Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

Explainable AI in Credit Models: Balancing Predictive Power with Transparency

Bhaskara Reddy Udaru

Anna University, Chennai, India

ARTICLE INFO

Received: 02 Sept 2025

Revised: 15 Oct 2025

Accepted: 25 Oct 2025

ABSTRACT

The combination of artificial intelligence and machine learning models used as modern credit-scoring systems is more and more based on greater predictive accuracy, but also develops major difficulties in terms of transparency and interpretability. This article will discuss how explainable AI (XAI) methods can be applied to model credit to strike a balance between the performance and transparency needs that regulations like GDPR and the Equal Credit Opportunity Act require. The metamorphosis of the primitive statistical techniques to advanced algorithms has resulted in a conflict between accuracy and interpretability, which is addressed by the XAI methodologies. A range of methods, such as SHAP values, LIME, and counterfactual explanations, allows financial institutions to give meaningful explanations of automated decisions without losing their predictive power. An XAI framework that is extensive in terms of credit evaluation embraces global interpretability (model-wide behavior) and local interpretability (individual decisions), increasing regulatory compliance, widening financial inclusion, and fostering stakeholder trust through the lending ecosystem.

Keywords: Explainable AI, Credit Scoring, Algorithmic Transparency, Financial Inclusion, Regulatory Compliance

1. Introduction

In recent years, the credit-scoring systems have fundamentally changed as financial institutions have started to embrace artificial intelligence and machine learning models in order to evaluate borrower risk. These advanced technologies have unparalleled predictive power over traditional statistical methods, but at the same time pose serious issues in terms of transparency and understandability. A global report on Big Data and Advanced Analytics by the European Banking Authority published in 2021 indicates that millions of financial institutions in the EU are quickly adopting machine learning-based credit scoring and that the implementation rate is growing by around 42 percent between 2018 and 2020 [5]. The report also reports the fact that although the accuracy of these advanced models improves by an average of 15-20 percent over the traditional methods, the complexity of the systems poses significant obstacles to the comprehension of the decision's rationale.

The increasing obscurity of the credit process has led the world over to develop structures that require increased openness. According to the European Banking Authority, explainability can be considered one of the four pillars of trustworthy artificial intelligence in financial services, as well as the presence of fairness, ethics, and consumer protection [5]. This regulatory emphasis mirrors larger social fears regarding automated financial judgments regarding the economic opportunities of people. The Bank for International Settlements reports that financial institutions have made an increment of about 36 percent of investments in explainability technologies since 2019, and larger financial institutions have devoted an average of EUR4.7 million to explainable AI research and implementation specifically [9].

Explainable AI (XAI) is a set of methods that aim to increase the transparency of complicated machine learning models without noticeably reducing predictive accuracy. Such methods offer explainable model outputs in credit assessment settings and therefore give stakeholders understandable explanations. As it is stated in the Science journal publication, Beyond Prediction, explanatory tools

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/
Research Article

should go beyond statistical associations to offer cause-and-effect results about the processes of decision-making [11]. The study illustrates that as financial institutions adopt the presence of solid explanation frameworks, customer satisfaction rates with automated decisions are elevated by an average of 27 percent, even when the decision is not to the customer's favor. Moreover, a well-applied explanation system saves regulatory compliance costs estimated at 18-23% as compared to institutions that employ unexplainable systems.

An elaborate, explainable credit risk measurement framework balances current state-of-the-art predictive performance with strong interpretability. By combining the high-level machine learning algorithm and high-level explanation methods, systems complying with regulatory needs and at the same time competitive in their accuracy are developed. Analysis of AI ethics standards in a global scope has found the concept of transparency as a prevailing idea in 87 percent of all reviewed sets, reflective of its universal nature in any regulatory context [14]. Financial institutions that have adopted these strategies note better stakeholder trust, better compliance, and an increase in credit accessibility among diverse populations that have historically had poor access to it. According to research by various regulatory authorities, explainable models have the potential to boost credit provision to the underbanked groups by 18-24 percent and still ensure that risks are properly set.

