International Journal of
Business Strategies
(IJBS)

Customer Engagement Prediction in Retail Banking with Explainable AI

Tuan Ngoc Nguyen, Quoc Giang Nguyen

International Journal of Business Strategies
ISSN 2519-0857 (online)
Vol.11, Issue 2, pp 1- 11, 2025

www.ajpojournals.org

Customer Engagement Prediction in Retail Banking with Explainable AI

Tuan Ngoc Nguyen[1*], Quoc Giang Nguyen[2]

.

ABSTRACT

Purpose: In the highly competitive retail banking environment, accurately predicting customer engagement is critical to enhancing customer satisfaction, ensuring retention, and ultimately boosting profitability.

Materials and Methods: This study embarks on a comprehensive exploration of advanced computational finance techniques by integrating Explainable Ensemble Learning (EEL) with a suite of Trustworthy Open AI tools. Utilizing state-of-the-art methods including Evidently AI for rigorous testing and LIME for interpretability this research leverages the publicly available Berka dataset.

Findings: After an intensive phase of feature engineering, extraction, and deep clustering, customers are segmented into distinct engagement categories. Four ensemble models are constructed and meticulously evaluated, with the blending model emerging as the most effective approach by achieving an impeccable AUC score of 1.000 along with outstanding accuracy, precision, recall, and F1 scores.

Implications to Theory, Practice and Policy: Only six out of 2000 data points were misclassified, underscoring the model's robustness. This paper not only highlights significant advancements in using computational finance techniques to predict customer engagement but also emphasizes the crucial role of transparency and interpretability in fostering trust within AI-based decision systems.

Keywords: *Explainable AI, Machine Learning, Ensemble Learning, Retail Banking*

INTRODUCTION

In today's dynamic banking ecosystem, where customer satisfaction and unwavering loyalty are critical components for sustaining a competitive edge, the capability to predict customer engagement with high precision has emerged as an essential requirement. This research underscores the profound impact of customer engagement, recognizing that not only do engaged customers exhibit steadfast loyalty, but they also tend to utilize an array of supplementary banking services, drive higher transaction volumes, and serve as influential brand advocates through positive word-of-mouth endorsements. Against this multifaceted backdrop, our study harnesses sophisticated computational finance methodologies to forecast customer behavior with unparalleled accuracy.

In recent years, the strategic incorporation of artificial intelligence (AI) and machine learning (ML) techniques has revolutionized the banking industry, particularly in addressing complex challenges such as customer churn. Among these techniques, ensemble learning has distinguished itself as a formidable strategy, enhancing prediction accuracy by amalgamating the diverse outputs of multiple models. Despite these advancements, conventional ML models frequently encounter difficulties in delivering interpretability a shortfall that is especially problematic in privacy-sensitive sectors like banking, where understanding the rationale behind predictions is paramount. Current conventional AI models perform well in prediction, however, such models serve as black-box models and do not provide much explanations of how a decision/prediction has been made. To address this critical gap, our research introduces an innovative framework known as Explainable Ensemble Learning (EEL), which synergizes the strengths of ensemble methods with open AI testing tools such as Evidently AI and LIME. This integration not only bolsters predictive accuracy but also provides lucid and transparent justifications for the outcomes, thereby fostering trust and yielding actionable insights for financial institutions.

The structure of this article is organized to provide a comprehensive narrative: Section 2 delineates the background and underlying motivation for the study, while Section 3 offers an in-depth review of existing literature in customer engagement prediction. Section 4 details the unique contributions of our research, and Section 5 elaborates on the employed methodologies including advanced feature engineering, deep clustering, and the construction of ensemble models coupled with a discussion of the evaluation techniques. Section 6 presents a thorough analysis of the experimental results, highlighting key performance metrics, and Section 7 provides an overarching synthesis and discussion of the findings. The paper concludes with Section 8, summarizing the research insights, followed by Section 9, which outlines prospective avenues for future investigation.

