

Agentic Artificial Intelligence For Automated Banking Risk Assessment And Decision Support: A Human-In-The-Loop Framework For Credit, Compliance, And Regulatory Governance

Jayadeep Pakala

Independent Researcher, USA.

Mail Id: jaydeep.pakala@gmail.com

Article Received 07-07-2026, Revised 21-08-2026, Accepted 06-09-2026

Author(s) Retains the Copyrights of This Article

Abstract

In this paper, a human-in-the-loop Agentic AI model of banking risk assessment, compliance and decision support is developed and evaluated. They were evaluated using 8,000 synthetic records, logistic regression, random forests, XGBoost and Isolation Forest. XGBoost had an accuracy of 61.25% and a ROC-AUC of 0.631 which has a weak autonomous reliability. Accuracy is raised to 85.83% and F1-score to 87.94% in case of human supervision, whereas, 97.17% escalation coverage signified operational dependency. The design synthesizes out the explainability, fairness, auditability, cybersecurity and regulatory controls. Findings indicate that real-time checks by supervised agents can become more effective in making responsible choices, whereas calibrated and continuous checks and balances are indispensable.

Keywords: *Agentic AI, Banking Risk, Human-In-The-Loop, Xgboost, Explainability, Compliance, Governance, Fairness, Decision Support, Auditability.*

I. Introduction

Background and Context

Artificial intelligence has evolved beyond banking systems that are controlled by rules to machine-learning frameworks and intelligent agents that understand data and organize tasks and prescriptions. Banks are using the technologies more to score credit, detect fraud, monitor compliance and manage operational risks [1]. This development facilitates quicker, ongoing and more adaptive decision-making in sophisticated financial conditions throughout the globe at present operations.

Research Problem

The traditional banking risk assessment of the risk is usually based on manual reviews, fixed rules and disintegrated information systems. These requirements stall down decision making, limit real time monitoring and give unequal evaluations [2]. Human judgement can create cognitive or demographic bias, and the dynamic nature of financial threats and regulation causes the human controls which are currently in place to be hard to change, audit and make consistent across institutions and jurisdictions.

Research Aim and Objectives

Research Aim

This study aims to create a multi-faceted human-in-the-loop Agentic AI architecture to support automated banking risk assessment, decision making, compliance

with regulations, transparency, accountability and responsible governance.

Research Objectives

- To analyze the Agentic AI use in banking risk, credit, fraud and compliance.
- To evaluate the advantages as well as risks of autonomous banking decisions.
- To explore human-in-the-loop controls to help transparency, accountability and responsible AI.
- To create a unified Agentic AI system that would comply with banking rules and requirements in decision support.

Significance of the Study

This research has academic value in that it builds upon the understanding of data on Agentic AI governance in banking. At the organizational level, it assists in quicker and evidence-based risk choices and effective compliance. Regulators get a guideline on regulation, transparency and responsibility [3]. It encourages explainable and safe, human-controlled agents that are technologically feasible. Socially, it promotes fair credit practices, enhanced consumer protection and increased trust by the people towards financial automation.

Scope and Limitations

The study discusses credit rating, fraud identification, operational risk and regulatory adherence in both commercial and digital banking. It explores

responsible-AI principles of autonomous and human-supervised AI agents and data-protection and financial governance [4]. It is based on secondary academic, regulatory and industry data, which is mainly related to the banking markets [5]. Generalizability is restricted by proprietary constraints of data, cross country differences and fast changing technology.

Novel Contribution

The research suggests a merged agentic AI system and links between autonomous credit, fraud, risk and compliance agents and human approval, explainability and regulatory measures [6]. It is innovative in its fusion of operational independence with risk-thus-escalation, ongoing interpretation of regulations and audit histories of decision making [7]. This offers a governance-based approach to accountable, open and transparent banking choices in intricate settings.

