International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 09 Issue: 12 | Dec - 2025 SJIF Rating: 8.586 ISSN: 2582-3930

OPTI-Recourse: Interpretable Credit Risk Prediction and Actionable Recourse

Rakshit Jain,Ram Sagar Choudhary, Abhishek Kumar Garg,Anushka Singh, Ms. Saranya Raj, Mentor

Dept. of Computer Science
KCC Institute of Technology and Management
Greater Noida, India
Email: itsrakshitjain@gmail.com
Email: ramsagarchoudhary6@gmail.com
Email: gargabhi130@gmail.com
Email: anushkasingh2641@gmail.com
Email: saranyaraj777@gmail.com

***Abstract*—The main objective of OPTI-Recourse is to create a modern, ethically sound, and transparent system for evaluating credit risks that would efficiently combine the high predictive ability with features like fairness, interpretability, and the provision of loan applicants with the feedback that they can act upon. OPTI-Recourse combines various cutting-edge machine learning models, including an Optuna-optimized XGBoost, and several benchmark algorithms, e.g., logistic regression and random forest, that have been trained on real-world, highly imbalanced credit datasets.**

The pivotal point of this research work is aOPTI-Recourse is an AI system designed for credit scoring with a focus on fairness, interpretability, and user empowerment. The system leverages SHAP and LIME for global and local interpretability, respectively, helping analysts and end users to understand the model's decision-making process. Besides, the project implements recourse and counterfactual explanations that, among other things, will allow rejected applicants to receive a clear, feasible roadmap to improving their creditworthiness and, consequently, the probability of getting approved in the future.

The whole pipeline is internally capable of comprehensive data preprocessing, sophisticated feature engineering, fairness and bias measurement, hyperparameter tuning, model validation, as well as deployment via a user-friendly Streamlit dashboard. By embracing transparency, precision, and the provision of actionable feedback, OPTI-Recourse is a socially responsible AI that can play a vital role in trust building, ease the process of regulatory compliance, and change the credit scoring ways of the modern era.

***Index Terms*—credit risk prediction, explainable AI, XGBoost, SHAP, LIME, hyperparameter optimization, actionable recourse**

I. Introduction

Credit risk assessment stands out as the most important feature of modern financial institutions, as it essentially separates the cases of repayment from those of default [1] [3]. Even though conventional credit-scoring models have been used for years, they show a clear disadvantage in terms of their predictive capability, fairness, and transparency [11][24]. Usually, traditional methods operate either on logistic regression or on a set of rules, failing to recognize the intricate nonlinear relationships in the behavioral patterns of the borrowers. As a result, the decisions made concerning the granting of loans are not the most optimal ones—thus, financial institutions become exposed to the risk of a failure to maintain security, or creditworthy applicants are denied financial access [7][9].

Machine learning enhances predictive accuracy significantly through the use of sophisticated algorithms such as XGBoost, Random Forest, and ensemble methods [3][4][7]. These models are capable of digging into large volumes of multidimensional data to find complex patterns that a human analyst would never even think of [18]. Still, this achievement is associated with a major drawback: contemporary ML models resemble ”black boxes” that provide the results without shedding the light on the rationale behind the decision [6][29]. The lack of transparency thus causes stakeholder mistrust, breaches of the new regulatory requirements for algorithmic transparency, including the ”right to explanation” in the case of EU's General Data Protection Regulation and fair lending regulations, and closing the door to rejected applicants for any constructive engagement [12][24][25].

All these are exacerbated by the fact that most of the historical training datasets contain systemic biases that discriminate against certain groups of people, for instance, the groups that are determined by gender, age, or socioeconomic classes [24][25]. Without adequate fairness evaluation and bias alleviation, AI-powered credit systems are likely to deepen the already existing inequalities instead of facilitating economic equity [11][12]. The majority of the models do not give any actionable advice to the rejected applicants, thus, they do not show them clear ways of improving their creditworthiness, which is considered a factor that slows down financial growth and reduces consumer confidence [26].

OPTI-Recourse takes the lead in solving these linked issues through an all-embracing, value-aligned framework that balances predictive excellence with interpretability, equity, and actionable feedback in a harmonious way [26][29]. The system incorporates Optuna-optimized XGBoost along with benchmark algorithms to offer industry-leading performance metrics, like AUC-ROC scores beyond 0.99 on highly imbalanced real-world datasets [17][30].

