
Knowledge Graph–Enhanced Decision Intelligence for Regulatory Compliance Reasoning and Cross-Domain Financial Risk Analysis

Gopichand Talluri

Overland Park, Kansas, USA, 66213

Gopichand.bigdtech@gmail.com

Article Received:10-09-2025 Article Modified:25-10-2025

Article Accepted:26-10-2025 Article Published:28-10-2025

DOI:10.37854/UIJMR.2025.3.4.121

Abstract

Financial institutions operate in complex regulatory environments where compliance violations, fraudulent transactions, and interconnected financial risks require accurate and explainable decision-making. Traditional rule-based compliance systems and conventional machine-learning models often struggle to represent complex relationships across regulatory documents, financial entities, and transaction networks, limiting cross-domain reasoning and risk propagation analysis. To address these limitations, this paper proposes Knowledge Graph–Enhanced Decision Intelligence (KG-EDI), an integrated framework for regulatory compliance reasoning and cross-domain financial risk analysis. The proposed model constructs a knowledge graph connecting regulations, compliance requirements, financial entities, transactions, and risk indicators. A hybrid reasoning engine combines graph-based relationship analysis, rule-based compliance verification, and machine-learning-based risk scoring to identify potential violations, detect suspicious relationships, and assess interconnected risks. An explainable decision module generates traceable evidence paths to support transparent and auditable decisions. The framework is evaluated using compliance detection accuracy, precision, recall, F1-score, risk classification performance, and computational latency. Since experimental validation has not yet been provided, quantitative performance improvements cannot be reported as established results. For a preliminary evaluation, the model will be compared against conventional rule-based systems and machine-learning baselines using identical datasets and evaluation conditions. The expected outcome is an integrated decision-support framework that improves regulatory traceability, relationship-aware risk assessment, and explainability across multiple financial domains.

Keywords: Knowledge Graph, Decision Intelligence, Regulatory Compliance, Financial Risk Analysis, Graph Reasoning, Fraud Detection, Explainable Artificial Intelligence, Risk Propagation.

1. Introduction

With the rise in digital financial services comes an increase in the complexity in compliance and risk management processes. Financial organizations need to process large numbers of transactions, interconnections between entities, regulatory compliance, and detect any fraud or financial anomalies. Conventional compliance systems mostly focus on rule-based systems, separate databases, and machine learning algorithms. While these techniques are useful in finding out violations of the rules and classifying financial risks, they lack the ability to capture more complex relations among regulations, financial organizations, transactions, and financial risk factors.

Knowledge graphs are structures of the entities that are interconnected with each other. Knowledge graphs are capable of analyzing the interconnected financial data and regulations as one entity. Regulations, compliance responsibilities, organizations, transactions, and risk factors can all be represented as entities and links on a knowledge graph and thus enable semantic query, relationship analysis, and reasoning through evidence. However, it should be noted that only knowledge graphs are not enough for the whole decision making process to work.

To address the aforementioned issues, this paper suggests Knowledge Graph-Enhanced Decision Intelligence (KG-EDI) that consists of four modules: knowledge graph construction, compliance verification using rules, relationship analysis in graphs, and risk scoring using machine learning algorithms. This framework is designed to identify possible compliance violations, detect suspicious relationships between financial entities, assess interlinked risks, and produce evidence to explain decision-making. Using the combination of structured knowledge and intelligent reasoning, KG-EDI attempts to build a unified approach to analyze regulatory requirements and financial risks in different domains.

The suggested framework can be measured by accuracy, precision, recall, F1-score, risk classification accuracy, and computational time. Evaluating the framework in comparison with rule-based and machine learning techniques will allow assessing its performance in compliance verification and risk detection tasks. The principal contribution of this research is an integrated decision intelligence framework that links regulatory knowledge and financial relationships for transparent, evidence-based, and scalable risk management. The quantitative benefits will be established experimentally rather than assumed.

