Aditi Rajput, 2026, 14:3
ISSN (Online): 2348-4098
ISSN (Print): 2395-4752

International Journal of Science, Engineering and Technology
An Open Access Journal

AI Based Personal Finance Advisor

Aditi Rajput, Avani Malviya, Dhruv Dekhane, Mahi Ishwe, Prof. Manish Vyas
IT Department Acropolis Institute of Technology & Research, Indore.

Abstract- Artificial intelligence is reshaping how individuals access, receive, and act on financial advice. Personal finance— covering budgeting, credit management, investment planning, and retirement preparation—has long been gatekept by cost and expertise. AI-driven systems, from simple calculators to adaptive robo-planners, are changing that. But expanded access does not automatically mean better outcomes. Many platforms reproduce the same incentive misalignments and opacity prob- lems that made traditional human advisors unreliable in the first place, only now at scale. This paper examines the current state of AI in personal finance advice across three angles: (1) a review of machine learning techniques used in credit risk assessment and portfolio optimization, (2) an analysis of trust, personalization, and behavioral risks in deployed platforms, and (3) a proposed five-principle framework—fiduciary duty, adaptive personalization, technical robustness, ethical fairness, and auditability—to evaluate whether these systems genuinely serve users. Experimental results from a supervised ML classifier trained on a Kaggle credit risk dataset show Random Forest achieving 89.25% accuracy and AUC of 0.77. The broader argument is that technological sophistication, while necessary, does not resolve the governance gaps that undermine user trust in AI financial advice.

Keywords: Personal Finance, Robo-Advisors, Credit Risk, Machine Learning, Trust, AI Ethics, Financial Planning, Person- alization

I. INTRODUCTION

Most people manage money without professional help. Not by choice—professional financial advice is expensive, often inaccessible, and frequently entangled with commission struc- tures that do not obviously favor the client. A 2023 Statista survey of 75,000 bank customers across 32 countries found that satisfaction with banking services correlates strongly with how well institutions meet individual needs, not just how reputable they appear [1].

AI promises to fix this. Automated systems can screen loan applicants at scale, recommend portfolios based on risk tolerance, flag suspicious transactions, and answer routine financial questions at any hour. The question worth asking is whether they actually do so in ways that benefit the people using them.

The answer so far is: sometimes. Platforms like Betterment and Wealthfront have genuinely democratized portfolio man- agement by making index-fund diversification accessible to retail investors at low cost [2]. But adjacent platforms—notably Robinhood—have demonstrated that design choices intended to widen access can simultaneously encourage speculative behavior among inexperienced users. Meanwhile, systems built for credit risk assessment routinely inherit biases from historical training data, producing disparate outcomes across demographic groups without any explicit intent to discriminate [3].

This paper does not argue that AI in personal finance is net harmful. It argues that the field has not yet solved the core problem: how to build systems that give individuals genuinely useful, unbiased, and accountable financial guidance rather than efficient-looking machinery that serves platform economics.

We address this through three contributions:

- A review of ML techniques applied to personal finance tasks—credit scoring, portfolio optimization, anomaly detection, and behavioral analysis.
- An experimental evaluation of four supervised classifiers on a credit risk dataset, establishing a performance baseline and discussing what accuracy metrics do and do not tell us about real-world suitability.
- A normative framework of five principles for responsible AI financial advice, grounded in economic theory and AI governance literature.

The remainder of the paper is structured as follows. Section II surveys the background. Section III describes the experi- mental methodology. Section IV presents results. Section V develops the five-principle framework. Section VI discusses limitations and future directions. Section VII concludes.

II. BACKGROUND AND RELATED WORK

A. Financial Advice and Its Structural Problems

Human financial advisors exist because financial markets are not transparent or frictionless. Borrowers and lenders do not share the same information. Contracts are costly to write and enforce. Retail clients face complex product decisions they lack the expertise to evaluate [2]. Advisors, in theory, bridge these gaps.

