# Enhancing Targeting in CRM Campaigns through Explainable AI

Teoman Berkay Ayaz[1][0000-0003-4318-7357], Muhammet Furkan Özara[1][0009-0006-2652-7179], Emrah Sezer[1][0000-0002-5078-9463], Ahmet Erkan Çelik[1][0000-0001-5462-698X], and Akhan Akbulut[1,2][0000-0001-9789-5012]

[1] R&D Center of Next4biz, Sahrayicedit Mah, Pakpen Plaza, No:40/4, 34734 Istanbul, Turkey
{teoman.ayaz,furkan.ozara,emrah.sezer,erkan.celik,akhan.akbulut}@next4biz.com
[2] Department of Computer Engineering, Istanbul Kültür University, 34536 Istanbul, Turkey
a.akbulut@iku.edu.tr

**Abstract.** Modern customer relationship management (CRM) solutions are vital to firms because they streamline the administration of customer interactions, sales processes, and marketing initiatives. To fully exploit the potential of massive volumes of customer data, these platforms need help from AI techniques to quickly evaluate and extract useful insights, personalize customer experiences, and optimize decision-making to improve business outcomes. This study delves into the use of explainable AI methods like SHAP, LIME, and ELI5 to analyze CRM campaign outcomes. The purpose of this research is to discover essential traits that serve as indications for successful targeting by analyzing a dataset that captures the results of customers' interactions with campaign content as responder or non-responder. Using these methods improves interpretability and closes the gap between AI-driven decision-making and human understanding. The findings add to the field by offering clear rationales for consumer actions, which in turn helps companies fine-tune their targeting tactics and boost the efficiency of their campaigns. This study emphasizes the value of AI systems being transparent and interpretable in order to promote trust and enable data-driven decision-making in CRM contexts.

**Keywords:** Explainable AI · XAI · SHAP · LIME · ELI5 · Transparency · Interpretability · Campaign Effectiveness.

## 1 Introduction

In today's competitive business environment, customer relationship management (CRM) platforms play a pivotal role in monitoring and enhancing marketing, sales, and contact optimization. However, organizations have the difficulty of extracting valuable insights and maximizing the value of this information as the volume of client data continues to rise exponentially. For CRM platforms to succeed now, they must use cutting-edge AI techniques. Effectively analyzing enormous volumes of data, discovering previously unseen trends, and providing meaningful recommendations are all made possible by AI in CRM systems [1]. Businesses may now provide individualized experiences for a large number of customers thanks to AI-powered algorithms that can learn individual tastes, anticipate purchases, and customize interactions. By streamlining customer interactions and accessing previously unattainable insights, new CRM technologies are enabling organizations to drive the effectiveness of their marketing campaigns.

The reasoning processes of traditional AI models, like deep neural networks, are often difficult to decipher due to their opaque operation. The broad use of AI in vital areas where transparency and responsibility are paramount may be impeded by this inability to be understood. Conversely, XAI approaches aim to shed light on the reasoning behind AI decisions in order to promote transparency and aid humans in understanding the underlying dynamics. By letting users comprehend and evaluate AI-generated results, XAI bolsters ethical decision-making, promotes trust, and assists with regulatory compliance [2].

XAI's insights on customer behavior, segmentation, and personalized recommendations can have a substantial impact on businesses in the CRM space. For instance, XAI methods can be used in customer churn prediction to help explain why a given client is at risk of leaving [3]. Businesses can take specific measures to keep clients whose engagement or sentiment has dropped after learning the variables that contributed to the projection. Lead scoring is another area where XAI might be useful; this is the process by which prospective clients are given a score depending on how likely they are to become paying customers [4]. Businesses can better comprehend the decision-making process and improve their lead qualifying tactics if they are provided with an explanation of the aspects and patterns contributing to the scoring. With this additional context, firms are better able to make decisions that boost customer relationship management and the efficiency of their CRM initiatives.

