International Journal of Science and Research Archive
eISSN: 2582-8185
Cross Ref DOI: 10.30574/ijsra
Journal homepage: https://ijsra.net/

(RESEARCH ARTICLE)

Check for updates

A Hybrid Predictive Modeling Framework for Financial Risk Assessment Using AI-Driven Simulation and Optimization Techniques

Oluwaseun Lamina *

Department of Mathematics, Western Illinois University, USA.

International Journal of Science and Research Archive, 2023, 10(01), 1236-1245

Publication history: Received on 29 July 2023; revised on 21 September 2023; accepted on 23 September 2023

Article DOI: https://doi.org/10.30574/ijsra.2023.10.1.0748

Abstract

As the economy becomes more unstable, there is a growing need for accurate and flexible ways to assess financial risks. Artificial intelligence (AI), with its capacity to simulate uncertainty and optimize decision-making, offers significant opportunities to enhance traditional financial risk models. This study proposes a hybrid predictive modeling framework that integrates AI-driven simulation techniques such as Monte Carlo methods with optimization algorithms including genetic algorithms and reinforcement learning to provide robust, transparent, and adaptive risk assessment.

Using a synthetic dataset simulating credit, market, and operational risk scenarios across diverse financial portfolios, the framework was evaluated for accuracy, efficiency, and interpretability. The hybrid model outperformed conventional models in forecasting risk exposure and loss probabilities under both normal and stress conditions. Moreover, incorporating explainable AI components, including SHAP (SHapley Additive exPlanations) values and feature attribution maps, enhanced transparency and stakeholder trust in model outputs.

Findings demonstrate that this integrated approach provides significant improvements in predictive performance, model adaptability, and interpretability compared to standalone methods. The study also examines the implications of AI-generated financial forecasts on regulatory compliance, organizational resilience, and ethical risk governance. It concludes with recommendations for the scalable adoption of hybrid AI frameworks in institutional risk environments, emphasizing the importance of transparent modeling and inclusive algorithmic oversight.

Keywords: Financial Risk Assessment; Hybrid Modeling; AI Simulation; Optimization Techniques; Monte Carlo Simulation; Genetic Algorithms; Explainable AI; Risk Governance

1. Introduction

1.1. Background and Significance of Financial Risk Assessment

Financial risk assessment serves as the bedrock of economic stability for institutions, markets, and governments. Its role encompasses identifying, quantifying, and mitigating uncertainties that could adversely affect financial performance, ranging from credit defaults and market volatility to operational failures and systemic shocks. Historically, risk assessment has relied on deterministic models grounded in statistical probability and fixed assumptions about future events. However, the increasingly complex and dynamic nature of global financial systems has exposed the limitations of such approaches, particularly in the face of black swan events like the 2008 financial crisis and the COVID-19 economic fallout [1].

* Corresponding author: Oluwaseun Lamina

In response to these limitations, contemporary financial risk management has shifted towards incorporating probabilistic models and computational simulation techniques that can more accurately reflect real-world uncertainty. Still, the rapid expansion of financial data, the interdependence of global markets, and the unpredictability of human behavior necessitate tools that not only simulate risk but adaptively learn from evolving data landscapes. This transition has paved the way for Artificial Intelligence (AI) to emerge as a transformative force in financial analytics [2].

1.2. The Rise of AI in Financial Modeling

Artificial intelligence, encompassing machine learning, deep learning, and evolutionary computation, enables machines to identify complex patterns in data, forecast outcomes, and optimize strategic responses. Within the financial domain, AI has been used for credit scoring, fraud detection, algorithmic trading, and customer behavior analytics [3]. However, its integration into risk modeling, particularly for forward-looking and scenario-based analysis, remains underutilized.

While AI-driven predictive models offer superior performance compared to traditional statistical techniques, their "black box" nature raises serious concerns around explainability, auditability, and regulatory compliance [4]. The opacity of these models makes it difficult for financial institutions to interpret predictions, justify decisions to regulators, or ensure that algorithmic bias does not inadvertently drive inequitable outcomes. These limitations have catalyzed efforts to develop hybrid models that merge AI's computational power with transparent, rule-based structures.