In practice, they often replicate them. Commission-driven compensation structures incentivize advisors to recommend products that pay out rather than products that fit [9]. Advisors also imprint their own preferences onto client portfolios: research on Canadian households found that advisor charac- teristics explained more variation in client asset allocation than the clients' own stated risk tolerances [8]. The result is a service that is both expensive—advised portfolios carry average annual costs roughly 1.5 percentage points higher than comparable lifecycle funds—and often poorly personalized.

B. Machine Learning in Personal Finance

AI has been applied to a range of personal finance prob- lems. Credit risk assessment is the most studied. Shi et al. [4] reviewed ML-driven credit risk models across multi- ple datasets and found that ensemble methods—particularly gradient-boosted trees and random forests—consistently out- perform both traditional statistical models and most deep learning approaches on standard accuracy metrics.

Portfolio optimization using genetic algorithms and rein- forcement learning has demonstrated potential for generating risk-adjusted returns that adapt dynamically to market condi- tions [5]. Recommendation systems modeled on collaborative filtering techniques have been applied to investment product selection, though their adoption by retail platforms remains limited [6].

Fraud detection is an area where AI has seen unambiguous operational benefit. Systems that screen transaction data using anomaly detection and graph-based pattern recognition can identify suspicious activity at a speed and scale no human compliance team could match [7].

The common challenge across all of these applications is the same: ML models perform well on the data they were trained on, and personal finance data is historically unequal, geographically specific, and subject to distributional shifts during economic stress [2].

C. Trust and Personalization in Digital Finance

Trust in digital financial services is not simply a matter of platform security. It involves users believing that the system is designed for their benefit rather than for the provider's. Kanaparthi [1] identifies five research gaps in this area: explainability, trustworthiness, data privacy, ethical alignment, and credit risk detection. These are not independent. Opacity in credit scoring directly reduces perceived trustworthiness, which in turn reduces willingness to engage with personalized recommendations.

Personalization itself is a layered problem. Early robo- advisors used short questionnaires to classify users into broad risk buckets (conservative, moderate, aggressive) and held them there regardless of changing circumstances. During the 2018 market downturn, users classified as aggressive and left in high-equity allocations suffered losses that arguably no responsible human advisor would have permitted [2].

More recent platforms integrate behavioral monitoring— spending patterns, income volatility, engagement signals— to update recommendations dynamically. These systems can, in principle, produce advice that is genuinely responsive to individual circumstances. The risk is that the same behavioral data used for personalization can also be used for engage- ment optimization, which is not the same thing as financial optimization.

D. Regulatory Landscape

Regulatory responses to AI in financial services vary widely across jurisdictions. The EU AI Act classifies credit scoring and risk assessment as high-risk applications requiring trans- parency, auditability, and human oversight provisions. The US SEC has sanctioned platforms—Robinhood most visibly—for misleading disclosures around payment for order flow and for approving users for complex products without adequate suitability checks [10]. In Australia, ASIC sued eToro in 2023 for offering contracts-for-difference products to users whose risk profiles were demonstrably unsuitable [2]. These cases confirm that the governance gap is real and consequential.

III. METHODOLOGY

A. Research Framework

This paper combines a Systematic Literature Review (SLR) following the PRISMA protocol [11] with an experimental evaluation of supervised ML classifiers on a credit risk dataset. The SLR informed the identification of key research gaps and the design of the five-principle framework. The experiment provides a concrete performance benchmark and illustrates the gap between classification accuracy and real-world advisory adequacy.

B. Dataset

The dataset is the Credit Risk Dataset from Kaggle [12], which contains 32,581 loan records with 11 features and a binary target variable (default/no default). Features include borrower age, annual income, home ownership status, em- ployment length, loan grade (A–G), loan intent, loan amount, interest rate, loan-to-income ratio, prior default history, and credit history length.

