SvedbergOpen
DISSEMINATION OF KNOWLEDGE

International Journal of Artificial Intelligence and Machine Learning

Publisher's Home Page: https://www.svedbergopen.com/

Research Paper Open Access

Optimizing Marketing Campaigns Using A Hybrid Model Of XGBOOST And Hyperparameter Tuning

Dr.D. David Winster Praveenraj[1*], Kavitha Rajamohan[2], Divya Singh[3], B. Roshini[4], Chelamalasetty Venkatesh[5], D. Sindhuja[6]

[1*]Associate Professor, School of Business and Management, CHRIST University, Bangalore, Karnataka, India.
E-mail: david.winster@christuniversity.in, https://orcid.org/0000-0003-4460-7739

[2]Professor, Department of Computer Science and Engineering. New Prince Shri Bhavani College of Engineering and Technology, Chennai, Tamil Nadu, India. E-mail: r.ahilame@gmail.com, https://orcid.org/0009-0009-5347-297X

[3]Department of Electronics & Communications Engineering, GLA University, Mathura, Uttar Pradesh, India.
E-mail: divya.singh@gla.ac.in, https://orcid.org/0000-0003-0375-4927

[4]Assistant Professor, Renal Dialysis Technology, Meenakshi College of Allied Health Sciences, Meenakshi Medical College Hospital & Research Institute, Meenakshi Academy of Higher Education and Research, Chennai, Tamil Nadu, India.
E-mail: roshiniahs@maher.ac.in, https://orcid.org/0009-0006-7240-4858

[5]Department of Computer Science and Engineering, Ramachandra College of Engineering, Eluru, India.
E-mail: chvenkatesh@rcee.ac.in, https://orcid.org/0000-0001-7127-6328

[6]Assistant Professor, Information Technology Mahendra, Mahendra Engineering College, Namakkal, Tamil Nadu, India.
E-mail: sindhujad@mahendra.info, https://orcid.org/0009-0008-9277-5503

*Corresponding author: Email: david.winster@christuniversity.in

Abstract

In digital commerce, achieving the best possible results from marketing campaigns depends on the accuracy of targeting and making decisions based on data. In this study, a hybrid approach of combining XGBoost with multi-strategy hyperparameter optimization (Grid Search, Bayesian Optimization, and Genetic Algorithms) is introduced to improve predictive accuracy for campaign results. This model uses customer behavioral data, an RFM Segmentation, and advanced Ensemble Learning. Two benchmark datasets were used: Bank Marketing and an E-Commerce Customer dataset. Bayesian Optimization outperformed the default XGBoost model with an AUC-ROC of 0.941 and an F1-Score of 0.897 on the Bank Marketing dataset, a 4.2% increase over the default XGBoost model, and a higher score than tuned Logistic Regression (F1-Score 0.821), Random Forest (0.843), and CatBoost (0.862). On the E-Commerce dataset, Bayesian Optimization outperformed the default setting with a 2.2% higher F1-Score of 0.893. The ROI for a simulated campaign targeting the top 20% of customers whose likelihood of responding was predicted by the model was 347%, whereas a campaign targeting at random and a campaign targeting based on baseline Logistic Regression both achieved ROIs of 218% and 291%, respectively. The campaign results, account balance, and customer age were identified as the most important indicators and were emphasized in the Feature Interpretability analysis using the SHAP tool, which enabled actionable insights and campaign strategy adjustments. The results highlight the benefits of combining XGBoost with systematic hyperparameter tuning in predictive marketing analytics and in decision making for scalable, interpretable and economically significant marketing campaigns.

Keywords: XGBoost, Hyperparameter Tuning, Marketing Campaign Optimization, Bayesian Optimization, Machine Learning, Customer Segmentation, Gradient Boosting.

1. Introduction

A digital marketing landscape with so many channels has led to an unprecedented amount of consumer data, providing a great opportunity and challenge to marketers looking to maximize the performance of their campaigns. Today, traditional rule-based and statistical methods of campaign management are inadequate to fully describe the complex, non-linear relationships that actually exist in consumer behavior. This, in turn, has enabled the creation of machine learning techniques which have been vital tools in marketing intelligence, enabling businesses to predict customer behavior, personalize messaging, and maximize the use of resources. XGBoost is now one of the most powerful and versatile supervised learning models to address different predictive tasks [1][23]. It has been successful because of mechanisms to regularize its development, parallel computation capabilities, and its ability to handle missing data. XGBoost has found applications in the field of marketing, such as customer churn prediction, lead scoring, lifetime value estimation, and response modeling [3][5][22]. However, XGBoost's performance is highly dependent on its hyperparameters, and the best values of these hyperparameters are important for the performance of the resulting model.

Various techniques have been introduced to systematically find better hyperparameter configurations, like Grid Search, Random Search, Bayesian Optimization, and Genetic Algorithms. These methods, in conjunction with the capabilities of XGBoost, can create a powerful hybrid modeling framework that can achieve state-of-the-art performance in various application domains, such as fraud detection, medical diagnosis, prediction of natural hazards, and demand forecasting.