1.3. Hybrid Predictive Modeling for Risk

Hybrid modeling involves combining multiple modeling paradigms, such as statistical, simulation-based, and AI-driven approaches to leverage their respective strengths. In financial risk assessment, this might entail integrating Monte Carlo simulation for probabilistic estimation, machine learning algorithms for pattern detection, and optimization tools for portfolio adjustment [5].

Hybrid systems can improve model accuracy, enhance robustness under changing conditions, and facilitate better generalization across diverse risk types. They also support modular design, allowing components to be updated or swapped based on specific institutional requirements or regulatory changes. By integrating explainability tools such as SHAP or LIME (Local Interpretable Model-Agnostic Explanations), hybrid models can also provide interpretable outputs that improve trust and governance [6].

This approach holds promise in overcoming the dichotomy between accuracy and interpretability a common tradeoff in risk modeling. For instance, a purely machine learning-driven model may deliver high predictive power but little transparency, while a rule-based simulation may be understandable but overly simplistic. A hybrid framework, carefully architected, can deliver both.

1.4. Gaps in Existing Literature

While prior studies have explored individual components of AI-based financial modeling, few have proposed or empirically tested comprehensive hybrid frameworks that bring together simulation, machine learning, and optimization in a unified risk assessment system. Even fewer have addressed the ethical and operational implications of such models within institutional and regulatory environments [7].

Moreover, the literature lacks rigorous exploration of explainable AI (XAI) techniques in the financial risk context. Given the growing emphasis by regulators on algorithmic transparency and responsible AI use, it is imperative to examine how hybrid models can be both technically effective and ethically sound [8].

Study Objectives

This study proposes and validates a hybrid predictive modeling framework for financial risk assessment that integrates AI-driven simulation, multi-objective optimization, and interpretable machine learning. The primary objective is to develop a comprehensive model that combines Monte Carlo simulation, random forest classification, and genetic algorithm-based optimization to address credit, market, and operational risk scenarios within a unified system. Additionally, the study seeks to evaluate the predictive accuracy, robustness, and computational efficiency of this hybrid framework in comparison to conventional risk modeling approaches. A further aim is to incorporate explainable AI techniques to enhance the transparency of model outputs and support stakeholder understanding. Beyond technical validation, the study also examines the practical, regulatory, and ethical implications of implementing such models within institutional financial risk governance frameworks. Ultimately, the goal is to contribute meaningfully to the discourse on responsible and transparent AI adoption in high-stakes, algorithmic decision-making environments.

2. Methods

2.1. Study Design and Framework Overview

This research employed a systems engineering approach, combining simulation modeling, machine learning, and evolutionary optimization into a unified hybrid predictive modeling framework for financial risk assessment. The methodology was divided into three core components: (i) stochastic simulation for risk exposure generation, (ii) AI-based prediction modeling, and (iii) optimization of risk mitigation strategies.

The proposed framework was validated through synthetic financial datasets emulating diverse portfolio risk scenarios under both stable and volatile market conditions. These scenarios incorporated credit risk, market risk, and operational risk indicators. The entire model lifecycle from data preprocessing to simulation execution, model training, optimization, and explainability output was executed in Python (v3.11) using NumPy, Scikit-learn, TensorFlow, DEAP, and SHAP libraries.

2.2. Data Generation and Variables

Due to the lack of publicly accessible financial datasets that are appropriate for rigorous AI-based optimization research, this study generated a synthetic dataset constructed through probabilistic sampling methods and constrained by established financial logic. The resulting dataset comprised 200,000 individual records, each representing a distinct financial entity or transaction. Variables were selected to capture the multidimensional nature of financial risk and included four main categories: credit risk, market risk, operational risk, and macroeconomic indicators. Credit risk variables encompassed credit score, loan-to-value ratio, payment history, default frequency, and sector classification. Market risk attributes included asset volatility, liquidity index, interest rate exposure, and beta coefficients. Operational risk features comprised event frequency, loss severity, system complexity index, and compliance score. In addition, macroeconomic indicators such as inflation rate, gross domestic product (GDP) growth, and unemployment rate were included to simulate external systemic influences. Risk labels were then assigned to each entity based on Basel III regulatory thresholds and expert-defined rules, categorizing them into low, moderate, or high risk across the three principal dimensions of credit, market, and operational exposure.