The contributions of this study are two-fold. Starting with a CRM campaign dataset, it analyzes three explainable AI methods: SHAP, LIME, and ELI5. The objective of this paper is to utilize these methods in order to clarify the factors that impact customers' choices to interact with campaign content. In order to shed light on the elements that contribute to accurate targeting, this investigation dives further than typical black-box AI models into the underlying features' interpretability. Additionally, the study contributes to CRM and marketing by offering a logical structure for understanding customer behavior. To bridge the knowledge gap between humans and AI-driven decisions, tools like SHAP, LIME, and ELI5 are helpful. Organisations can enhance their targeting strategies and the overall success of their campaigns by gaining a better understanding of the variables driving customer engagement with explanations that are easy to digest.

## 2 Related Work

The emergence of XAI in recent times has provided businesses with novel opportunities to improve the transparency and effectiveness of CRM systems. This enables them to gain actionable insights and establish customer confidence through decision-making processes that are comprehensible and interpretable. By integrating XAI techniques, CRM-based solutions empower organizations to not only generate accurate predictions of customer preferences and behavior but also to provide rationales for those predictions that are easily understood. Haag et al. [5] provide a rationale for their investigation into the use of XAI for cross-selling in the energy retail industry. They highlight the potential of machine learning to improve business outcomes with an accuracy rate of 86% (AUC), which would result in significantly greater comprehension and implementation. They conduct a case study of the SHAP XAI method using data from an energy provider in order to assess how well it explains ML predictions. The results of their research illustrate the soundness and consistency of SHAP's explanations, thereby potentially enhancing the confidence of sales personnel and facilitating more accurate cross-purchase predictions.

A different initiative sought to improve the predictability and understandability of Credit Risk Models proposed by Torrent et al. [6]. In order to train an explainable ML model for use with databases related to banking, this project will make use of open data. After adjusting the hyperparameters, the authors found that, out of all the proven models, CatBoost performed the best, with an implementation producing a GINI of 0.68. The 20 most important features were identified using the SHAP software, and a complete model that is easy for humans to understand was presented, illustrating the global and local relationships between an individual's attributes and the model's prediction. One such initiative that aims to apply XAI methods in commercial contexts is the XMANAI project [7]. The authors of this study endorse the XAI methodology and the points of view expressed in it, and they also provide a groundbreaking XAI Platform that caters to the specific needs and peculiarities of the manufacturing sector. Following ethical principles, the XMANAI approach advances the state-of-the-art by offering practical and transparent methods for creating, training, and validating a catalog of interoperable XAI models that are hybrid and graph-based. These models can be applied to a variety of manufacturing problems. The major contribution of this research is a trustworthy, human-centered approach to XAI in manufacturing. This approach improves human cognition, builds confidence in the technology, and provides value-based explanations for the data, models, and outcomes.

Eswarapu et al. [8] introduced a framework for integrated customer analytics in telecommunications, which utilizes explainability and AutoML. The study's overarching goal is to provide a flexible and intuitive machine-learning paradigm for consumer data analysis. These predictions and explanations are made using XAI, H2O AutoML, and K-Means segmentation. In terms of lowering customer turnover and increasing retention rates, the results show that the framework is useful. The main contributions of the study are the thorough consumer analysis framework and the use of state-of-the-art tools such as AutoML and XAI to generate accurate predictions and explanations. A recent systematic review paper authored by Owens et al. [9] describes the use of XAI in the insurance industry.

Discussed are related publications on the subject as well as the effects of XAI on the insurance value chain (IVC) and related taxonomies. The methodology used to gather and analyze relevant data on XAI use along the IVC is detailed, along with the results of the literature review that focused on AI approaches in a systematically selected sample, according to the stated XAI criteria. In this work, the authors concentrate on the explainability of AI applications along the IVC and provide a novel evaluation of the review's results regarding the frequency of XAI along the IVC. Following an exhaustive review of the related literature, the article concludes by outlining critical factors to keep in mind when planning for the future of XAI in insurance operations.