Background and Motivation

Banks are increasingly leveraging machine learning techniques, such as ensemble models, to predict customer engagement with heightened precision. Ensemble learning, a highly regarded approach in computational finance, applies to a suite of sophisticated computational methods to analyze financial data and generate robust predictions. By adopting ensemble techniques, banks are able to significantly enhance their engagement prediction capabilities—an imperative for maintaining a competitive advantage in the fast-evolving financial industry. These models are capable of processing and synthesizing vast quantities of data in an efficient manner, thereby uncovering valuable insights into customer behavior. In addition, the integration of open-source tools like Evidently AI provides a layer of transparency and control over model performance, as these tools offer rigorous testing, reporting, and monitoring functionalities. Moreover, the incorporation of explainable techniques further demystifies the model outcomes, ensuring that the results are not only accurate but also readily interpretable by stakeholders.

One of the powerful algorithms for prediction is ensemble learning which involves combining predictions from multiple models to achieve an overall improvement in accuracy and a reduction in prediction variance. Among the myriad ensemble approaches, popular methods include Stacking, Blending, and Voting. Stacking utilizes a meta-learner to integrate the diverse predictions generated by various base models, thereby synthesizing a final output that is more precise than any individual model's prediction. Blending, on the other hand, trains the meta-learner on a distinct held-out validation set, ensuring that the ensemble is finely tuned to generalize well on unseen data. Meanwhile, Voting amalgamates the predictions through techniques such as majority or weighted averaging, which further stabilizes the final prediction. This powerful ensemble strategy is not only instrumental in financial forecasting but also finds application in image recognition, natural language processing, and numerous other domains. The primary motivation behind employing ensemble learning in our study is to develop a model that significantly enhances the reliability and accuracy of customer engagement predictions in the banking sector.

Another standard practice among the researchers and practitioners of AI involves utilizing artificial intelligence systems that are developed, deployed, and utilized in an ethical, equitable, and transparent manner, ensuring that these systems remain open to scrutiny and contributions from a diverse array of stakeholders. Such systems are expected to be reliable, accurate, and devoid of bias, while simultaneously aligning with core social values and ethical principles. Evidently AI is an open-source AI which exemplifies trustworthiness through its commitment to transparency, reliability, and ethical operation. Evidently AI is an open-source tool designed to monitor and evaluate machine learning models by detecting issues related to data drift, model performance degradation, and data integrity. It plays a crucial role in MLOps by providing automated reports and dashboards for continuous monitoring. By clearly articulating the mechanisms of its machine learning solutions, openly sharing datasets and documentation, and implementing safeguards to eliminate bias and discrimination, Evidently AI fosters a trustworthy environment. Its robust framework comprises three key components—Reports, Tests, and Monitors—which collectively provide a comprehensive toolkit for assessing and ensuring the quality of data and machine learning models throughout their lifecycle. These interfaces are versatile enough to be integrated into various machine learning stacks, supporting scenarios ranging from real-time monitoring to automated pipeline testing and visual analysis. With its intuitive API and extensive array of metrics, tests, and visualizations, Evidently AI plays a pivotal role in the development of reliable and effective machine learning solutions that cater to the complex needs of modern businesses and society at large, particularly within the realms of tabular and textual data.

Most of the advanced AI algorithms work as black box models that do not offer much information about how a decision has been made within the model. Explainable AI (XAI) solves this problem and represents an active area of research dedicated to enhancing the transparency and interpretability of AI models, thereby bolstering user trust. This discipline enables individuals to understand, interpret, and communicate the decision-making processes of AI systems, which is crucial as these systems increasingly influence high-stakes decisions. The imperative for XAI stems from the need to elucidate how AI models arrive at their conclusions, particularly as these systems grow more accurate and are integrated into critical operational contexts. One notable XAI technique is Local Interpretable Model-Agnostic Explanations (LIME), which facilitates the interpretation of complex machine learning models by approximating their behavior with simpler, local linear models. LIME operates by perturbing input data around a specific instance and observing the resultant changes in predictions, thereby identifying the most influential features that contribute to the final decision. This approach allows data scientists and specialists to discern the relative importance of factors such as age, income, and historical transaction history in predicting customer behavior. By employing techniques like LIME, the interpretability and transparency of AI models are significantly enhanced, enabling banks to make more informed decisions regarding customer engagement strategies. Moreover, the integration of monitoring tools like Evidently AI further supports data scientists in rigorously testing and analyzing model construction processes, ensuring that the models are both robust and interpretable. Overall, the study seeks to determine whether the combination of ensemble models and automated machine learning techniques can reliably predict bank customer engagement by revealing the most pertinent variables, thereby transforming traditional, error-prone approaches into agile, data-driven solutions.