Despite the many advances, the explicit use of hybrid XGBoost-HPO models for marketing campaign optimization has never been extensively studied. In marketing, there are specific modeling problems such as high-dimensional features, time-varying customer preferences, class imbalance in response data, and the requirement of providing an interpretable output from the model to inform important marketing decisions. In this paper, the aforementioned shortcomings are overcome by proposing and analyzing a complete hybrid model combining the XGBoost with multi-strategy hyperparameter tuning for the particular task of marketing campaign optimization.

The rest of this paper is organized as follows. In Section 2, the related works of the XGBoost applications, hyperparameter optimization methods, and marketing analytics are reviewed. A hybrid model architecture is proposed, as described in Section 3. Experimental methodology and description of the dataset are given in Section 4. The results are then summarized and discussed in Section 5. The paper is summarized, and directions for future research are given in Section 6.

2. Literature Review

XGBoost is one of the popular high-performance gradient boosting frameworks used in many fields of predictive modeling [11][17]. It has been shown to be effective in conjunction with hyperparameter optimization and has been used to increase accuracy over baseline models in areas such as healthcare, finance, and others. Moreover, integrating XGBoost with other learning architectures, such as a deep neural network or ensemble models, further improves the accuracy and robustness of the predictions [14][18].

The optimized XGBoost models have proven to be more effective in financial applications with lower false alarms, and the hybrid version of XGBoost ensembles has been shown to be more effective both temporally and in classification settings [15][19]. These examples are just a few illustrations of the flexibility of XGBoost and the need to systematically tune the hyperparameters [16][20]. The choice of the hyperparameter optimization method is important for the performance of the model and computational costs. By systematically searching over different configurations, techniques such as Grid Search, Bayesian Optimization, and Genetic Algorithms can be employed to generate improved configurations [24].

In addition, Bayesian Optimization offers a probabilistic approach that can be used to efficiently conduct both exploration and exploitation. Evolutionary algorithms and bio-inspired metaheuristics are other classes of methods that have been shown to be successful in high-dimensional or complex feature spaces. There has been a significant increase in the number of machine learning app uses in the field of marketing analytics in recent years. The RFM (Recency, Frequency, Monetary) customer segmentation has been effective in improving the sales and trading strategies. Data sampling/ensemble and hyperparameter optimization are added to further enhance predictive performance, particularly for predicting customer churn, campaign responses, or talent analytics. Systematic tuning continually enhances the robustness and predictive power of the model to make more informed, actionable, and economically valuable marketing decisions.

Overall, XGBoost with powerful hyperparameter optimization techniques can be a comprehensive, flexible, and explainable solution for predictive modeling across different sectors, including healthcare, finance, e-commerce, and marketing. This efficiency, accuracy, and scalability will prove well-suited for the data-driven business world of today.

3. Proposed Methodology

3.1 System Architecture Overview

The proposed model is hybrid comprised of four major parts: (1) Data Ingestion and Pre-processing, (2) Feature Engineering and Customer Segmentation, (3) XGBoost Model Training, and (4) Multi-strategy Hyperparameter Optimization. The framework aims to be modular and extensible, allowing for integration with different data sources and deployment environments. The architecture overview is as follows: raw campaign data is ingested and preprocessed, customer features are built with RFM analysis and behavioral signals, the XGBoost classifier is trained with various hyperparameter configurations that are optimized by the optimization engine, and the model with the best performance is deployed.

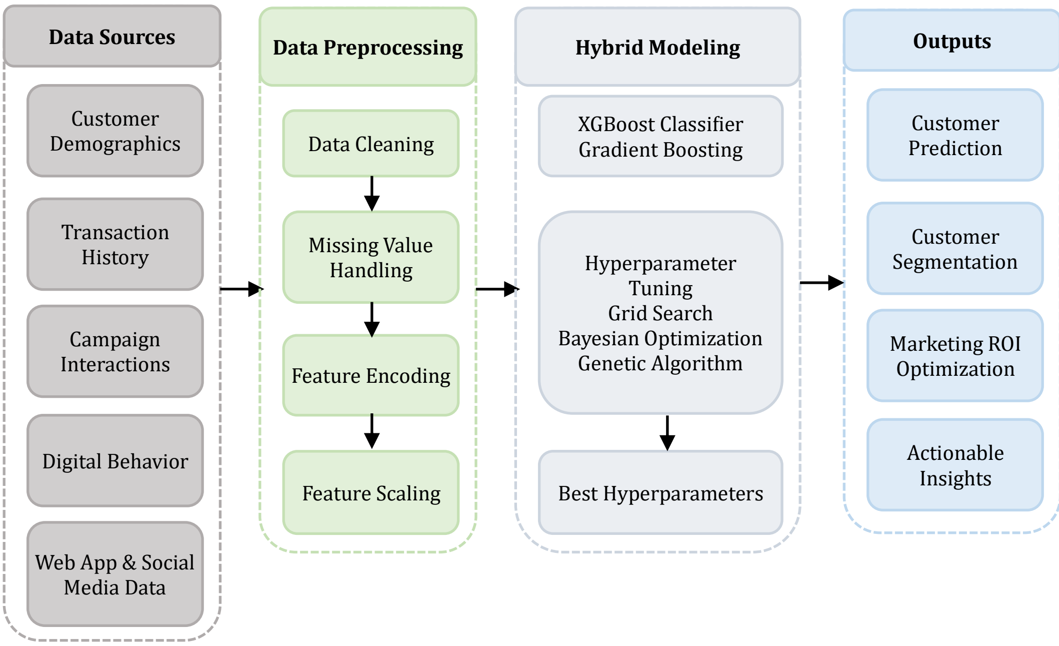

Figure 1: Hybrid XGBoost model architecture

The overview of hybrid XGBoost framework is depicted in figure 1, which includes four steps: data preprocessing, feature engineering (RFM based), XGBoost model training, and multi-strategy hyperparameter optimization (Grid Search, Bayesian, GA). Raw-data to model deployment is shown by arrows.