Marín Díaz et al.'s research [10] emphasizes the importance of predicting client attrition in a Business-to-Business (B2B) scenario using machine learning models. The demand for more accurate and interpretable models is what prompted this investigation of client turnover in business-to-business settings. When interpreting the choices produced by black box algorithms, the authors use an agnostic interpretability methodology that is then applied to the problem at hand. The RFID model and the client's history of contact with the company are used to form an evaluation of the consumer. The study's primary contributions are the implementation of ML models dedicated to customer churn rate prediction in B2B settings and the deployment of an agnostic interpretability technique that is applicable to any predictive model in B2B or B2C settings. Using a one-of-a-kind dataset of unsecured consumer loans offered by a Norwegian bank, de Lange et al. propose an XAI model for bank credit evaluation [11]. The authors test their Gradient Boosting Decision Trees utilizing LightGBM XAI model against a classic logistic regression model and compare the results. They also analyze each variable's importance in predicting default and the possible financial benefits of a more accurate credit scoring system. Because it allows models to analyze larger datasets and produce more visible and interpretable findings, the scientists believe their XAI model can improve credit scoring. The outcomes favor the XAI model over the LR model for both prediction accuracy and clarity of results. The authors do admit, however, that further work is required to fine-tune the LightGBM models' calibration and to test out alternative tree-based models using XAI on a wide range of bank data.

A different approach is taken by another XAI study [12], which applies community discovery methods and graph theory to medical recommender systems. By taking into account the views of the most powerful users and applying this data to address the cold start problem, the authors aim to enhance the precision of medical advice. After applying two community detection algorithms to the user and physician graphs, they utilize the overlapping community graph to determine the final recommendation's weight. When compared to employing each community detection algorithm separately, the suggested method achieves better accuracy (93.06) and precision (88.34).

## 3 Methodology

This study's methodology delves into the application of explainable AI methods such as SHAP, LIME, and ELI5 to the analysis of CRM campaign results. Examining a dataset that records customers' interactions with campaign content as interested or not is the goal in order to discover key characteristics that show successfully targeted ads.

### 3.1 Dataset

The experiments' dataset was provided by Next4Biz, the market leader in the Turkish CSM, CRM, and BPM software industries. This dataset contains 315,121 records and 19 attributes related to commercial clients. Examples of such variables include customers' names, email addresses, job titles, and the dates of their businesses' founding. Duplicate column removal was the initial step in cleaning up the dataset. We next divided the incorporation dates into five sets, with each set representing a twenty-year period. We used this method to classify every attribute in the dataset. Upon completion of the encodings, data imputation was employed to fill in any missing values in the dataset. To prepare the data for the techniques we were going to use, one-hot encoding was finally applied.

Out of a total of 253 features, the target variable "is-clicked" indicates whether a client has clicked on a campaign link or not. A custom parser script was added to the dataset, allowing us to raise the feature count to 6,734. We aimed to determine how the campaign performance was affected by the values in the special fields utilized to define CRM customers. The ability to identify users based on their domain-specific information is what makes these additional fields possible in the Next4biz CRM system. We broadened our focus by incorporating a large amount of data from the public domain into the dataset. Many instances of the target variable were null, suggesting that customers did not click on the link to engage with the campaign. The dataset orignally contained 5,421,542 samples, or about 98 times more occurrences of 0 than occurrences of 1. To address this class disparity, we employed undersampling by dividing the dataset in half, with 131,650 samples in each half. Then, using supplementary consumer data as customer segments, we further divided the entire client base into eleven groups using the k-means clustering algorithm.

### 3.2 Predictive model

In this research, we use a Random Forest Ensemble model as our machine learning model, and we adjust it by utilizing a grid search to exhaustively search for the most effective values for each of the algorithm's parameters. Among machine learning algorithms, the Random Forest approach is widely employed for such research since it provides respectable performance relative to other options, and the model's performance may be considerably improved through fine-tuning. There are 105320 rows used in the training process, and 26330 rows used in the validation. A data distribution of 8:2 between training and testing sets is typical and allows for fair comparisons of model abilities.

# 4 Results

## 4.1 Predictive performance

The initial step is to validate a black-box prediction model. The model provides a respectable cross-validated accuracy of 87%, making it a reasonably usable model for such a study. Precision and f1-score of 87% show that the model does not discriminate between the two distinct classes and the recall of 88% alongside the ROC curve support this even further. Despite the fact that this performance was enough for the campaign prediction evaluation, we did not work with the complete dataset due to its underestimation for explainability.

