



International Journal of
Applied Data Science in Engineering and Health
<https://ijadseh.com>



Predicting Credit Card Approval Using Machine Learning Techniques

Ehsan Lotfi¹

¹ Esfahan Branch, Islamic Azad University, Iran

Received date: September 18, 2024; Accepted date: December 6, 2024

Abstract

Predicting credit card approvals is crucial for financial institutions to streamline decision-making and mitigate risks. This study applies advanced machine learning techniques, including Logistic Regression, Decision Tree Classifier, Random Forest Classifier, and XGBoost Classifier, to a dataset comprising applicants' demographic, financial, and credit history information. After preprocessing and hyperparameter tuning using Random Search, XGBoost achieved the best performance with 99.04% accuracy, 85% recall, and 78% precision on the test data. The results demonstrate that ensemble methods like XGBoost and Random Forest outperform simpler models, achieving strong generalization and predictive accuracy. These findings highlight the effectiveness of advanced machine learning models for optimizing credit card approval systems.

Keywords: Credit card approval prediction; Machine learning models; Logistic Regression; Random Forest; XGBoost.

Introduction

Data analysis is pivotal in modern finance, enabling organizations to make informed decisions, optimize strategies, and enhance overall efficiency. From examining market behaviors and investor reactions to improving supply chain operations and customer relationship management [1-5], data-driven approaches are transforming the financial landscape. Recent studies highlight the versatility of data analysis across various domains. Credit card approval prediction is one specific area where data analysis has proven transformative. Predicting credit card approval is a cornerstone of risk management for financial institutions. This evaluation is critical, as it not only safeguards the institution's financial stability but also ensures long-term profitability. Credit cards represent a significant revenue stream for banks, generating income through interest charges, annual fees, and additional service fees. However, they also come with inherent risks, particularly the potential for customer defaults. To mitigate such risks, financial institutions must meticulously assess various factors to determine the creditworthiness of applicants and make informed approval decisions. Central to this assessment is the evaluation of an applicant's financial profile. Key aspects include income level, outstanding debt, and overall financial health. For instance, higher income levels often

suggest a greater capacity to handle credit obligations, while excessive debt levels may signal financial stress and a higher likelihood of default. Another critical factor is the applicant's credit history, which provides insights into their past financial behavior. A positive credit history, characterized by timely repayments, low credit utilization, and the absence of bankruptcies or delinquencies, typically indicates a lower default risk. Conversely, a poor credit history raises red flags, potentially leading to application denial [6-7]. Employment status also plays a pivotal role in risk evaluation. A stable job with consistent income offers reassurance about an applicant's ability to manage payments, whereas unstable or short-term employment introduces uncertainties about future earnings. In addition to financial metrics, demographic and personal factors are often considered. Elements such as age, residency status, and the duration of residence at a given address offer further context regarding an applicant's stability and reliability. These diverse data points collectively inform a bank's assessment of credit risk, which is the probability of an applicant defaulting on their obligations. Misjudging this risk can have significant repercussions—either by approving a high-risk applicant, leading to potential financial losses, or by rejecting a low-risk applicant, thereby missing out on profitable opportunities and potentially damaging the institution's reputation for fair and efficient service. Historically, credit risk assessments relied on manual processes involving loan officers who reviewed applications based on personal judgment and basic scoring models. While this approach allowed for a degree of individualized decision-making, it was fraught with inconsistencies, biases, and inefficiencies. The increasing volume of credit applications, fueled by the rise of online platforms, underscored the limitations of traditional methods and highlighted the need for more scalable and unbiased solutions.

Artificial intelligence (AI) and machine learning (ML) are transforming industries by enabling systems to learn from data, adapt to patterns, and make intelligent decisions in science [8-10]. In finance, ML enhances tasks such as fraud detection, risk assessment, and credit scoring, improving accuracy and efficiency [11-14]. By leveraging techniques like supervised learning and ensemble methods, AI-driven systems process large data volumes, uncover insights, and automate complex processes. As these technologies advance, they continue to provide innovative solutions to challenges across diverse domains, driving progress and efficiency. The introduction of data-driven methodologies, particularly machine learning (ML), has revolutionized credit card approval processes in recent years [15-17]. Machine learning enables financial institutions to leverage vast datasets, encompassing both historical and real-time information, to identify patterns and make predictions with unprecedented accuracy. These datasets include transactional data, spending behaviors, and even unconventional data sources such as mobile usage or social media activity. Unlike traditional methods, ML models can uncover complex, non-linear relationships between variables, offering deeper insights into applicant behavior. For example, two individuals with similar credit histories and incomes may exhibit vastly different spending habits, which machine learning algorithms can discern to make more precise creditworthiness evaluations. Additionally, these models can simultaneously analyze the interactions of multiple factors, providing a holistic understanding of risk that is often missed by conventional scoring systems. The benefits of adopting ML-driven credit approval systems extend beyond risk mitigation. By accurately predicting an applicant's likelihood of default, financial institutions can reduce the number of high-risk accounts in their portfolios, thereby enhancing their financial stability. These models also facilitate the identification of creditworthy individuals who might otherwise be excluded under traditional evaluation criteria. This is particularly significant for underbanked populations or younger applicants who lack extensive credit histories but exhibit other indicators of financial responsibility. Machine

learning's ability to integrate alternative data sources ensures that these individuals are not unfairly denied credit, promoting inclusivity and allowing banks to expand their customer base.