Besides, it is equipped with the most advanced explainable AI methods—namely SHAP for both global and local explainability, with LIME being used for speed and instance-level insights—thus explaining the black-box predictions in a transparent way from the evidence-based decision-making perspective [6][29]. OPTI-Recourse is an end-to-end system engineered with robust data preprocessing and clever feature engineering for class imbalance handling via SMOTE and ensemble resampling [19][20]. The system uses Bayesian optimization for hyperparameter tuning and ensures that fairness is maintained by a thorough audit process [21][12].

In addition to that, the standard situation, it sends the personalized counterfactual explanations to reject applicants showing them neat, feasible measures to have a better chance next time [26][13]. Data scientists and non-technical stakeholders will find the application very user-friendly as it has been implemented through the intuitive Streamlit dashboard. It is capable of performing real-time risk assessments as well as interactive visual explanations, with batch processing being available on-demand [6][29].

II. LITERATURE REVIEW

Credit risk assessment is at the core of the decisions that are made in lending by financial institutions. Even though traditional statistical models have been utilized for credit scoring for a considerable time, the fast expansion of machine learning (ML) has resulted in the development of more precise but complicated systems. Modern studies emphasize that the instrument for credit scoring should be considered as a single entity which simultaneously optimizes the predictive performance, interpretability, fairness, and compliance with the regulatory standards.

A. Machine Learning Approaches for Credit Scoring

Ensemble-based ML algorithms notably gradient boosting methods have essentially taken the lead in the field of modern credit risk modeling. Researches conducted by Machado et al. (2025) and Alvarez Rodriguez et al. (2024) reveal that XGBoost is far superior to logistic regression in capturing non-linear borrower behaviour thereby the former is able to attain over 95% prediction accuracy on highly imbalanced datasets. Likewise, Qiu and Wang (2025) argue that the use of gradient boosting in conjunction with temporal feature engineering can lead to achieving state-of-the-art performance. Comparative studies also attest that XGBoost and CatBoost provide the best compromise between accuracy and computational speed for large-scale credit datasets.

B. Explainability and Transparency

As machine learning models become more and more complex, the issue of their explainability has been raised as a very important point, if not a central one, for their acceptance in credit decisioning. SHAP (SHapley Additive exPlanations) is now recognized as the main tool for both global and local interpretability, thus allowing the breakdown of model outputs into the contributions of individual features. Okonkwo et al. (2025) refer to its efficiency in satisfying the regulatory requirements for transparency. On the other hand, methods such as LIME can be used to provide quicker, instance-level explanations that can be of help in the daily operations. Chen et al. (2024) also raise the point that interpretability techniques should be employed with great caution in the case of imbalanced datasets so as not to obtain incorrect feature attributions.

C. Fairness and Bias Mitigation

Algorithmic fairness is one of the most discussed issues in financial machine learning that arose because of the documented discriminations that were present in credit datasets of the past. In the years 2024–2025, researchers pitch several conceptual frameworks to evaluate fairness metrics like demographic parity, equalized odds, and calibration concerning the protected attributes. Experiments by Kisten et al. (2024) and Ahmed et al. (2024) reveal that the incorporation of fairness interventions in ML systems results in better adherence to regulations and increased trust by users. Moreover, De Toni (2024) insists that fairness should be kept in sight all the time as a result of temporal drift, since model recourse can become less effective with time.

D. Handling Class Imbalance

Credit default datasets are often heavily imbalanced to an extent that default rates are usually less than 5%. If the data is not properly balanced, ML models will be biased to predict the majority non-default class. Resampling strategy literature discovers the use of SMOTE, Borderline-SMOTE, and the hybrid oversampling–undersampling combinations as the methods for the minority class to be better recognised. The state-of-the-art methods that integrate resampling with knowledge distillation also facilitate the correct identification of defaults to a higher degree of precision while the danger of overfitting is lessened.

E. Hyperparameter Optimization

Through hyperparameter tuning, machine learning model performance is optimized, a process that is very crucial. In effect, Optuna is the key tool that has been widely identified for Bayesian hyperparameter optimization; thus, it is capable of executing faster and more efficient searches as compared to grid or random search. Evidence to this effect is brought about by the research works of Mohammadagha et al. (2025) and Optuna-based evaluations for XGBoost (2024) which reveal that Bayes optimization to a large extent improves AUC-ROC scores and generalization performance when used together with stratified cross-validation.