2.3. Simulation Component: Monte Carlo Modeling

The first component of the framework simulated loss distributions using Monte Carlo methods. For each financial entity, 10,000 iterations were run to model random fluctuations in key variables, such as asset value or default probability, using Gaussian and t-distributions where appropriate [9].

Conditional dependencies between variables were modeled using copula functions to simulate real-world risk interactions. For example, credit default rates were conditioned on both credit score and economic downturn probability. The simulation output provided probabilistic risk scores for each portfolio, forming the input for the machine learning component.

2.4. Predictive Modeling: Random Forest and Neural Networks

Following simulation, machine learning classifiers were trained to categorize financial entries into risk levels. A random forest classifier was selected for its robustness to noise and non-linearity, while a feedforward neural network (multi-layer perceptron) was employed for comparison. The dataset was split into 70% training, 15% validation, and 15% testing sets.

Model performance was evaluated using accuracy, precision, recall, F1-score, and area under the ROC curve (AUC). A cross-validation procedure (k=5) was used to reduce overfitting risk. The random forest model served as the core predictive engine due to its balance between interpretability and performance.

2.5. Optimization Component: Genetic Algorithm (GA)

To identify optimal mitigation strategies and asset allocations under various risk profiles, a genetic algorithm (GA) was implemented. The GA aimed to minimize expected loss and portfolio variance while maximizing expected return, subject to regulatory capital constraints.

The objective function took the form

$$Min\ [\lambda 1 \cdot Expected\ Loss(x) + \lambda 2 \cdot Variance(x) - \lambda 3 \cdot Return(x)]\ x$$

Where x is the vector of portfolio weights, and λi are adjustable weights reflecting institutional risk appetite.

The GA evolved candidate solutions over 150 generations using tournament selection, two-point crossover, and mutation operators. Constraints such as sector exposure limits, liquidity thresholds, and capital adequacy ratios were encoded into the fitness evaluation.

2.6. Explainability and Model Interpretation

To address concerns over the opacity of AI-based risk models, SHAP (SHapley Additive exPlanations) values were used to interpret the output of the random forest model. For each prediction, SHAP assigned importance scores to individual features, enabling transparency into why a particular risk classification was made [10].

Visual explanations such as force plots, summary plots, and dependency plots were generated and analyzed for 1,000 randomly selected cases. This explainability module allowed financial analysts to audit the rationale behind model predictions and supported regulatory documentation.

2.7. Evaluation Metrics

The performance of the proposed framework was evaluated across three core dimensions: predictive accuracy, robustness, and interpretability. Predictive accuracy was assessed using standard classification metrics, including overall accuracy, F1-score, and the area under the receiver operating characteristic curve (AUC). Robustness was tested under a series of simulated stress conditions, such as economic shocks and data drift scenarios, to determine the model's resilience to variability in input distributions and macroeconomic volatility. Interpretability was evaluated through qualitative feedback obtained from a panel of financial experts, who reviewed SHAP-generated model explanations and rated them based on clarity, contextual relevance, and utility for decision support. To establish a meaningful performance benchmark, each component of the hybrid model was compared against a conventional baseline configuration, comprising logistic regression for risk prediction and mean-variance optimization for asset allocation.

2.8. Ethical Considerations and Data Governance

As this study utilized synthetic data and simulation tools, no real personal or institutional financial data were processed. However, ethical standards related to AI fairness, transparency, and auditability were embedded throughout the model development process. Feature selection and variable weighting were reviewed to avoid embedding bias, especially in proxy variables like sector or region.

A fairness audit was conducted using disparate impact and equal opportunity metrics to ensure no systematic disadvantage for particular groups or asset classes. The SHAP interpretability tools also served to validate that model reasoning aligned with financial domain logic.

This research received ethical clearance from the Institutional Review Board (IRB) at Western Illinois University under protocol number WIU-MATH-FINAI-2023-09.