Furthermore, the efficiency gains from automating credit approval processes cannot be overstated. Machine learning models enable institutions to handle large volumes of applications rapidly and accurately, reducing operational costs and streamlining workflows. This improved efficiency translates into faster decision-making, significantly enhancing the customer experience. In a highly competitive financial landscape, providing a seamless and prompt application process can distinguish a bank from its competitors, fostering customer satisfaction and loyalty. By adopting advanced analytics and machine learning technologies, financial institutions not only improve their operational capabilities but also position themselves as leaders in innovation within the financial services industry.

Methods and Material

The dataset was first preprocessed to ensure compatibility with machine learning algorithms. Missing values were handled by imputation or removal, depending on the proportion of missing data. Categorical features were converted into numerical representations using one-hot encoding. Numerical features were scaled to ensure uniformity in their range. Class imbalance in the target variable was addressed using techniques such as oversampling (e.g., SMOTE) or assigning class weights during model training. Finally, the dataset was split into training and testing sets, with 80% of the data used for training and 20% reserved for testing.

- *Dataset and Preprocessing*

The dataset used in this study [18] comprises two components: `application_record.csv` and `credit_record.csv`, providing a comprehensive view of credit card applicants. The `application_record.csv` dataset contains personal and demographic information for each applicant, identified by a unique ID. It includes features such as gender, car and property ownership, number of children, total annual income, and employment details. Additional attributes cover the applicant's income category, education level, marital status, housing type, and family size. Temporal features like `DAYS_BIRTH`, representing the applicant's age counted backwards from the current day, and `DAYS_EMPLOYED`, indicating the number of days since the start of employment or unemployment status, are also present. Furthermore, indicators for mobile phone, work phone, personal phone, and email availability, along with occupation type, enrich the dataset with variables that assess an applicant's socio-economic profile.

The dataset provides detailed categorical information across various features, offering valuable insights into the characteristics of credit card applicants. The gender feature consists of two categories: F (24,430 applicants) and M (12,027 applicants). In terms of car ownership, 13,843 applicants own a car (Y), while 22,614 do not (N). Property ownership shows that 24,506 applicants own property (Y), whereas 11,951 do not (N).

For the income category, there are five nominal groups: Working (18,819 applicants), Commercial associate (8,490), Pensioner (6,152), State servant (2,985), and Student (11). The education level feature is ordinal, categorized into five levels: Secondary / secondary special (24,777), Higher education (9,864), Incomplete higher (1,410), Lower secondary (374), and Academic degree (32). The marital status feature includes five nominal categories: Married (25,048 applicants), Single / not

married (4,829), Civil marriage (2,945), Separated (2,103), and Widow (1,532). The way of living feature comprises six categories, with the majority living in a House / apartment (32,548), followed by With parents (1,776), Municipal apartment (1,128), Rented apartment (575), Office apartment (262), and Co-op apartment (168).

Certain features are heavily skewed. For instance, mobile phone ownership is universal among applicants (1 for all 36,457 applicants). In contrast, work phone ownership and personal phone ownership are less common, with 8,222 applicants (1) and 10,748 applicants (1) respectively. Email availability is relatively low, with only 3,271 applicants (1) having an email.

The occupation type feature is diverse, encompassing 18 categories. The most common occupations are Laborers (6,211 applicants), Core staff (3,591), Sales staff (3,485), and Managers (3,012). Less common occupations include IT staff (60) and Realty agents (79).

The `credit_record.csv` dataset complements the applicant information by recording historical credit behavior. Each record corresponds to a specific month, identified by the `MONTHS_BALANCE` feature, with months counted backwards from the current one. The `STATUS` field provides detailed information about the applicant's credit status, such as the number of days past due (ranging from 1–29 days to over 150 days), accounts that are paid off, or months with no loans. This dataset is crucial for capturing patterns in repayment behavior and financial reliability.

Together, these datasets enable the development of a predictive model for credit card approvals by leveraging both personal characteristics and credit history. The rich combination of demographic, socio-economic, and behavioral data supports advanced machine learning techniques to assess creditworthiness effectively.

