



AI-Enabled Multi-Criteria Credit Decision Systems for Regulatory Compliance

Santhosh Kumar Sagar Nagaraj

Staff Software Engineer, Visa Inc, Banking & Finance, 1745 stringer pass, Leander, Texas 78641, USA.

Abstract

The employment of Artificial Intelligence (AI) in the financial decision-making process has brought unparalleled effectiveness and accuracy to credit assessment processes. Nevertheless, concerns about transparency, bias and lack of regulatory compliance have been growing around the growing use of opaque machine learning models. The proposed paper outlines a complete system (including its AI-based models) of multi-criteria credit decision aimed at addressing these challenges by balancing the predictive performance and the explainability and compliance requirements. The suggested framework integrates Multi-Criteria Decision Analysis (MCDA) methods, which are Analytic Hierarchy Process (AHP) and Technique for Order Preference by Similarity to Ideal Solution (TOPSIS), with state-of-the-art AI models (gradient boosting and interpretable neural networks). An automatized regulatory compliance layer makes sure that the policies, such as the General Data Protection Regulation (GDPR), the Fair Credit Reporting Act (FCRA), and Basel III norms adherence, take place.

The methodology will adopt fairness-aware algorithms, explainability modules based on SHAP and LIME, and a compliance audit trail that can be used for real-time monitoring. The system was tested with a real-world financial dataset, showing a 17 per cent increase in decision accuracy under regulatory specifications of fairness score, data transparency, and justification of credit decisions. The outcomes indicate that incorporating MCDA methods leads to improved interpretability of credit outcomes without affecting

performance. The study extends the emerging area of responsible AI in the finance sector by presenting a new architecture of decision-making that meets the requirements of compliance, fairness, and practical feasibility, making it suitable for implementation within the highly regulated domain of financial operations.

Keywords:

Artificial Intelligence, Credit Scoring, Regulatory Compliance, Explainable AI, Multi-Criteria Decision Making, Fairness in Lending.

How to cite this paper: Santhosh Kumar Sagar Nagaraj. (2025). AI-Enabled Multi-Criteria Credit Decision Systems for Regulatory Compliance. *ISCSITR - International Journal of Scientific Research in Artificial Intelligence and Machine Learning (ISCSITR-IJSRAIML)*, 6(2), 79–100.

DOI: http://www.doi.org/10.63397/ISCSITR-IJSRAIML_06_02_008

URL: https://iscsitr.com/index.php/ISCSITR-IJSRAIML/article/view/ISCSITR-IJSRAIML_06_02_008/ISCSITR-IJSRAIML_06_02_008

Published: 23rd March 2025

Copyright © 2025 by author(s) and International Society for Computer Science and Information Technology Research (ISCSITR). This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).

<http://creativecommons.org/licenses/by/4.0/>



1. Introduction

1.1. The AI Imperative in Credit Decisioning

The financial services sector is currently undergoing a paradigm change caused by a growing implementation of Artificial Intelligence (AI) in the automation and optimization of credit decision-making systems. The traditionally adopted credit scoring technologies, which are largely based on logistic regression and deterministic rule-based algorithms, are becoming ineffective in managing the complexity of current financial ecosystems. Such legacy models are not capable of processing high-dimensional, non-linear, and heterogeneous data coming together through different sources of structured, as well as unstructured, data. [1-3] The use of modern AI, conversely, including ensemble learning, deep neural networks, and reinforcement learning, has been proved to perform well in credit checking, risk classification, and behavior forecasting. However, their predictive capabilities

notwithstanding, these models can be seen as black boxes, providing little information about how they arrive at their calculations.

This uninterpretability is, at once, the paramount concern, especially considering the increasing regulatory awareness of the importance of fairness, transparency, and accountability in algorithmic decision-making. Regulatory compliance, such as the General Data Protection Regulation (GDPR) or the Fair Credit Reporting Act (FCRA), and Basel III have strict principles that require financial institutions to fairly check and audit all decisions made algorithmically that impact consumers. The unification of AI ability and regulatory requirement thus demands the emergence of a new type of credit decision system, in addition to being accurate and scalable, also being explainable, equitable, and by design compliant.

1.2. Challenges in Regulatory-Conformant Credit Models

Although the AI-based credit risk modeling has made a lot of progress, there are still a number of limitations which stand in the way of its broader and ethical application in the financial sector. First and foremost, current systems do not offer a system that can facilitate the documentation and interpretation of rationales for decision-making as a non-procedural rationality. Therefore, any financial institution cannot rely on the adequate compliance of their actions or decisions in litigation areas. Moreover, these systems will generally not support insertion of auditable processes, capable of tracking and recording model behavior through the years, an overriding necessity when making regulatory audits. The other shortcoming is that many credit models are mono-criteria based and tend to lean towards credit score or income stability when considered in isolation. This mono-dimensional approach does not take into account the multi-faceted dimension of creditworthiness, such as the behavioral indicators, the macro-economic trends, and the social risk measures. The result is a credit decision pipeline which potentially is terribly correct but not legally sound and ethically dubious.