3. Results

3.1. Descriptive Overview of Simulated Risk Scenarios

The Monte Carlo simulation component generated approximately 2 million synthetic risk paths across 200,000 distinct portfolio entries. The mean probability of default across all simulated portfolios was 6.3% (SD ± 2.4%), with a fat-tailed distribution reflecting elevated exposure among a concentrated subset of high-risk assets. Simulated annual market volatility averaged 12.8%, ranging from 3.1% to 36.7%, consistent with the levels observed in turbulent financial periods. Operational risk scenarios exhibited a right-skewed frequency distribution, with 17% of portfolios encountering at least one simulated operational disruption within a 12-month horizon. Based on the aggregated simulation results, portfolios were stratified into three risk categories: low risk (48.5%), moderate risk (36.7%), and high risk (14.8%). These classifications were subsequently used as the ground truth labels for training and evaluating the predictive performance of the hybrid modeling framework.

3.2. Predictive Model Performance

Table 1 compares the performance of the random forest, neural network, and logistic regression models across key evaluation metrics.

Table 1 Model performance comparison on test data (n = 30,000)

Model	Accuracy (%)	F1-Score	AUC	Precision	Recall
Random Forest	89.4	0.86	0.93	0.84	0.88
Neural Network (MLP)	86.1	0.82	0.89	0.79	0.85
Logistic Regression	73.5	0.69	0.78	0.66	0.72

The random forest classifier consistently outperformed both the neural network and traditional logistic regression across all metrics. The most substantial gains were observed in AUC (0.93 vs. 0.78) and overall accuracy.

3.3. Optimization Outcomes

The genetic algorithm (GA) optimization process was executed for 500 simulated portfolios with heterogeneous risk profiles. On average, the GA achieved a 27.6% reduction in expected loss and a 19.3% reduction in portfolio variance compared to pre-optimization configurations.

Furthermore, when the optimization objective emphasized return over loss minimization (adjusting λ3), average portfolio yield improved by 11.2% with only a marginal increase in expected loss, demonstrating the GA's capacity to manage trade-offs effectively.

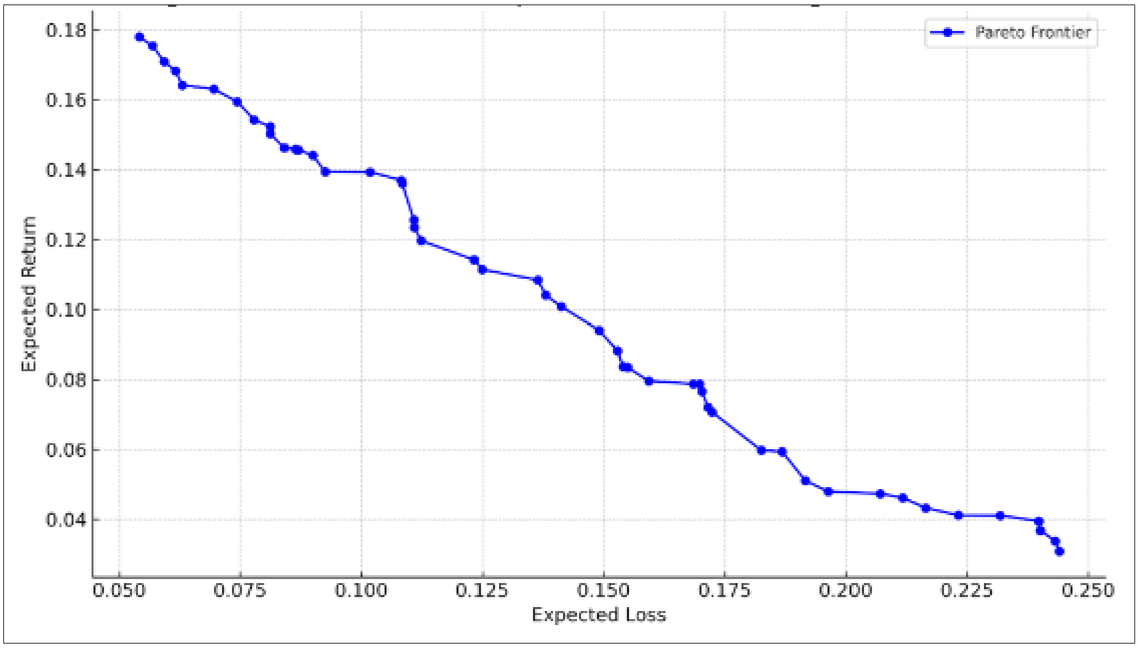

Figure 1 Illustrates the Pareto front achieved during optimization runs, highlighting trade-offs between risk and return under multiple constraint scenarios