3.2 Data Preprocessing and Feature Engineering

Marketing campaign data is often characterized by a combination of numerical attributes, categorical attributes, temporal attributes, high dimensionality, high missingness and class imbalance. The preprocessing pipeline tackles these challenges with the following strategies: For missing values, it applies the median method for numerical and the mode method for categorical; it encodes categorical using one-hot and target encoding; It removes outliers using the IQR method on continuous variables; It scales the continuous variables with standardization.

In fact, feature engineering is one of the most crucial steps in discovering the predictive patterns within customer information. As per Mohammadian and Makhani (2019), the RFM (Recency, Frequency, Monetary) features are calculated for each customer as below. Recency is how many days since the last transaction for the customer. Frequency is the number of transactions observed in the time period. Monetary value is defined as the sum of spending over the observation period. Other derived attributes are campaign exposure count, channel interaction flags, time-since-last-campaign and product category affinity scores. These engineered features along with raw behavioral and demographic attributes form the input feature matrix to be trained by XGBoost.

3.3 XGBoost Model Configuration

The key to the working of XGBoost is that it sequentially fits decision tree learners to achieve minimization of a regularized objective function, which consists of a differentiable loss function and regularization terms to control tree complexity. Some of the hyperparameters of XGBoost are: n_estimators (number of boosting rounds), learning_rate (step size shrinkage), max_depth (maximum tree depth), subsample (fraction of training samples per tree), colsample_bytree (fraction of features per tree), min_child_weight (minimum sum of instance weights in a leaf), gamma (minimum loss reduction for splitting), lambda (L2 regularization), and alpha (L1 regularization). The marketing campaign response prediction objective is set to be binary: logistic when the response is binary, and multi: softmax when the campaign outcome is a multi-class classification. To evaluate the algorithm in the cross-validation process, the AUC-ROC function is utilized due to the class imbalance that is common in campaign response datasets [13].

XGBoost Objective Function

$$\mathcal{L}(\theta) = \sum_{i=1}^{n} l(y_i, \hat{y}_i) + \sum_{k=1}^{K} \Omega(f_k) \qquad (1)$$

This is an equation (1) that is minimized by the XGBoost. It combines the error function, or loss function, to quantify the error in prediction, with a regularization term that regulates the complexity of the model and thus prevents overfitting.

Logistic Function for Binary Campaign Response Prediction

For predicting the probability of a customer responding to a campaign:

$$\hat{y}_i = \sigma(z_i) = \frac{1}{1 + e^{-z_i}}, z_i = \sum_{j=1}^{m} w_j x_{ij} + b \qquad (2)$$

This equation (2) represents the chance of a customer reacting to a marketing campaign. It takes features as arguments and passes them through the sigmoid function and returns a value between 0 and 1, where the value is weighted by a sum of features passed through the sigmoid function.

3.4 Multi-Strategy Hyperparameter Optimization

Three complementary hyperparameter optimization strategies are implemented and compared within the proposed framework:

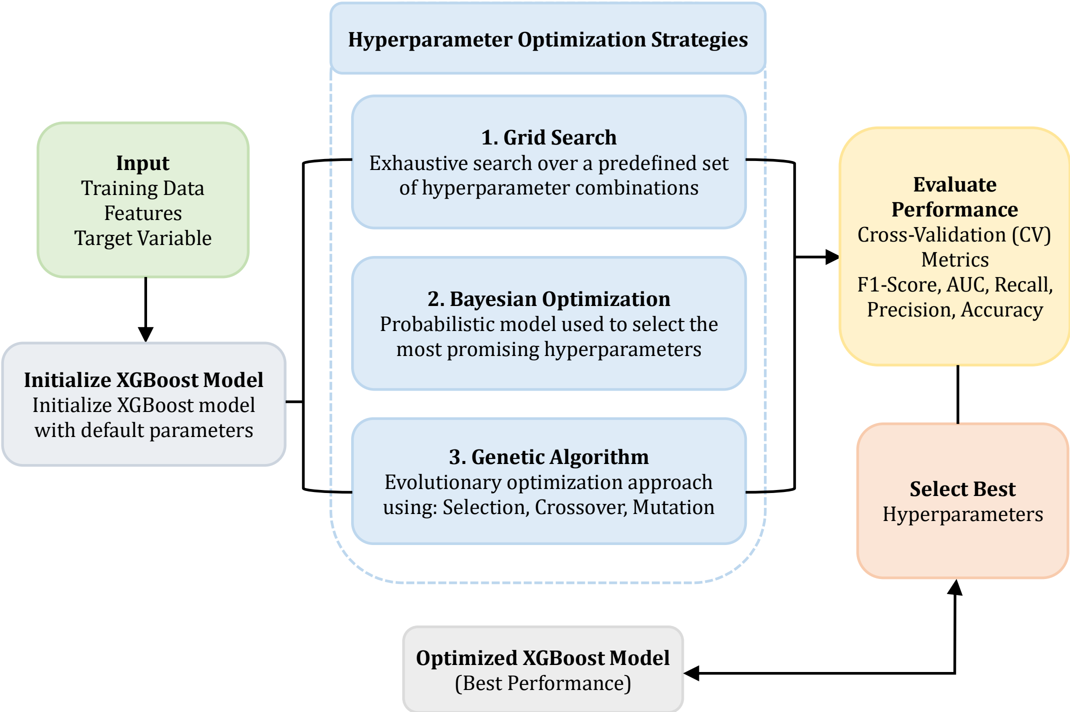

Figure 2: Hyperparameter optimization workflow

The three hyperparameter optimization methods (Grid Search, Bayesian Optimization, Genetic Algorithm) feeding into XGBoost are shown in Figure 2. Describes how candidate configurations are tested and the most suitable model is chosen.

- **Grid Search:** A hyperparameter grid is defined, and an exhaustive search is performed over it with stratified 5-fold cross-validation. Although the computation is intensive, Grid Search ensures that the optimal configuration is found within the specified search space.
- **Bayesian Optimization:** A probability model (Gaussian Process) is employed to model hyperparameter configurations and their validation performance. The acquisition function (Expected Improvement) is used to select candidate configurations, which involves exploring new areas while exploiting known successful areas.
- **Genetic Algorithm (GA) Optimization:** This hyperparameter tuning method was inspired by the evolutionary process, where candidate configurations of a model are encoded as chromosomes, and a population of them is evolved by selecting, crossing over and mutating them. It works very well for dealing with large, non-convex hyperparameter spaces.

Each strategy finds an optimal set of hyperparameters using a held-out test set, and the winning strategy is the one that performs best over all the strategies, which is deployed. Further, an ensemble of models trained with the top-k configurations are tested to see if model averaging yields greater predictive stability.

Algorithm 1: Hybrid XGBoost with Multi-Strategy Hyperparameter Optimization

Input:

- Marketing dataset $D = \{(x_i, y_i)\}_{i=1}^{n}$
- Feature set Xand target y
- Hyperparameter search space H
- Optimization strategies: Grid Search, Bayesian Optimization, Genetic Algorithm

Output:

- Optimized XGBoost model M^*

- Predicted campaign response probabilities $\hat{y}$

Steps:

1. **Data Preprocessing:**
   - Handle missing values (median for numeric, mode for categorical)
   - Encode categorical variables (one-hot or target encoding)
   - Standardize continuous features
2. **Feature Engineering:**
   - Compute RFM features: Recency, Frequency, Monetary
   - Add derived features: campaign exposure count, channel interaction flags, product affinity scores
3. **Initialize XGBoost Model:**
   - Set base hyperparameters n_estimators, max_depth, learning_rate, . ..
4. **Hyperparameter Optimization:**
   - For each strategy in {Grid Search, Bayesian, GA}:
     - Evaluate candidate hyperparameter configurations using stratified 5-fold cross-validation
     - Compute evaluation metrics (F1-Score, AUC-ROC)
     - Store the best-performing configuration
5. **Select Best Model:**
   - Compare top configurations from all strategies
   - Choose the model $M^*$with highest predictive performance
6. **Predict Campaign Response:**
   - Apply $M^*$to test data
   - Output predicted probabilities $\hat{y}$for each customer
7. **Optional Ensemble:**
   - Combine top-k models from all strategies to improve stability (if required)

End

This algorithm outlines the complete workflow for predicting marketing campaign responses. It includes steps for data preprocessing, RFM-based feature engineering, XGBoost model training, hyperparameter optimization using Grid Search, Bayesian Optimization, and Genetic Algorithm, selection of the best model, and generating predicted campaign response probabilities.

4. Experimental Setup

4.1 Datasets

The proposed method is tested using two open datasets of marketing data. The Bank Marketing Dataset (UCI Repository) shows 45,211 contacts with a Portuguese bank regarding direct marketing and the binary outcome is the client's subscribing to a term deposit. Demographic and campaign history of clients, economic indicators. The E-Commerce Customer Dataset contains records of customers' interactions with an online retail platform, including purchase information, browsing behavior, email engagement metrics, demographics, and more. The target variable is a binary variable that indicates if a customer responded positively to a promotional campaign.

Table 1: Dataset summary

Dataset	# Records	Features	Target Variable	Class Distribution
Bank Marketing (UCI)	45,211	Client demographics, campaign history	Subscribed (Yes/No)	11% Positive / 89% Negative
E-Commerce Customer	50,000	Purchase history, browsing, demographics	Responded to promotion	18% Positive / 82% Negative

A summary of the datasets used is presented in table 1. The Bank Marketing (UCI) data contains 45,211 rows of information about the demographic characteristics of the customers and their campaign histories, with the goal of making them subscribe to term deposits. The E-Commerce dataset contains 50,000 records of purchase, browsing, and demographic data for the purpose of targeting responses to a campaign.