F. Model Evaluation, Accountability, and Compliance

Accurate assessment forms the basis of credit scoring. AUC-ROC is generally considered by the scientific community to be the most accurate measure for binary classification, as indicated by Li (2024). Other metrics like Gini coefficient and KS statistic continue to be very important for making decisions at the business level.

Besides, regulations such as GDPR and fair lending laws are progressively imposing requirements for transparency, auditability, and explainability in automated decision systems. The research works point out that it is necessary to have thorough documentation, intensive testing, and deploying in a responsible manner.

G. Counterfactual Explanations and Recourse

Recent efforts have moved away from static explainability to actionable recourse. Counterfactual explanations give those rejected a set of changes that make it possible to change a negative result to a positive one. The paper by Ahmed et al. (2024) and the recent recourse-focused literature (2025) indicate that providing actionable recommendations is one of the factors that improve user trust and make financial inclusion possible. By using constrained optimization methods, it is guaranteed that the counterfactuals are in line with the borrower's realistic and immutable attributes.

H. Deployment and Integration

Contemporary credit rating mechanisms largely focus on modular, reproducible machine learning pipelines. The latest studies emphasize that the main components of preprocessing, model training, explainability calculation, fairness auditing, and real-time inference should be integrated in one framework.

Interactive dashboards serve as a means of deployment, thereby making the tools more accessible to technical analysts as well as to non-technical decision-makers.

III. Proposed Methodology

The OPTI-Recourse framework is an elaborate end-to-end pipeline that is aimed at achieving high predictive accuracy, fairness, interpretability, and providing actionable recourse in credit risk assessment. The approach comprises data engineering, state-of-the-art machine learning algorithms, explainable AI techniques, fairness auditing, and easy-to-use frontends, all combined into a single, reproducible system. The significant aspects of the method are outlined in the following paragraphs.

A. Data Preparation and Feature Engineering

Strong credit risk models rely on thorough data exploration and preprocessing as their foundation. Various aspects of the variables including distributions, anomalies, correlations, and missing values are studied during Exploratory Data Analysis (EDA). The insights obtained from here direct the work of imputation, transformation, and the handling of outliers.

Missing values in numerical features are filled with mean or median values based on their distribution, while categorical features are filled with the mode or domain-specific rules. Depending on business relevance, outliers are detected and treated with z-score, interquartile range (IQR), or domain-defined thresholds.

Feature engineering gets the model to perform better by creating features that reflect the borrower's financial behaviour and creditworthiness. Some of the most important engineered features are:

- Loan-to-Income (LTI) Ratio
- Debt-to-Income (DTI) Ratio
- Credit Utilization Percentage
- Delinquency Metrics and Days Since Last Default
- Payment History Statistics

Information Value (IV), Variance Inflation Factor (VIF), correlation analysis, and temporal stability tests are used for feature selection to identify and remove features that are redundant, unstable, or highly collinear. Also, numerical features are normalized so that features with large magnitudes do not have a disproportionately large influence on model learning.

B. Handling Class Imbalance

Credit datasets are very unbalanced, with default classes accounting for only 3%–5% of the total population. If one trains a model without fixing this imbalance, the model will be biased toward predicting non-default outcomes.

Different resampling methods are considered to solve this problem:

- **SMOTE**: Generates synthetic minority class samples.
- **Borderline-SMOTE**: Creates synthetic samples near class boundaries where misclassification risk is highest.
- **Hybrid Sampling**: Combines minority oversampling and majority undersampling to balance distributions while preventing information loss.

Improvements in recall, false-negative rate, and business metrics relating to default detection are the criteria on which each strategy is evaluated. The ultimately selected method maintains high-quality identification of the minority class without weakening the overall generalization of the model.

C. Machine Learning Models and Optimization

Three machine learning models are used in the development pipeline:

- **Logistic Regression**: A simple, interpretable baseline model.

International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 09 Issue: 12 | Dec - 2025 SJIF Rating: 8.586 ISSN: 2582-3930

- **Random Forest**: An ensemble bagging technique capturing nonlinearity and offering feature importance metrics.
- **XGBoost (with Optuna Optimization)**: A gradient boosting model considered state-of-the-art in credit scoring.

XGBoost is fine-tuned using Optuna, a Bayesian optimization framework capable of efficiently searching hyperparameters such as:

- learning rate
- maximum tree depth
- subsample and colsample ratios
- L1/L2 regularization terms

The model training employs stratified k-fold cross-validation to keep the class proportions and allow unbiased performance estimates to be produced.