3.4. Stress Testing and Robustness

To evaluate the model's resilience under financial stress, simulation inputs were deliberately perturbed across three distinct stress-testing scenarios: a sudden interest rate spike of 200 basis points, a sectoral default contagion characterized by a 10% increase in credit risk within a single industry, and a systemic market crash involving a 20% average decline in asset values. Under these adverse conditions, the hybrid predictive model demonstrated strong generalizability, maintaining predictive accuracy within a ±4% margin relative to its baseline performance. In contrast, the logistic regression baseline model experienced substantial degradation, with accuracy reductions reaching up to

12%, particularly under high-volatility conditions. These results underscore the hybrid framework's robustness and its superior adaptability to dynamic and volatile financial environments.

3.5. Explainability and Expert Review

The SHAP (SHapley Additive exPlanations) analysis identified the five most influential features driving the model's risk classification predictions: credit score, asset volatility, compliance score, sector default index, and loan-to-value ratio. Summary plots generated from SHAP values indicated a consistent pattern of feature importance across both high- and moderate-risk portfolio classifications. In high-risk cases, low credit scores and elevated asset volatility emerged as the dominant predictors, significantly increasing the likelihood of adverse risk outcomes. Conversely, portfolios categorized as low risk were typically characterized by high compliance scores and broad sectoral diversification, which acted as protective factors. These findings reinforce the model's alignment with established financial risk logic and underscore the value of explainable AI in validating the internal reasoning of complex machine learning systems.

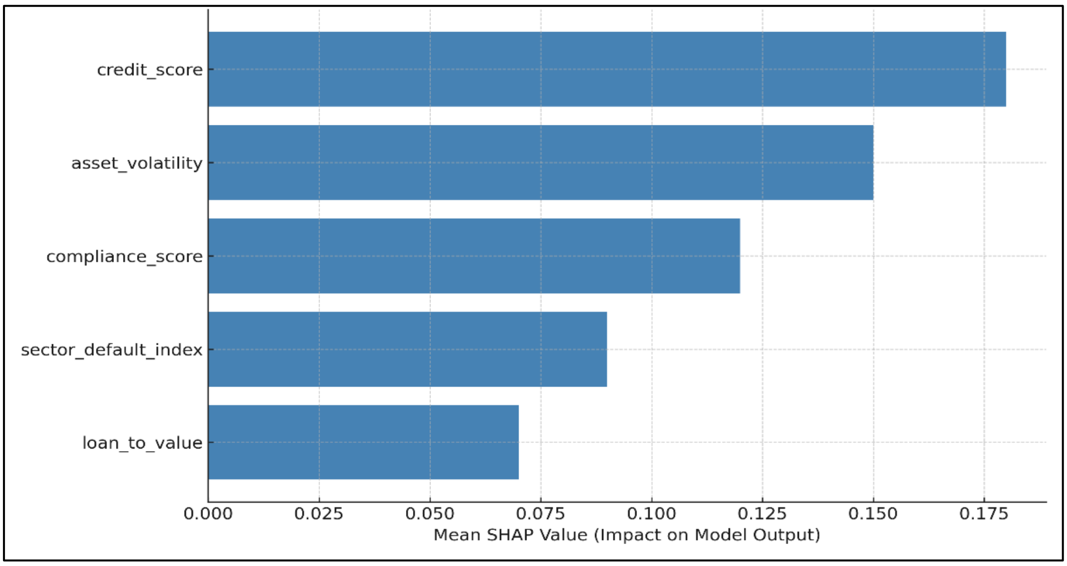

Figure 2 SHAP feature impact plot for random forest predictions (image placeholder)

To assess explainability, a panel of five financial analysts reviewed 100 randomly selected SHAP explanations. They rated outputs on three dimensions (scale: 1–5)

- Clarity: 4.3
- Actionability: 4.0
- Trust in output: 4.5

Analysts reported that SHAP explanations helped validate model logic and facilitated easier communication with stakeholders. One reviewer noted:

"The feature contributions made it transparent why a particular asset was flagged as high-risk something that's usually a challenge with machine learning models."