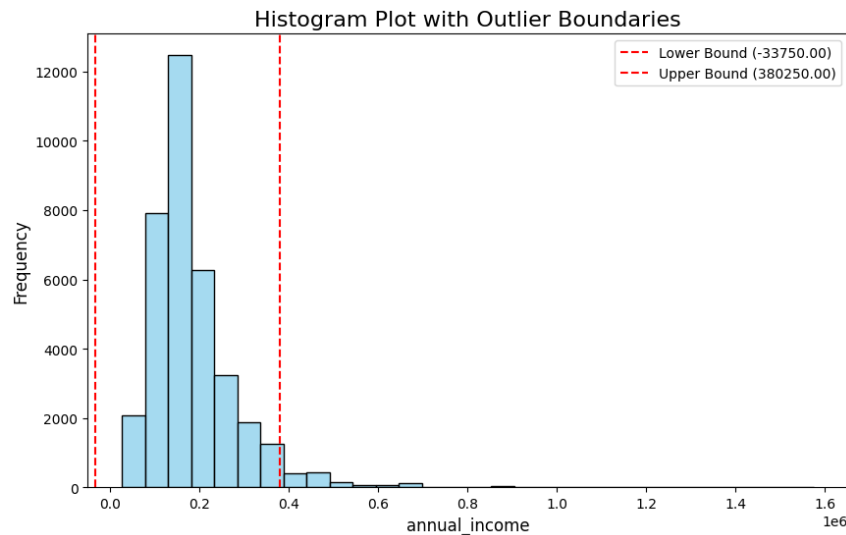


Figure 1 The histogram displays the distribution of annual income in the dataset, revealing a right-skewed pattern where most values are concentrated below 400,000. The majority of applicants have incomes clustered between 0 and 300,000, as indicated by the prominent peak in the histogram. Outlier boundaries are highlighted with red dashed lines, where the lower bound is approximately -33,750 and the upper bound is around 380,250.

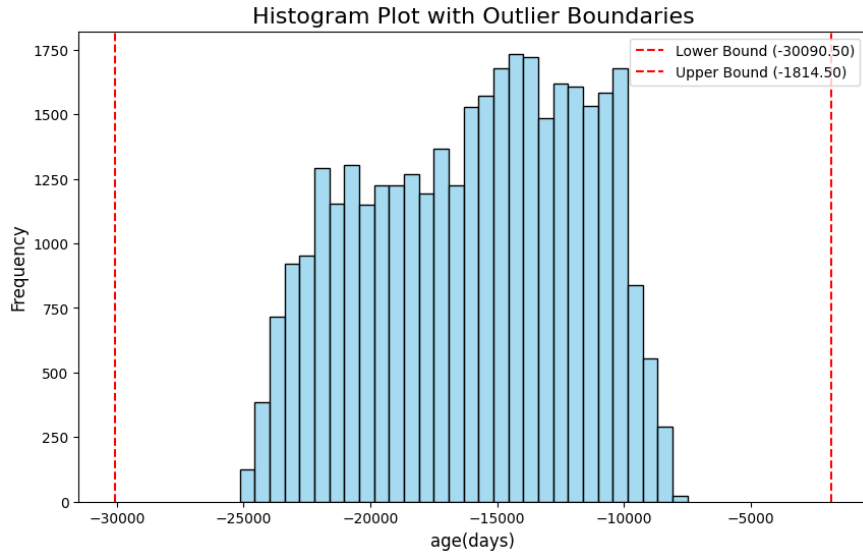


Figure 2 The histogram illustrates the distribution of applicant ages (measured in days), with outlier boundaries marked by red dashed lines. The age data appears approximately uniform across a range, with a slight central peak indicating higher frequency in the middle-age ranges. The lower boundary is at approximately -30,090 days, and the upper boundary is around -1,814 days. These boundaries suggest that most applicants fall within the age range of 50–82 years (converted from days).