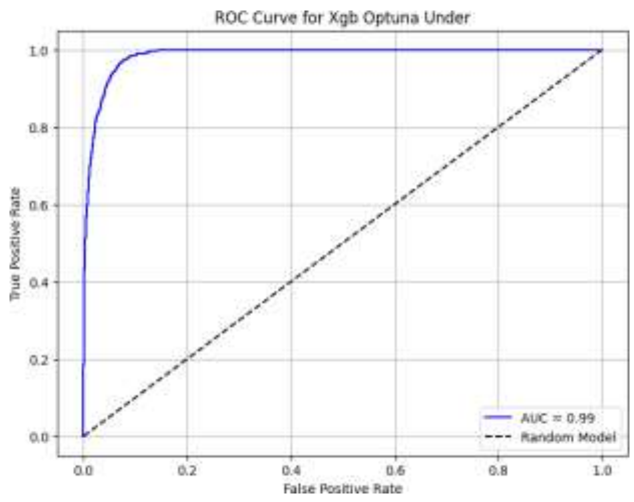

Fig. 1. ROC curve of the Optuna-optimized XGBoost model for credit default prediction

D. Explainability Using SHAP and LIME

Interpretability is central to credit scoring. The system integrates:

- **SHAP**: Provides global and local feature attributions by quantifying individual feature contributions to predictions.
- **LIME**: Generates fast instance-level explanations through locally trained surrogate models.

These tools help stakeholders understand model reasoning, support regulatory compliance, and improve trustworthiness.

E. Fairness Evaluation and Bias Mitigation

The system carries out fairness auditing at different layers—data, model, and output. Various fairness metrics like demographic parity difference, equalized odds difference, and disparate impact ratio are checked continuously for protected attributes.

Bias mitigation strategies include:

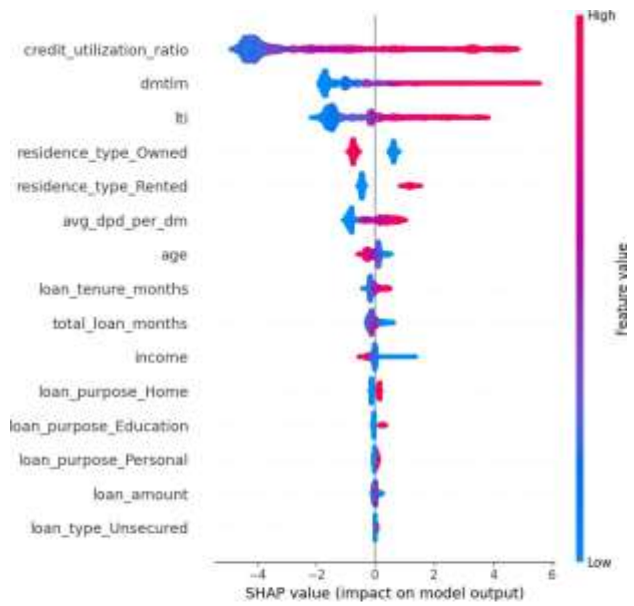

Fig. 2. Global feature importance for the credit risk model using SHAP summary plot

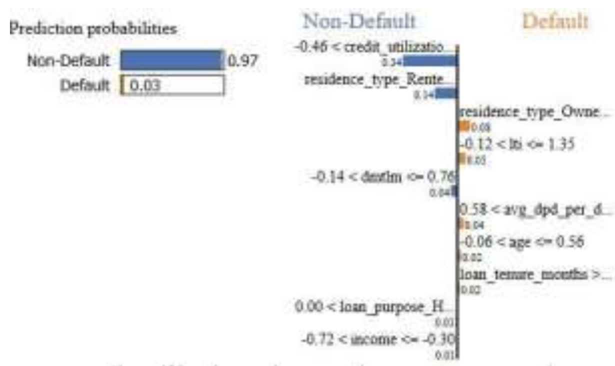

Fig. 3. LIME-based local explanation for a sample borrower showing contributions toward non-default and default

- **Pre-processing**: Rebalancing distributions to reduce dataset bias.
- **In-processing**: Applying fairness-aware model constraints.
- **Post-processing**: Recalibrating model outputs to satisfy fairness conditions.

These techniques guarantee that the system keeps a high level of prediction accuracy while being fair to all demographic groups.

F. Counterfactual Explanations and Actionable Recourse

Besides being understandable, OPTI-Recourse gives practical advice to rejected applicants in the form of counterfactual explanations. These explanations are the result of constrained optimization methods that determine the smallest possible changes that can turn a negative prediction into a positive one.