II. Related Work

Agentic AI Applications in Banking Risk, Credit, Fraud and Compliance




Fig.1. Application of Agentic AI in Banking Sector
Literature introduces Agentic AI as a next stage of predictive analytics as agents are able to interpret information, plan actions and use banking systems [8]. Within credit assessment, borrower risk can be determined by using transactional data to speed up the underwriting process and detect evolving borrower risk [9]. They are also able to respond to anomalies and coordinate in fraud detection [10]. The compliance applications can aid in monitoring transactions, customer due diligence, and reporting [11]. The autonomy, however, can enhance biased information sources in history, provide inconsistent results, or unequivocally act on incomplete data [12]. Therefore, success is reliant on information quality, control, and well-defined decision-making power [13].

Benefits, Risks and Limitations of Autonomous Banking Decisions




Fig.2. Advantages and disadvantages of AI in banking

A study reveals efficiency and scalability, consistency, and quicker decision making to be positive sides of banking systems [14]. Against the context of enhancing responses to credit and fraud risks, agentic AI is able to handle information continuously and enhance response to these risks [15]. However, freedom comes with risks. When agents amalgamate models, tools and sequential actions, decisions might be hard to explain [16]. Mistakes have the potential to spread through intertwined processes before they are corrected [17]. The problem of cybersecurity, model drift, privacy, and automation bias, when employees rely too much on machine recommendations, are also mentioned in literature [18]. These constraints suggest that deployment must be continued to ensure validation, contextualized judgement, and protective measures that can curb malicious or otherwise unwarranted automated results in practice in a variety of banking settings [19]. Thus, accuracy and resilience are to be considered as the determinants of autonomy, and explainability and the outcomes of making wrong choices.

Human-in-the-Loop Controls for Transparent, Accountable and Responsible AI

Technical Approaches Framework for Human-in-the-Loop AI





Fig.3. Human-in-the-Loop AI Framework

Human-in-the-loop governance is considered significant in high impact banking decisions [20]. Recommendations can be challenged, action approved that may be sensitive, exceptions investigated and correction made by human oversight [21]. This enhances accountability as responsibility is able to be recognized rather than being generalized to the self-sustaining systems [22]. Nevertheless, the

involvement of humanity is not necessarily effective. Too many alerts will cause review fatigue and bad interfaces will promote rubber-stamping [23]. Sophisticated control should have intervention thresholds and escalation procedures, audit trail, role-based authority and competent reviewers [24]. It has been recommended in the literature that it be overseen: all routine low risk operations can be automated, but consequential credit, fraud, and compliance choices must have greater human involvement.

Integrated Agentic AI Frameworks for Banking Regulation and Decision Support



Note: Synthetic fsQCA is interpreted as theory-development and forecasting, not as evidence of actual banking-sector prevalence.

Fig.4. Agentic AI implementation in banking sector

Responsible AI is linked to principles in terms of fairness, transparency, accountability, security, privacy, as well as human oversight, as presented in recent scholarship [25]. There should be an integrated system linking the design of the agent, the model governance, data controls, human review and within the decision lifecycle [26]. Aligning with the regulations is a challenge as one in one jurisdiction may require one thing whilst in another jurisdiction may require another thing. The framework that is purely on technical performance is not sufficient [27]. Nonetheless, it is politically endorsed in literature that autonomy should be proportional to risk, that decisions should be traceable, that controls should be constantly observed, and that humans should be able to intervene or override agents [28]. This integration can make Agentic AI an ability to make decisions in a more controlled fashion and not an efficiency tool.

Research Gap

The existing literature illustrates that banking risk, credit, fraud, and compliance can be effectively handled by the mechanisms of Agentic AI, but few studies combine autonomous decision-making with the functioning of effective human supervision and regulatory needs. There are still serious gaps in regard to explainability and accountability, intervention mechanisms, and proportional autonomy [29]. Hence, more studies are necessary to come up with an intertwined, regulation-consistent system of responsible Agentic AI implementation.