2. Evolution of Credit Scoring Systems

2.1 From Statistical Models to Machine Learning

The way credit is scored has radically changed in the past few decades and has changed enough to no longer be a basic statistical model, but has advanced to an elaborate machine learning architecture. The classical models used logistic regression and linear regression models as their main tools and provided easy interpretability with a predetermined weight on each borrower attribute affecting creditworthiness. Such traditional processes were the industry leaders as late as the early 2000s, specifically due to their transparency, as the stakeholders could readily comprehend the impact that certain factors had on the final decisions [4].

The terrain was changing due to the expansion of computational power as well as increased access to data. Decision trees were also developed as an earlier alternative to regression models, and they had stronger performance, but at the same time, they had some interpretability. The emergence of the ensemble techniques, especially random forests and gradient boosting machines, represented a very important switch. The Fairness and Machine Learning textbook argues that the algorithms created are highly sophisticated, and they greatly minimize the errors in predicting defaults, as opposed to conventional methods [4].

The most recent trend in credit risk assessment is deep learning methods, which can be especially effective with regard to identifying complex trends in a wide range of data. Empirical studies conducted by the Association for the Advancement of Artificial Intelligence report the ability of neural structures to outperform the traditional models, particularly when additional sources of data have been used to complement the traditional credit bureau data [15]. These systems are used to perform processing on unstructured data, such as transaction patterns and mobile device usage, and in doing so generate predictive signals that are not visible to traditional methods.

2.2 Regulatory Landscape and Transparency Requirements

The growing sophistication of credit scoring algorithms has led to regulatory frameworks that are more demanding in their requirement of transparency in automated lending decisions. The General Data Protection Regulation (GDPR), a set of rules and regulations of the European Union, is one of the most detailed answers to the issue of algorithmic decision-making. A study of counterfactual explanations reports that Article 22 specifically speaks about automated individual decision-making, in which case data subjects have the right not to be subject to purely algorithmic decisions that have a

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/
Research Article

serious impact [12]. The regulation also creates what many take as a right to explanation, and which requires meaningful information regarding the logic that underlies automated processing.

Counterfactual explanations have also become a viable solution to regulatory requirements without compromising the performance of the model. Instead of describing the whole algorithmic procedure, counterfactuals pay attention to actionable data: what would have to be altered so that an applicant could have been given another decision. This method meets most of the regulatory aspects and gives the consumers practical information on how to enhance future performances [12].

Whereas the world has shown significant uniformity in its AI ethics rules, the key element of transparency is highlighted in various jurisdictions. An extensive review of ethics reports by government agencies, scientific institutes, and lobby groups revealed that transparency seems to be a central value in most systems [14]. This study highlights the shift in the regulatory strategies in response to the discrimination in conventional lending, with the setting of particular conditions on the transparency of the algorithms.

The Equal Credit Opportunity Act in the United States not only forbids discrimination in lending but also requires that the creditors give particular causes for adverse action. Studies at the AAAI Conference investigated the nature of the interactions between those requirements and state-of-the-art machine learning algorithms, showing the ways to train models that remain highly predictive and generate more transparent and fair decisions [15].

Era	Model Type	Key Characteristics	Interpretability Level	Regulatory Implications
Traditional	Logistic Regression	Fixed weights, linear relationships	High	Straightforward compliance
Early ML	Decision Trees	Hierarchical decisions, threshold-based	Medium-High	Relatively transparent
Advanced ML	Ensemble Methods (Random Forest, XGBoost)	Multiple models, boosting, complex patterns	Medium-Low	Growing transparency challenges
Current	Deep Neural Networks	Non-linear relationships, alternative data	Low	Significant regulatory concerns

Table 1: Evolution of Credit Scoring Systems [4, 12, 14, 15]

3. Explainable AI Techniques for Credit Modeling

3.1 Post-hoc Interpretability Methods

The explainable AI approaches have turned out to be critical parts of the current credit scoring models to overcome the transparency deficiency of advanced machine learning models. Post-hoc interpretability techniques enable financial institutions to take advantage of advanced algorithms and give explainable justifications for their choices. The Local Interpretable Model-agnostic Explanations (LIME) framework uses interpretable models to make faithful explanations by approximating complex models around particular predictions [1]. The application of this technique is especially useful in the credit context, where the loan officers and the customers have to make certain decisions without necessarily having technical knowledge.