2. Literature Survey

Abdul-Rahman et al. [1] have created a model to study the construction readiness parameters for abandoned housing projects based on partial least square structural equation modeling (PLS-SEM) method in Malaysia. They studied the factors that affected the readiness for construction and helped to implement housing projects. The authors used PLS-SEM to study relationships between the relevant parameters, and gave insights on what parameters should be taken into

account to minimize the risk of project abandonment. They concluded that it was important to consider project construction readiness systematically in project planning and implementation.

Abid et al. [2] studied the link between risk governance and bank risk-taking behaviour in the context of Asian banks. The authors looked at the effect of governance mechanisms on banks' risk exposure and financial stability decisions. The authors looked at how risk governance influences banking practices, and how good governance structures can be used to mitigate financial risks. Their research has helped to understand the role of appropriate risk oversight and governance mechanisms in helping banks balance activities that entail taking risks with financial stability.

Aburumman et al. [3] discussed how the researcher can interpret and reporting the results of the research using the technique of Partial Least-Squares Structure Equation Modeling (PLS-SEM). Their study included an explanation of suitable procedures for testing measurement models and structural models, and the interpretation of the statistical results from PLS-SEM analyses. The reliability and validity of the instruments, as well as the relationships between the constructs were emphasized by the authors before making conclusions. Their research methodology gave methodological directions to researchers who started using PLS-SEM in business and in technology research in order to analyze complex relationships among the variables.

Afjal et al. [4] studied financial fraud and credit risk by analyzing the aspects and fraudulent financial practices and their consequences for banking stability. Their research revealed the impact of fraudulent activities and credit risk management weaknesses on financial institutions, and how this exposure to potential losses can be increased. The authors spoke about the need to detect financial irregularities and put in place adequate risk assessment and monitoring systems. Their investigations highlighted the importance of better anti-fraud measures and the management of credit risk to strengthen the stability of banking systems and mitigate vulnerabilities in financial systems.

Alaassar et al. [5] conducted a review of the factors that enabled innovations in financial technology (FinTech) and offered a research agenda for future research. They explored the impact of technology advancements, institutional settings, cooperation and organizational competence on innovation in the FinTech industry. The authors compiled the relevant literature to understand the major themes, problems and opportunities of FinTech innovation. Their results have pointed to the importance of grasping the dynamics of technology, regulation and financial institutions to facilitate innovation and the development of digital financial services.

Alfadil et al. [6] performed a systematic literature review to observe the exposure of the construction industry towards force majeure events and environmental risks, and how these challenges compared with the effects of the COVID-19 outbreak. The impact of unforeseen events and environmental uncertainties on construction operations and activities, project

schedules, construction resources and overall project performance were analyzed. The authors outlined the relevance of risk assessment, contingency planning and the right way to manage risks when the situation is not as expected. They highlighted that good risk management and preparedness practices can increase the resilience of construction organisations and reduce the negative impacts of unexpected disruptions.

Allen et al. [7] carried out a literature review of the current literature on financial technology and examined policy debates. They looked at the evolution of FinTech and its impact on financial institutions, competition, regulation and financial services provision. The writers spoke of the impact of new technology on the traditional banking business and the opportunities it provides to enhance financial accessibility and efficiency. They also noted the potential regulatory issues with the new financial technologies, such as the tension between innovation and consumer protection and financial stability. Their survey resulted in a broad overview of the major research themes and policy considerations in the context of FinTech.

Based on comparative analysis using the combination of PLS-SEM and machine learning algorithms, Alshurideh et al. [8] examined the factors that affect the real use of mobile learning systems (m-learning). They reviewed the potential of using statistical modeling and predictive techniques to analyze and predict users' acceptance of m-learning systems. The authors applied PLS-SEM to explore the relationships between the factors and machine learning algorithms to evaluate the predictive performance. Their research showed that using explainable statistical techniques alongside predictive modeling can yield a broader understanding of technology adoption and real-world system usage.

Al-Suwaidi and Nobanee [9] surveyed the existing literature on anti-money laundering (AML) and counter-terrorism financing (CTF). They studied all of the approaches, regulatory measures, and challenges that are encountered in combating money laundering and terrorist financing activities. The authors highlighted that regulatory compliance, financial monitoring and institutional cooperation are all important in the context of detecting and preventing illicit financial transactions. Also, they defined the areas in need of further study and suggested a research agenda for the future. Their findings underscored the importance of having improved financial controls and integrated regulatory action to react to threats to financial integrity that are now occurring.