Feature selection used correlation analysis against the target variable. Features with $|r| < 0.1$ were excluded. This elimi- nated age ($r = -0.022$), employment length ($r = -0.087$), loan intent ($r = -0.059$), and credit history length ($r = -0.016$). Annual income was subsequently removed based on exploratory data analysis: it showed near-zero variance, severely unbalanced class distributions, and inconsistent distri- butional behavior compared to other features. The final feature set comprised six variables: home ownership, loan grade, loan amount, interest rate, loan-to-income ratio, and prior default history.

C. Classifiers

Four supervised classification algorithms were evaluated:

Support Vector Machine (SVM): Locates a maximum-margin hyperplane in feature space. Effective on high- dimensional data with clear class separation, but computation- ally intensive and sensitive to kernel choice [13].

Random Forest: An ensemble of decision trees trained on bootstrapped subsamples with random feature subsets. Resis- tant to overfitting and produces feature importance estimates, which matter for explainability in credit contexts [14].

Decision Tree: A single flowchart-structured classifier using recursive attribute splits. Transparent and interpretable, but prone to overfitting without pruning [15].

Logistic Regression: A linear probabilistic classifier that estimates default probability as a logistic function of weighted features. Interpretable, well-calibrated, and straightforward to audit [16]. All classifiers used a 75/25 train-test split with no temporal ordering, since the dataset does not include origination dates.

D. Evaluation Metrics

$$\text{Accuracy}= \frac{TP+TN}{TP+TN+FP+FN} \quad (1)$$

$$\text{Precision} = \frac{TP}{TP+FP} \quad (2)$$

$$\text{Recall} = \frac{TP}{TP+FN} \quad (3)$$

$$\text{Specificity} = \frac{TN}{TN+FP} \quad (4)$$

$$\text{F1}=\frac{2\cdot\text{Precision}\cdot\text{Recall}}{\text{Precision}+\text{Recall}} \quad (5)$$

Performance was measured using accuracy, precision, recall, specificity, F1-score, and area under the ROC curve (AUC). Let TP, TN, FP, and FN denote true positives, true negatives, false positives, and false negatives respectively.

The AUC measures the probability that the classifier ranksa randomly chosen positive instance above a randomly chosen negative one. Values above 0.70 are generally considered acceptable for credit risk screening.

IV. EXPERIMENTAL RESULTS

A. Classifier Performance

Table I shows the performance metrics across all four classifiers. Random Forest achieves the best results across every metric except AUC, where it ties with SVM. Deci- sion Tree performs comparably to SVM on most measures, which is consistent with the relatively simple feature space after selection. Logistic Regression underperforms but remains interpretable and useful as a baseline.

TABLE I
CLASSIFIER PERFORMANCE METRICS

Classifier	Acc.	Prec.	Rec.	Spec.	F1	AUC
SVM	0.885	0.880	0.885	0.885	0.878	0.77
Random Forest	**0.893**	**0.889**	**0.893**	**0.893**	**0.890**	0.77
Decision Tree	0.879	0.874	0.879	0.879	0.871	0.70
Logistic Regr.	0.840	0.828	0.840	0.840	0.826	0.71

B. Confusion Matrix Analysis

Table II shows the confusion matrices for each classifier. The most practically relevant figures here are the false negative and false positive rates. In credit risk assessment, a false negative (actual default classified as no-risk) carries higher institutional cost than a false positive (non-defaulter denied credit). However, from an equity standpoint, false positives are the higher-stakes error: they deny access to credit for individuals who would have repaid.

TABLE II
CONFUSION MATRICES (TEST SET, n = 8,146)

Classifier	TN	FP	FN	TP
SVM	6252	704	230	960
RandomForest	6144	537	338	1127
DecisionTree	6132	751	237	1026
LogisticRegr.	6040	976	329	801

Random Forest reduces false positives substantially com- pared to SVM (537 vs. 704), at the cost of slightly more false negatives (338 vs. 230). For a deployment prioritizing financial inclusion, this trade-off may be worth accepting.

Table III compares this model's accuracy against published results on related credit datasets. The present model matches or exceeds the Lending Club denoising autoencoder result (0.875) and outperforms several deep learning approaches on smaller or domain-specific datasets. The AUC of 0.77 trails only the Lending Club autoencoder model when AUC is reported.