## 4.2 SHAP explanation of predictions

SHAP (SHapley Additive exPlanations) is a technique used to explain why a machine learning model makes certain predictions. SHAP is an approach that provides the explanation of black box models, such as deep learning and ensemble methods. The SHAP approach can produce local and global explanations and feature importance outputs using the Shapley value [13]. In Figures 1.a) and 1.b), we present dot plots of our binary classes 0 and 1 (is_clicked), derived from SHAP values, to better illustrate the local values of the calculated according to the SHAP explanation. By allowing dots to stack up when they don't align, the pictures show a combination of a scatter plot and density estimate. One data point from the provided dataset is represented by each dot. The implication of the feature on the category is displayed along the horizontal X axis. The length of the tail has a greater effect on the model's classification if it is interpretable.

Comparison of the two figures is an appropriate initial step before looking at the global scores. The main distinction between the two is that Figure Figure 1.a) shows clients who did not make a decision, whereas Figure Figure 1.b) depicts the inverse scenario. At first glance, these figures follow a similar trend, since they represent the top two attributes with the highest global measurement values. Upon further examination, however, it is clear that the classes share some of the most important traits, yet in a different order of importance and with different traits visible on the graphs' spectra. Even though the two competing classes' feature importance levels are usually different, this might be seen as a natural convergence of those levels. The campaign's participants were divided into three groups: the chemical, financial, and machinery sectors.

Notably, the Chemistry sector has the highest significance with a remarkable value of 7.45E+26, indicating that it has a significant impact on the model's decision-making process (see Fig. 2. This is followed by terms with importance values of 6.42E+26 and 5.70E+26, respectively: machinery and customs. The importance values for succeeding features, such as packing and publishing, decrease progressively. These feature importances help us comprehend which factors are most influential in predicting outcomes and can guide subsequent analysis and decision-making. By incorporating these insights, we can improve the interpretability and explainability of our model, paving the way for better-informed and more effective decision-making across multiple domains.

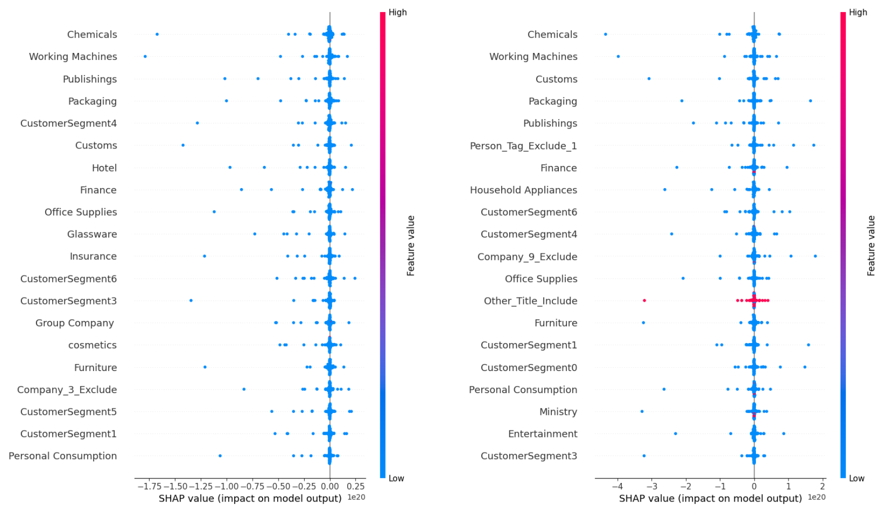


a) b)

**Fig. 1.** Bee swarm plot of SHAP value magnitude for a) Campaign Non-responders and b) Campaign Responders

### 4.3 LIME explanation of predictions

Local Interpretable Model-agnostic Explanations (LIME) is a simple and accessible XAI. Ribeiro et al. [14] proposed the LIME approach to explain supervised learning models. This approach is advantageous as it is suitable to work with all supervised learning algorithms. The LIME approach allows overfitting to be detected based on data quality issues, thanks to the local prediction explanation for each instance [15]. It provides model agnosticism where it treats a given model as a separate black box allowing it to handle any given available model. It also gives local explanations where each individual explanation is centered around the individual instance. Despite LIME explaining each instance individually, to understand the explanations and feature importance scores that comes from the data alongside understanding the overall influence and the implications the said influence has on the models prediction making process, it is necessary to look at not each individual explanation but the cumulative explanation of each individual instance combined.