Brandis et al. [10] studied governance, risk and compliance (GRC) in cloud computing systems. Their research explored the issues that organizations encountered with the governance requirements, security risks, and compliance requirements that they had to manage in cloud-based infrastructure. The authors highlighted the need to set up appropriate governance structures, put in place risk management mechanisms, and ensure adherence to relevant policies and regulations. Their research pointed out that an integrated approach to GRC is needed to

enhance accountability, enhance security, and enable the reliable use of cloud computing services.

3. Proposed Model

The proposed Knowledge Graph–Enhanced Decision Intelligence (KG-EDI) framework integrates knowledge graph construction, regulatory compliance reasoning, cross-domain financial risk analysis, and explainable decision support. The methodology is designed to transform heterogeneous regulatory and financial data into structured knowledge, identify compliance violations, estimate interconnected risks, and generate evidence-based decisions. The framework consists of six major stages: data collection and preprocessing, knowledge graph construction, regulatory rule representation, graph-based risk analysis, decision intelligence, and performance evaluation

A. Data Collection and Preprocessing

The first stage collects heterogeneous data from regulatory documents, compliance policies, financial statements, transaction records, organizational information, and historical risk reports. These sources provide complementary information about regulatory obligations, financial entities, transactions, and potential risk indicators. Regulatory documents are processed to identify obligations, restrictions, thresholds, and exceptions, while financial records are analyzed to extract entities, transaction attributes, and risk-related features.

Data preprocessing includes duplicate removal, missing-value handling, entity normalization, timestamp standardization, and identification of inconsistent records. Entity resolution is applied to determine whether different references correspond to the same organization, account, transaction, or regulatory requirement. The resulting structured dataset provides a consistent foundation for knowledge graph construction and downstream analysis.

Let the collected heterogeneous dataset be represented as:

$$\mathcal{D} = \{\mathcal{D}_r, \mathcal{D}_f, \mathcal{D}_t, \mathcal{D}_e\}$$

where $\mathcal{D}_r$ represents regulatory documents, $\mathcal{D}_f$ represents financial records, $\mathcal{D}_t$ represents transaction data, and $\mathcal{D}_e$ represents entity-related information.

B. Knowledge Graph Construction

After preprocessing, the framework constructs a knowledge graph to represent relationships among regulations, compliance requirements, financial institutions, customers, accounts, transactions, and risk indicators. Each entity is represented as a node, while relationships such as *governed by*, *owns*, *transacts with*, *violates*, and *associated with risk* are represented as edges. Node and edge attributes store relevant information, including transaction amounts, timestamps, regulatory thresholds, jurisdiction, and risk categories.

The graph is designed to support cross-domain analysis by connecting regulatory knowledge with financial activities. For example, a transaction may be linked to an account, an organization, a jurisdiction, and one or more regulatory obligations. Such relationships enable the system to examine indirect associations that may not be visible when the underlying records are analyzed independently. Graph validation checks entity consistency, relationship validity, and missing connections before the graph is used for reasoning.

The knowledge graph is formally represented as:

$$G = (V, E, A_V, A_E)$$

where V is the set of entities, E is the set of relationships, A_V represents node attributes, and A_E represents edge attributes.

A relationship between two entities can be represented as a knowledge triple:

$$(h, r, t)$$

where h is the head entity, r is the relationship, and t is the tail entity.

C. Regulatory Compliance Reasoning

The compliance reasoning module converts applicable regulatory requirements into machine-evaluable rules. Each rule contains a set of conditions, an expected compliance state, and an associated regulatory reference. These rules may address transaction limits, mandatory reporting, documentation requirements, jurisdiction-specific restrictions, or other explicitly represented obligations. The rules are applied to relevant entities and activities in the knowledge graph to determine whether the available evidence indicates compliance, a potential violation, or insufficient information.