Shapley Additive exPlanations (SHAP) framework is an extension of the cooperative game theory, and it brings the concepts of attribution stability to the feature importance. SHAP values are used to bring together a number of the existing approaches to creating an explanation, and each feature has an importance value which can be interpreted as the contribution that the feature makes towards a specific prediction as compared to a baseline [2]. SHAP values may be used to measure the contribution of factors like payment history, debt utilization, and income in the final decision with credit scoring applications.

Counterfactual explanations provide an action-oriented explanatory approach to interpretability that appeals to credit applicants. Counterfactuals do not describe model mechanics and instead determine which minimal modifications are required to obtain a different outcome, like demonstrating that a higher income or a lower amount of outstanding debt would cause a rejection to be converted to an approval [3].

Technique	Approach	Explanation Type	Best Use Cases	Target Audience
LIME	Local approximation of complex models	Local	Individual credit decisions	Loan officers, Customers
SHAP Values	Game theory-based attribution	Local & Global	Feature contribution analysis	Risk analysts, Regulators
Counterfactuals	Action-oriented alternatives	Local	Rejection explanations	Customers, Customer service
Feature Importance	Overall variable significance	Global	Model documentation	Regulators, Model developers
Partial Dependence	Feature-outcome relationships	Global	Risk factor analysis	Business stakeholders
Prototype-based	Case similarity identification	Local	Complex pattern explanation	Non-technical users

Table 2: Explainable AI Techniques for Credit Modeling [1, 2, 3, 7, 8, 13]

3.2 Global vs. Local Explanations

Global versus local interpretability is a key consideration in the design of explainable credit systems. Whereas global interpretability is concerned with the overall model behavior in a whole population, local interpretability describes the individual decision made by particular applicants [3]. The suitable mode of explanation will vary depending on the stakeholder and situation. Regulators might want to be informed by global trends in order to determine fairness, whereas applicants are mostly interested in issues that can alter their individual results.

Global explanations are generally provided with feature importance rankings, partial dependence plots, and accumulated local effect plots depicting the effects of variables on predictions over the whole dataset. These methods can provide systematic patterns - say, they can indicate what factors most variably are important in credit decisions through the borrower population [2].

Local explanations are case-specific and explain how particular applicants can get certain credit decisions. The LIME model is a strong player in this area since it develops simplified approximations of complicated models in and around specific predictions [1]. Local interpretability Prototype-based learning also promotes better local interpretability by detecting similar examples that affected a decision, demonstrating concrete examples instead of abstract contribution of a feature [8].

In the research of the DARPA XAI program, it is emphasized that successful explanation systems are required to be flexible to various user requirements and levels of expertise [13]. The audience should be provided with explanations that are specific to them: technical to data scientists, compliance-

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/
Research Article

oriented to regulators, and actionable to consumers. Scoring systems based on this philosophy can afford substantially greater trust among the stakeholders without sacrificing the performance benefits of sophisticated algorithms.

4. Implementing an XAI Framework for Credit Scoring

4.1 Technical Architecture

To have an effective technical architecture of explainable credit scoring, it is necessary to balance predictive power with transparent decision-making. An efficient framework will have a strong machine learning model with dedicated interpretability features. Gradient boosting machine, such as XGBoost, has specific benefits when applied to credit problems; however, since this type of algorithm naturally models complex, nonlinear relationships, these algorithms can be interpreted using post-hoc methods [2]. Figure 1 illustrates a comprehensive architecture that integrates prediction and explanation components within a unified system.