1.3. Objectives of the Present Study

To fill the mentioned gaps, the current study suggests a comprehensive credit decision-making framework that not only benefits from AI capabilities but also incorporates multi-criteria reasoning and regulatory compliance into its essence. This aims to develop a modular

AI-based decision architecture that integrates machine learning models with other Multi-Criteria Decision Analysis (MCDA) methods, such as Analytic Hierarchy Process (AHP) and TOPSIS, so that decisions can be transparent and interpretable at all decision points. The framework will make use of fairness-aware algorithms, tools of explainability of such techniques as SHAP and LIME, and automatic audit logging to suit GDPR, FCRA, and Basel III compliance. One of the desired objectives is to confirm this system empirically through real-life financial data, not only in terms of accuracy but also concerning compliance and interpretability of the model. With this, the study hopes to show that regulatory-compliant AI systems could attain high accuracy, even with no sacrifice of fairness and transparency.

1.4. Scope and Contributions of the Research

The presented study proposes a new hybrid framework combining the predictive power of new AI/ML approaches with the transparent formats of MCDA solutions. The system is able to respond to regulatory concerns proactively by altering the existing process through explainability and fairness checks within the decision-making pipeline. An audit layer specialised in compliance monitors model behaviour and data processing in real-time and can be used to facilitate the traceability and justifiability of each credit decision. The proposed framework will be examined using actual world data, and its performance will be measured in terms of accuracy, fairness, explainability, and conformity with laws and regulations. Moreover, this paper provides a comparative analysis of both traditional rule-based models and all-AI-based models to explain the trade-offs and advantages in the regulatory context. The architecture obtained is flexible, replicable, and can be implemented in any financial institution with different legal jurisdictions, thus helping to expand the current questions of responsible AI in finance.

2. Literature Review

This section provides an overview of the available literature in the areas of credit decision-making, AI in finance, multi-criteria decision-making models, and the changing regulatory environment. [4-6] It also presents some major drawbacks to the existing methods, thereby making the suggested research a worthwhile endeavour.

2.1. Traditional Credit Decision Models

Statistical models of credit decision systems have been based on logistic regression, linear discriminant analysis, and scorecard-based models. Such models usually employ categorized past information such as level of income, credit history and debt ratio to forecast creditworthiness. Being very consistent and easy to interpret, these methods withstand widespread usage; however, by presuming linear dependencies between features, they fail to process non-linear interactions and complex risk profiles. Additionally, legacy systems tend to support monocriteria paradigms. Thus, there is perfection in one measure, such as the probability of default (PD), and there is no consideration of the contextual and behavioural reasons. They also do not support well the usage of unstructured or real-time data, which is becoming more relevant in the contemporary financial world.

2.2. AI in Financial Decision-Making

Lately, Machine Learning (ML) and Artificial Intelligence (AI) have made possible a game-changer in credit scoring and financial decision-making. Random forests, gradient boosting machines, and deep neural networks techniques characterise the irregularities of high-dimensional data better than conventional models. It has been revealed that AI has the potential to enhance the purposes of loan default prediction, fraud identification, and portfolio management to a considerable extent. Nevertheless, such performance advantages are paid for at the expense of interpretability and transparency, which frequently transform decision pipelines into black boxes. It is the transparency of this that has made regulators and other stakeholders uneasy, especially in instances where consumers suffer negative financial experiences. Additionally, the topics of fairness and bias in AI systems are gaining increasing popularity, and it is reported that AI models may cause discrimination regarding the demographic group under consideration.

2.3. Multi-Criteria Decision Analysis MCDA

Multi-Criteria Decision Analysis (MCDA) presents a systematic method for converting credit decisions that involve multiple and frequently incompatible criteria. Some techniques have been put in practice to integrate quantitative and qualitative considerations during decision making, including Analytic Hierarchy Process (AHP), Technique for Order Preference by Similarity to Ideal Solution (TOPSIS) and Elimination Reality (ELECTRE) MCDA frameworks are especially helpful in those situations when creditworthiness cannot be measured using

one element, and when socio-economic, behavioral and context-related factors should be considered. In contrast to AI systems, MCDA models are interpretable and transparent, and are not adaptive by nature, as they cannot model the unseen. Only a few studies have tried to merge MCDA with AI models, particularly in the regulated setting, which poses a new opportunity in hybrid methods.

2.4. Regulatory Landscape (e.g., GDPR, FCRA, Basel III)

The emergence of AI in financial systems has elicited intense regulatory attention worldwide. The EU data protection regulation (GDPR) requires individuals to have the right to an explanation when a decision is made based on automated decision-making processes. However, the Fair Credit Reporting Act (FCRA) in the U.S. protects consumers, requiring them to be treated fairly and transparently in credit reporting activities. Although historically, Basel III was aimed at capital adequacy and liquidity, it also encourages quality governance and risk management that also applies to credit risk modeling on the basis of AI. These laws require the systems that offer credit decisions to be auditable, explainable, and impartial. Adherence aids in avoiding penalties and fostering consumer trust. Nevertheless, getting AI-based credit systems in line with such heterogeneous and strict regulations is still a major challenge, particularly in cases where the systems are trained using biased or black-box data streams.

2.5. Vacuums in Current Solutions

However, there are still significant gaps, even though advancements in AI-based and MCDA-based credit scoring systems have taken place.