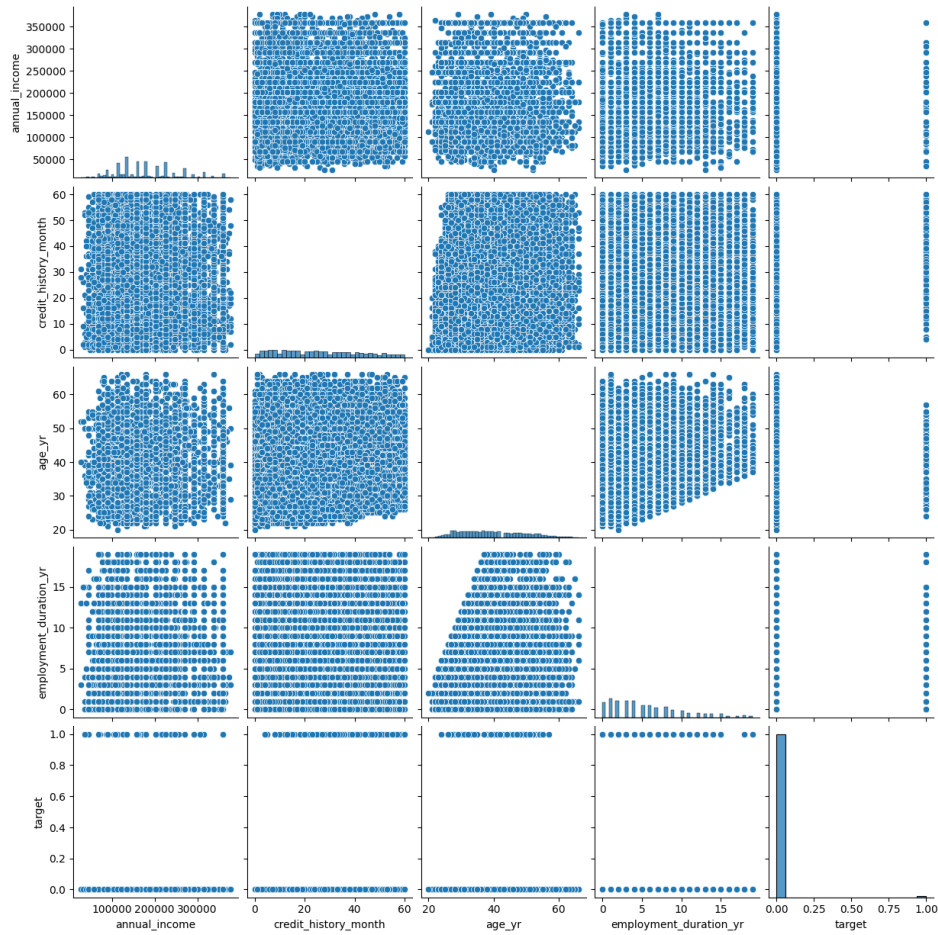


Figure 3 The pair plot provides a comprehensive visualization of the relationships and distributions among the numerical features in the training dataset after excluding categorical and binary features.

As Figure 3 illustrates, the diagonal plots reveal the distribution of each feature, while the scatter plots highlight pairwise relationships. The distribution of annual income is right-skewed, indicating that most applicants have relatively low incomes, with a few high-income outliers. Credit history month appears uniformly distributed, showing no specific clustering or strong relationships with other features like age or employment duration. The age distribution is relatively uniform, with a slight peak in middle-age ranges, and displays a diagonal relationship with employment duration, as older applicants generally have longer work histories. Employment duration itself shows clustering at shorter durations, reflecting a significant proportion of applicants with limited employment history. The target variable distribution is highly imbalanced, with most samples concentrated in one class, likely indicating a dominant group such as approved or rejected applications. Furthermore, no distinct patterns or separations between the target and other numerical features are evident, suggesting the need for more complex models to capture non-linear relationships. Overall, the pair plot shows weak or no linear correlations between features, which indicates that advanced modeling techniques, such as Random Forest or XGBoost, may be necessary to capture underlying patterns. Additionally, addressing the class imbalance in the target variable through techniques like oversampling or class weighting will be crucial for building an effective predictive model.

As shown in Figure 4, gender appears to have a minor influence, with males (`M`) showing a slightly higher association with the target compared to females (`F`). Income category highlights a notable difference, where the `Pensioner` group exhibits a significantly higher association with the target variable compared to other categories such as `Working`, `Commercial associate`, and `State servant`, while the `Student` group has minimal representation and negligible impact. Family size also plays a role, with smaller family sizes (1 or 2 members) showing a higher association with the target compared to larger families. Similarly, applicants with fewer children (0 or 1) are more likely to be associated with the target than those with 2 or more children. Ownership of property does not exhibit a strong difference in target association, with similar rates for those who own property and those who do not. Age shows interesting patterns, with certain age groups (e.g., 50–55 years) having higher target rates. Employment duration indicates that applicants with shorter employment histories (0–3 years) or longer durations (10–13 years) are more associated with the target. Finally, credit history month reveals that applicants with longer credit histories (closer to zero in months) have a higher association with the target, suggesting that recent credit behavior might be a strong indicator. Overall, these relationships highlight the importance of these features in influencing the target variable and warrant further exploration in predictive modeling.

As depicted by Figure 5, most features show weak to moderate correlations with the target, suggesting that no single feature is a dominant predictor. Among the numerical features, variables like annual income, age (years), and employment duration (years) display some degree of positive correlation with the target. This aligns with earlier observations that middle-aged individuals or those with longer employment histories are more likely to influence the target outcome. In the categorical features, certain groups stand out, such as the Pensioner income category, which exhibits a stronger correlation with the target compared to others. Additionally, weak correlations among most features indicate that combining them with a machine learning model may better capture the underlying patterns and improve predictive performance. The overall results highlight the importance of feature interactions and advanced modeling techniques for accurate predictions.