The reasoning process combines graph traversal with rule evaluation. Graph traversal identifies the entities and activities relevant to a regulatory obligation, while rule evaluation determines whether the observed attributes satisfy the specified conditions. When a potential violation is identified, the system records the applicable rule, the affected entity or transaction, and the supporting evidence path. If the available data are incomplete, the framework marks the case for further review instead of automatically treating missing evidence as a confirmed violation.

Let $R = \{r_1, r_2, \dots, r_m\}$ denote the set of regulatory rules. The compliance status for entity v under rule r_j is defined as:

$$C(v, r_j) = \begin{cases} 1, & \text{if the rule is satisfied,} \\ 0, & \text{if the rule is violated,} \\ \perp, & \text{if evidence is insufficient.} \end{cases}$$

Here, $\perp$ denotes an unknown or indeterminate compliance state.

D. Cross-Domain Financial Risk Analysis

The financial risk analysis module examines relationships among transactions, accounts, organizations, counterparties, and regulatory indicators. Rather than evaluating each financial record independently, the module uses graph connectivity and relationship attributes to identify potentially suspicious patterns, repeated interactions, unusual transaction flows, and associations with previously identified risk indicators. Risk analysis can span multiple domains, such as transaction monitoring, corporate financial reporting, counterparty relationships, and regulatory compliance.

The framework derives risk features from node attributes, edge attributes, graph neighborhoods, transaction patterns, and compliance outcomes. These features are supplied to a risk-scoring model that estimates the likelihood of a predefined risk event. A graph-based analysis may also identify paths connecting a suspicious transaction to a high-risk entity or a relevant regulatory requirement. Such paths provide contextual evidence for subsequent decision-making.

Let x_v represent the feature vector for entity v . A risk-scoring model estimates the probability that the entity is associated with a risk event:

$$P(y_v = 1 \mid x_v, G) = \sigma(f_\theta(x_v, G))$$

where $y_v = 1$ indicates the presence of the defined risk event, f_θ is a learned prediction function, G is the knowledge graph, and $\sigma(\cdot)$ is the sigmoid function. The model parameters θ are learned from labeled training data when such data are available.

E. Knowledge Graph–Enhanced Decision Intelligence

The decision-intelligence module combines compliance outcomes, financial risk scores, and graph-derived evidence to generate a unified assessment. Compliance status indicates whether a specified obligation is satisfied, violated, or unresolved, while risk scores estimate the likelihood of a defined financial risk. These signals are kept distinguishable because a high financial risk score does not, by itself, establish a regulatory violation.

For prioritization, the system combines normalized risk and compliance indicators using configurable weights. The resulting priority score helps rank cases for investigation, escalation, or routine monitoring. The system also records the regulatory rule involved, the relevant entities, the supporting graph paths, and the model output so that analysts can inspect the basis of each recommendation. High-priority or uncertain cases may be referred for human review.

A general decision-priority function is expressed as:

$$S(v) = \alpha R_v + \beta V_v + \gamma U_v$$

subject to:

$$\alpha + \beta + \gamma = 1, \quad \alpha, \beta, \gamma \geq 0$$

where R_v is the normalized financial risk score, V_v is a normalized compliance-violation indicator, U_v represents uncertainty or evidence insufficiency, and α, β, γ are configurable weights. The priority score is a proposed ranking mechanism and must be calibrated and validated using appropriate evaluation data.

F. Performance Evaluation

The proposed framework should be evaluated using labeled compliance and financial-risk datasets, with training, validation, and test partitions separated to reduce data leakage. Where transactions are time-dependent, a chronological split is preferable. Baseline methods should include a rule-based compliance system, a conventional machine-learning classifier, and a graph-based model without integrated compliance reasoning. All approaches should use equivalent test data and comparable evaluation conditions.

Compliance and risk-detection performance can be measured using accuracy, precision, recall, and F1-score. Computational latency, graph-construction time, and the quality of evidence supporting model decisions should also be assessed. Where the dataset contains imbalanced classes, precision, recall, and F1-score should receive particular attention. Explainability can be assessed through evidence-path validity, traceability to regulatory sources, and expert review of generated explanations. No performance improvement should be claimed until these experiments have been conducted.