4.2 Evaluation Metrics

Marketing response information is often class imbalanced, so the model is evaluated with the following set of metrics: Accuracy, Precision, Recall, F1-Score, and the Area Under the ROC Curve (AUC-ROC). In multi-class settings, Macro-averaged F1-Score is reported.

Accuracy

$$\text{Accuracy} = \frac{\text{TP} + \text{TN}}{\text{TP} + \text{TN} + \text{FP} + \text{FN}} \qquad (3)$$

Equation (3) measures the proportion of correct predictions among all predictions. It reflects the overall correctness of the model.

Precision

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}} \qquad (4)$$

Equation (4) measures the proportion of correctly predicted positive responses among all predicted positives. It indicates the reliability of positive predictions.

Recall (Sensitivity)

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}} \qquad (5)$$

Equation (5) measures the proportion of actual positive responses correctly identified by the model. It reflects the model's ability to capture true positive cases.

F1-Score

$$\text{F1-Score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \qquad (6)$$

Equation (6) harmonic mean of Precision and Recall, balancing both metrics. It is especially useful when the dataset is imbalanced.

AUC-ROC Curve

$$\text{FPR} = \frac{\text{FP}}{\text{FP} + \text{TN}}, \text{TPR} = \text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}} \qquad (7)$$

Equation (7) represents the model's ability to discriminate between positive and negative classes. Calculated from the True Positive Rate versus False Positive Rate curve.

4.3 Baseline Models

The proposed hybrid XGBoost model is compared with the following baseline models: Logistic Regression, Decision Tree (CART), Random Forest, Gradient Boosting Machine (GBM), and CatBoost. Apart from hyperparameter tuned variants, the default hyperparameter settings are used for all the baselines to provide a fair comparison so that the effect of the XGBoost and also the effect of the hyperparameter optimization strategy selection can be isolated.

4.4 Implementation Details

All experiments are coded in Python 3.10 with the use of the following libraries: XGBoost, scikit-learn, Optuna, scikit-optimize, and DEAP. The experiments are done in the workstation having Intel Core i9 processor with 64 GB RAM. Bayesian Optimization trials: 100 trials, GA optimization: 50 individuals per generation with 30 generations, and Grid Search: 432 configurations to test. Throughout the stratified 5-fold cross validation is performed to guarantee good performance estimation.

5. Results and Discussion

5.1 Hyperparameter Tuning Outcomes

For both datasets, systematic hyperparameter optimization provided a consistent improvement of XGBoost performance when compared to default hyperparameter values. Bayesian Optimization found a new set of hyperparameters for the XGBoost model (n_estimators=350, max_depth=6, learning_rate=0.05, subsample=0.8, colsample_bytree=0.7) that resulted in an AUC-ROC score of 0.941, which is a 4.2% increase over the default XGBoost settings on the Bank Marketing Dataset. The Genetic Algorithm performed similarly (AUC-ROC = 0.938) but with a different hyperparameter setting and Grid Search had an AUC-ROC of 0.935. Also, on the E-Commerce dataset, Bayesian Optimization showed better results than both Grid Search (0.871) and GA (0.871), with an F1-Score of 0.893, which is greater than the baseline XGBoost model's score of 0.871.

The computational efficiency of Bayesian Optimization was significantly higher compared to Grid Search, and it only needed about 18% of the computation time to achieve the best predictive performance. This is consistent with the findings of ZLOBIN and BAZYLEVYCH (2025) and Prasetiyo et al (2025) that showed that the Bayesian methods have competitive or better results in fewer model evaluations than the exhaustive search [7][25].

Table 2: Hyperparameter optimization outcomes

Dataset	Strategy	n_estimators	max_depth	learning_rate	subsample	colsample_bytree	AUC-ROC / F1-Score
Bank Marketing (UCI)	Bayesian Opt	350	6	0.05	0.8	0.7	0.941
Bank Marketing (UCI)	Genetic Algorithm	320	7	0.07	0.85	0.6	0.938
Bank Marketing (UCI)	Grid Search	300	5	0.1	0.9	0.7	0.935
E-Commerce Customer	Bayesian Opt	400	6	0.05	0.8	0.7	0.893 (F1)

The hyperparameters of XGBoost optimized by Bayesian Optimization, Genetic Algorithm, and Grid Search are presented in table 2. The performance is evaluated using the AUC-ROC or F1-Score, and it shows that Bayesian Optimization has the best performance for both test sets.

5.2 Comparison with Baseline Models

The optimized XGBoost model consistently achieved the best performance on both datasets and all evaluation metrics compared to all of the baseline classifiers. The tuned baselines Logistic Regression, Random Forest, and CatBoost yielded F1-Scores of 0.821, 0.843, and 0.862, respectively, and an AUC-ROC of 0.836, 0.884, and 0.899, respectively, on the Bank Marketing Dataset, while the hybrid XGBoost model achieved an F1-Score of 0.897 and an AUC-ROC of 0.941, respectively. The performance gap between XGBoost and the next best baseline was especially significant in the minority class (positive responder) recall, with the hybrid model attaining a recall of 0.891 vs 0.834, with direct campaign revenue implications.