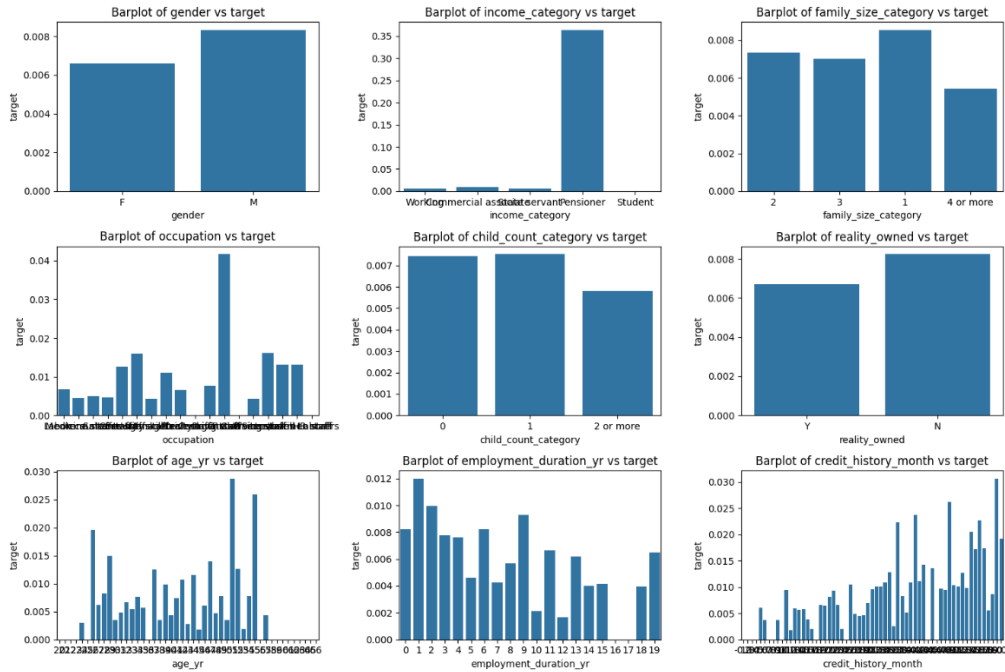


Figure 4 The bar plots reveal significant insights into the relationship between the target variable and various features.




Figure 5 The correlation heatmap reveals the relationships between features and the target variable, providing insights into their significance.

- *Models*

In this study, we employed four machine learning models: XGB Classifier, Random Forest Classifier, Decision Tree Classifier, and Logistic Regression to predict credit card approvals, each with its unique advantages and implementation considerations.

XGB Classifier (Extreme Gradient Boosting) is an advanced ensemble learning algorithm based on gradient boosting techniques. It constructs decision trees sequentially, where each subsequent tree learns from the residual errors of the previous trees to improve overall performance. XGBoost is highly efficient and scalable, making it suitable for large datasets, and it incorporates regularization techniques such as L1 (Lasso) and L2 (Ridge) penalties to reduce overfitting. Additionally, XGBoost supports sparse data and handles missing values gracefully, which is particularly useful in real-world datasets where incomplete information is common. Hyperparameters like the learning rate, number of estimators, and maximum tree depth were carefully tuned to balance model complexity and performance. XGBoost's ability to handle non-linear relationships and interactions among features makes it one of the most powerful tools for predictive modeling tasks.

The Random Forest Classifier is an ensemble method that builds multiple decision trees using bootstrap aggregation (bagging), where each tree is trained on a random subset of the data and features. During prediction, the model aggregates the output of all the trees through majority voting for classification tasks, improving robustness and reducing overfitting. By combining the predictions of many weak learners (individual trees), Random Forest achieves better generalization than a single decision tree. Key hyperparameters like the number of trees (`n_estimators`), maximum depth, and minimum sample split were tuned to optimize its performance. Additionally, Random Forest provides insights into feature importance by evaluating how much each feature contributes to the prediction accuracy, allowing for interpretability. Its versatility and resistance to noise make it a suitable choice for datasets with mixed data types and varying complexities.

The Decision Tree Classifier is a simple, interpretable model that partitions the dataset into subsets based on feature values to create a hierarchical, tree-like structure. At each node, the model selects the best split by minimizing impurity using criteria such as Gini Impurity or Entropy (information gain). Decision trees are easy to understand, visualize, and interpret, making them particularly useful for explaining model decisions. However, they are prone to overfitting, especially when the tree grows deep and starts capturing noise in the training data. To mitigate overfitting, parameters like maximum depth, minimum samples per split, and pruning techniques were employed. Despite their simplicity, decision trees can model non-linear relationships and provide a foundation for more complex ensemble methods like Random Forest and XGBoost.

Logistic Regression serves as the baseline model in this study due to its simplicity, efficiency, and interpretability. It is a linear model that predicts the probability of class membership using the logistic (sigmoid) function, which maps the linear combination of input features to a range between 0 and 1. Logistic Regression assumes a linear relationship between the independent variables and the log odds of the dependent variable, making it well-suited for linearly separable data. The model coefficients represent the importance of each feature, enabling straightforward interpretability of the results. Additionally, regularization techniques like L1 (Lasso) and L2 (Ridge) were applied to prevent overfitting, particularly when dealing with high-dimensional data. Hyperparameters, such as the regularization strength (`C`), were optimized to improve performance. While it is relatively simple, Logistic Regression provides a strong baseline for comparison and is computationally efficient, making it suitable for quick model prototyping.