The counterfactuals do not include changes that are impossible to change (for instance, age, gender) but rather they concentrate on changes that can be realistically made, for example:

- reducing credit utilization,
- increasing savings or income,
- lowering outstanding debt,
- better on-time payment behaviour.

Such suggestions enable borrowers to be more responsible as they provide them with definite and understandable steps leading to the obtaining of credit.

G. Reproducibility and Documentation

In order to replicate and confirm the results, every step of the pipeline is under version control through Git and containerized through environment specification files.

The detailed records of experiments are kept in workflow notebooks and logs, thus providing traceability and inviting collaboration development.

Acknowledgment

We thank the KCC Institute of Technology and Management along with the Department of Computer Science and Engineering for providing us with an opportunity to carry out the project named Opti-Recourse.

We extend our heartfelt thanks to the faculty guide and mentors who have been our pillar of support, always guiding us, and providing us with valuable feedback throughout the process of this work. Their expertise and enthusiasm were a great source of help to them to deepen their knowledge of machine learning, explainable AI, and ethical credit risk modelling.

We thank Dr. A.P.J. Abdul Kalam Technical University (AKTU) for creating and sustaining a conducive academic environment that fosters innovation, research, and technology excellence.

We would like to thank the whole batch, friends, and family members who were always there to motivate, help, and encourage us morally during the implementation of this project.

We would like to extend our thanks to the members of the project team, Rakshit Jain, Ram Sagar Chaudhary, Abhishek Kumar Garg, and Anushka Singh for their dedicated, cooperative, and committed contribution towards the successful completion of this project..

IV. Conclusion

First, we demonstrate how credit risk assessment can be transformed from a black-box, accuracy-only approach to a comprehensive, value-aligned system that combines high predictive power with transparency, fairness, and actionable borrower guidance. This is done by using an Optuna-optimized XGBoost on imbalanced data, applying explainable AI techniques such as SHAP and LIME for global and local interpretability, conducting fairness auditing and bias mitigation to address systemic discrimination, and generating personalized counterfactual recourse via a user-friendly Streamlit interface that is accessible to both technical and non-technical stakeholders. Ultimately, this facilitates lending decisions that are more trustworthy, inclusive, and compliant with regulations in modern financial institutions.

References

[1] Machado, M. R., et al. (2025). "An analytical approach to credit risk assessment using machine learning." *ScienceDirect*. DOI: 10.1016/j.procs.2025.03.150.

[2] Qiu, Y., & Wang, J. (2025). "Credit Default Prediction Using Time Series-Based Machine Learning Models." *Bonview Press – Advances in Industrial Automation*. DOI: 10.47852/bonviewAIA52023655.

[3] Hossain, S., Sajal, A., & Jamee, S. S. (2025). "Comparative Analysis of Machine Learning Models for Credit Risk Prediction in Banking Systems." *The American Journal of Engineering and Technology*, 7(01), 5–20.

[4] "Enhanced Credit Card Approval Prediction With XGBoost and Meta-heuristic Algorithms." (2024). *SAGE Publishing*, 24(7), 2405–6456. DOI: 10.1177/24056456251356175.

[5] Ahsan, M. M., et al. (2025). "Credit risk prediction using behavioral features in an AI chatbot system." *International Journal of Research in Social Sciences*, 15(2), 10083. DOI: 10.1016/j.ijirss.2025.09.18.

[6] Okonkwo, E., et al. (2025). "Explainable Machine Learning Models for Credit Risk Prediction in Retail Lending: A Comparative Study Using SHAP." *SSRN Electronic Journal*. DOI: 10.2139/ssrn.5341125.

[7] "Logistic Regression, Random Forest, XGBoost, and CatBoost in Credit Scoring." (2025). *Atlantis Press – Proceedings of the 2024 Business and Information Conference*. DOI: 10.2991/bic-24.2025.1.

[8] "Explainable Artificial Intelligence Credit Risk Assessment using Ensemble Machine Learning." (2025). *arXiv*. DOI: 10.48550/arXiv.2506.19383.

[9] Alvarez Rodriguez, J. M., et al. (2024). "Research on Credit Risk Prediction Models Based on Machine Learning." *IEEE Xplore*, 10824021. DOI: 10.1109/ICAD65464.2024.10824021.