Related Work

In recent years, the imperative to comprehend consumer behavior and the overall experience within the banking sector has grown exponentially, as these factors directly influence customer engagement, attrition, and loyalty. To address this complex issue, researchers have extensively focused on developing quantitative models and sophisticated algorithms that meticulously analyze client data to inform judicious investment decisions. One particularly dynamic area of investigation is the application of computational finance methods within banking. These methods entail the creation of mathematical models and algorithms for analyzing financial data, such as consumer transaction data, and formulating well-informed investment choices. For instance, academics have created a technology to measure and operationalize consumer interaction in the Indian retail banking market [21]. With the help of this application, it is possible to examine consumer transaction data and spot behavioral trends that point to increased levels of involvement. Meanwhile, [11] found that experience predicts loyalty and referrals better than customer satisfaction. [13] focused on controlling risks from customer fluctuations in retail banks in Iran, applying a churn prediction model. In [19], the research recommended a course of action to minimize portfolio buying and switching actions by enhancing experience surveillance. The effect of customers' perception of corporate social responsibility on their participation intention was investigated in [19]. [16] sought to comprehend greater client involvement and experience strengthening banks' relationships with them. The study examined a range of customer engagement and customer experience problems, as well as their effects. Using unstructured data in [22], the proposed customer churn prediction model can generate meaningful insights to assist managers in creating retention plans that are tailored for various client categories. Finally, [7] investigated the impact of RQ on relational outcomes and the role of customer satisfaction in relational quality in the banking industry by contrasting clients with and without a specialized account manager.

Despite these numerous contributions to the field of customer engagement prediction, there remains a notable gap in research pertaining to the integration of Trustworthy Open AI and Explainable AI techniques. Trustworthy Open AI encompasses systems that are developed, deployed, and maintained in a transparent, accountable, and ethical manner, ensuring that they deliver reliable and unbiased outcomes. This is particularly critical in sensitive areas such as banking, where the ramifications of AI-driven decisions can profoundly affect both customers and financial institutions. On the other hand, Explainable AI is dedicated to crafting models and algorithms that offer clear and interpretable explanations of their predictions—a feature essential for cultivating trust and ensuring that decision-making processes are fair and free from discrimination. In light of these considerations, the present study is committed to leveraging Trustworthy Open AI and Explainable AI methodologies, along with an open-source real-world dataset, to significantly enhance the accuracy, transparency, and fairness of customer engagement predictions in the banking industry.

Research Contributions

This study makes several significant contributions to the domain of customer engagement prediction:

- **Data Synthesis via Deep Clustering:** By applying deep clustering techniques alongside comprehensive feature engineering on the Berka dataset, we synthesized a refined dataset that accurately represents customer behaviors. This process enabled the segmentation of customers into active and inactive clusters.
- **Advanced Ensemble Modeling:** We developed and compared four distinct ensemble models stacking, blending, voting, and rotation forest—to identify the most effective approach for predicting customer engagement.
- **Transparent Model Evaluation:** By integrating Evidently AI and employing rigorous testing protocols, we ensured that our models were evaluated transparently, with clear metrics and visualizations that validated their performance.
- **Enhanced Interpretability with LIME:** The use of LIME as an explainable AI method allowed us to unpack the model's predictions, providing granular insights into how specific features such as transaction history and credit usage influence customer engagement outcomes.