III. Research Methodology

Research Philosophy

The study embraces pragmatism since it considers how well the performance of Agentic AI and governance are measurable and suitable. Pragmatism advocates integration of computational experimentation with regulatory interpretation, so the framework can be assessed based upon technical correctness, human control and explainability and practical applicability of banking value as opposed to a single epistemological standpoint.

Research Approach

Using a deductive method is applied in this study that tests propositions based on the banking risk literature, literature on responsible AI and human-in-the-loop governance. Predicted links amongst agent autonomy, accurateness of decisions, compliance and human intervention are turned into measurable variables [30]. Results subsequently verify, perfect or contravene these expectations as it imparts information to the proposed integrated framework.

Research Design

The study has a quantitative experimental and design-science design. A simulated banking environment is an assessment of autonomous credit, fraud and compliance agents with the contingency of the risk condition. Design science assists in construction of frameworks, demonstration and evaluation and controlled experiments comparing agent-only, human-only decisions and human-supervised decisions with the same synthetic cases.

Data Sources

The data used in the study is the synthetic banking information which is a reflection of applicants, credit histories and transactions, fraud indicators and compliance status. Variable distributions, and relationships, are estimated using peer-reviewed literature and regulatory reports and public banking statistics. No actual customer history is taken, thus ensuring the reproducibility whilst minimizing privacy risks.

Sampling Strategy

The purposive scenario sampling technique picks representative low, medium and high-risk cases of credit, fraud, operational and compliance decisions. Stratified generation ensures a sufficient representation of the classes, and diversity of demographics. A statistical power analysis and sensitivity analysis take place to determine the final sample size and then model evaluation and comparative testing occur.

Data-Collection Procedure

Synthetic records are created based on existing distributions, dependency and controlled risk parameter distributions. Missing values, unreasonable

combinations, duplication and class imbalance are screened off and data is divided into the training and validation and test sets. Procedural reproducibility in all of experimentation is maintained using versioned generation scripts and fixed random seeds.

Data-Analysis Methods

Descriptive statistics, confusion matrices, and subgroup comparisons are used to analyze the agent-only, human-only and human-in-the-loop results. The synthesis of the framework links empirical outcomes and regulatory controls. The improvement of the decision is said to be:

$$\Delta P = P_{HITL} - P_{agent}$$

Here P represents performance and positive values mean better-supervised decision quality with similar experimentation conditions.

Machine-Learning Model Selection

Logistic regression, Random Forest and XGBoost can be used to score credit and compliance risk, and Isolation Forest can be used to monitor transactions that are anomalous to monitor fraud. Their productions provide scores to the agents that produce recommendations and cause human inspection due to breached confidence or regulatory limits.

Model-Performance Measures

Accuracy, precision, recall, F1-score, area under receiver-operating-characteristic curve, false-positive rate, fairness disparity, explanation quality, response time and compliance rate are used to determine performance. Accuracy is

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN}$$

where TP, TN, FP and FN denote True Positive, True Negative, False Positive and False Negative values for the classification outcomes. Using these values, the Precision and Recall values are calculated.

$$Precision = \frac{TP}{TP + FP}$$

$$Recall = \frac{TP}{TP + FN}$$

$$F1 = 2 \times \frac{Precision \times Recall}{Precision + Recall}$$

F1 Score is solved based on this precision and recall values.

Reliability and Validity

Repeated simulation, constant parameter, documented procedure and sensitivity analysis enhances reliability. Construct validity is evaluated via mapping synthetic variables to accepted notions of banking risk and criterion validity compares outputs with pre-existing ground-truth labels. Intersection of metrics, objects and regulatory standards facilitates credibility of the framework.

Ethical and Legal Considerations

Though there are no identifiable customers in the synthetic data, the study adheres to principles of privacy, fairness, cybersecurity and research-integrity. Only bias testing is done using demographic variables. Data, agent prompts and outputs are captured and the decisions are experimental. Cited sources that are under copyright are accompanied by interpretations of the regulations that are made transparent.