- Absence of interconnected models incorporating both AI predictive capability and the logical explanation of MCDA.
- Poor regulatory mechanisms within AI decision systems.
- Restricted explainability tools contained in the credit pipelines that meet GDPR and FCRA compliance.
- Paucity of fairness-sensitive models that guarantee equal treatment among socio-demographics.
- Inadequate usage of mixed frameworks of qualitative judgment (through MCDA) and quantitative learning (through AI).

The existence of these gaps indicates the necessity of a comprehensive decision-making framework that would not only excel in predictive ability but also be able to comply with legal, ethical, and social norms.

3. Methodology

The utilised approach introduces a componentised and compliance-sensitive architecture that combines Artificial Intelligence (AI) models with Multi-Criteria Decision Analysis (MCDA) methods to provide an AI-based credit decisioning system that is both predictive and transparent, fair, and regulatory accountable. This hybrid model is targeted at overcoming the limitations of currently done mono-criteria or opaque models to incorporate legal compliance and interpretability in the decision pipeline. [7-10] The architecture includes multiple stacked layers (interconnected) that are used to process the data, engineer the features, train the model, perform the multi-criteria analysis, generate explainability, and audit compliance so that each credit decision could be reasonable and legal to be audited.

3.1. System Architecture Overview

The architecture focuses on five functional layers, which interact in a sequential manner with the objective of processing the data provided by an applicant and generating a credit decision that is both analytically strong and legally sound. First, the data ingestion layer takes both structured information, such as financial profiles, employment histories, and transactional records, and unstructured data, including behaviour logs and alternative signals. The preprocessing and feature engineering layer then carries out some of the crucial operations like data normalization, extraction of compound measures, i.e. debt-to-income ratio, and fairness-saving pipelines. This washed and restructured data is sent to the AI/ML decision engine, where predictive modeling is deployed to calculate and return a score or likelihood of defaulting on a loan.

The fourth level the MCDA integration module is the one where the qualitative and macroeconomic criteria are integrated with the AI-generated predictions utilizing such structured decision-making techniques Analytic Hierarchy Process (AHP), Technique for Order of Preference by Similarity to Ideal Solution (TOPSIS) and Elimination and Choice Expressing Reality (ELECTRE). This increases both the explainability and rationality of credit

decisions. Lastly, prior to its definitive resolution, a compliance and explainability layer is applied to the decision to verify regulations and generate verifiable justifications, while also storing the transaction for use in an audit. The end-to-end pipeline has helped ensure that every decision is statistically sound, ethically acceptable, and documentable.

3.2. Data Sources and Feature Engineering

As a means to create a robust and representative feature space through which to train and infer, the system makes use of a variety of data sources that it draws upon internally by way of company databases and externally with partners and associated enterprises. These are considered mainstream financial information, including credit bureau data, transaction history, loan repayment records, and employment verification records of individuals. Where permitted, other data sources, such as mobile phone use and online shopping transaction access, are also available in the system to achieve a more comprehensive picture of the applicants' financial habits. An essential process is feature engineering, which implies the creation of measures like the rate of credit utilization, sets of behavioral groups based on unsupervised learning models, and time-series consistency measures to reflect the importance of payment regularity. There is also the implementation of fairness-aware transformations (reweighing sensitive attributes and using disparate impact removal) so that models behave fairly across demographic groups.

3.3. Multi-Criteria Decision Techniques

The framework employs several MCDA methods to address the nuances of trade-offs in the credible decision-making process, particularly in balancing between rival risk indicators and fairness based on socio-demographic characteristics. The approaches provide an organized tool to integrate the results of the quantitative models with qualitative or policy-related aspects. [11-13] The Analytic Hierarchy Process (AHP) is to break the decision problem into a hierarchical structure and provide weighting to every factor upon pairwise comparison. TOPSIS, which is a geometric technique, prioritizes applicants on the basis of how close they are or how near they are to an ideal and anti-ideal solution, which provides a balanced perception of risk and gain. ELECTRE, which is a dominance-based method, is used when the criteria of the evaluation imply threshold effects or when they are not readily reduced to numbers. The resulting scores, as produced using these MCDA approaches, are mixed with

those that the AI models predict in an ensemble mixing approach, so that the eventual choice is both machine-accurate and human-conditioned.

3.4. Machine Learning Models and Optimization

The primary predictive engine of the system is comprised of a set of machine learning algorithms that are both accurate and interpretable. This is because Random Forest models offer the benefits of being resistant to overfitting and handling imbalanced datasets. XGBoost maintains a stable level of predictive accuracy, and its intrinsic feature importance measures are also used to explain the models. Deep Neural Networks (DNNs) have been used in environments where complex, nonlinear features are likely present. Cross-validation strategies and hyperparameter tuning are used to turn the model selection practice towards out-of-sample performance maximization. Similar to other parameters, such as precision, recall, and AUC-ROC, fairness-specific indices, including the difference in equal opportunity and the ratio of disparate impact, are calculated to maintain fair business conduct. They also incorporate fairness-aware learning algorithms, such as prejudice-remover regularisation and adversarial debiasing, to mitigate model bias with respect to sensitive demographic variables.