The evaluation metrics are defined as follows:

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

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

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

$$F_1 = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}$$

where TP , TN , FP , and FN represent true positives, true negatives, false positives, and false negatives, respectively.




Fig 1. Proposed Model Architecture

G. Framework Workflow

Algorithm 1: KG-EDI for Regulatory Compliance and Financial Risk Analysis

Input:

- Regulatory documents and compliance policies
- Financial statements and transaction records
- Entity and organizational information
- Historical risk labels, if available
- Regulatory rules and model parameters

Output:

- Knowledge graph G
- Compliance status for each evaluated case
- Financial risk scores
- Ranked decision-priority list with traceable evidence

Steps:

1. Initialize the KG-EDI framework and load the input datasets.
2. Preprocess the datasets by normalizing entities, handling missing values, removing duplicates, and standardizing attributes.
3. Extract regulatory obligations, financial entities, transactions, and risk indicators.
4. Construct the knowledge graph $G = (V, E, A_V, A_E)$.
5. Validate entities and relationships and resolve inconsistent or duplicate references.
6. Convert applicable regulatory obligations into machine-evaluable rules.
7. Evaluate each relevant entity or transaction against the applicable rules.
8. Assign a compliance state of compliant, violated, or indeterminate.
9. Extract graph-based features and relationship paths for financial risk analysis.
10. Estimate financial risk scores using the selected risk model.
11. Combine risk scores, compliance indicators, and uncertainty measures to calculate decision-priority scores.
12. Generate evidence-linked recommendations and flag uncertain or high-priority cases for review.
13. Evaluate the framework using the selected test dataset and performance metrics.
14. **Termination condition:** Stop when all eligible cases have been processed and evaluation is complete, or when a predefined computational budget or stopping criterion is reached. If new data arrive continuously, the framework operates iteratively rather than terminating permanently.




Fig 2: Workflow of the Proposed Model

Key Implementation

1. Data Preprocessing: Clean regulatory documents, financial statements, transaction records, and entity data; resolve duplicate and inconsistent entities.
2. Knowledge Graph Construction: Represent regulations, organizations, accounts, transactions, and risk indicators as interconnected nodes and relationships.
3. Regulatory Reasoning: Convert regulatory requirements into rules and evaluate cases as compliant, violated, or uncertain.
4. Financial Risk Analysis: Extract graph-based features and use machine learning to identify suspicious patterns and estimate risk scores.
5. Decision Intelligence: Combine compliance signals, risk scores, and graph evidence to prioritize cases and generate recommendations.
6. Evaluation and Validation: Compare against rule-based, traditional ML, and graph-based baselines using classification metrics, latency, and expert review.

Key Highlights

- **Unified Framework:** Integrates regulatory compliance reasoning and cross-domain financial risk analysis.
- **Knowledge Graph Integration:** Captures relationships among regulations, entities, accounts, and transactions.
- **Explainable Decisions:** Provides traceable evidence paths to support compliance findings and risk alerts.
- **Risk-Based Prioritization:** Ranks cases according to risk scores, violation severity, and uncertainty.
- **Modular and Scalable Design:** Supports updates to data, regulations, graph relationships, and analytical models.
- **Performance Evaluation:** Assesses accuracy, precision, recall, F1-score, AUC, and processing latency.

4. Results and Discussions

The proposed Knowledge Graph–Enhanced Decision Intelligence (KG-EDI) framework is designed to integrate regulatory compliance reasoning with cross-domain financial risk analysis. To illustrate its evaluation, five baseline approaches are considered: a rule-based compliance system, a traditional machine learning model, graph-based risk analysis, a knowledge graph combined with machine learning (KG + ML), and the proposed KG-EDI framework. The comparison considers classification performance, risk detection capability, processing latency, and the contribution of

individual framework components. The plots generated for this study use randomly generated illustrative values rather than results obtained from trained models or a real-world dataset.

Therefore, the following observations demonstrate how the proposed framework can be evaluated; they do not establish empirical superiority.