IV. Proposed Agentic Ai Framework Framework Principles and Data Integration




Fig. 5. Banking Data Integration and Agentic AI Framework Architecture

Controlled autonomy, accountability, transparency, security, fairness, adaptability and regulatory compliance are applied in the framework. Its data-integration layer is an amalgamation of its synthetic customer, transaction, credit, market, operational, and regulatory data in a regulated repository. The Standardization, validation, access controls and encryption provides a similar input, consistently and reliably to agent decisions and human reviewing as represented in this below figure.

Autonomous Agents and Orchestration

Specialized agents assess credit risk, detect fraud, monitor operational risks, check compliance, understand regulations and make recommendations. An orchestration mechanism allows sharing evidence between the agents, assigning tasks and solving conflicting recommendations based on priority lists. Permissions blocks out available tools and actions and coordination logs maintain traceability at every stage of the workflow.

Risk Assessment and Human Control




Fig. 6. Human-in-the-Loop Risk Assessment and Escalation Workflow

The engine risk-assessment integrates the risk scoring, anomaly detection, predictive assessment, scenario analysis and continuous monitoring. The cases that are

very high in cases of confidence, impact or fairness are sent to human experts. Reviewers are able to accept, refuse or counter-meet suggestions and leave suggestions to be controlled. The figure below depicts such a guided pathology in decision making in the banking operations.

Explainability, Regulation and Decision Support




Fig. 7. Regulatory Governance and Explainable Decision-Support Interface

The explainability and audit layer captures the decision rationales, evidence, agent action logs, versions and human interventions. Banking, data-protection, model-risk and responsible-AI requirements are mapped to the regulatory controls. The interface presents scores of risks, detailing, warnings, suggested measures and options of approval, which allows making responsible decisions. The integrated governance and decision-support interface are shown in the below figure.

V. RESULTS AND FINDINGS




Fig. 8. Distribution of Synthetic Banking-Risk Classes

The figure indicates that there are 8,000 cases under each category. In credit risk, there are 2,800 cases of high risks and in fraud, there are almost 7,850. The compliance and operational risks are mainly low-risk at about 7,100 and 7,000, whereas the total cases of high-risk account to approximately 4,500, showing class imbalance.




Fig. 9. Correlation Structure of Synthetic Banking Variables

The heatmap reveals credit risk to be the key factor making up the overall high risk with $r = 0.61$. The correlation of debt-to-income and credit risk is 0.22 and that of credit score is negative at -0.19. The majority of relationships tend towards 0, which validates that they have low linear relationships.

	Model	Accuracy	Precision	Recall	F1-Score	ROC-AUC
0	Logistic Regression	0.5950	0.6471	0.6204	0.6335	0.6312
1	Random Forest	0.6058	0.6594	0.6233	0.6489	0.6278
2	XGBoost	0.6125	0.6299	0.7592	0.6885	0.6314

	False-Positive Rate
0	0.4379
1	0.4168
2	0.5774

	precision	recall	f1-score	support
0	0.53	0.56	0.55	523
1	0.65	0.62	0.63	677
accuracy			0.59	1200
macro avg	0.59	0.59	0.59	1200
weighted avg	0.60	0.59	0.60	1200

	precision	recall	f1-score	support
0	0.54	0.58	0.56	523
1	0.66	0.62	0.64	677
accuracy			0.61	1200
macro avg	0.60	0.60	0.60	1200
weighted avg	0.61	0.61	0.61	1200

	precision	recall	f1-score	support
0	0.58	0.42	0.49	523
1	0.63	0.76	0.69	677
accuracy			0.61	1200
macro avg	0.60	0.59	0.59	1200
weighted avg	0.61	0.61	0.60	1200

Fig. 10. Classification Reports for Banking-Risk Models

The findings indicate that there is a weak discrimination among all models. XGBoost has the highest accuracy of 61.25% and F1-score of 68.85% and the lowest being 59.50% in logistic regression. Random forest reaches 60.58%. The values of ROC-AUC are close to 0.63 and the performance is only modestly better than random classification.