When all features are considered, the feature that has the greatest impact on the model's prediction making is being in the healthcare sector, with a remarkably high importance of 3.718866. Following the most important feature, the scores for its significance fall predictably, falling within the range of the top 20 most important features with a standard deviation of 0.626. Aside from the primary function, the exclusion of certain businesses from the intended users also plays a role. We have anonymized these company names within the scope of the article. In terms of the workforce, it's clear that the Construction, Technology,

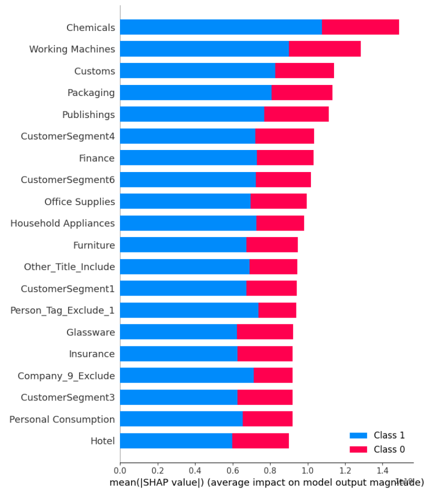


**Fig. 2.** Rich summaries of entire model based on SHAP

and Education target audiences impact the final product. Similarly, success is impacted by the target audience, which does not possess certain personal traits. To exclude customers with a specific attribute, use the Person_Tag_X_Exclude attribute. For the purposes of this paper, these traits are considered personally identifiable information. The impact of occupational groups is addressed in the last point. The inclusion of an Account Manager, for example, slightly enhances the model. This is what separates different types of professions. The feature importance values mentioned above are clearly the most important ones available shown in Figure 3. These values serve as a helpful guide for investigating the model's decision-making process and conducting additional research on the model's predictions.

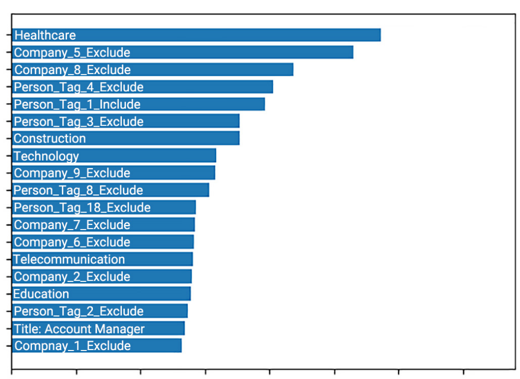


**Fig. 3.** Bar chart of the average LIME value magnitude of feature importance

### 4.4 ELI5 explanation of predictions

ELI5 [16] is a Python XAI Framework licensed by MIT to aid in the debugging of machine learning classifiers and provide clear explanations for their predictions. It is important to mention a shortcoming of the ELI5 model: it does not provide assistance for model-agnostic interpretations. It is only effective for linear/parametric and tree-based models. In addition to the traditional tree models in the sci-kit-learn library, it can also explain current tree models such as XGBoost, LightGBM, and CatBoost. ELI5 can produce ready-made outputs for simple cases and is revisable for specific and complex problems other than simple cases. ELI5 deploys a feature-based weight assignment technique to create a tree map explaining holistic and individual predictions [17]. The ELI5 algorithm showed that only 11 features contributed significantly to the model. Among these 11 features, 9 contain the sector information of the receiver to which the notification is sent. Food, Logistics, Banking, Construction, Machines, Cables, Telecommunication, Packaging, and Textile are the industries ranked from most important to least according to the ELI5 feature importance value. Two other features are an option to exclude individuals based on their responses to previously posted content and an option to exclude employees of a company. The ELI5 algorithm's explanations highlight the significance of industry tags in submitting successful campaigns.