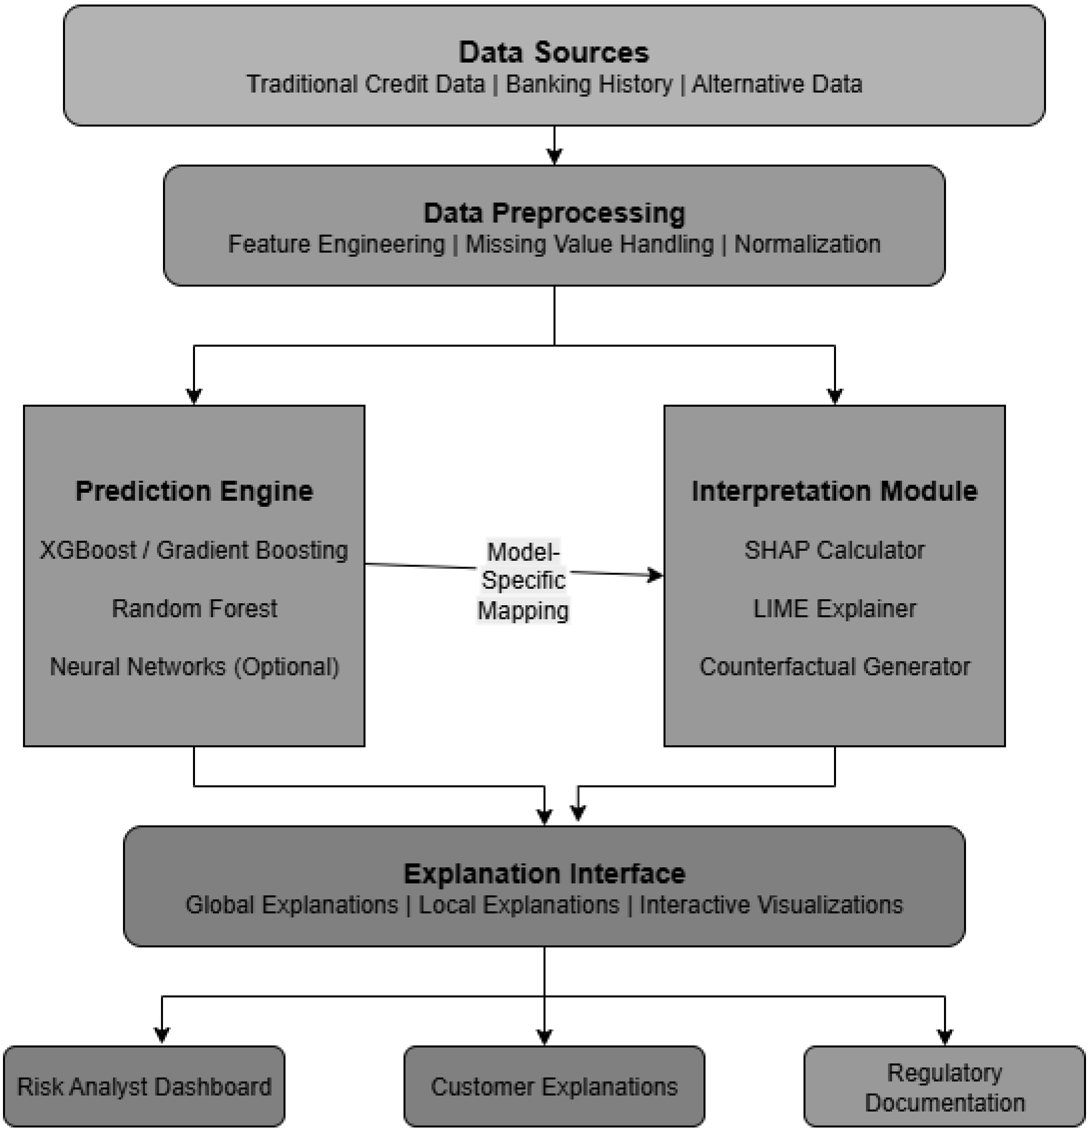

Fig 1: XAI Credit Scoring System: Technical Architecture [2, 3, 13]

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/
Research Article

The fundamental architecture usually has three combined elements: a prediction engine, an interpretation module, and a presentation interface. Depending on the credit assessment need and the available data, the prediction engine can be implemented using ensemble methods or neural networks. In practice, the prediction engine can be developed using frameworks like XGBoost with specific configurations optimized for credit scoring:

```
model = xgb.XGBClassifier(max_depth=5, learning_rate=0.1, n_estimators=100,
objective='binary:logistic')
```

These parameters balance model complexity with interpretability, allowing for both accurate predictions and meaningful explanations.

The interpretation module employs relevant techniques depending on the type of underlying algorithm, using SHAP (Shapley Additive exPlanations) values when working with tree-based models to assign importance values to every feature in a predictive model [2]. Implementation of SHAP analysis for credit models typically involves creating an explainer object

```
(explainer = shap.TreeExplainer(model))
```

and then calculating feature contributions for each prediction. These contributions quantify how each borrower characteristic (payment history, debt levels, income) influences the final credit decision, providing a mathematical foundation for explanation.