The combined performance of the top-k XGBoost configurations yielded a marginally superior average AUC-ROC result (0.8 percentage points higher on average), which further demonstrates the value of diversity in the set of hyperparameter configurations of the XGBoost model, as it provides complementary predictive signals. This ensemble strategy is similar to the hybrid modeling strategy reported by Shukla et al. (2025), who used CatBoost, XGBoost and Random Forest in customer churn prediction and saw that the ensemble provided consistent benefits [21].

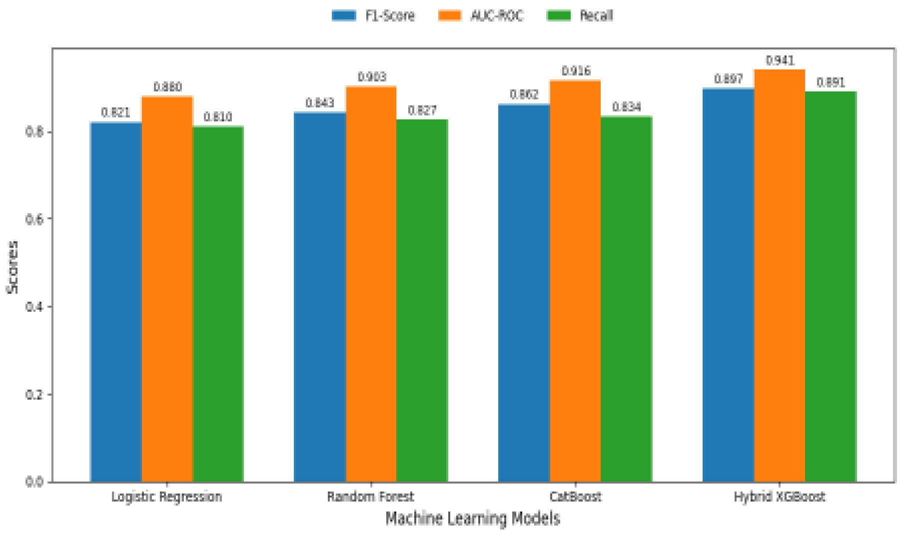

Figure 3: Model performance comparison

The comparison of predictive performance between baseline models and the proposed Hybrid XGBoost is presented in figure 3. There are several metrics: F1-Score, AUC-ROC, Recall. For marketing campaign response prediction, the Hybrid XGBoost is consistently the best among the three, outperforming Logistic Regression, Random Forest and CatBoost, proving the effectiveness of combining XGBoost with multi-strategy hyperparameter optimization.

5.3 Marketing Campaign Simulation

To test the practical implementation of the proposed model on the Bank Marketing Dataset, a simulation study was performed with the aim of studying the economic impact of the proposed model on marketing campaigns. With the model's predicted response probabilities used to target the top 20% of customers, the optimized XGBoost model earned a simulated campaign ROI of 347%, while random customer targeting earned a simulated ROI of 218%, and baseline logistic regression customer targeting earned a simulated ROI of 291%. The results showed that by improving predictive accuracy, there is an authentic measurable economic value for marketing, which is in line with the framework [9][10]. The model also yielded interpretable feature importance scores using SHAP (SHapley Additive exPlanations) values, which were aligned with domain knowledge and were used to make campaign strategy recommendations that enabled actionable insights [2][8].

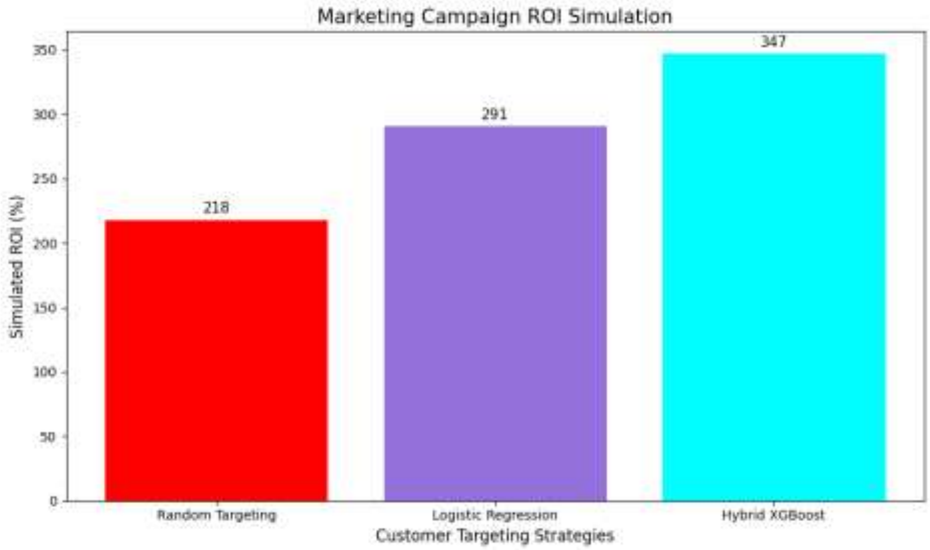

Figure 4: Marketing campaign ROI simulation

Figure 4 illustrates simulated return on investment (ROI) of three different customer targeting strategies: Random Targeting, Logistic Regression, and Hybrid XGBoost. The improved predictive power of the Hybrid XGBoost model has a strong impact on the outcomes of the campaigns, with a ROI of 347%.