TABLE III
ACCURACY COMPARISON ACROSS CREDIT RISK STUDIES

Dataset	Model	Acc.	AUC
Balanced FICO	Two-Layer	0.740	–
	Additive		
Italian Bank (76 firms)	Feedforward Net	0.870	–
NYU Salomon Center	Boosting	0.863	–
Bondora (European p2p)	Neural Network	0.744	–
Singaporean Credit Firm	Neural Network	0.840	–
Chinese p2p Platform	AM-LSTM	–	0.669
Lending Club	Denoising AE	0.875	–
Chinese Dataset	Relief-CNN	–	0.699
Present Work	**Rnd. Forest**	**0.893**	**0.77**

D. What These Results Do and Do Not Show

An accuracy of 89.25% on a held-out test set is a reasonable result. It does not mean the model is ready for deployment in a real lending or financial planning context. The dataset contains no temporal ordering, so there is no test of generalization across market cycles. Feature importance from the Random Forest model (not shown for space) weights loan grade and loan-to-income ratio most heavily—both of which may cor- relate with demographic variables not present in the feature set.

V. A FIVE-PRINCIPLE FRAMEWORK FOR RESPONSIBLE AI FINANCIAL ADVICE

ML performance metrics alone cannot determine whether an AI system gives good financial advice. This section proposes five principles that, taken together, describe the conditions un- der which AI-driven personal finance systems can be trusted. The framework draws from two traditions: financial services fiduciary law, which imposes duties of care and loyalty on ad- visors, and AI ethics research, which addresses explainability, robustness, and bias. Neither tradition alone is sufficient. A system can be technically excellent—accurate, fast, reliable— and still systematically harm users through opaque design or misaligned incentives.

A. Principle 1: Fiduciary Duty

The core question for any financial advice system is: whose interests does it serve? Human advisors answer this through fiduciary law and reputational incentives, imperfectly. AI sys- tems must answer it through design.

Fiduciary alignment requires three specific things. First, the system must form a genuine model of the individual user's goals—not just their stated risk tolerance in a five- question onboarding survey, but their actual financial situa- tion, constraints, and priorities, updated over time. Second, incentive compatibility: the system's optimization objective must reflect client welfare, not engagement or transaction volume. Payment-for-order-flow models, proprietary product prioritization, and upselling logic all create conflicts that compromise fiduciary duty. Third, equitable treatment: users with smaller balances or lower financial literacy should receive the same analytical rigor as wealthier users.

B. Principle 2: Adaptive Personalization

Financial circumstances change. Job loss, health expenses, family expansion, and market volatility can all materially alter what constitutes appropriate financial advice. A sys- tem that classifies someone as "aggressive" at onboarding and holds them in that category for years is not providing personalization—it is providing an initial segmentation.

Genuine adaptive personalization requires continuous moni- toring of behavioral and contextual signals, with a mechanism to distinguish genuine regime changes (a client's income drops by 40%) from noise (a client's spending spikes during the holidays). Overly reactive systems churn portfolios unneces- sarily; overly static ones ignore material changes. The balance requires both algorithmic guardrails and explainable output so that clients can understand why their recommendations changed [2].

The risks of anthropomorphic design in this context de- serve mention. Chatbots and virtual assistants designed to feel personable can create misplaced trust—users may follow advice from a friendly interface they would question from a formal one. Conversational design should enhance clarity, not substitute for it.

C. Principle 3: Technical Robustness and Resilience

AI financial systems fail in ways human advisors do not. A human advisor who encounters an unusual situation can pause, reflect, and ask for help. An ML model extrapolates from its training distribution—and during market stress, the distribution of inputs often departs significantly from what the model was trained on.