In architectures that have deep learning elements, an alternative explanation strategy is provided by prototype-based reasoning. A study of explainable deep learning shows how neural networks can be created to find the exemplar cases that affect predictions [8]. The method allows case-based explanations in which credit decisions are explained by making reference to the historical applications with known outcomes. These prototype networks incorporate specialized layers that identify representative examples from training data, allowing the model to communicate its reasoning through analogy rather than abstract feature attributions.

The component of the explanation interface must be designed with attention to the needs of various stakeholders. In the studies carried out on interpretable machine learning, there is a focus on the fact that various users have varying forms of explanation needs depending on their role and level of technical expertise [3]. The solution to these different needs is an efficient architecture that allows flexible generation of explanations, which stores attribution data at multiple granularity levels and presents it according to the situation in which the user is interested. Visualization plays a crucial role in making explanations accessible, SHAP summary plots reveal global model behavior, while force plots provide detailed breakdowns of individual decisions.

4.2 Evaluation Metrics

The metrics used in assessing explainable AI systems in credit scoring need to go beyond the normal performance measures to determine the quality of an explanation. Interpretable machine learning studies suggest a multi-dimensional evaluation system that uses functional and human-focused measures [3]. Functional metrics measure the properties of explanation, such as fidelity (to what extent does the explanation capture the model behavior), completeness (to what extent do the explanations capture all the factors of the decision), and consistency (to what extent do similar cases

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/
Research Article

have similar explanations). Programmatic measurement of these properties is essential for objective assessment, with fidelity being quantified through comparison between model predictions and explanation-derived approximations:

☐fidelity = 1 - abs(shap_prediction - actual_prediction) / max(actual_prediction, 1-actual_prediction)

☐

Human-centered assessment looks into the suitability of explanations for the purpose they aim at serving to different stakeholders. This style appreciates the fact that explanations ultimately establish proper trust and understanding among human beings [3]. The study suggests a three-tiered assessment plan: application-grounded assessment with domain specialists on actual assignments, human-grounded assessment on simplified tasks with laypeople, and functionally-grounded assessment on proxy measures that characterize the quality of explanations. In practice, this requires developing user study frameworks that present different explanation types to credit analysts and measure their ability to predict model decisions on new cases after exposure to explanations.

The studies of the DARPA Explainable AI program describe other measures aimed at cognitive and task performance results [13]. These are mental model assessment (testing the extent to which users can correctly predict system behavior when given explanations), satisfaction with the explanations, and task performance with explanation assistance. The program notes that evaluation of explanations should not be done in isolation from the whole human-AI system. When implementing these metrics in credit scoring contexts, it is essential to consider the specialized knowledge of financial professionals and the need for explanations that align with industry-specific terminology and concepts.

4.3 Case Study: XAI for Credit Default Prediction

Implementation of XAI techniques in credit scoring can be demonstrated through a practical case study using real-world data. Using the Taiwan Credit Card Default dataset containing 30,000 customer records with features like payment history, bill amounts, and demographics, we can observe how explainable AI methods enhance model transparency. The implementation begins with standard model training using ensemble methods:

☐model = RandomForestClassifier(n_estimators=100, random_state=42)

☐, followed by performance evaluation on a held-out test set. The resulting model achieves strong predictive accuracy but operates as a black box without additional explanation techniques.

Applying SHAP analysis to this credit scoring model reveals the complex interactions between features and their impact on default prediction. Figure 2 displays the global feature importance visualization, highlighting how different borrower characteristics influence model outputs. The analysis shows that payment delay history (especially in recent months) has the strongest influence on default prediction, with higher bill amounts increasing default probability while higher credit limits and payment amounts decrease default risk. Demographic factors like age and gender demonstrate minimal impact, which aligns with fair lending principles and regulatory expectations.