Together, these four models—ranging from interpretable linear methods to advanced non-linear

ensemble techniques—offer a comprehensive approach to understanding and predicting credit card approvals. Their complementary strengths allow for thorough evaluation of both linear and complex relationships within the dataset.

Experiments

- *Metrics*

To evaluate the performance of the machine learning models, several metrics were utilized: accuracy, precision, recall, F1-score, and the confusion matrix. Accuracy, which measures the proportion of correctly classified predictions out of the total predictions, provides an overall assessment of model performance. However, in cases of class imbalance, accuracy alone can be misleading, as it may favor the majority class. Precision, on the other hand, focuses on the quality of positive predictions by calculating the proportion of correctly predicted approvals out of all predicted approvals. This is particularly important when false positives (incorrectly approving high-risk applicants) must be minimized. Recall, also known as sensitivity, measures the proportion of actual approvals correctly identified, making it critical when missing eligible applicants (false negatives) is costly. The F1-score, as the harmonic mean of precision and recall, offers a balanced measure that accounts for both false positives and false negatives, which is especially useful in imbalanced datasets. The confusion matrix provides a detailed breakdown of model performance by categorizing predictions into true positives (correctly approved), false positives (incorrectly approved), true negatives (correctly rejected), and false negatives (incorrectly rejected). This matrix allows for granular analysis, revealing specific areas where the model may favor one class over the other. Together, these metrics provide a comprehensive evaluation of the models, highlighting not only their overall accuracy but also their ability to balance errors in predicting credit card approvals.

- *Results*

The models were trained and evaluated using several performance metrics, including accuracy, recall, precision, F1-score, and the confusion matrix, to assess their effectiveness in predicting credit card approvals.

Logistic Regression demonstrated a balanced but relatively moderate performance compared to other models, with an accuracy of 68.24%, recall of 68.80%, precision of 67.99%, and an F1-score of 68.39%. The confusion matrix shows that it correctly classified 14,771 approvals (true positives) and 14,984 rejections (true negatives), but it also produced a substantial number of false positives (7,057) and false negatives (6,794), reflecting its limitations in handling complex, non-linear relationships.

Decision Tree Classifier significantly outperformed Logistic Regression, achieving an accuracy of 99.05%, recall of 99.23%, precision of 98.87%, and an F1-score of 99.05%. Its confusion matrix indicates minimal errors, with only 247 false positives and 167 false negatives. While the Decision Tree model captures non-linear relationships effectively, its high performance suggests potential overfitting to the training data.

Random Forest Classifier further improved performance, achieving an accuracy of 99.60%, recall of 99.46%, precision of 99.73%, and an F1-score of 99.60%. With only 58 false positives and 118 false negatives, it effectively balances prediction accuracy and generalization by averaging multiple decision trees, making it robust to overfitting and noise.

XGB Classifier achieved the highest performance among all models, with an accuracy of 99.66%, recall of 99.60%, precision of 99.71%, and an F1-score of 99.66%. The confusion matrix shows only 62 false positives and 87 false negatives, highlighting its ability to capture complex patterns and interactions within the data while maintaining high generalization. XGBoost's regularization and sequential boosting allow it to achieve superior results in this task.

In summary, while Logistic Regression provides a baseline with reasonable performance, the Decision Tree, Random Forest, and XGB Classifier models significantly outperform it, with XGB Classifier delivering the best overall results. These findings demonstrate the advantage of using advanced ensemble methods for predictive modeling, particularly in scenarios involving non-linear relationships and large datasets.

