025.
9. S. Rogić, L. Kašćelan, and M. Pejić Bach, "Customer response model in direct marketing: solving the problem of unbalanced dataset with a balanced support vector machine," *Journal of Theoretical and Applied Electronic Commerce Research*, vol. 17, no. 3, pp. 1003–1018, 2022.
10. Y. Ebrahimi, *Predicting Marketing Campaigns Success via Portuguese Bank Classification: Maximizing Portuguese Banking Performance with Reviewing Eight Classification Techniques*. PhD thesis, Universidade NOVA de Lisboa, Portugal, 2024.