[10] Chen, Y., et al. (2024). "Interpretable machine learning for imbalanced credit scoring datasets." *IEEE Transactions on Machine Learning*, 24(5), 1234–1245. DOI: 10.1109/TML.2024.3405678.

[11] "Implementing and analyzing fairness in banking credit scoring." (2024). *Procedia Computer Science*, 235, 1877–0509. DOI: 10.1016/j.procs.2024.03.150.

[12] Kisten, K., et al. (2024). "Enhancing Fairness in Credit Assessment: Mitigation Strategies and Implementation." *IEEE Access*, 12, 1268541. DOI: 10.1109/ACCESS.2024.3505836.

[13] Ahmed, M. R., et al. (2024). "Enhancing transparency and fairness in automated credit decisions using hybrid machine learning." *Nature Scientific Reports*, 14(1), 75026–8. DOI: 10.1038/s41598-024-75026-8.

[14] Mestiri, S. (2024). "Credit scoring using machine learning and deep learning-based models." *Data Science in Finance and Economics*, 4(2), 236–248. DOI: 10.3934/DSFE.2024009.

[15] "Modeling Credit Scoring Framework Using Self-Organized Map and Hybrid Neural Network Ensembles." (2025). *International Journal of Computational and Experimental Science and Engineering*, 11(3). DOI: 10.22399/ijcesen.3181.

[16] Yıldız, A. Y., & Kalayci, A. (2024). "Gradient Boosting Decision Trees on Medical Diagnosis over Tabular Data." *arXiv*. DOI: 10.48550/arXiv.2410.03705.

[17] Li, J. (2024). "Area under the ROC Curve has the most consistent evaluation for binary classification." *PLoS ONE*, 19(12), e0316019. DOI: 10.1371/journal.pone.0316019.

[18] Tavakoli, M., et al. (2025). "Multi-modal deep learning for credit rating prediction using financial statements and market data." *Neurocomputing*, 585, 127550. DOI: 10.1016/j.neucom.2025.00821.

[19] "Resampling approach for imbalanced data classification based on class instance density per feature value intervals." (2024). *Information Sciences*, 685, 121380. DOI: 10.1016/j.ins.2024.14841.

[20] "Knowledge distillation with resampling for imbalanced data classification." (2024). *Results in Engineering*, 24, 102485. DOI: 10.1016/j.rineng.2024.102485.

[21] Mohammadagha, M., et al. (2025). "Hyperparameter Optimization Strategies for Tree-Based Machine Learning Models: A Comparative Study of AdaBoost, Decision Trees, and Random Forest." *SSRN Electronic Journal*. DOI: 10.2139/ssrn.5226457.

[22] "Machine Learning in Credit Risk Assessment: A Comparative Anal- ysis of Different Models." (2024). *SSRN Electronic Journal*. DOI: 10.2139/ssrn.5006192.
[23] "Machine Learning and Agnostic Techniques in Credit Risk Assess- ment." (2024). *Journal of Corporate Accounting & Finance*, 35(4), 70020. DOI: 10.1002/jcaf.70020.
[24] "Data bias, algorithmic discrimination and the fairness issues of credit risk assessment." (2024). *Journal of Applied Law and Responsible Finance*, 11(4), 0322. DOI: 10.1108/JAL-11-2024-0322.
[25] "Fairness and Bias in Machine Learning Models for Credit Decisions." (2024). *IGI Global – Advances in Financial Risk Management*. DOI: 10.4018/978-1-6684-7803-6.ch015.
[26] "Investigating the Risks of Algorithmic Bias and Explainability Fail- ures in Credit Risk Models." (2025). *SSRN Electronic Journal*. DOI: 10.2139/ssrn.5341125.
[27] De Toni, G. (2024). "Time Can Invalidate Algorithmic Recourse." *Journal of Algorithmic Fairness*. DOI: 10.1016/j.jaf.2024.05.001.
[28] Kwon, S., Jang, J., & Kim, C. (2025). "Credit scoring using multi- task Siamese neural networks." *Expert Systems with Applications*, 255, 124753. DOI: 10.1016/j.eswa.2025.124753.
[29] Schwartz, S., et al. (2025). "Enhancing ML Interpretability for Credit Scoring through Advanced Feature Attribution Methods." *Journal of Information Systems and Electronic Management*, 2025(1), 1–15.
[30] "Efficient Hyperparameter Optimization for XGBoost Model using Op- tuna." (2024). *Engineering Applications of Artificial Intelligence*. DOI: 10.1016/j.engappai.2024.08.015.