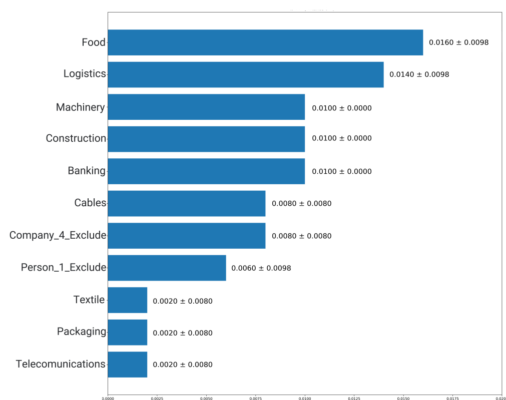


**Fig. 4.** Bar chart of the average ELI5 value magnitude of feature importance

## 5 Conclusion

Explanatory AI approaches can, in conclusion, enhance the targeting strategies of CRM campaigns, as demonstrated by this research. Establish precise justifications for these findings while distinguishing the fundamental attributes that serve as indicators of effective targeting. By utilizing SHAP, LIME, and ELI5 techniques, this research advances the interpretability of CRM campaign data beyond the capabilities of conventional black-box AI models. A practical recommendation was provided to organizations seeking to enhance their marketing strategies based on the identification of influential consumer engagement elements from a real-world dataset using SHAP, LIME, and ELI5. Our findings indicate that several factors significantly influence a customer's decision to engage with a campaign: the content of an email, demographic information, and purchasing history.

These results may have significant implications for businesses who want to enhance the efficiency of their CRM initiatives by better targeting their audiences. Companies may better analyze client behavior and adjust their marketing strategies by using explainable AI. The findings shed light on the determinants of click-through rates, allowing businesses to better serve their customers. These features' interpretability allows businesses to fine-tune their targeting tactics and raise the efficiency of their campaigns. There are many potential directions for research in the future. The use of alternative explainable AI approaches and a comparison of their efficacy in assessing CRM campaign data is a potential direction to investigate. Research on the effect of various email content types on customer engagement and the identification of best practices for creating successful campaigns could also be pursued. This research has proven that explainable AI has the potential to revolutionize the way firms handle customer relationship management. Businesses can get ahead in today's dynamic market by adopting these approaches and using the information they yield. Companies that are looking to improve their CRM tactics might benefit from this research. XAI aproaches provide clear explanations that enable businesses to fine-tune their targeting strategies, enhance client interaction, and boost marketing efficiency.

Future investigations in this field may focus on various approaches to augment our comprehension and application of XAI techniques in scrutinizing CRM campaign data. Firstly, by examining the integration of various XAI methodologies, a more thorough comprehension of the fundamental elements that impact the effectiveness of a campaign may be achieved. This approach has the ability to reveal intricate connections between distinct indicators. In addition, performing longitudinal studies to examine the evolution of these indicators over time and their influence on campaign performance could provide significant insights for creating flexible and responsive campaign tactics. Moreover, doing research on the scalability and generalizability of XAI techniques across various industries or marketing settings could provide evidence to support their practicality beyond individual case studies. Furthermore, exploring the advancement of automated decision-making systems utilizing XAI data could simplify campaign optimization processes and improve efficiency in resource allocation. Finally, integrating specialized knowledge in a certain field and input from individuals with a vested interest in the outcome of explainable artificial intelligence (XAI) findings could enhance the congruence between analytical discoveries and practical marketing tactics. In conclusion, future research in this field shows potential for enhancing both the theoretical comprehension and practical utilization of XAI in CRM campaign optimization.

# References


[1] Max Fatouretchi. *The The Art of CRM: Proven strategies for modern customer relationship management*. Packt Publishing Ltd, 2019.

[2] David Gunning, Mark Stefik, Jaesik Choi, Timothy Miller, Simone Stumpf, and Guang-Zhong Yang. Xai—explainable artificial intelligence. *Science robotics*, 4(37):eaay7120, 2019.

[3] Jianrong Tao, Yu Xiong, Shiwei Zhao, Runze Wu, Xudong Shen, Tangjie Lyu, Changjie Fan, Zhipeng Hu, Sha Zhao, and Gang Pan. Explainable ai for cheating detection and churn prediction in online games. *IEEE Transactions on Games*, 2022.