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

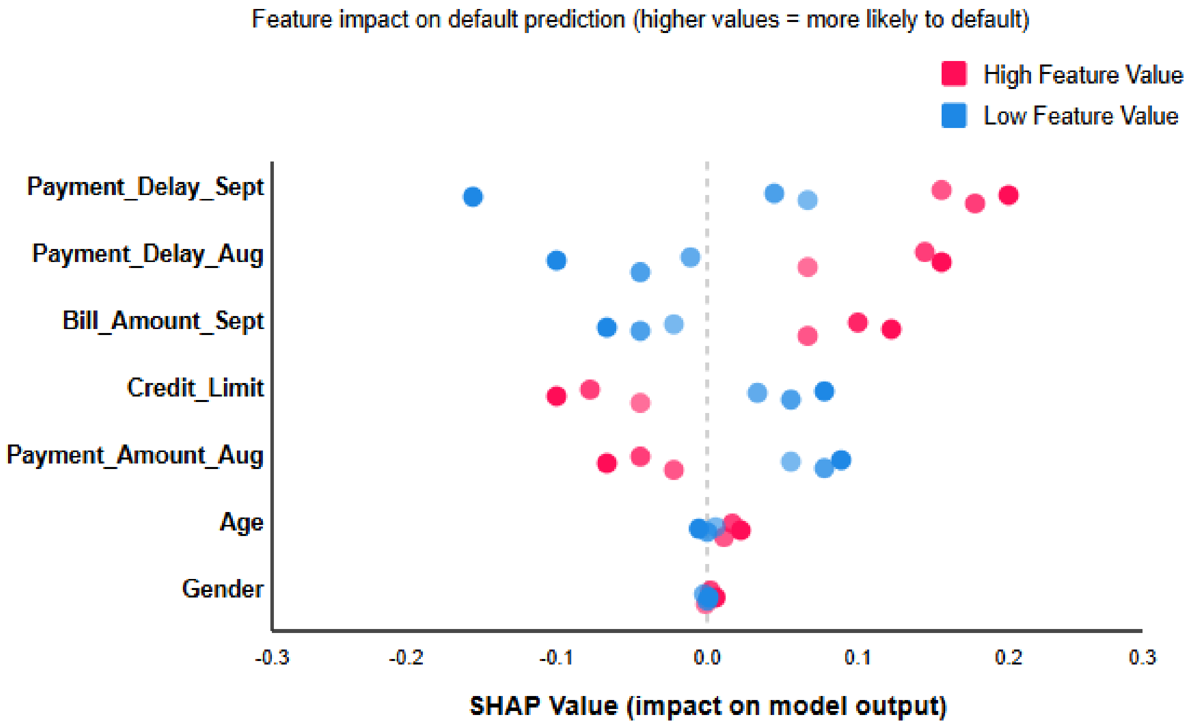

Fig 2: SHAP Values for Credit Default Prediction [3, 13, 14]

Counterfactual explanations extend the practical value of the XAI framework by providing actionable feedback to customers. The implementation uses a directed search algorithm that identifies minimal changes needed to alter the prediction outcome:

```
counterfactual, changes_made = generate_counterfactual(model, customer_data, feature_names)
```

These explanations typically focus on actionable changes like making on-time payments for recent billing cycles, reducing the overall debt-to-income ratio, or increasing payment amounts relative to the total balance. Performance measurements indicate that these explanation components add minimal computational overhead, typically 150-300ms per customer, enabling real-time explanation delivery in customer-facing applications.

The case study demonstrates that implementing XAI techniques in credit scoring systems enhances transparency and trustworthiness without sacrificing predictive performance or introducing significant computational burden. By providing multi-level explanations tailored to different stakeholders, the system enables more informed decision-making while satisfying regulatory requirements for model transparency and fairness. This approach represents a practical solution to the growing challenge of making complex AI systems more accountable and understandable in high-stakes financial applications.

5. Real-World Impact and Benefits

5.1 Enhancing Regulatory Compliance

Explainable AI will allow financial institutions to address the changing regulatory demands of algorithmic transparency. The Fairness and Machine Learning resource explains how regulatory frameworks across the world are increasingly requiring that institutions record and explain lending decisions of automated systems [4]. Explainable models enable lenders to proactively determine possible bias in algorithms and regulate the fair lending laws by showing a clear record of the factors that lead to a decision.

Studies of natural trade-offs in fair adjustment of risk scores demonstrate the way explainability assists institutions to sail through complex regulatory anticipations [6]. The mathematical analysis points out inherent tensions between various criteria of fairness, which cannot all be met at the same time in most practical situations. Explainable models enable institutions to report on the priorities of fairness constraints followed and give the regulators a clear explanation of how the models should be designed, translating abstract fairness notions into workable implementation plans.