Methodologies

Our approach is structured around the development of a robust machine learning pipeline that is both predictive and interpretable. The methodology is organized into several sequential steps:

Data Acquisition and Preprocessing

We began with the publicly available Berka dataset, originally compiled by the Czech banking institution Česká spořitelna. This dataset comprises extensive transactional, demographic, and behavioral data that has been anonymized to protect customer privacy. Given the inherent complexity of the dataset, which includes multiple interrelated tables (e.g., loan, account, transaction, and demographic information), we performed rigorous data cleaning and exploratory data analysis (EDA) to assess data quality, identify missing values, and detect outliers.

Feature Engineering and Deep Clustering

To better capture the nuances of customer behavior, we engaged in extensive feature engineering and extraction. This process involved synthesizing a multitude of raw features into meaningful attributes that could be used to predict engagement. Subsequently, deep clustering techniques implemented via neural networks with multiple hidden layers and ReLU activation functions were applied. This allowed us to partition the customers into two distinct clusters (active versus inactive), thereby laying the groundwork for more targeted predictive modeling.

Ensemble Model Construction

We developed four ensemble models stacking, blending, voting, and rotation forest each of which combines the outputs of multiple base classifiers. The base models, including decision tree classifiers, random forest classifiers, and XGB classifiers, were selected based on their individual performance. Each ensemble technique was designed to integrate these base model predictions in a unique manner:

- **Stacking:** Utilizes a meta-learner trained on the predictions generated by the base models.
- **Blending:** Involves training a meta-model on a held-out validation set that incorporates both original features and base model predictions.
- **Voting:** Aggregates predictions from each base model through either majority or weighted voting.
- **Rotation Forest:** Enhances prediction accuracy by applying Principal Component Analysis (PCA) on randomly selected subsets of features, thereby reducing feature redundancy.

Model Evaluation with Evidently AI and LIME

To ensure the robustness and transparency of our models, we employed Evidently AI—a tool that offers both reporting and testing modules to evaluate model performance across several metrics, including accuracy, precision, recall, F1-score, and AUC. Moreover, we utilized LIME to dissect individual predictions, enabling us to understand the contribution of each feature in driving the model's decisions. This dual approach not only validated the overall performance of the models but also provided actionable insights into their internal mechanics.

Model Findings and Evaluations

Our ensemble models were rigorously evaluated using a comprehensive set of performance metrics. The blending model emerged as the most effective, achieving an exceptional AUC score of 1.000 along with consistently high accuracy, precision, recall, and F1 scores. The confusion matrices revealed that only a minimal number of instances 6 out of 2000 in the current dataset and 19 out of 2000 in the reference dataset were misclassified, highlighting the model's precision.

Table 1: Ensemble Learning Models Results

Models	Accuracy	Prediction	Recall	F-1 Score	AUC
Stacking	0.99	0.99	0.99	0.99	0.997
Blending	**0.991**	**0.99**	**0.99**	**0.99**	**1.00**
Voting	0.992	0.99	0.99	0.99	0.999
Rotation Forest	0.992	0.98	0.98	0.98	0.996

Table 2: Model Findings after Testing

Models	Accuracy	Prediction	Recall	F-1 Score
Stacking	0.992	0.979	0.997	0.988
Blending	**0.99**	**0.996**	**0.978**	**0.987**
Voting	0.994	0.99	0.991	0.991
Rotation Forest	0.988	0.993	0.974	0.983

A detailed performance analysis demonstrated that:

- **Stacking Model:** Delivered robust performance with an AUC of 0.997.
- **Voting Model:** Achieved an AUC of 0.999, showcasing a balanced integration of base model predictions.
- **Rotation Forest Model:** Offered competitive performance with an AUC of 0.996.

By integrating both the testing and reporting components of Evidently AI into our evaluation framework, our models demonstrated exceptional performance, achieving accuracy, precision, recall, and F-1 scores of 99%, 99%, 97%, and 98% respectively, as detailed in Table 2. These impressive metrics serve as a strong testament to the effectiveness and reliability of our predictive model. Moreover, the confusion matrices presented in Figures 5 and 6 further underscore the superior performance of our blending model; it misclassified only 6 out of 2000 data points in the current dataset and 19 out of 2000 in the reference dataset. Such results clearly illustrate the robustness and high accuracy of our AI model in classifying customer engagement.