Fig. 11. Comparative Performance of Banking-Risk Models

The comparison proves that there is no model that corresponds to the level of 80%. XGBoost offers the best recall of about 0.76 and F1-score of 0.69, nevertheless, the accuracy is 0.61. The highest precision of random forest is around 0.66 which shows poor overall predictive reliability to be deployed.




Fig. 12. ROC Curves for Overall Banking-Risk Classification

The ROC curves are near to the random-classifier diagonal. XG Boost, when used alone, has the lowest AUC, which is 0.631, but this is only slightly higher than that of the random forest (0.628). These values indicate poor ranking capability and lack of evidence to rely on automated decisions.




Fig. 13. Confusion Matrices for Banking-Risk Models

On 1,200 cases, the logistic regression correctly classifies on 294 and 420 low and high-risk cases respectively; random forest correctly classifies on 305 and 422. The XGBoost correctly classifies in 514 high-risk cases but falsely classifies 302 low-risk

cases, having a high sensitivity but low specificity and having costly false alarms.




Fig. 14. Key Risk Drivers Identified by XGBoost
XGBoost considers credit score to be the most predictive with around 0.033 mean F1-score decrease in permutation. After it comes debt-to-income of almost 0.020, followed by operational incidents of about 0.007. The rest of the variables are kept at less than 0.005, and as a result capture dedicated credit-risk oversight and explainability regulations in use.

```
*** Isolation Forest fraud precision: 0.9948
Isolation Forest fraud recall: 0.1216
Isolation Forest fraud F1-score: 0.2167
```

Fig. 15. Isolation Forest Fraud-Detection Performance

Isolation Forest has extremely high accuracy of 99.48%, such that nearly all flagged anomalies are fraud. Yet, recall is only 12.16% and F1-score 21.67% indicating that about 87.84% of frauds go undetected. It, therefore, cannot be used, but we can use it to achieve a specific steep.




Fig. 16. Isolation-Forest Detection of Fraud Anomalies

The scatterplot indicates that the anomalies are located among larger transactions and the range of foreign-transactions ratios though red and blue observations have a lot of overlap. Point size represents zero to six failures of logins. This overlap is supporting the 12.16% recall as it justifying more supervised fraud features and models.

Agent-only accuracy: 61.25%
HITL accuracy: 85.83%
Decision improvement, ΔP : 24.58%
Human-escalation rate: 97.17%

Fig. 17. Human-in-the-Loop Decision Improvement Summary

Human-in-the-loop control raises accuracy from 61.25% to 85.83%, producing $\Delta P = 24.58\%$. Nonetheless, 97.17% of cases can only be escalated by humans. Despite enhancing reliability, the fact that supervision effectively is near-universal undermines automation efficiency, and necessitates a stronger level of confidence and achieves optimal model calibration before deployment.



Fig. 18. Performance Improvement Through Human Oversight

The chart confirms that HITL makes accuracy go up by 0.613 to 0.858 and F1-score go up by 0.689 to 0.879 is verified by the chart. A 0.245 and 0.190 increase are strong reasons to consider human oversight but compensation should be weighed against the 97.17% increase overhead seen in operation in general.

Gender	Agent High-Risk Rate	HITL High-Risk Rate	Agent Accuracy
0 Female	0.6938	0.6380	0.6364
1 Male	0.6649	0.5812	0.5864

HITL Accuracy
0 0.8644
1 0.8517

Demographic parity difference: 0.0568

Fig. 19. Gender-Based Fairness Results

The accuracy of the female agents is at 63.64% compared to 58.64% of the male and HITL increases the two to 86.44% and 85.17% respectively. HITL high-risk rates fall to 63.80% and 58.12%. A difference of 5.68 in demographic parity suggests that there is still disparity which has to be checked and ameliorated prior to the approval of the deployment.




Fig. 20. Gender-Based Fairness and Accuracy Assessment

The graph visually confirms increased female scores on all measures. Accuracy of HITL is around 0.864 is for females and 0.852 is for males and the HITL high-risk rate is 0.638 and 0.581 respectively. Continuous discrepancies between subgroups are reasons to justify threshold auditing and validation procedures that are fair.