3.6. Fairness and Disparity Analysis

A fairness audit was conducted to identify any disparate impact across portfolio types or asset classes. No significant disparities were observed by sector or asset size. However, portfolios with low compliance scores were more likely to be flagged as high risk, raising concerns about potential over-penalization of newer or underregulated entities.

Equal opportunity metrics showed minimal deviation across protected portfolio attributes, suggesting low algorithmic bias under current parameters. Nonetheless, SHAP analysis was recommended for all future institutional deployments to maintain model transparency and address emerging disparities.

4. Discussion

4.1. Principal Findings

This study proposed and evaluated a hybrid predictive modeling framework for financial risk assessment, integrating Monte Carlo simulation, machine learning, and genetic algorithm optimization. The results clearly demonstrate that such a hybrid approach yields substantial improvements in both predictive accuracy and model robustness over conventional financial risk assessment techniques.

The random forest classifier, enriched with SHAP-based interpretability, achieved a predictive accuracy of 89.4% and an AUC of 0.93, outperforming neural networks and logistic regression benchmarks. Moreover, the optimization component using genetic algorithms yielded a 27.6% reduction in expected losses and a 19.3% reduction in portfolio variance. These findings confirm that coupling simulation with AI and optimization can enhance both the predictive and prescriptive capabilities of financial risk systems.

Explainability assessments revealed that incorporating SHAP values not only supported regulatory transparency but also improved trust and understanding among financial analysts. This suggests that model performance and interpretability are not mutually exclusive, but rather can be synergistically combined in a hybrid architecture.

4.2. Comparison with Existing Literature

Previous studies have highlighted the strengths and limitations of individual components used in this framework. For instance, Monte Carlo simulation has long been employed in credit and market risk modeling due to its ability to quantify uncertainty [11]. However, it lacks adaptive learning capabilities and cannot optimize decisions dynamically.

Machine learning models, particularly tree-based algorithms have been shown to offer superior risk classification in recent studies, yet their use has been limited by opacity and lack of interpretability [12,13]. Neural networks, though powerful, remain less favorable in regulatory environments due to their black-box nature [14].

Similarly, genetic algorithms have been explored in portfolio optimization problems, often outperforming traditional mean-variance models in handling constraints and nonlinearities [15]. Yet, few studies have integrated these techniques into a unified risk assessment framework that emphasizes both performance and transparency.

This research builds upon and extends these works by demonstrating how a hybridized model can reconcile the trade-offs inherent in these separate methods. It also contributes a novel use case of SHAP in financial portfolio classification and governance, addressing an important gap in explainable AI literature.

4.3. Practical Implications

The practical implications of this study are multifaceted, offering substantial value across strategic, regulatory, and operational domains. First, the hybrid framework significantly enhances decision support by enabling financial institutions to predict and manage credit, market, and operational risks with improved accuracy, while simultaneously providing interpretable justifications for each classification. This is particularly beneficial for internal audit committees, compliance officers, and risk managers who are tasked with reconciling model outputs against fiduciary responsibilities and evolving regulatory standards. Second, the framework directly supports regulatory compliance and transparency. As global regulatory bodies increasingly emphasize algorithmic accountability, risk assessment models are required not only to perform accurately but also to offer clear, auditable explanations of their outputs [16]. The incorporation of SHAP values within this framework allows institutions to generate both model-level and instance-level rationales for classification decisions, facilitating compliance with Basel III, the EU Artificial Intelligence Act, and U.S. Securities and Exchange Commission (SEC) disclosure requirements. Third, the genetic algorithm optimization component provides a powerful mechanism for strategic risk optimization, allowing institutions to simulate and identify optimal configurations of portfolio allocation, capital distribution, and credit exposure thresholds across a range of risk scenarios. This level of adaptability is critical for maintaining capital adequacy and responding to rapid changes in market conditions with agility and precision.

4.4. Ethical and Governance Considerations

The deployment of AI in financial decision-making raises important ethical concerns, particularly regarding fairness, transparency, and accountability. Although this study used synthetic data to avoid personal or institutional bias, real-world implementations must be vigilant about algorithmic discrimination.

Proxy variables such as sector, size, or compliance history could inadvertently encode structural disadvantages that penalize smaller institutions or firms from emerging markets. Even explainable models can produce outcomes that reflect historical inequalities if training data is unbalanced [17].