Robustness here means three things: consistency (similar users get similar advice; advice does not shift arbitrarily between sessions), resilience (the system degrades gracefully under anomalous inputs or infrastructure stress rather than silently misfiring), and accuracy under distribution shift (the model's

predictions remain calibrated as market conditions evolve). Each requires specific engineering: version-controlled models, continuous validation pipelines, adversarial stress test- ing, and fallback procedures.
The 2020–2022 period illustrated these risks concretely. Robo-advisors trained on historical equity-bond correlations performed poorly when correlations broke down during COVID-era volatility and the subsequent rate cycle. Systems that lacked stress-testing infrastructure gave users no warning that their models were operating outside their reliable range.

D. Principle 4: Ethical Fairness

Machine learning amplifies bias. A model trained on his- torical lending data learns historical lending patterns—patterns that reflect decades of discriminatory practice. Unless fairness is explicitly engineered into the training and evaluation pro- cess, the model will reproduce that discrimination at scale.

In personal finance, the fairness problem has three compo- nents. First, privacy: financial data is sensitive, and users need to know how it is collected, stored, and used. Second, bias detection: models should be routinely audited for disparate impact across demographic groups. Accuracy metrics averaged across the full test set can mask systematic underperformance on minority subgroups. Third, inclusive design: advisory qual- ity should not degrade for users with lower balances, lower financial literacy, or access only through mobile interfaces.

Research on credit markets has confirmed that ML-based lending models can produce racially disparate outcomes even when protected attributes are excluded from the feature set, because correlated proxies—zip code, loan-to-income ratio, home ownership status—carry the same discriminatory infor- mation [3].

E. Principle 5: Auditability and Accountability

When AI-driven financial advice produces bad outcomes, who is responsible? In practice, accountability is often diffuse: data providers, model developers, deploying institutions, and regulators each share some responsibility but none fully owns it. This diffusion benefits no one except those who would prefer not to be held accountable.

Auditability is the technical precondition for accountability. Systems must log the inputs and model versions that produced each recommendation. Recommendations must be explainable in terms clients and regulators can understand—not necessarily showing model weights, but at minimum explaining which factors drove the output and in which direction. Compliance logic should be embedded in the system architecture rather than bolted on afterward.

German regulators' 2021 shutdown of an AI investment platform that provided unsuitable tax optimization advice without adequate disclosures illustrates what happens when accountability structures are absent [2]. The practical cost to clients—millions of euros returned after the fact—confirms that accountability is not optional.

Table IV summarizes how the five principles apply differ- ently to human and AI advisors, drawing on the comparison developed by Feng et al. [2].

TABLE IV
COMPARISON OF HUMAN AND AI ADVISORS
UNDER FIVE PRINCIPLES

Principle	HumanAdvisor	AISystem
FiduciaryDuty	Professionalnormspro- vide some protection,but commission struc-tures frequently over-ride them	Can be programmed toavoidcommissioncon -flicts,buthiddenincen- tives and engagement-optimizedobjectivesin- troducenewdistortions
Personalization	Periodicreviews;lim- ited by memory andavailability	Continuous monitoringpossible, but risksnoise-chasing and opaque adjustments
Robustness	Inconsistent;subjectto fatigue and cognitivebias	Consistent in normalconditions,brittl eunderdistribution shift
Fairness	Susceptible to individ- ualbias,constrainedbypr ofessional norms	Can reproduce struc- tural bias at scale with- out any discriminatoryintent
Accountability	Traceable;professional liabilityisclear	Oftenopaque;responsi- bility is diffuse acrossdevelopers, deployers,and regulators

VI. DISCUSSION AND LIMITATIONS

A. What the Experiment Shows About Real-World Deployment

The Random Forest classifier achieved 89.25% accuracy with a 0.77 AUC on held-out test data. These are competitive results by the standards of the literature. But several things matter for deployment that this experiment does not test.

First, the dataset has no temporal structure, so we cannot test generalization across economic cycles. Credit risk models notoriously degrade during recessions—precisely when good risk assessment matters most.

Second, we did not audit the model for demographic dispar- ities. Home ownership, loan grade, and loan-to-income ratio are plausible proxies for race and socioeconomic status in the US context. A model that achieves 89% accuracy on the full test set may achieve substantially lower accuracy—or higher false positive rates—on demographic subgroups.