Furthermore, the LIME-based explanations provided critical insights into feature importance. For instance, the analysis revealed that customers with an existing loan tended to exhibit lower engagement levels, whereas those with a history of frequent transactions and longer tenure were more actively engaged. These insights are invaluable for tailoring customer relationship management strategies and refining engagement models.

Description and Analysis of Findings

The comprehensive evaluation of our ensemble models underscores the effectiveness of blending multiple predictive techniques to forecast customer engagement in retail banking. By integrating deep clustering with feature engineering, we successfully synthesized a highly informative dataset that captures the subtleties of customer behavior. The resulting ensemble models, particularly the blending model, not only demonstrate superior predictive performance but also deliver meaningful interpretability through the use of LIME.

Our analysis indicates that the integration of Evidently AI enhances model transparency, enabling us to validate model performance through both quantitative metrics and qualitative visualizations. This dual approach provides a robust framework for ensuring that the predictions are reliable, unbiased, and actionable. The exceptional performance of the blending model, as evidenced by its near-perfect AUC score and minimal misclassification rate, validates the potential of explainable ensemble learning in transforming customer engagement prediction.

.

CONCLUSION

This study makes a significant contribution to the field of customer engagement prediction in retail banking by integrating Explainable Ensemble Learning with Trustworthy Open AI tools. Our innovative approach—centered on deep clustering, advanced ensemble modeling, and rigorous evaluation using Evidently AI and LIME—has resulted in a highly accurate and interpretable predictive framework. The blending model, in particular, has achieved an outstanding accuracy of 99% and an AUC of 1.000, affirming its capability to distinguish between engaged and non-engaged customers with exceptional precision.

The insights derived from the LIME analyses further reinforce the importance of key behavioral features, such as transaction count and loan status, in predicting customer engagement. These findings have practical implications for the banking industry, offering a reliable method for anticipating customer behavior and proactively addressing engagement challenges. Overall, our research advances the state-of-the-art in computational finance by presenting a novel, transparent, and effective strategy for customer engagement prediction.

Recommendation

Looking ahead, future research will focus on integrating Responsible AI principles to further enhance the accuracy, interpretability, and accountability of the computational finance pipeline. One key limitation of this study is the reliance on the dated Berka dataset; hence, acquiring up-to-date customer activity data from modern banking institutions will be imperative. In the case of privacy concerns and regulatory restrictions, future endeavors can also rely on synthetic data or federated learning approach to address data accessibility issues. In addition, future endeavors will explore the integration of more advanced deep learning techniques alongside ensemble methods, as well as the implementation of other interpretability tools—such as SHAP and Partial Dependence Plots—to provide even richer insights into model behavior. Future studies can also explore the influence of external factors to improve prediction accuracy. Incorporating real-time monitoring capabilities via the Evidently AI framework will also be a priority, ensuring the model's adaptability to evolving data landscapes and contributing to a more responsible and sustainable financial ecosystem.

.

REFERENCES

Aaleya Rasool, Farooq Ahmad Shah, and Muhammad Tanveer. 2021. Relational dynamics between customer engagement, brand experience, and customer loyalty: An empirical investigation. Journal of Internet Commerce 20, 3 (2021), 273–292.

Anna Omarini. 2022. Retail Banks' Challenges and Opportunities from Vision and Strategy to Managing People, Processes and Capital. Current Aspects in Business, Economics and Finance Vol. 3 (2022), 74–103.

Bo Li, Peng Qi, Bo Liu, Shuai Di, Jingen Liu, Jiquan Pei, Jinfeng Yi, and Bowen Zhou. 2023. Trustworthy ai: From principles to practices. Comput. Surveys 55, 9 (2023), 1–46.

Dhruti Dheda, Ling Cheng, and Adnan M Abu-Mahfouz. 2022. Long short term memory water quality predictive model discrepancy mitigation through genetic algorithm optimisation and ensemble modeling. IEEE Access 10 (2022), 24638– 24658.

Diego Monferrer, Miguel Angel Moliner, and Marta Estrada. 2019. Increasing customer loyalty through customer engagement in the retail banking industry. Spanish Journal of Marketing-ESIC 23, 3 (2019), 461–484.