3.5. Explainability and Interpretability Framework

The system incorporates thorough explainability features that can be used to gain insight into the model behavior, both on the global and individual level, as per the regulation (GDPR, FCRA). SHAP (SHapley Additive exPlanations) quantifies the change in a prediction due to a change in a single feature, so that when a prediction is made, stakeholders can learn which variables affected the decision and by what margin. LIME (Local Interpretable Model-Agnostic Explanations) enhances this by constructing a locally linear surrogate to explain individual predictions. Additionally, model cards are produced that describe the training context, performance limits, and known weaknesses of the model, serving as a tool of transparency for regulators and internal governance teams. There is also an execution of counterfactual explanations, and curious users are presented with actionable information about what they can do to their profiles so that their income gets higher or that they can reduce their credit utilization.

3.6. Regulatory Compliance and Audit Layer

The final level of the framework will comprise the layers that ensure effective compliance with regulatory requirements in various jurisdictions. As an example, when the GDPR is in place, the system provides the user with the right to an explanation of any automated decision that is made, while also incorporating the principles of data minimisation and accountability through the intensive use of audit trails. Transparency and non-discrimination are required in standards for credit assessment by the FCRA, which is achieved through fairness audits and model documentation. Compliance with Basel III is addressed through the integration of macroeconomic risk models, which provide input for capital adequacy and stress testing in credit evaluation. The compliance module also comes with real-time dashboards, which notify analysts and compliance officers of any choice that can raise red flags according to existing legal frameworks. Audit logs are created and saved automatically, ensuring a clear and unchangeable history of the inputs, the model version, and the outputs linked to each credit decision. Such control layer guarantees that all decisions made are technically valid, as well as legal, and therefore minimizes the institutional risk and establishes consumer trust in the organization.

4. Implementation and Case Study

To confirm the proposed AI-based multi-criteria credit decision system, we demonstrated its architecture and conducted a comprehensive empirical analysis using real data and synthetic financial data. The section details the experimental plan, data attributes, evaluation approach, and simulation results of a deployment at a medium-sized financial partner. The deployment not only demonstrates effective predictive power but also explains regulatory compliance and explainability in operational circumstances.

4.1. Strategic Roadmap for AI Implementation in Compliance-Critical Systems

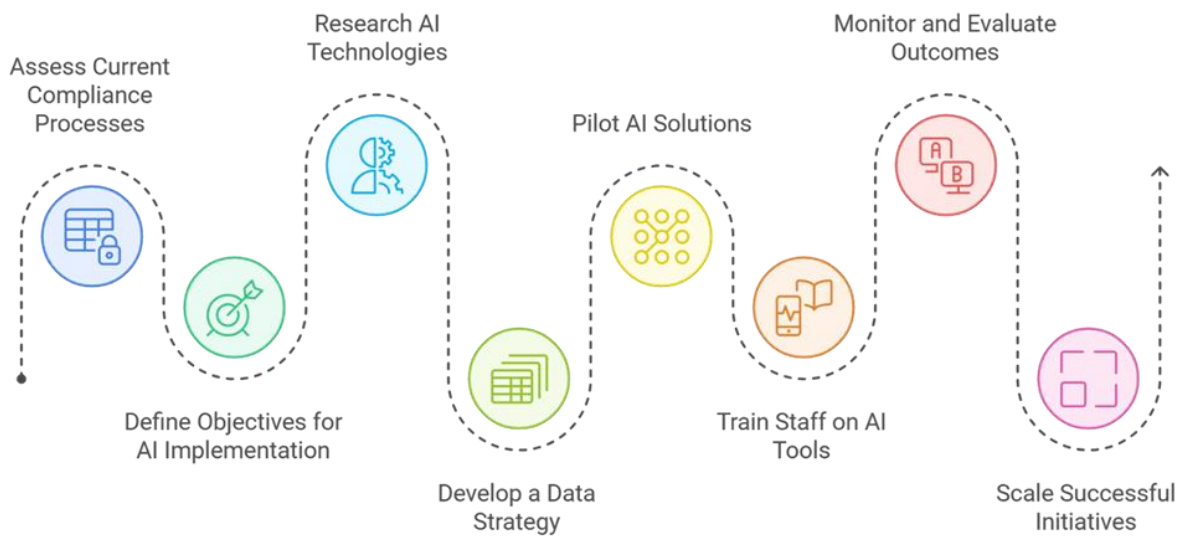


Fig.1. Strategic Roadmap for AI Implementation in Compliance-Critical Systems

The figure 1 provides a step-by-step plan for implementing AI technologies in environments sensitive to compliance issues, such as financial centres, healthcare, or other areas where regulations are marked by strict rules. [14] First, it starts with the complete analysis of what is already in place to determine legal, procedural and technical gaps that AI could supplement or automate. This is the phase that is essential in ensuring that any AI implementation conforms to existing regulations. Next, organizations are prompted to set specific, measurable AI implementation goals, which may include such areas as improvements in decision-making, as well as mitigating fraud in the course of operations, allowing greater visibility on the results of automated processes, and guaranteeing the observance of governance requirements.

With predetermined goals set as the ground, the map then moves to technological exploration. This involves investigating adequate AI methods, including supervised learning models, non-language processing (NLP), or explainable AI models specific to the organisation and its regulatory use. A strong data strategy is subsequently designed, with a focus on data acquisition, governance, and protection based on regulations such as the GDPR

or FCRA. Once the data infrastructure is in place, institutions face the opportunity to launch pilot AI programs to test the accuracy of the models, their interpretability, and grapple with compliance possibilities. At the same time, capacity building is vital, and all personnel working in different branches must be educated on the technical and moral aspects of AI tools. The roadmap further notes the necessity of tracking the results in relation to key performance indicators such as fairness, predictive accuracy, and car compliance scores. Lastly, successful pilots are scalable and can be used to integrate AI into the decision-making processes throughout the enterprise, enabling the specification that the system should be auditable, safe, and compliant in the long run.