Third, accuracy is not the right metric for evaluating whether the model gives users good financial advice. Credit scoring is one input into a lending decision; it is not itself financial advice. Translating a risk score into a recommendation for an individual borrower requires additional context—income stability, family obligations, alternative financing options—that no classifier captures.

B. The Explainability Gap

Random Forest and SVM are harder to explain than Logistic Regression or Decision Tree. This matters because credit decisions in many jurisdictions require an explanation the borrower can understand and contest. A model that achieves 2% better accuracy at the cost of explainability may not be the right choice in a regulated context.

Feature importance measures from Random Forest (SHAP values or mean decrease in impurity) can partially address this. But feature importance is not the same as a causal explanation, and users who receive adverse credit decisions deserve something closer to causation than correlation.

C. Limitations of the Framework

The five-principle framework is normative—it describes what responsible AI financial advice should look like, not what any existing system achieves. Level 4 and Level 5 systems in Feng et al.'s roadmap [2] remain largely aspirational. Most deployed platforms sit at Level 2 (chatbot) or Level 3 (basic robo-advisor), where the principles provide a checklist for minimum compliance rather than a description of current practice.

Operationalizing fairness diagnostics at scale remains tech- nically challenging. Defining what counts as a "similar" client for the purposes of consistency audits requires choices that are not value-neutral. And the governance structures needed for real accountability—independent audits, public model version logs, clear liability allocation—are not yet standard industry practice in most markets.

VII. FUTURE WORK

Several directions extend this work. On the technical side:temporal train-test splits to test generalization across market cycles; fairnessauditsusing demographicproxies; and SHAPbased explanations for the Random Forest model to assess whether the most predictive features are also the most equitable. On the advisory system side: evaluation frameworks that measureoutcomes—actual default rates, client financial health over time—rather than just classification accuracy. And governance research: how to design liability structures and audit
Requirements that createreal accountability with out stifling the
Genuine benefits AIcande liverin expanding financial access.

VIII. CONCLUSION

AI is doing useful things in personal finance. Credit risk models trained on behavioral and financial data can identify like lydefaults with reasonableaccuracy. Portfoliooptimization tools have made low-cost diversification accessible to millions of retail investors. Conversational interfaces have reduced barriers to basic financial information.

The problem is not that these systems exist. The problem is that the field has not yet established robust standards for when they can be trusted—by the individuals who use them, by the regulators who oversee them, or by the researchers who evaluate them. High accuracy on a held-out test set is not trust. Broad deployment is not validation. A friendly chatbot interface is not fiduciary alignment.

The five-principle framework proposed here—fiduciary duty, adaptive personalization, technical robustness, ethical fairness, and auditability—offers a structured way to evaluate whether a given system actually serves users or merely ap- pears to. These principles are not independent. Personalization without fairness audits can reproduce inequality. Robustness without accountability makes it harder to catch and correct failures. Fiduciary intent without incentive-compatible archi- tecture collapses the moment the business model changes.

The experimental results in this paper confirm that ML- based credit risk classification can achieve competitive per- formance on standard benchmarks. They also illustrate, con- cretely, the limits of what benchmark performance tells us. A 89.25% accurate model that produces racially disparate false positive rates, cannot explain its decisions to a borrower, and degrades during recessions is not good financial advice. Get- ting from good benchmark performance to genuinely trustwor- thy AI in personal finance requires engineering, governance, and sustained regulatory attention—all three, and probably in that order.

Acknowledgment
The authors thank the research teams whose systematic reviews and empirical studies formed the foundation for this analysis, and the open data community for making the credit risk dataset publicly available.