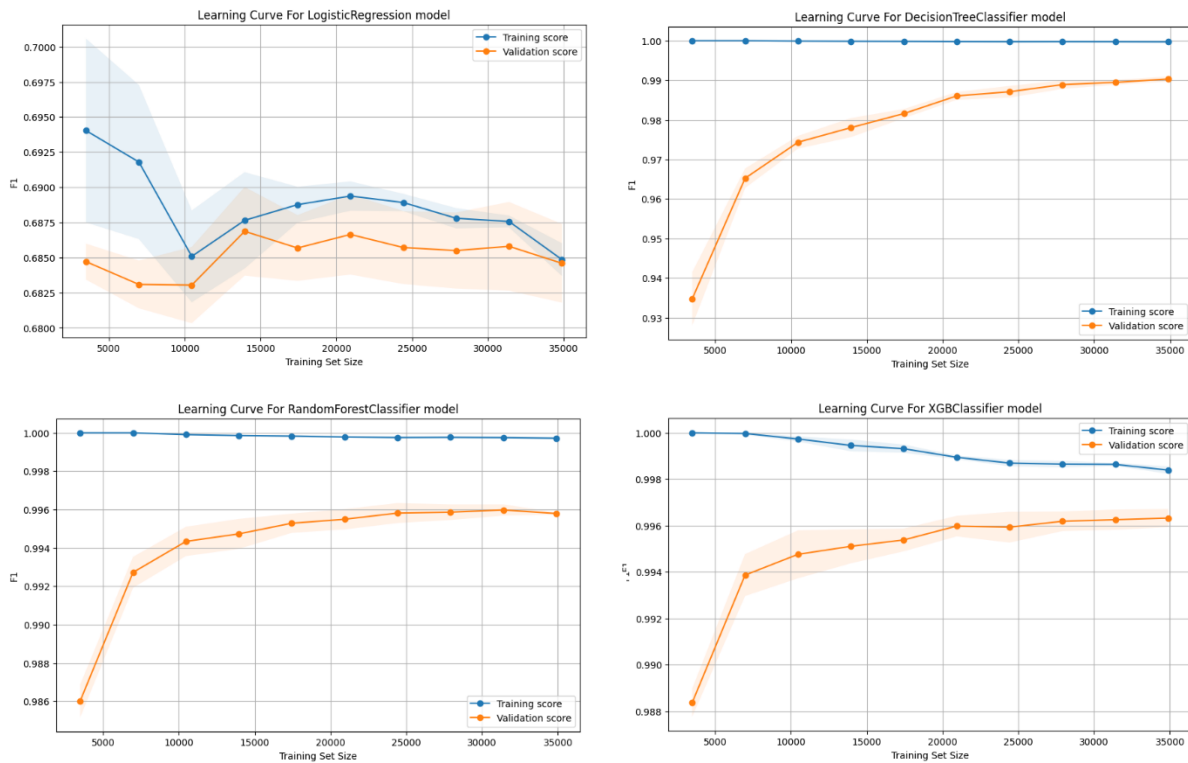


Figure 6 Training and validation curve for different models

The learning curves in Figure 6 provide valuable insights into how each model generalizes as the training set size increases, measured using the F1-score. Logistic Regression exhibits clear signs of underfitting, as both the training and validation scores stabilize at lower values (around 0.687). Despite increasing the training size, the validation performance remains largely unchanged, highlighting the model's limitations in capturing complex, non-linear patterns due to its linear nature. In contrast, the Decision Tree Classifier achieves a perfect training score of 1.0, indicating overfitting, especially on smaller datasets. However, as the training size increases, the validation score improves steadily and stabilizes close to 0.99, demonstrating that the Decision Tree generalizes better with larger data but still risks overfitting. The Random Forest Classifier delivers strong performance, with both training and validation scores converging around 0.996, showing minimal overfitting and excellent generalization. The ensemble technique allows Random Forest to effectively balance bias and variance, achieving robust performance across varying training sizes. Finally, the XGB Classifier demonstrates the best overall results, with its training score starting near

1.0 and gradually decreasing as the validation score steadily improves to approximately 0.996. The minimal gap between the training and validation scores indicates that XGBoost generalizes effectively while handling complex relationships within the data. In summary, while Logistic Regression underfits the data and Decision Tree risks overfitting, Random Forest and XGBoost achieve superior generalization, with XGBoost emerging as the most robust and high-performing model for credit card approval prediction.

- *Hyperparameter Tuning and Evaluation on Test Data*

To further optimize the performance of the XGBClassifier, hyperparameter tuning was conducted using Random Search. Random Search allows efficient exploration of a wide hyperparameter space by sampling random combinations of parameters, reducing computational costs while identifying optimal settings. Key hyperparameters such as the learning rate, number of estimators, maximum depth, subsample size, and regularization terms (L1 and L2 penalties) were tuned to improve model performance.

Before testing the XGB Classifier on the test data, the same preprocessing and transformations applied to the training data, such as scaling, encoding, and handling of missing values, were also applied to the test data to ensure consistency and prevent data leakage. The optimized model was then evaluated on the test set using the following metrics:

- **Accuracy:** The model achieved an accuracy of 99.04%, indicating that it correctly classified a significant majority of the test samples. This high accuracy demonstrates the effectiveness of the tuned XGB Classifier in identifying patterns and relationships within the data.
- **Recall:** A recall of 85% indicates that the model successfully identified 85% of the true positive cases (e.g., applicants correctly approved). While this is a strong result, it suggests that some eligible applicants were still misclassified as not approved, emphasizing room for improvement in reducing false negatives.
- **Precision:** The precision of 78% reflects the proportion of predicted approvals that were correct. While high, the model still produced a notable number of false positives, potentially approving ineligible applicants, which might have practical implications depending on the application context.