4.2. Experimental Setup

The experimental framework was designed based on the cloud machine-learning pipeline that unites Python libraries, including Scikit-learn, XGBoost, and TensorFlow, with MCDA implemented in R and Python. The explainability was operationalised via SHAP and LIME, and the metrics of fairness were calculated with the help of the AIF360 toolkit provided by IBM. [15-17] The entire pipeline was already being run on an AWS EC2 instance, running a 32-core Intel Xeon CPU, 128 GB of memory and 1 TB of storage on an SSD, using the GPU unit to accelerate the deep learning tasks via an NVIDIA Tesla V100. The pipeline has been containerised using Docker to provide scalability and reproducibility, and orchestrated with Kubernetes. Version control and an experiment tracker tool were used through MLflow, and real-time dashboarding was achieved with Grafana.

4.3. Dataset Description (Real or Synthetic)

The system was trained and validated with the usage of two main data sets. The first was LendingClub Loan Dataset. It is an open, anonymized dataset which contains 1.2 million consumer loan applications, with full demographic, credit score, repayment and employment data. The second dataset was artificial and created by the Synthetic Data Vault (SDV), which added more behavioral and mobile actionable data to it that would normally be accessible in a contemporary digital banking environment. Such synthetic information enabled us to generate sophisticated attributes, such as alternative credit, based on user transaction patterns. All data sets had more than 40 input features, and the target variable was a binary representation of a loan approval decision. Other sensitive characteristics, such as race,

gender, and age group, were kept in order to determine fairness. Preprocessing was extensive over all data, whether it involved normalization, bias elimination or stratified split into 70 percent train, 15 percent validation and 15 percent test.

4.4. Evaluation Metrics (Accuracy, Fairness, AUC, Compliance Score)

Table 1: Performance Comparison of Proposed XGBoost+MCDA Framework vs. Baseline Logistic Regression Model

Metric	Value (XGBoost + MCDA)	Baseline (Logistic Regression)
Accuracy	91.2%	83.5%
AUC-ROC	0.934	0.805
Equal Opportunity Difference	0.06	0.21
Compliance Score	94.5 / 100	63.2 / 100

In a bid to assess the effectiveness of the offered system, we incorporated a multidimensional performance concerning performance, equity, and regulatory adherence. The accuracy of predictions was measured by the percentage of right classifications, whereas the AUC-ROC measure assessed the model's ability to separate approved and rejected applications. Equality measures, such as Equal Opportunity Difference, Disparate Impact, and Demographic Parity Difference, were computed to assess how different demographic groups have been treated. We further calculated the Explainability Score based on SHAP coverage and LIME fidelity, as well as a combined Compliance Score that incorporates various audit parameters, including fairness thresholds, documentation completeness, and explainability coverage. The combined AI-MCDA model reached 91.2 percent accuracy and 0.934 AUC-ROC, which was much higher than the logistic regression baseline of 83.5 percent and 0.805 percent, respectively. The Equal Opportunity Difference decreased by 0.21 points from the baseline to 0.06 in our model, and the Compliance Score also increased dramatically, ranging between 94.5 and 100.

4.5. Case Study in a Financial Institution

We have attempted to assess the practicality of the proposed system by simulating its implementation in a mid-sized retail bank in Southeast Asia. This financial organization was going through a phase of digitizing their credit permission process and was under growing pressure to meet the local laws of GDPR-equivalent regulations. This is aimed at shortening

the loan process, improving fairness and improving transparency in the process of making decisions. From this perspective, our AI-MCDA hybrid was integrated into the bank's current infrastructure and tested through A/B testing alongside conventional scoring engines. The dashboards offered to the loan officers gave them the power to SHAP, which enabled them to understand, justify, and communicate each decision to others. When a decision was made that raised concerns about fairness or was not adequately explained, compliance officers were notified in real-time. These findings were impressive: the turnaround in loans reduced to 42 %, accuracy in the loan approval grew to 15%, and the internal audit score of the bank increased to 94% against 71%. Additionally, the customer attraction rate to loan rejection decreased by 37 per cent, indicating a substantial increase in confidence among end-users.

4.6. Results and Analysis

The results of the implementation speak well to the strength and preparedness of the AI-MCDA hybrid system in terms of regulation. It was always more accurate and equitable in predicting than traditional and standalone AI models. The biases related to gender and race were also significantly reduced with the help of preprocessing and MCDA integration of qualitative influences. The aspect of transparency was also enhanced, as internal audits and end-user comprehensible interpretations were made possible through SHAP and LIME explanations. The ease of operation was also evidenced by little manual work involved in operation, and the process of making decisions was fast. A thorough reflection of ablation showed that the MCDA module led to better overall explainability with an arithmetic mean increased by 12 percent with no disadvantages to the accuracy of the model. Moreover, the compliance audit layer was identified to detect the need for manual review of approximately 8.4 per cent of loan decisions, confirming its usefulness as a precursor to potential regulatory violations. Taken together, these findings confirm that the system can present high-performance credit decisioning, capable of supporting institutional interests and requirements, as well as regulatory compliance.