A. Comparative Model Performance

The first evaluation compares the models using accuracy, precision, recall, F1-score, and area under the receiver operating characteristic curve (AUC). These metrics represent complementary aspects of regulatory compliance classification and financial risk identification. Accuracy measures overall classification correctness, precision measures the proportion of predicted positive cases that are positive, and recall measures the proportion of actual positive cases detected. The F1-score balances precision and recall, while AUC evaluates discrimination across classification thresholds. Fig. 3 presents the illustrative comparison of these metrics across the five approaches. The synthetic values show an increasing performance trend for the integrated KG-EDI configuration. However, this trend is based on predefined example values and must be verified using identical test data, consistent labels, and appropriately tuned baseline models before drawing conclusions about actual performance.




Fig. 3. Performance comparison of compliance and financial risk analysis models.

B. Precision, Recall, and F1-Score Analysis

Precision and recall are particularly important in regulatory compliance applications because false-positive alerts can increase manual investigation costs, whereas false negatives may allow potential violations or financial risks to remain undetected. The F1-score provides a combined measure of these two objectives. Fig. 4 compares these metrics for the selected baseline approaches and KG-EDI. In the illustrative data, the proposed configuration is assigned higher values across all three measures, representing a possible benefit of combining graph relationships, compliance rules, and machine learning. In a real evaluation, the results should be calculated from predictions on a held-out test set, with particular attention to recall for high-risk cases and precision for actionable alerts. Class imbalance should also be considered because a model can achieve high accuracy while failing to identify rare regulatory violations.




Fig. 4. Comparative precision, recall, and F1-score of the evaluated approaches.

C. Receiver Operating Characteristic Analysis

The receiver operating characteristic (ROC) curve illustrates the trade-off between the true-positive rate and the false-positive rate at different classification thresholds. This analysis is relevant to financial risk assessment because decision-makers may require different alert thresholds depending on the severity of a potential violation, the cost of investigation, and the acceptable level of missed risk. Fig. 5 displays illustrative ROC curves for the comparison approaches. The curves are generated using synthetic classification scores and therefore demonstrate the intended evaluation procedure rather than measured model behavior. In the actual experiment, ROC curves should be generated from each model's continuous prediction scores, and AUC values should be computed from the same test set. Where positive cases are rare, ROC analysis should be supplemented with precision–recall analysis.




Fig. 5. Receiver operating characteristic (ROC) curves for risk identification.

D. Precision–Recall Curve Analysis

Precision–recall analysis provides additional insight into performance when the dataset contains substantially more compliant or low-risk records than violation or high-risk records. Unlike accuracy alone, the precision–recall curve directly highlights the relationship between detecting positive cases and maintaining reliable positive predictions. Fig. 6 illustrates the comparison of precision and recall across classification thresholds. For the KG-EDI framework, this evaluation can help determine whether combining regulatory evidence with graph-derived features improves the identification of entities, transactions, or organizations that require further investigation. Nevertheless, the displayed curves are based on synthetic scores and should not be interpreted as evidence that the proposed framework achieves a particular level of risk-detection

performance. Actual assessment should report average precision or area under the precision–recall curve, along with the positive-class prevalence.




Fig. 6. Precision–recall curves for financial risk classification.

E. Confusion Matrix Analysis

A confusion matrix provides a detailed view of classification errors by separating true positives, true negatives, false positives, and false negatives. For the proposed framework, these outcomes can represent correctly identified high-risk cases, correctly identified low-risk cases, unnecessary alerts, and missed risks, respectively. Fig. 7 presents an illustrative confusion matrix for the KG-EDI classification example. This visualization can be used to identify whether the system requires improved sensitivity to violations or better suppression of false alerts. In practical deployment, the relative costs of these errors should inform the decision threshold, particularly when missed high-severity violations have greater consequences than additional manual reviews. The matrix shown here is produced from randomly generated labels and simulated predictions; it is not a validated confusion matrix for KG-EDI.




Fig. 7. Confusion matrix for the illustrative KG-EDI classification example.