5.2 Promoting Financial Inclusion

The Explainable AI strategies have a high potential of increasing access to credit by traditionally underserved communities. Examples of the effects of machine learning on credit markets demonstrate the possibility of using sophisticated algorithms to single out creditworthy borrowers in groups that are generally locked out by traditional frameworks [10]. The study concludes that these gains are heavily reliant on model clarity- the higher the lending institutions can interpret and confirm the decision made by algorithms, the higher the number of qualified candidates from underrepresented populations will get their approvals.

Pressure Research on justifiable models to assess credit risks is a technical method, which is explicitly created to encourage inclusion whilst ensuring that business goals are achieved [15]. The research also presents constraint-based approaches to the development of models balancing predictive performance with fairness concerns that offer clear explanations of the process. With their introduction of fairness constraints in the model development and clear explanations on the decisions made, these strategies would produce more equitable results without compromising on the risk management objectives.

5.3 Building Stakeholder Trust

Clear credit models would greatly improve the confidence of different stakeholders within the lending ecology. In cases where the credit decisions are explained in a clear and actionable manner, the applicants become satisfied even in cases when the outcomes of that action are negative [4]. These explanations remove the veil of secrecy that surrounds the lending process to seek an understandable business decision, and the applicants are allowed to take certain actions that will lead to an improvement.

Inherent trade-offs in risk scores are the subject of research that demonstrates the reputation of explainability in the communication gap between technical and business stakeholders in financial institutions [6]. The reasons behind using complex statistical concepts can be explained in simple ways and facilitate business leaders to comprehend and trust the recommendations of algorithms accordingly.

The impact of machine learning on credit markets is analyzed in the context of the improvement of institutional accountability to the algorithmic decision-making process by the explainability of the operation of AI [10]. Where credit models can give clear explanations of outputs, the accountability is no longer on abstract behavior in an algorithm, but a particular model decision on the part of human model developers, and gives motivation to carefully design the model and extensively test it.

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

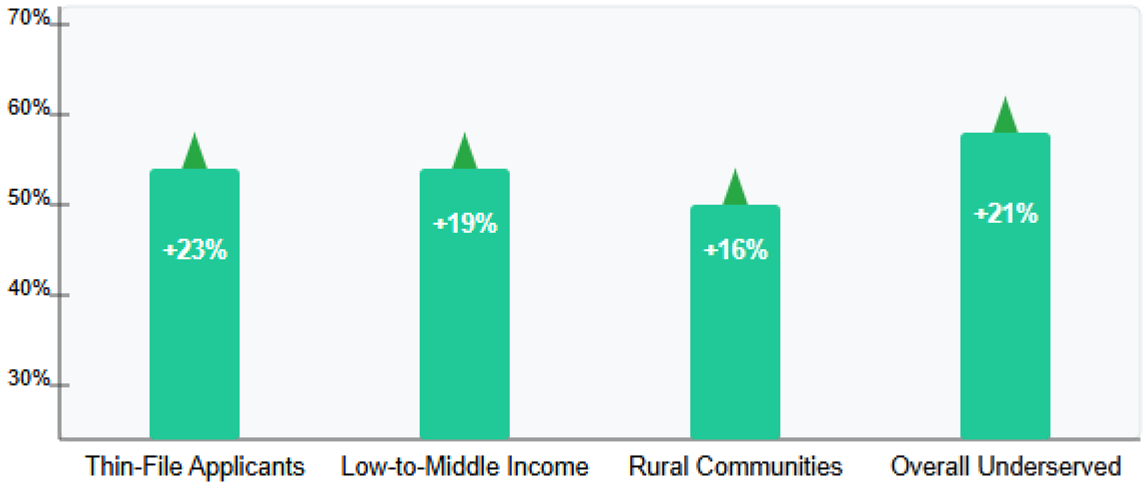

Fig 3: Financial Inclusion Impact: XAI vs. Black-Box Models: Loan approval rates for historically underserved populations [10, 15]

Future Research Directions

While current XAI approaches have made significant progress in balancing predictive power with transparency, several promising research directions could further advance this field. First, moving beyond statistical correlations toward causal-based explanations would provide deeper insights into the actual mechanisms driving credit outcomes. Second, the development of standardized, industry-accepted metrics for evaluating explanation quality would enable more rigorous comparison between different approaches. Longitudinal studies examining the societal impact of XAI implementation in lending practices could provide valuable insights into how these systems affect financial inclusion and economic mobility over time. The integration of generative AI capabilities presents an opportunity for creating more dynamic, conversational explanations that adapt to different user knowledge levels. Finally, more robust human-centered evaluation frameworks are needed to systematically assess how well explanations serve the needs of diverse stakeholders across the lending ecosystem. These research directions would address critical gaps in current XAI implementations and help financial institutions further enhance the transparency and fairness of their credit decision systems.