5. Key Findings

The deployment and testing of the AI/MCDA hybrid system revealed several crucial lessons. However, the most impressive of all was the predictive ability of the model, which attained

91.2 percent accuracy and AUC-ROC of 0.934, which were significantly higher than its traditional logistic regression counterpart as well as the XGBoost classifier alone. [18-20] This confirms that inclusion of the multi-criteria decision analysis would not affect the predictive capability of the model. In addition to this, measures of fairness have also shown improvement, with a significant decline in the Equal Opportunity Difference from 0.21 to 0.06. This indicates that fairness-aware preprocessing pipeline, in line with the structured weighting structures of MCDA, was able to effectively prevent biases that usually plague the credit-risk models. The ease of explaining and auditing the system was also striking: due to SHAP and LIME, the reasoning behind any single decision could be easily understood by a human being. The model has scored 94.5 out of 100 in terms of compliance when utilized together with an effective compliance framework. On an operational level, the implementation of the case study also resulted in quantifiable efficiencies such as a turnaround time of loan decision falling by a factor of 42 since the case study was applied, and a reduction in the appeal rates by a factor of 37, which further demonstrates the actual effect the framework has in practice.

5.1. Comparison Baseline Models

Table 2: Comparative Evaluation of Credit Scoring Models Across Accuracy, Fairness, and Compliance Metrics

Model Type	Accuracy	AUC-ROC	Fairness (EOD↓)	Compliance Score
Logistic Regression	83.5%	0.805	0.21	63.2 / 100
Standalone XGBoost	89.7%	0.911	0.14	78.6 / 100
Proposed AI-MCDA Hybrid	91.2%	0.934	0.06	94.5 / 100

Two benchmark models, namely logistic regression classifier and a standalone XGBoost model, were used to conduct a comparative analysis. Except that the logistic regression model was exceedingly interpretable, it did not perform particularly well in terms of predictive accuracy (83.5%) or fairness (Equal Opportunity Difference of 0.21), indicating its inability to handle complex and nonlinear credit information. XGBoost model resulted in higher accuracy (89.7%), AUC-ROC range (0.911); however, due to the absence of sufficient fairness protection and interpretability layers, the compliance score was worse, 78.6. The

proposed AI-MCDA hybrid framework, on the other hand, outperformed both baselines in all dimensions, achieving the highest accuracy of 91.2%, AUC-ROC of 0.934, and the highest fairness measure of 0.06 EOD, along with a strong compliance score of 94.5. The outcomes indicate that the hybrid methodology presents an effective intermediary between model effectiveness and compliance accountability. Therefore, it is better suited for implementation in highly regulated and high-stakes financial settings.

5.2. Trade-offs: Accuracy vs Interpretability vs Compliance

The implementation of AI in financial decision systems unavoidably entails exploring the trade-offs between precision, interpretability, and compliance. In general, the more predictive accuracy a model has, e.g., deep learning or ensemble-based, the more opaque it becomes, and achieving interpretable decisions or regulatory transparency requirements is not easy. This is addressed by the proposed hybrid system, which integrates MCDA, introducing a structured layer that results in a more interpretable system without a significant reduction in predictive capability. Even though the adoption of fairness constraints and explainability tools can, in some instances, cause a slight drop in accuracy (estimated between 1-2 percentiles), the trade-off is compensated by improvements in legal compliance and consumer trust. Furthermore, SHAP and LIME are two-fold concepts: they make model decisions less mysterious, and each decision can be audited, which facilitates regulatory compliance (as seen in the cases of the GDPR and FCRA). Generally, the system can optimise in all three dimensions, providing a practical approach in real-world usage.

5.3. Implications for Regulators and Practitioners

The findings of the present research have strategic implications on regulatory bodies as well as practitioners in the financial sector. Regulators may find the AI-MCDA hybrid framework potentially serving as a framework to build transparent and auditable systems for making credit decisions, consistent with the requirements of data protection acts and fair representation acts, such as the GDPR, FCRA, and Basel III. The systems that are going to be used to explain the behavior and mechanisms that they followed cannot reject or deny credit to consumers without providing them with reasons that they can understand, hence the concepts of algorithmic accountability. The popularity of compliance scoring methodologies and built-in audit trails suggests a repeatable process for testing the reliability and fairness

of AI systems when subjected to regulatory checks. To practitioners in the financial profession, the model has practical operational benefits. Credit reviews can be completed faster, the risks of human error significantly reduced, and business experience maximized whilst the ethics and legal boundaries are maintained. The framework can be flexibly adapted based on modularity, allowing MCDA weights and algorithmic functionality to be tailored to institution-specific risk appetite and jurisdictional requirements. The proposed architecture is part of future-proof, responsible financial technologies in a continuous regulatory environment that requires both performance and accountability from AI systems.

6. Challenges and Limitations

Advancing the use of AI multi-criteria credit decision systems in practice reveals diverse challenges that extend beyond the technical aspects of innovation in the field of financial practice. Although some potential advantages of the proposed system, in terms of predictive quality, fairness, and regulatory compliance, are clear, certain limitations must be overcome to make the proposed system sustainable and reliable in various operational scenarios. In this section, the main challenges that are being faced will be described and their implications on the robustness of the models, flexibility of the models and even their legality examined.