Forbes Makudza. 2020. Augmenting customer loyalty through customer experience management in the banking industry. Journal of Asian Business and Economic Studies 28, 3 (2020), 191–203.

Ggaliwango Marvin and Md. Golam Rabiul Alarm. 2021. An Explainable Lattice based Fertility Treatment Outcome Prediction Model for TeleFertility. In 2021 IEEE International Conference on Biomedical Engineering, Computer and Infor- mation Technology for Health (BECITHCON). 64–68. https://doi.org/10.1109/ BECITHCON54710.2021.9893623

Ibomoiye Domor Mienye and Yanxia Sun. 2022. A survey of ensemble learning: Concepts, algorithms, applications, and prospects. IEEE Access 10 (2022), 99129– 99149.

Joyee Chatterjee. 2018. Role of Services Marketing Mix and Customer-Company Identification in building Engaged Customers. Amity Journal of Management Research 3, 1 (2018), 112–122.

Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. " Why should I trust you?" Explaining the predictions of any classifier. In Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining. 1135–1144.

Mudasir A Ganaie, Minghui Hu, AK Malik, M Tanveer, and PN Suganthan. 2022. Ensemble deep learning: A review. Engineering Applications of Artificial Intelli- gence 115 (2022), 105151.

Neena Sondhi, Baldev R Sharma, and Supriya M Kalla. 2017. Customer engagement in the Indian retail banking sector: an exploratory study. International Journal of Business Innovation and Research 12, 1 (2017), 41–61.

Nguyen, Q. G., Nguyen, L. H., Hosen, M. M., Rasel, M., Shorna, J. F., Mia, M. S., & Khan, S. I. (2025). Enhancing Credit Risk Management with Machine Learning: A Comparative Study of Predictive Models for Credit Default Prediction. *The American Journal of Applied sciences*, *7*(01), 21-30.

Nhi NY Vo, Shaowu Liu, Xitong Li, and Guandong Xu. 2021. Leveraging unstructured call log data for customer churn prediction. Knowledge-Based Systems 212 (2021), 106586

Petr Berka et al. 2000. Guide to the financial data set. PKDD2000 discovery challenge (2000).

Philipp 'Phil' Klaus and Stan Maklan. 2013. Towards a better measure of customer experience. International journal of market research 55, 2 (2013), 227–246.

R Salah Khairy, A Hussein, and HTHS ALRikabi. 2021. The detection of counterfeit banknotes using ensemble learning techniques of AdaBoost and voting. International Journal of Intelligent Engineering and Systems 14, 1 (2021), 326–339.

Rudresh Dwivedi, Devam Dave, Het Naik, Smiti Singhal, Rana Omer, Pankesh Patel, Bin Qian, Zhenyu Wen, Tejal Shah, Graham Morgan, et al. 2023. Explainable AI (XAI): Core ideas, techniques, and solutions. Comput. Surveys 55, 9 (2023), 1–33.

Ruiting Hao, Xiaoqian Xia, Siyi Shen, and Xiaorong Yang. 2020. Bank Direct Marketing Analysis Based on Ensemble Learning. Journal of Physics: Conference Series 1627, 1 (aug 2020), 012026. https://doi.org/10.1088/1742-6596/1627/1/

Soumi De, P Prabu, and Joy Paulose. 2021. Effective ML techniques to predict customer churn. In 2021 Third international conference on inventive research in computing applications (ICIRCA). IEEE, 895–902.

Sudi Murindanyi, Ben Wycliff Mugalu, Joyce Nakatumba-Nabende, and Ggaliwango Marvin. 2023. Interpretable Machine Learning for Predicting Customer Churn in Retail Banking. In 2023 7th International Conference on Trends in Elec- tronics and Informatics (ICOEI). 967–974. https://doi.org/10.1109/ICOEI56765. 2023.10125859

Teresa Fernandes and Teresa Pinto. 2019. Relationship quality determinants and outcomes in retail banking services: The role of customer experience. Journal of Retailing and Consumer Services 50 (2019), 30–41.