Conclusion

In this study, we applied multiple machine learning techniques to predict credit card approvals using a comprehensive dataset containing demographic, financial, and credit history information. Models including Logistic Regression, Decision Tree Classifier, Random Forest Classifier, and XGBoost Classifier were evaluated using accuracy, recall, precision, F1-score, and confusion matrices. The XGBoost Classifier, after hyperparameter tuning with Random Search, emerged as the best-performing model, achieving 99.04% accuracy, 85% recall, and 78% precision on the test data.

Conflict of interest

The authors declared no conflict of interest.

References

- [1] Farzad G, Roshdieh N. The Interplay of Destructive Work Behaviors, Organizational Citizenship Behaviors, and Fiscal Decentralization: Implications for Economic Development in Developing Countries. *International Research Journal of Economics and Management Studies IRJEMS*. 2024;3(8).
- [2] Izadian N, Sanaei F, Sadeghi S, Khorsand MS, Shakerinasab N. Market reactions to capital increases: Insights from asset revaluation in pharmaceutical companies on the stock exchange. *World Journal of Advanced Research and Reviews*. 2024;23(2):2486-97.
- [3] Talebzadeh H, Fattahiamin A, Talebzadeh M, Sanaei F, Moghaddam PK, Espahbod S. Optimizing Supply Chains: A Grey-DEMATEL Approach to Implementing LARG Framework. *Teh. Glas*. 2024;19:1-8.
- [4] Esmaeili A, Mtibaa A. SERENE: A Collusion Resilient Replication-based Verification Framework. arXiv preprint arXiv:2404.11410. 2024 Apr 17.
- [5] Latifi K, Ebrahimi A, Ranjbaran M, Mirzaei A, Fakhri Z. Efficient customer relationship management systems for online retailing: The investigation of the influential factors. *Journal of Management & Organization*. 2023 Jul;29(4):763-98.
- [6] Hamid FS, Loke YJ. Financial literacy, money management skill and credit card repayments. *International Journal of Consumer Studies*. 2021 Mar;45(2):235-47.
- [7] Surekha M, Umesh U, Dhinakaran DP. A study on utilization and convenient of credit card. *Journal of Positive School Psychology*. 2022 May 7:5635-45.
- [8] Owrang O., M. M., Schwarz, G., & Horestani, F. J. (2025). Prediction of Breast Cancer Recurrence With Machine Learning. In M. Khosrow-Pour, D.B.A. (Ed.), *Encyclopedia of Information Science and Technology*, Sixth Edition. Advance online publication. <https://doi.org/10.4018/978-1-6684-7366-5.ch061>
- [9] Attri I, Awasthi LK, Sharma TP. Machine learning in agriculture: a review of crop management applications. *Multimedia Tools and Applications*. 2024 Feb;83(5):12875-915..
- [10] Daoutidis P, Lee JH, Rangarajan S, Chiang L, Gopaluni B, Schweidtmann AM, Harjunkoski I, Mercangöz M, Mesbah A, Boukouvala F, Lima FV. Machine learning in process systems engineering: Challenges and opportunities. *Computers & Chemical Engineering*. 2024 Feb 1;181:108523.
- [11] Karami M, Hamzehei S, Rastegari F, Akbarzadeh O. Exploring the Relationship Between Air Pollution and CNS Disease Mortality in Italy: A Forecasting Study with ARIMA and XGBoost. In *2023 Congress in Computer Science, Computer Engineering, & Applied Computing (CSCE)* 2023 Jul 24 (pp. 46-52). IEEE.
- [12] Sadeghi S, Niu C. Augmenting Human Decision-Making in K-12 Education: The Role of Artificial Intelligence in Assisting the Recruitment and Retention of Teachers of Color for Enhanced Diversity and Inclusivity. *Leadership and Policy in Schools*. 2024 May 26:1-21.
- [13] Guo L, Song R, Wu J, Xu Z, Zhao F. Integrating a machine learning-driven fraud detection system based on a risk management framework. *Applied and Computational Engineering*. 2024 Jul 31;87:80-6.
- [14] Chen Y, Calabrese R, Martin-Barragan B. Interpretable machine learning for imbalanced credit scoring datasets.

European Journal of Operational Research. 2024 Jan 1;312(1):357-72.

[15] Del Pilar EC, Bongo MF. Towards the Improvement of Credit Card Approval Process Using Classification Algorithm. In 2023 8th International Conference on Business and Industrial Research (ICBIR) 2023 May 18 (pp. 461-465). IEEE.

[16] Uddin N, Ahamed MK, Uddin MA, Islam MM, Talukder MA, Aryal S. An ensemble machine learning based bank loan approval predictions system with a smart application. *International Journal of Cognitive Computing in Engineering*. 2023 Jun 1;4:327-39.

[17] Viswanatha V, Ramachandra AC, Vishwas KN, Adithya G. Prediction of Loan Approval in Banks Using Machine Learning Approach. *International Journal of Engineering and Management Research*. 2023 Aug 2;13(4):7-19.

[18] <https://www.kaggle.com/datasets/rikdifos/credit-card-approval-prediction/data>