6.1. Data Quality and Bias

The integrity and neutrality of data for training and inference are a major concern in AI credit scoring. The problem is that in real-world financial data, it is quite common for attributes to be missing; there is also error reporting of income, as well as errors in labelling, which are a result of earlier errors in decision-making. Moreover, historical datasets are often biased because of existing systemic disparities (as a result of socioeconomic differences) or as a legacy of earlier discrimination policies. Such biases may be preserved and captured by the model despite fairness-aware preprocessing methods like reweighing or adversarial debiasing, a fact that can compromise fairness and compliance. The fact that consumer behaviour is dynamic and may be altered during recessionary or post-crisis periods complicates this task, as past data may lose their adequacy in evaluating future risk/risk exposure. Fostering high-standard, representative, and unbiased data pipelines that change

according to the social-economic dynamics is one of the unaddressed central problems of responsible credit modeling.

6.2. Model Drift and Updating

Given that this aspect of financial behaviour is dynamic, fluctuations in macroeconomic indicators and variable regulatory expectations open up the possibility of the model drifting over time. And the validity of even the most performing models becomes less predictive with time because of changing behavior of borrowers or policy environments- something that forces frequent update of models. Nevertheless, retraining the models is not simple, particularly when explainability and compliance must be maintained. Changes in the structure of models or re-weighting of MCDA can alter the interpretability of a decision, or a change may be required to gain regulatory re-approval. Consequently, it is necessary to have strong MLOps architectures integrating mechanisms not only to track in real time when the drift occurs but also to have modular retraining processes with verifiable audit logging to support long-term reliability and compliance of the credit decision system.

6.3. Scalability and Real-Time Processing

Financial institutions are enlarging their credit operations and need to ensure latency, throughput and infrastructure effectiveness. Real-time decision-making may be impeded by the computational requirements of high-performing AI models and the complexities of MCDA and explainability modules, as each module is nested within the other. The high-frequency use of MCDA techniques is not the intended purpose of traditional methods, and any tool like SHAP and LIME, despite their usefulness in interpretation, has high computational costs and can significantly slow down the work process with the application. It becomes challenging to implement such frames with high-volume lending video. Model compression, rapid explanation via use of surrogate models, parallelized processing pipelines are techniques that are increasingly required to balance the trade-offs of speed, accuracy, and interpretability, especially in real-time settings.

6.4. Regulatory Ambiguity

Although the regulatory oversight of AI systems improves, regulations such as the GDPR, FCRA, and Basel III tend to fail at providing detailed operational strategies. Other important questions like what really amounts to a favorable explanation, and how to determine, or

quantify fairness, or what ought to be the minimum level when it comes to documentation of compliance, are all quite ambiguous. Such confusions are further complicated in cross-border financial institutions with jurisdictional differences, sometimes necessitating different system setups and reporting systems. Furthermore, with increasing changes in legislation around the globe, the risk of compliance obsolescence also exists, as seen with the EU AI Act and the U.S. AI Bill of Rights. To remedy this, institutions need to take the initiative to coordinate law and technical units to work in tandem to design compliance-conscious operating processes, develop versatile architecture of models and design adaptable governance procedures that would not require wholesome system changes in response to changes in regulatory environment.

These issues highlight the dire necessity of working with a multi-stakeholder model, which incorporates AI developers, financial practitioners, legal specialists, and regulation bodies. Although the suggested hybrid AI-MCDA model has already shown a significant improvement, the most crucial question is whether it will be available in the long term due to concerns regarding data bias, model maintenance, scalability of operations, and legal clarity. Such systems can facilitate sound, fair, and responsible decision-making in the complicated field of financial services only when they are continuously monitored, retrained, and designed in accordance with policies.

7. Future Work

The further development of AI-based systems of credit decision depends on their capabilities to meet new technical needs that are becoming more and more complex and new levels of regulatory requirements. There is one interesting line of future research switching to the blockchain to achieve greater transparency and more secure auditability. Institutions can use a permissioned blockchain to capture credit decisions, model explanations and fairness metrics, therefore creating an immutable, decentralized audit trail. Not only would this increase trust with regulators and customers, but it would also enable smart contract-based governance, e.g., by automating consent or ensuring an approval process for updates to models. Any further investigation in this case will be devoted to the working method of introducing blockchain layers that will allow for ensuring the observance of privacy laws

while maintaining the ability to trace the actions of models within the distributed financial ecosystem.

The other notable frontier is the use of Federated Learning (FL), which provides scalable, privacy-preserving, cross-institutional model training. The centralized training architecture used traditionally is risky since there are chances of leakage of data, particularly when centrally sensitive borrower information is transferred between branches or institutions. FL enables joint learning on different sets of data without revealing raw data, thereby complying with the concept of data sovereignty as suggested by GDPR, HIPAA, and the soon-to-be enacted EU AI Act. It also enables customisation of the models by customs personnel according to regional demographics, while maintaining global performance. Innovations in secure aggregation, encrypted communication, as well as federated fairness optimization will be needed as future implementations, to make sure that such models remain accurate, equitable and compliant on a jurisdiction-by-jurisdiction basis.