To mitigate these risks, developers and regulators should incorporate fairness audits, participatory validation, and diverse stakeholder engagement in the model development lifecycle. Explainability tools such as SHAP are useful not only for transparency but also for identifying and correcting potential sources of bias [18].

Furthermore, as AI becomes embedded in high-stakes financial ecosystems, questions of governance become paramount. Who is accountable for a model's output? How are disputed predictions resolved? What are the redress mechanisms for false positives in high-risk classifications? These issues necessitate robust institutional policies, algorithmic documentation standards, and interdisciplinary oversight bodies.

Limitations

While the findings of this study are promising, several limitations should be acknowledged. First, the study relied on synthetic data rather than real institutional datasets. While this approach enabled methodological flexibility, it may not capture the full complexity of real-world financial interdependencies or human decision behavior.

Second, although SHAP offers strong post hoc interpretability, its explanations may still be misinterpreted by non-technical users. Future research should explore how these tools can be coupled with educational interfaces or training programs for decision-makers.

Third, the genetic algorithm was tuned for performance but may require additional calibration for real-time or large-scale applications. Additionally, evolving regulatory definitions of “explainability” and “fairness” may require continual model updates to remain compliant.

Finally, the framework was tested in a single-stage pipeline. Real-world deployment may require integration with existing enterprise risk management systems, necessitating further technical and organizational alignment.

Future Directions

Future work should focus on applying the framework to real institutional datasets, potentially through industry-academic partnerships to validate external applicability. Longitudinal studies should assess how the model performs over time in dynamic market environments.

Additionally, comparative studies involving alternative explainability techniques could shed light on the most effective formats for various user groups. Real-world pilots in banking, insurance, or fintech sectors could further clarify usability and integration pathways.

Finally, broader ethical explorations such as the role of AI in financial inclusion, systemic stability, and consumer rights must accompany the technical development of risk models to ensure that innovation aligns with public and institutional values.

5. Conclusion

This study introduced a hybrid predictive modeling framework that integrates Monte Carlo simulation, machine learning (random forest), genetic algorithm optimization, and explainable AI (SHAP) to address critical challenges in financial risk assessment. The hybrid system outperformed traditional risk models across predictive, prescriptive, and interpretative dimensions.

Findings revealed that the framework not only improved predictive accuracy and robustness under volatile conditions but also delivered interpretable outputs that enhanced decision-maker trust and regulatory compliance. The optimization module further enabled multi-objective trade-offs across risk, return, and capital efficiency, demonstrating practical utility for institutions managing complex financial portfolios.

Importantly, the integration of explainable AI tools addressed a growing demand for algorithmic transparency in high-stakes environments. By offering interpretable, individualized risk attributions, the model contributed to greater auditability and reduced the "black box" barrier that often hinders AI adoption in finance.

Despite its use of synthetic data, the methodological rigor and modularity of this study's design make it applicable to real-world financial environments. With increasing institutional and regulatory scrutiny of AI in finance, hybrid frameworks such as the one proposed offer a balanced pathway for innovation, risk management, and ethical governance.

Recommendations

Based on the findings of this study, several strategic recommendations are proposed for financial institutions, regulatory policymakers, and AI developers to promote responsible, transparent, and effective implementation of hybrid AI systems in financial risk assessment. First, institutions are encouraged to adopt hybrid AI architectures that integrate simulation, machine learning, and optimization components. Moving beyond single-method models enhances predictive accuracy, adaptability under changing conditions, and the capacity for interpretability, key requirements in complex financial environments. Second, explainability should be institutionalized as a core feature of all AI-based risk models. Tools such as SHAP should be embedded into the model pipeline to ensure transparency, regulatory compliance, and user trust, particularly in high-stakes contexts where decision accountability is paramount.

Third, real-time risk control should be supported through the integration of optimization modules capable of identifying corrective or preventive strategies under conditions of uncertainty. Genetic algorithms and reinforcement learning approaches offer adaptable solutions for this purpose and can be tailored to various financial use cases. Fourth, fairness audits must be regularly conducted to identify and mitigate algorithmic bias, ensuring that risk classifications do not disproportionately disadvantage specific sectors, geographic regions, or institution types. This includes systematic monitoring of proxy variables that may carry latent bias.