F. Processing Latency Comparison

In addition to predictive performance, operational latency is important when the framework processes large regulatory documents, financial transactions, and interconnected entity records. The processing time may include data preprocessing, graph construction or retrieval, regulatory rule evaluation, risk scoring, and evidence-path generation. Fig. 8 compares illustrative latency distributions for the five approaches using synthetic measurements. The plot demonstrates how median latency, variability, and outliers can be examined when assessing operational feasibility. A knowledge graph may introduce additional graph-traversal and maintenance overhead, even when it improves the availability of relationship-based evidence. Consequently, the integrated approach should be evaluated for both computational cost and decision quality. Actual latency measurements must be collected under controlled hardware, data-volume, graph-size, and concurrency conditions before scalability claims can be made.




Fig. 8. Processing latency comparison among alternative model architectures.

G. Ablation Study of KG-EDI Components

An ablation study can help determine the contribution of the individual components of KG-EDI. The configurations considered are machine learning alone, machine learning with knowledge graph features, knowledge graph reasoning with regulatory rules, a knowledge graph combined with machine learning, and the complete KG-EDI configuration. Fig. 9 presents an illustrative comparison of accuracy, F1-score, and AUC across these configurations. The synthetic example assigns higher values to the complete configuration to demonstrate how the contribution of integrated components may be visualized. To establish whether each component provides a genuine benefit, all configurations must be trained and evaluated under consistent data splits, feature availability, and hyperparameter-selection procedures. An actual ablation study should also consider the impact of rule reasoning, evidence-path extraction, and graph construction on both predictive performance and processing cost.




Fig. 8. Ablation analysis of knowledge graph, machine learning, and rule-based components.

5. Conclusion

This paper proposes Knowledge Graph–Enhanced Decision Intelligence (KG-EDI), an integrated framework for regulatory compliance reasoning and cross-domain financial risk analysis. The framework combines data preprocessing, knowledge graph construction, rule-based compliance evaluation, graph-based relationship analysis, and machine learning-based risk scoring to identify potential violations and assess interconnected financial risks. By integrating compliance signals, risk scores, and traceable evidence paths, the proposed approach aims to support explainable decisions and prioritize cases for further investigation. Its modular architecture also enables updates to regulatory requirements, financial data, and analytical models. Future work will focus on implementation using real-world or validated datasets, comparison with established baseline methods, and evaluation using accuracy, precision, recall, F1-score, AUC, processing latency, and evidence-path validity. These experiments are necessary to verify the framework’s effectiveness, scalability, and practical applicability in financial compliance systems.

References

- [1]. Abdul-Rahman, N. Aishah, Rahimi A. Rahman, and Ahmad Rizal Alias. 2024. Modeling the construction readiness parameters for abandoned housing projects in Malaysia: PLS-SEM approach. *Open House International* 49: 904–26.
- [2]. Abid, Ammar, Ammar Ali Gull, Nazim Hussain, and Duc Khuong Nguyen. 2021. Risk governance and bank risk-taking behavior: Evidence from Asian banks. *Journal of International Financial Markets, Institutions and Money* 75: 101466.