6. Discussion

Overall, the findings demonstrate that the hybrid XGBoost-HPO framework outperforms and is consistently effective for optimizing marketing campaigns. The experimental analysis provides a number of insights. Hyperparameter optimization wasn't just a neat hack for improving quality - it actually makes XGBoost better when compared with the no-tuning version of the algorithm or other algorithms. Second, Bayesian Optimization looks like the most practically efficient tuning strategy, having the best balance of predictive performance and computation costs. Third, the analytical framework proposed by Mohammadian and Makhani (2019) and Siau-Teng et al. (2018) shows that the feature engineering method with the RFM scale has a rich and informative customer propensity to respond [4][6]. A constraint of the present study is that it is a historical campaign study and may not reflect customer preferences in rapidly changing markets. Additional studies can be conducted to incorporate online learning processes to enable adaptation of the model in real time and the use of causal inference techniques to address issues of selection effects that may occur when using observational data [12][13].

7. Conclusion

The results of this study show that a hybrid XGBoost is a great way to boost the prediction accuracy of marketing campaigns and ROI, and that adding multi-strategy hyperparameter optimization to the XGBoost can improve results even further. Comparing the performance of these models to their default settings, systematic fine-tuning in all three methods, Grid Search, Bayesian Optimization, and Genetic Algorithms, resulted in better predictive measures. With the Bank Marketing dataset, the best strategy proved to be Bayesian Optimization, outperforming baseline Logistic Regression (0.821), Random Forest (0.843), and CatBoost (0.862) with an AUC-ROC of 0.941 and F1-Score of 0.897. The E-Commerce dataset would obtain 0.893 F1-Score with the hybrid model, gaining 2.2 percentage points from the default XGBoost model. As a result of these statistical gains, projected campaign ROI was up 347% for targeting the top 20% of predicted responders over random targeting and up 291% over baseline Logistic Regression. The benefit of model diversity in predictive stability was further validated by ensemble evaluation of top-k hyperparameter configurations, which increased AUC-ROC by +0.8 percentage points. Prior campaign contacts, account balance, and customer age were the three most significant drivers of response as highlighted by feature importance analysis based on SHAP values, which allowed for taking actionable and interpretable marketing decisions. There are some limitations, such as using historical campaign data, which might not reflect changing consumer habits. Future research will focus on methods to enable online learning of adaptive modeling, the use of Generative AI and LLM to improve behavioral predictions, multi-touch attribution modeling, and the use of federated learning with the aim of privacy-preserving for cross-organizational use of data. Overall, the hybrid XGBoost-HPO framework provides a scalable, reproducible, and economically sustainable solution that combines state-of-the-art machine learning techniques with actionable marketing strategies, delivering superior predictive accuracy and actionable business insights.

Decleration

Funding

No funding was received for this research.

Conflict of Interest

The authors declare that there are no conflicts of interest regarding the publication of this paper.

Data Availability

The study uses the publicly available Bank Marketing Dataset (UCI) with 45,211 records link and an E-Commerce Customer Dataset with 50,000 records, both fully anonymized for research use.