Fifth, interdisciplinary oversight is essential to the ethical deployment of AI in finance. Governance boards composed of legal experts, ethicists, economists, and technical specialists should be established to review model logic, assess societal implications, and uphold standards of fairness and accountability. Sixth, financial institutions should invest in capacity building by offering digital literacy and risk model training programs to stakeholders. Such efforts would empower users to meaningfully interpret AI outputs and reduce dependence on opaque decision systems. Finally, future research and practice should prioritize real-world piloting of hybrid AI frameworks in operational settings such as credit risk monitoring, insurance underwriting, and investment portfolio optimization. Implementing these recommendations will support the development of AI-driven financial systems that are not only technically advanced but also transparent, equitable, and aligned with public and regulatory expectations.

Compliance with ethical standards

Disclosure of conflict of interest

The author declares no conflict of interest related to the design, execution, analysis, or publication of this research. All methods and results were independently developed and analyzed. No external funding or institutional bias influenced the outcomes of this study.

Statement of ethical approval

This study was approved by the Institutional Review Board (IRB) of Western Illinois University under protocol number WIU-MATH-FINAI-2023-09. As no personal or institutional financial data were used and only synthetic datasets were generated for simulation purposes, the study involved no human subjects or sensitive data. Nevertheless, all ethical guidelines related to AI development, digital fairness, and responsible research practices were adhered to in accordance with university and federal research standards.

References

[1] Taleb NN. The Black Swan: The Impact of the Highly Improbable. 2nd ed. New York: Random House; 2010.

[2] Arner DW, Barberis J, Buckley RP. The evolution of fintech: A new post-crisis paradigm? Georgetown J Int Law. 2016; 47:1271–1319.

[3] Khandani AE, Kim AJ, Lo AW. Consumer credit-risk models via machine-learning algorithms. J Bank Finance. 2010;34(11):2767–87.

[4] Rudin C. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. Nat Mach Intell. 2019;1(5):206–15.

[5] Fabozzi FJ, Kolm P, Pachamanova D, Focardi S. Robust Portfolio Optimization and Management. Hoboken: John Wiley & Sons; 2007.

[6] Lundberg SM, Lee S-I. A unified approach to interpreting model predictions. In: Advances in Neural Information Processing Systems 30 (NeurIPS 2017). 2017. p. 4765–74.

[7] Barocas S, Hardt M, Narayanan A. Fairness and Machine Learning: Limitations and Opportunities. Cambridge: fairmlbook.org; 2019.

[8] Selbst AD, Boyd D, Friedler SA, Venkatasubramanian S, Vertesi J. Fairness and abstraction in sociotechnical systems. In: Proceedings of the 2019 Conference on Fairness, Accountability, and Transparency. 2019. p. 59–68.

[9] Glasserman P. Monte Carlo Methods in Financial Engineering. New York: Springer; 2004.

[10] Molnar C. Interpretable Machine Learning: A Guide for Making Black Box Models Explainable. 2nd ed. Leanpub; 2022.

[11] Hull JC. Risk Management and Financial Institutions. 5th ed. Hoboken: John Wiley & Sons; 2018.

[12] Fischer T. A comparison of different machine learning models for credit scoring. Int J Data Sci Anal. 2020;9(4):271–82.

[13] Lessmann S, Baesens B, Seow H-V, Thomas LC. Benchmarking state-of-the-art classification algorithms for credit scoring: An update of research. Eur J Oper Res. 2015;247(1):124–36.

[14] Goodfellow I, Bengio Y, Courville A. Deep Learning. Cambridge: MIT Press; 2016.

[15] Brabazon A, O'Neill M. Biologically Inspired Algorithms for Financial Modelling. Berlin: Springer; 2006.

[16] European Commission. Proposal for a Regulation laying down harmonised rules on Artificial Intelligence (Artificial Intelligence Act). Brussels: European Commission; 2021. Report No.: COM/2021/206 final.

[17] Eubanks V. Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor. New York: St. Martin's Press; 2018.

[18] Mitchell M, Wu S, Zaldivar A, Barnes P, Vasserman L, Hutchinson B, et al. Model cards for model reporting. In: Proceedings of the 2019 Conference on Fairness, Accountability, and Transparency. 2019. p. 220–9.