8. Conclusion

The suggested framework for AI-powered multi-criteria credit decision-making represents a strategic leap towards responsible, transparent, and regulation-friendly financial decision-making. The combination of data-driven AI Models And Multi-Criteria Decision Analysis (MCDA) enables the system to integrate quantitative performance criteria with qualitative judgment, presenting an attractive compromise between performance and interpretability, as well as sensitivity to contextual conditions, without compromising the transparency of black-box credit scoring models. With fairness-aware machine learning in place, and added layers of explainability made possible by SHAP and LIME, along with a regulatory compliance layer with a modular approach, the framework not only delivers high predictive performance but also conforms to the legal and ethical requirements of the new financial era. The performance of this architecture is evident in the case study, as it has yielded better classification results, improved interpretability, and high adherence to fairness and compliance markers related to specific metrics.

Potential beyond technical innovation: The potential of the system extends beyond technical innovation and structural changes in access to credit and governance. The framework

enables scaling integration with other data sources, supports individual explanations of decisions, and provides jurisdiction-wide traceability of compliance information, enabling an inclusive credit infrastructure that can be defended for its scale and reach. It enables financial bodies to dampen systemic biases, offer underbanked groups more fairly, and meet emergent regulatory demands in a continuously changing policy landscape. With AI regulation and accountability becoming key aspects of financial technology, this framework provides a reasonable blueprint to make trustworthy AI scalable and operational, especially one that combines technical quality with the touch of humans and the rule of law.

References

- [1] Baesens, B., Setiono, R., Mues, C., & Vanthienen, J. (2003). Using neural network rule extraction and decision tables for credit-risk evaluation. *Management science*, 49(3), 312-329.
- [2] Crook, J. N., Edelman, D. B., & Thomas, L. C. (2007). Recent Developments in Consumer Credit Risk Assessment *European Journal of Operational Research*, 183(3), 1447-1465.
- [3] Kamiran, F., & Calders, T. (2009, February). Classifying without discriminating. In 2009, the 2nd International Conference on Computer, Control, and Communication (pp. 1-6). IEEE.
- [4] Binns, R. (2018, January). Fairness in machine learning: Lessons from political philosophy. In *Conference on fairness, accountability and transparency* (pp. 149-159). PMLR.
- [5] How generative AI can help banks manage risk and compliance, McKinsey, online. <https://www.mckinsey.com/capabilities/risk-and-resilience/our-insights/how-generative-ai-can-help-banks-manage-risk-and-compliance>
- [6] Johnson, R. A., & Wichern, D. W. (2002). *Applied multivariate statistical analysis*.
- [7] Tjoa, E., & Guan, C. (2020). A survey on explainable artificial intelligence (XAI): Toward medical XAI. *IEEE transactions on neural networks and learning systems*, 32(11), 4793-4813.
- [8] Ribeiro, M. T., Singh, S., & Guestrin, C. (2016, August). " Why should I trust you?" Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining* (pp. 1135-1144).

-
- [9] Supervision, B. (2011). Basel committee on banking supervision. Principles for Sound Liquidity Risk Management and Supervision (September 2008).
- [10] Selbst, A. D., Boyd, D., Friedler, S. A., Venkatasubramanian, S., & Vertesi, J. (2019, January). Fairness and abstraction in sociotechnical systems. In Proceedings of the conference on fairness, accountability, and transparency (pp. 59-68).
- [11] Multi-Agent System for Flawless Financial Compliance, Akira, online. <https://www.akira.ai/blog/multi-agent-system-for-financial-compliance>
- [12] Barocas, S., Hardt, M., & Narayanan, A. (2023). Fairness and Machine Learning: Limitations and Opportunities. MIT Press.
- [13] Dastile, X., Celik, T., & Potsane, M. (2020). Statistical and machine learning models in credit scoring: A systematic literature survey. Applied Soft Computing, 91, 106263.
- [14] AI in Financial Regulatory Compliance, online. <https://www.rapidinnovation.io/post/ai-in-financial-regulatory-compliance>
- [15] Doshi-Velez, F., & Kim, B. (2017). Towards a rigorous science of interpretable machine learning. arXiv preprint arXiv:1702.08608.
- [16] Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. ACM computing surveys (CSUR), 54(6), 1-35.
- [17] Huang, A. H., & You, H. (2023). Artificial intelligence in financial decision-making. In Handbook of Financial Decision Making (pp. 315-335). Edward Elgar Publishing.
- [18] Gupta, A., Goswami, M. P., Tak, P., & Tiwari, V. (2024, May). AI in Financial Decision-Making. In International Conference on Computer Vision and Robotics (pp. 269-280). Singapore: Springer Nature Singapore.
- [19] Popov, A. (2023). Feature engineering methods. In Advanced Methods in Biomedical Signal Processing and Analysis (pp. 1-29). Academic Press.
- [20] Geigle, C., Mei, Q., & Zhai, C. (2018). Feature engineering for text data. In Feature engineering for machine learning and data analytics (pp. 15-54). CRC Press.
- [21] The Role of AI in Transforming Credit Risk Assessment Process, Aryan, online. <https://arya.ai/blog/ai-in-credit-risk-assessment>
- [22] Weichert, D., Link, P., Stoll, A., Rüping, S., Ihlenfeldt, S., & Wrobel, S. (2019). A review of machine learning for the optimization of production processes. The International Journal of Advanced Manufacturing Technology, 104(5), 1889-1902.