Measure	Score
0 Prediction Accuracy	0.8583
1 F1-Score	0.8794
2 ROC-AUC	0.6314
3 Explanation Quality	0.8600
4 Compliance Rate	0.9100
5 Fairness Score	0.9432
6 Human Oversight Coverage	0.9717

Fig. 21. Integrated Agentic AI Evaluation Summary

The evaluation records 85.83% accuracy, 87.94% F1-score and 63.14% ROC-AUC. The results of governance are also better with the explanation quality of 86.00%, compliance and fairness of 91.00% and 94.32% respectively. Nevertheless, 97.17% coverage of oversight means that human reliance is too important and the ability to scale the operations autonomously does not allow it to achieve increased reliability in decision making overall.




Fig. 22. Integrated Agentic AI Governance and Performance Outcomes

Six out of the seven measures in the chart had values above the 80% threshold with human oversight

topping the list at 0.972 and fairness at 0.943. ROC-AUC is still just 0.631 means of poor discrimination. Therefore, poor predictive capabilities cannot be fully offset in terms of governance power alone in an operative mode.




Fig.23. Agentic AI Risk-Level and Recommended-Action Distribution

The figure summarizes 1,200 audited cases: 29 low-risk, 852 medium-risk and 319 high-risk cases. Medium risk and high risk create 469 rejection/enhanced-diligence, expert-review create 383 and the high risk produces 259 and 60 rejection and enhanced-diligence respectively. This view point

justifies giving such cases of uncertain medium risks a priority in time of review, however; low-risk escalations have to be reviewed with threshold consideration.




Fig.24. Compliance Warnings Across Predicted Risk Levels

The figure indicates 150 compliance alerts in total: 86 out of 852 medium risk cases and 64 out of 319 high risk cases and low-risk cases have no records. Most cases have no warning, although there are higher proportional warnings at high risk, which provides evidence of the targeted, rather than universal, review.

Table .1. Summary Of Key Results

Measure	Result	Interpretation
Best autonomous model	XGBoost	Strongest classifier
Agent accuracy	61.25%	Limited reliability
Agent F1-score	68.85%	Moderate balance
ROC-AUC	0.631	Weak discrimination
HITL accuracy	85.83%	Strong improvement
HITL F1-score	87.94%	Reliable classification
Performance improvement	24.58%	Human oversight effective
Escalation rate	97.17%	Excessive human dependence
Fairness score	94.32%	Strong overall fairness
Gender parity gap	5.68%	Residual disparity
Compliance rate	91.00%	Strong alignment
Explanation quality	86.00%	Good transparency
Fraud precision	99.48%	Highly accurate alerts

Fraud recall	12.16%	Most frauds missed
Fraud F1-score	21.67%	Weak standalone detection

VI. Discussion

Results show that only autonomous models offer a low reliability of decisions: XGBoost with 76% of its recall value, with 61.25% accuracy and 0.631 ROC-AUC. The human-in-the-loop control achieved an accuracy of 85.83% and F1-score of 87.94% that is consistent with previous studies which favor supervised automation. Nonetheless, the 97.17% growth of the coverage slows down efficacy, whereas the 5.68 points of dissimilarity in the demography illustrate the unmitigated dangers of fairness. This results in banks having quicker prioritization and auditable recommendations although they need to have calibrated thresholds, expert overrides and continuous monitoring. Explainability, audit trails and subgroup testing should be required by regulators. The research relates to autonomy, oversight and compliance.

VII. Conclusion And Future Directions

Risk identification and decision support can be expedited with agentic AI, however autonomous performance is lacking. Trustful banking adoption requires human supervision, revealing justification, and data security, equity appraisal and coordination with regulations.

Recommendations

Use of Agentic AI by banks should be based on risk-based human approval, adjusted thresholds and overrides. Explainability, fairness testing, audit trails and responsible ownership should also be enforced by regulators. Stronger efforts should be made by the developers on calibration, cybersecurity and subgroup performance. Pilots should be implemented in a controlled manner.