REFERENCES

1. V. Kanaparthi, "AI-based Personalization and Trust in Digital Finance," IEEE International Conference on Digital Finance and AI, 2024.
2. R. Feng, H. Li, and M. Liu, "Robo-Advisors Beyond Automation: Prin- ciples and Roadmap for AI-Driven Financial Planning," arXiv preprint arXiv:2509.09922, 2025.
3. R. Bartlett, A. Morse, R. Stanton, and N. Wallace, "Consumer-Lending Discrimination in the Fintech Era," Journal of Financial Economics, vol. 143, no. 1, pp. 30–56, 2022.
4. S. Shi, R. Tse, W. Luo, et al., "Machine Learning-Driven Credit Risk: A Systemic Review," Neural Computing and Applications, vol. 34, pp. 14327–14339, 2022.
5. R. Meier and R. Danzinger, "Personalized Portfolio Optimization Using Genetic (AI) Algorithms," in Big Data and AI in Digital Finance, J. Soldatos and D. Kyriazis, Eds. Springer, Cham, 2022.
6. R. McCreadie et al., "Next-Generation Personalized Investment Recom- mendations," in Big Data and AI in Digital Finance, J. Soldatos and D. Kyriazis, Eds. Springer, Cham, 2022.
7. F. Koprivec et al., "Screening Tool for Anti-Money Laundering Supervi- sion," in Big Data and AI in Digital Finance, J. Soldatos and D. Kyriazis, Eds. Springer, Cham, 2022.
8. S. Foerster, J. T. Linnainmaa, B. T. Melzer, and A. Previtero, "Retail Financial Advice: Does One Size Fit All?" Journal of Finance, vol. 72, no. 4, pp. 1441–1482, 2017.
9. B. P. Edwards, "Conflicts and Capital Allocation," Ohio State Law Journal, vol. 78, p. 181, 2017.
10. U.S. Securities and Exchange Commission, "Robinhood Financial LLC—Order Instituting Administrative and Cease-and-Desist Proceed- ings, Release No. 33-10906," 2020.
11. D. Moher et al., "Preferred Reporting Items for Systematic Review and Meta-Analysis Protocols (PRISMA-P) 2015 Statement," Systematic Reviews, vol. 4, pp. 1–9, 2015.
12. L. Laotse, "Credit Risk Dataset," Kaggle, 2021. [Online]. Available: https://www.kaggle.com/datasets/laotse/credit-risk-dataset
13. M. E. Mavroforakis and S. Theodoridis, "Support Vector Machine (SVM) Classification through

Geometry," in Proc. 13th European Signal Processing Conf., 2005, pp. 1–4.

14. Y. Xiao, W. Huang, and J. Wang, "A Random Forest Classification Algorithm Based on Dichotomy Rule Fusion," in Proc. IEEE ICEIEC, Beijing, 2020, pp. 182–185.
15. A. Navada, A. N. Ansari, S. Patil, and B. A. Sonkamble, "Overview of Use of Decision Tree Algorithms in Machine Learning," in Proc. IEEE Control and System Graduate Research Colloquium, 2011, pp. 37–42.
16. A. Mahajan, V. S. Baghel, and R. Jayaraman, "Credit Card Fraud Detection using Logistic Regression with Imbalanced Dataset," in Proc. 10th Int. Conf. on Computing for Sustainable Global Development (INDIACom), New Delhi, 2023, pp. 339–342.
17. T. A. Bach et al., "A Systematic Literature Review of User Trust in AI-Enabled Systems: An HCI Perspective," International Journal of Human-Computer Interaction, 2022.
18. A. B. Arrieta et al., "Explainable Artificial Intelligence (XAI): Concepts, Taxonomies, Opportunities and Challenges Toward Responsible AI," Information Fusion, vol. 58, pp. 82–115, 2020.
19. A. Fuster, P. Goldsmith-Pinkham, T. Ramadorai, and A. Walther, "Pre- dictably Unequal? The Effects of Machine Learning on Credit Markets," Journal of Finance, vol. 77, no. 1, pp. 5–47, 2022.
20. H. Zhu, O. Vigren, and I.-L. So¨derberg, "Implementing Artificial Intel- ligence Empowered Financial Advisory Services: A Literature Review and Critical Research Agenda," Journal of Business Research, vol. 174, p. 114494, 2024.