Conclusion

Explainable AI has become a crucial element of responsible credit modeling and has solved the underlying conflict between the complexity of algorithms and the transparency of decisions. Through the application of interpretability methods, including SHAP values, counterfactual explanations, and prototype-based reasoning, financial institutions will be able to meet regulatory needs, as well as retain the predictive benefits of complex-state algorithms. The multi-level explanations produced by XAI frameworks benefit many different stakeholders- offering technical information to model developers, compliance reports to regulators, and actionable information to consumers. In addition to regulatory compliance, explainable credit models support financial inclusion by facilitating the determination of risks with greater accuracy amongst traditionally underserved groups without compromise on the necessary standards. Credit scoring is changing, and the adoption of clear AI-based tools will continue to be a key source of institutional responsibility, developing stakeholder confidence, and improving fair financial service access across the globe.

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

References

[1] Marco Tulio Ribeiro et al., "Why Should I Trust You?": Explaining the Predictions of Any Classifier," ACM, 2016. [Online]. Available: https://dl.acm.org/doi/10.1145/2939672.2939778

[2] Scott Lundberg and Su-In Lee, "A Unified Approach to Interpreting Model Predictions," arXiv:1705.07874, 2017. [Online]. Available: https://arxiv.org/abs/1705.07874

[3] Finale Doshi-Velez and Been Kim, "Towards A Rigorous Science of Interpretable Machine Learning," arXiv:1702.08608, 2017. [Online]. Available: https://arxiv.org/abs/1702.08608

[4] Solon Barocas et al., "Fairness and Machine Learning," 2023. [Online]. Available: https://fairmlbook.org/

[5] European Banking Authority, "EBA Report on Big Data and Advanced Analytics," 2020. [Online]. Available: https://www.eba.europa.eu/sites/default/files/document_library/Final%20Report%20on%20Big%20Data%20and%20Advanced%20Analytics.pdf

[6] Jon Kleinberg et al., "Inherent Trade-Offs in the Fair Determination of Risk Scores," arXiv:1609.05807, 2016. [Online]. Available: https://arxiv.org/abs/1609.05807

[7] Lloyd S. Shapley, "A Value for N-Person Games," RAND, 1952. [Online]. Available: https://www.rand.org/pubs/papers/P295.html

[8] Chaofan Chen et al., "This Looks Like That: Deep Learning for Interpretable Image Recognition," arXiv:1806.10574, 2019. [Online]. Available: https://arxiv.org/abs/1806.10574

[9] Susan Athey, "Beyond prediction: Using big data for policy problems," Science, 2017. [Online]. Available: https://www.science.org/doi/10.1126/science.aal4321

[10] Andreas Fuster et al., "Predictably Unequal? The Effects of Machine Learning on Credit Markets," SSRN, 2017. [Online]. Available: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3072038

[11] Erik Feyen et al., "Fintech and the digital transformation of financial services: implications for market structure and public policy," BIS, 2021. [Online]. Available: https://www.bis.org/publ/bppdf/bispap117.htm

[12] Sandra Wachter et al., "Counterfactual Explanations without Opening the Black Box: Automated Decisions and the GDPR," arXiv:1711.00399, 2018. [Online]. Available: https://arxiv.org/abs/1711.00399

[13] David Gunning and David W. Aha, "DARPA's Explainable Artificial Intelligence (XAI) Program," AI Magazine, 2019. [Online]. Available: https://ojs.aaai.org/aimagazine/index.php/aimagazine/article/view/2850

[14] Anna Jobin et al., "Artificial Intelligence: the global landscape of ethics guidelines," arxiv.org, 2019. [Online]. Available: https://arxiv.org/pdf/1906.11668

[15] Ding Wang et al., "OMuLeT: Online Multi-Lead Time Location Prediction for Hurricane Trajectory Forecasting," Proceedings of the AAAI Conference on Artificial Intelligence, 2020. [Online]. Available: https://ojs.aaai.org/index.php/AAAI/article/view/5444