Limitations and Future Research

The ability of synthetic data limits realism, and regulatory variation, model bias, limited explainability and technological change as constraints to generalizability. This should be followed in future studies with real life bank implementation tests, cross nation comparisons, longitudinal tracking and model testing.

References

[1]. Buckley, R.P., Zetzsche, D.A., Arner, D.W. and Tang, B.W., 2021. Regulating artificial intelligence in finance: Putting the human in the loop. *Sydney Law Review*, *The*, 43(1), pp.43-81.

[2]. Mthembu, E.C., 2022. Ethical AI-First Banking: Cloud-Powered Cyber Decision Infrastructure with Quantum Machine Learning and Agentic Negotiation Frameworks. *International Journal of Engineering & Extended Technologies Research (IJEETR)*, 4(4), pp.5017-5021.

[3]. Fletcher, G.G.S. and Le, M.M., 2021. The future of AI accountability in the financial markets. *Vand. J. Ent. & Tech. L.*, 24, p.289.

[4]. Johnson, K.N., 2019. Automating the risk of bias. *Geo. Wash. L. Rev.*, 87, p.1214.

[5]. Ruiz, C.I.S., 2021. AI-Driven Software Development for Scalable IoT Hybrid Fuzzy WPM and TOPSIS Integration with Deep Learning and Particle Swarm Optimization in Agentic Negotiation Frameworks. *International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)*, 4(5), pp.5570-5574.

[6]. Truby, J., Brown, R. and Dahdal, A., 2020. Banking on AI: mandating a proactive approach to AI regulation in the financial sector. *Law and Financial Markets Review*, 14(2), pp.110-120.

[7]. Iqbal, J. and Saleh, A., 2020. Artificial Intelligence-Driven Decision Support Systems For Smart Enterprises. *International Research Journal of Advanced Science*, 1(1), pp.1-11.

[8]. Zerilli, J., Knott, A., Maclaurin, J. and Gavaghan, C., 2019. Algorithmic Decision-Making and the Control Problem: J. Zerilli et al. *Minds and Machines*, 29(4), pp.555-578.

[9]. Ejeofobiri, C.K., Adelere, M.A. and Shonubi, J.A., 2022. Developing adaptive cybersecurity architectures using Zero Trust models and AI-powered threat detection algorithms. *Int J Comput Appl Technol Res*, 11(12), pp.607-621.

[10]. Kolyshkina, I. and Simoff, S., 2021. Interpretability of machine learning solutions in public healthcare: the CRISP-ML approach. *Frontiers in big data*, 4, p.660206.

[11]. Kumar, A., Sharma, S. and Mahdavi, M., 2021. Machine learning (ML) technologies for digital credit scoring in rural finance: a literature review. *Risks*, 9(11), p.192.

[12]. Zhang-Zhang, Y., Rohlfer, S. and Rajasekera, J., 2020. An eco-systematic view of cross-sector fintech: The case of Alibaba and Tencent. *Sustainability*, 12(21), p.8907.

[13]. Ullah, I., Rios, A., Gala, V. and Mckeever, S., 2021. Explaining deep learning models for tabular