[4] Lara Marie Demajo, Vince Vella, and Alexiei Dingli. An explanation framework for interpretable credit scoring. *International Journal of Artificial Intelligence and Applications (IJAIA)*, 12(1), 2021.

[5] Felix Haag, Konstantin Hopf, Pedro Menelau Vasconcelos, and Thorsten Staake. Augmented cross-selling through explainable ai–a case from energy retailing. *arXiv preprint arXiv:2208.11404*, 2022.

[6] Neus Llop Torrent, Giorgio Visani, and Enrico Bagli. Psd2 explainable ai model for credit scoring. *arXiv preprint arXiv:2011.10367*, 2020.

[7] Fenareti Lampathaki, Carlos Agostinho, Yuri Glikman, and Michele Sesana. Moving from 'black box'to 'glass box'artificial intelligence in manufacturing with xmanai. In *2021 IEEE International Conference on Engineering, Technology and Innovation (ICE/ITMC)*, pages 1–6. IEEE, 2021.

[8] Sai Tejashwin Eswarapu, S Sesharhri, Yashwanth Deshaboina, P Bhargawa, Ashly Ann Jo, and Ebin Deni Raj. Integrated customer analytics using explainability and automl for telecommunications. In *2023 2nd International Conference on Applied Artificial Intelligence and Computing (ICAAIC)*, pages 228–235. IEEE, 2023.

[9] Emer Owens, Barry Sheehan, Martin Mullins, Martin Cunneen, Juliane Ressel, and German Castignani. Explainable artificial intelligence (xai) in insurance. *Risks*, 10(12):230, 2022.

[10] Gabriel Marín Díaz, José Javier Galán, and Ramón Alberto Carrasco. Xai for churn prediction in b2b models: A use case in an enterprise software company. *Mathematics*, 10(20):3896, 2022.

[11] Petter Eilif De Lange, Borger Melsom, Christian Bakke Vennerød, and Sjur Westgaard. Explainable ai for credit assessment in banks. *Journal of Risk and Financial Management*, 15(12):556, 2022.

[12] Arun Kumar Sangaiah, Samira Rezaei, Amir Javadpour, and Weizhe Zhang. Explainable ai in big data intelligence of community detection for digitalization e-healthcare services. *Applied Soft Computing*, 136:110119, 2023.

[13] Scott M. Lundberg and Su-In Lee. A unified approach to interpreting model predictions. In I Guyon, UV Luxburg, S Bengio, H Wallach, R Fergus, S Vishwanathan, and R Garnett, editors, *Advances In Neural Information*

*Processing Systems 30 (NIPS 2017)*, volume 30 of *Advances in Neural Information Processing Systems*, 2017. 31st Annual Conference on Neural Information Processing Systems (NIPS), Long Beach, CA, DEC 04-09, 2017.

[14] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. "why should i trust you?" explaining the predictions of any classifier. In *KDD'16: Proceedings Of The 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1135–1144, 2016.

[15] Varada Vivek Khanna, Krishnaraj Chadaga, Niranajana Sampathila, Srikanth Prabhu, Venkatesh Bhandage, and Govardhan K. Hegde. A distinctive explainable machine learning framework for detection of polycystic ovary syndrome. *Applied System Innovation*, 6(2), APR 2023.

[16] Angela Fan, Yacine Jernite, Ethan Perez, David Grangier, Jason Weston, and Michael Auli. Eli5: Long form question answering. *arXiv preprint arXiv:1907.09190*, 2019.

[17] Aabhas Vij and Preethi Nanjundan. Comparing strategies for post-hoc explanations in machine learning models. In S Shakya, R Bestak, R Palanisamy, and KA Kamel, editors, *Mobile Computing and Sustainable Informatics.*, volume 68 of *Lecture Notes on Data Engineering and Communications Technologies*, pages 585–592, 2022. International Conference on Mobile Computing and Sustainable Informatics (ICMCSI), Tribhuvan Univ, Kirtipur, NEPAL, JAN 29-30, 2021.