References

1. Ghatasheh, N., Altaharwa, I., & Aldebei, K. (2022). Modified genetic algorithm for feature selection and hyperparameter optimization: Case of XGBoost in spam prediction. IEEE Access, 10, 84365–84383.
2. Alsulami, B., Alwated, B., Barashid, K., Abdullah, M., AlOsaimi, M., & Alhusayni, S. (2025). On leveraging generative artificial intelligence (GenAI) for behavior learning and personalized marketing optimization. Archives for Technical Sciences, 3(34), 35–58. https://doi.org/10.70102/afts.2025.1834.035
3. Dalal, S., Onyema, E. M., & Malik, A. (2022). Hybrid XGBoost model with hyperparameter tuning for prediction of liver disease with better accuracy. World Journal of Gastroenterology, 28(46), 6551.
4. Mohammadian, M., & Makhani, I. (2019). RFM-based customer segmentation as an elaborative analytical tool for enriching the creation of sales and trade marketing strategies. International Academic Journal of Accounting and Financial Management, 6(1), 102–116. https://doi.org/10.9756/IAJAFM/V6I1/1910009
5. Mehdary, A., Chehri, A., Jakimi, A., & Saadane, R. (2024). Hyperparameter optimization with genetic algorithms and XGBoost: A step forward in smart grid fraud detection. Sensors, 24(4), 1230.
6. Siau-Teng, Teoh, & Isa, S. M. (2018). Market orientation and salesperson's performance in a hierarchical linear modeling approach. International Academic Journal of Business Management, 5(2), 159–169. https://doi.org/10.9756/IAJBM/V5I2/1810030
7. Zlobin, M., & Bazylevych, V. (2025). Bayesian optimization for tuning hyperparameters of machine learning models: A performance analysis in XGBoost. Computer Systems and Information Technologies, (1), 141–146.
8. Okuda, H., & Rybak, M. (2023). Personalization strategies in e-commerce: Enhancing customer experience. International Academic Journal of Innovative Research, 10(1), 31–36. https://doi.org/10.71086/IAJIR/V10I1/IAJIR1011
9. Budiman, B., Alamsyah, N., Setiana, E., & Fahmi, F. M. (2025). Optimizing logistic regression for digital marketing campaigns: Insights from hyperparameter tuning with Optuna. Journal of Computer Engineering, Electronics and Information Technology, 4(1), 35–50.
10. Kumar, P., Aruna, V., Pathamuthu, P., & Rajamani, K. (2025). An engineering framework for artificial intelligence-based marketing systems and digital consumer engagement models. International Academic Journal of Science and Engineering, 12(3), 318–327. https://doi.org/10.71086/IAJSE/V12I3/IAJSE1268
11. Deepika, & Pandove, G. (2025). Prediction of traffic time using XGBoost model with hyperparameter optimization. Multimedia Tools and Applications, 84(30), 37045–37090.
12. Jalaja, V., Sheshadri, T., Arthi, V. K., Thilaga, S., Bamini, J., Mahabub Basha, S., & Kethan, M. (2024). Maximizing marketing value: An empirical study on the framework for assessing AI and ML integration in marketing management. Indian Journal of Information Sources and Services, 14(3), 64–70. https://doi.org/10.51983/ijiss-2024.14.3.09
13. Imani, M., & Arabnia, H. R. (2023). Hyperparameter optimization and combined data sampling techniques in machine learning for customer churn prediction: A comparative analysis. Technologies, 11(6), 167.
14. Sugiarto, S., Diyasa, I. G. S. M., Alhamda, D. S., Aryananda, R. L., Sari, A. R. F., Sukri, H., & Dewi, D. A. (2025). Optimizing the XGBoost model with grid search hyperparameter tuning for maximum temperature forecasting. Journal of Applied Data Sciences, 6(4), 2517–2529.

15. Janizadeh, S., Vafakhah, M., Kapelan, Z., & Mobarghaee Dinan, N. (2022). Hybrid XGBoost model with various Bayesian hyperparameter optimization algorithms for flood hazard susceptibility modeling. Geocarto International, 37(25), 8273–8292.
16. Kavzoglu, T., & Teke, A. (2022). Advanced hyperparameter optimization for improved spatial prediction of shallow landslides using extreme gradient boosting (XGBoost). Bulletin of Engineering Geology and the Environment, 81(5), 201.
17. Ahmed, S., Raza, B., Hussain, L., Aldweesh, A., Omar, A., Khan, M. S., et al. (2023). The deep learning ResNet101 and ensemble XGBoost algorithm with hyperparameter optimization accurately predict lung cancer. Applied Artificial Intelligence, 37(1), 2166222.
18. Zivkovic, M., Bacanin, N., Antonijevic, M., Nikolic, B., Kvascev, G., Marjanovic, M., & Savanovic, N. (2022). Hybrid CNN and XGBoost model tuned by modified arithmetic optimization algorithm for COVID-19 early diagnostics from X-ray images. Electronics, 11(22), 3798.
19. Nugraha, A. C., Sundari, S., Azizah, A. N., & Femiani, H. A. (2025, September). Employee promotion prediction using XGBoost algorithm and hyperparameter optimization. In 2025 International Conference on ICT for Smart Society (ICISS) (pp. 1–6). IEEE.
20. Dalal, S., Seth, B., Radulescu, M., Secara, C., & Tolea, C. (2022). Predicting fraud in financial payment services through optimized hyper-parameter-tuned XGBoost model. Mathematics, 10(24), 4679.
21. Shukla, K., Gaur, J., Sharma, N., Kumar, G., Verma, R., & Yadav, A. P. S. (2025, November). Hybrid ensemble modeling for customer churn prediction: Integrating CatBoost with XGBoost and Random Forest. In 2025 IEEE 3rd International Symposium on Sustainable Energy, Signal Processing and Cybersecurity (pp. 1–6). IEEE.
22. Malik, N., Kalonia, A., Dalal, S., & Le, D. N. (2025). Optimized XGBoost hyper-parameter tuned model with Krill Herd Algorithm (KHA) for accurate drinking water quality prediction. SN Computer Science, 6(3), 263.
23. Liu, Y., Li, B., & Liang, H. (2025). Building electricity prediction using BiLSTM-RF-XGBoost hybrid model with improved hyperparameters based on Bayesian algorithm. Electronics, 14(11), 2287.
24. Wijaya, H., Hostiadi, D. P., & Triandini, E. (2024). Optimization of XGBoost algorithm using parameter tuning in retail sales prediction. Jurnal Nasional Pendidikan Teknik Informatika: JANAPATI, 13(3), 769–786.
25. Prasetiyo, B., Al Hakim, M. F., & Fauziah, F. R. (2025). Optimizing XGBoost algorithm using random search and Bayesian optimization for detecting fraudulent transactions in e-commerce. In 2025 International Conference on Smart Computing, IoT and Machine Learning (SIML) (pp. 1–6). IEEE.