- data using layer-wise relevance propagation. *Applied Sciences*, 12(1), p.136.
- [14]. Bekamiri, H., Ghasempour Ganji, S.F., Simonetti, B. and Seno, S.A.H., 2021. A new model to identify the reliability and trust of internet banking users using fuzzy theory and data-mining. *Mathematics*, 9(9), p.916.
- [15]. Pu, R., Teresiene, D., Pieczulis, I., Kong, J. and Yue, X.G., 2021. The interaction between banking sector and financial technology companies: qualitative assessment—a case of Lithuania. *Risks*, 9(1), p.21.
- [16]. Ramos Montesdeoca, M., Sanchez Medina, A.J. and Blazquez Santana, F., 2019. Research topics in accounting fraud in the 21st century: A state of the art. *Sustainability*, 11(6), p.1570.
- [17]. Lakhani, M.J., Wang, S., Urbański, M. and Egorova, M., 2020. Sustainable B2B E-commerce and blockchain-based supply chain finance. *Sustainability*, 12(10), p.3968.
- [18]. Bayram, O., Talay, I. and Feridun, M., 2022. Can FinTech promote sustainable finance? Policy lessons from the case of Turkey. *Sustainability*, 14(19), p.12414.
- [19]. Zarrouk, H., El Ghak, T. and Bakhouch, A., 2021. Exploring economic and technological determinants of FinTech startups' success and growth in the United Arab Emirates. *Journal of Open Innovation: Technology, Market, and Complexity*, 7(1), p.50.
- [20]. Mao, D., Wang, F., Hao, Z. and Li, H., 2018. Credit evaluation system based on blockchain for multiple stakeholders in the food supply chain. *International journal of environmental research and public health*, 15(8), p.1627.
- [21]. Tsiakas, K. and Murray-Rust, D., 2022, June. Using human-in-the-loop and explainable AI to envisage new future work practices. In *Proceedings of the 15th international conference on Pervasive technologies related to assistive environments* (pp. 588-594).
- [22]. Methnani, L., Aler Tubella, A., Dignum, V. and Theodorou, A., 2021. Let me take over: Variable autonomy for meaningful human control. *Frontiers in Artificial Intelligence*, 4, p.737072.
- [23]. Enarsson, T., Enqvist, L. and Naarttijärvi, M., 2022. Approaching the human in the loop—legal perspectives on hybrid human/algorithmic decision-making in three contexts. *Information & Communications Technology Law*, 31(1), pp.123-153.
- [24]. Nakao, Y., Stumpf, S., Ahmed, S., Naseer, A. and Strappelli, L., 2022. Toward involving end-users in interactive human-in-the-loop AI fairness. *ACM Transactions on Interactive Intelligent Systems (TiiS)*, 12(3), pp.1-30.
- [25]. Hamrouni, B., Bourouis, A., Korichi, A. and Brahmi, M., 2021. Explainable ontology-based intelligent decision support system for business model design and sustainability. *Sustainability*, 13(17), p.9819.
- [26]. Cabrera, D., Cubillos, C., Urra, E. and Mellado, R., 2020. Framework for incorporating artificial somatic markers in the decision-making of autonomous agents. *Applied Sciences*, 10(20), p.7361.
- [27]. Reis, J., Santo, P. and Melão, N., 2020. Impact of artificial intelligence research on politics of the European Union member states: The case study of Portugal. *Sustainability*, 12(17), p.6708.
- [28]. Amoako, G., Omari, P., Kumi, D.K., Agbemabiase, G.C. and Asamoah, G., 2021. Conceptual framework—artificial intelligence and better entrepreneurial decision-making: the influence of customer preference, industry benchmark, and employee involvement in an emerging market. *Journal of Risk and Financial Management*, 14(12), p.604.
- [29]. Vakili, S.V., Ölçer, A.I. and Schönborn, A., 2021. Identification of shipyard priorities in a multi-criteria decision-making environment through a transdisciplinary energy management framework: a real case study for a Turkish shipyard. *Journal of Marine Science and Engineering*, 9(10), p.1132.
- [30]. Yoo, S.K. and Kim, B.Y., 2018. A decision-making model for adopting a cloud computing system. *Sustainability*, 10(8), p.2952.
- [31]. V. K. B. Parasaram, V. T. Bathini, S. K. Nalluri and A. H. Sannidhanam, "A Sustainable AI-Enabled Predictive Maintenance Framework for Smart Industrial Systems Using Industrial IoT Data," 2026 Third International Conference on Innovations in Cybersecurity and Data Science (ICICDS), Pathum Thani, Thailand, 2026, pp. 440-447, doi: 10.1109/ICICDS70526.2026.11604878.