- [3]. Aburumman, Omar Jaber, Khatijah Omar, Mohannad Al Shbail, and Mohammed Aldoghan. 2023. How to Deal with the Results of PLS-SEM? In *International Conference on Business and Technology*. Cham: Springer International Publishing, pp. 1196–206.
- [4]. Afjal, Mohd, Aidin Salamzadeh, and Léo-Paul Dana. 2023. Financial Fraud and Credit Risk: Illicit Practices and Their Impact on Banking Stability. *Journal of Risk and Financial Management* 16: 386.
- [5]. Alaassar, Ahmad, Anne-Laure Mention, and Tor Helge Aas. 2023. Facilitating innovation in FinTech: A review and research agenda. *Review of Managerial Science* 17: 33–66.
- [6]. Alfadil, Mohammad Omar, Mukhtar A. Kassem, Kherun Nita Ali, and Wael Alaghbari. 2022. Construction Industry from Perspective of Force Majeure and Environmental Risk Compared to the COVID-19 Outbreak: A Systematic Literature Review. *Sustainability* 14: 1135.
- [7]. Allen, Franklin, Xian Gu, and Julapa Jagtiani. 2021. A Survey of Fintech Research and Policy Discussion. *Review of Corporate Finance* 1: 259–339.
- [8]. Alshurideh, Muhammad, Barween Al Kurdi, Said A. Salloum, Ibrahim Arpaci, and Mostafa Al-Emran. 2023. Predicting the actual use of m-learning systems: A comparative approach using PLS-SEM and machine learning algorithms. *Interactive Learning Environments* 31: 1214–28.
- [9]. Al-Suwaidi, Noura Ahmed, and Haitham Nobanee. 2020. Anti-money laundering and anti-terrorism financing: A survey of the existing literature and a future research agenda. *Journal of Money Laundering Control* 24: 396–426.
- [10]. Brandis, Knud, Srdan Dzombeta, Ricardo Colomo-Palacios, and Vladimir Stantchev. 2019. Governance, Risk, and Compliance in Cloud Scenarios. *Applied Sciences* 9: 320.
- [11]. Buchak, Greg, Gregor Matvos, Tomasz Piskorski, and Amit Seru. 2018. Fintech, regulatory arbitrage, and the rise of shadow banks. *Journal of Financial Economics* 130: 453–83.
- [12]. Cardella, Giuseppina Maria, Brizeida Raquel Hernández-Sánchez, and José Carlos Sánchez-García. 2021. Development and validation of a scale to evaluate students' future impact perception related to the coronavirus pandemic (C-19FIPS). *PLoS ONE* 16: e0260248.
- [13]. Cernisevs, Olegs, Andrey Surmach, and Stanislavs Buka. 2022. Analysis of Aspects of the Regional Economy in the Digital Economy, Using the Example of Financial Services. *Review of Economics and Finance* 20: 203–7.
- [14]. Cernisevs, Olegs, and Yelena Popova. 2023. ICO as Crypto-Assets Manufacturing within a Smart City. *Smart Cities* 6: 40–56.
- [15]. Cernisevs, Olegs, Yelena Popova, and Dmitrijs Cernisevs. 2023a. Business KPIs Based on Compliance Risk Estimation. *Journal of Tourism and Services* 14: 222–48.
- [16]. Cernisevs, Olegs, Yelena Popova, and Dmitrijs Cernisevs. 2023b. Risk-based approach for selecting company Key performance indicator in an example of financial services. *Informatics* 10: 54.

-
- [17]. Chin, Wynne, Jun-Hwa Cheah, Yide Liu, Hiram Ting, Xin-Jean Lim, and Tat Huei Cham. 2020. Demystifying the role of causal-predictive modeling using partial least squares structural equation modeling in information systems research. *Industrial Management & Data Systems* 120: 2161–209.
- [18]. Faccia, Alessio, Narcisa Roxana Moşteanu, Luigi Pio Leonardo Cavaliere, and Leonardo Jose Mataruna-Dos-Santos. 2020. Electronic Money Laundering, The Dark Side of Fintech: An Overview of the Most Recent Cases. In *Proceedings of the 2020 12th International Conference on Information Management and Engineering*. New York: ACM, pp. 29–34.
- [19]. Fantazzini, Dean, and Raffaella Calabrese. 2021. Crypto Exchanges and Credit Risk: Modeling and Forecasting the Probability of Closure. *Journal of Risk and Financial Management* 14: 516.
- [20]. Gaviyau, William, and Athenia Bongani Sibindi. 2023. Customer Due Diligence in the FinTech Era: A Bibliometric Analysis. *Risks* 11: 11.
- [21]. Goecks, Lucas Schmidt, André Luis Korzenowski, Platão Gonçalves Terra Neto, Davenilcio Luiz de Souza, and Taciana Mareth. 2022. Anti-money laundering and financial fraud detection: A systematic literature review. *Intelligent Systems in Accounting, Finance and Management* 29: 71–85.
- [22]. Khan, Mr Ashraf, and Majid Malaika. 2021. *Central Bank Risk Management, Fintech, and Cybersecurity*. IMF Working Papers, WP/21/105. Washington, DC: International Monetary Fund.
- [23]. Kineber, Ahmed Farouk, Ayodeji Oke, Mohammed Magdy Hamed, Ashraf Alyanbaawi, Ali Elmansoury, and Ahmed Osama Daoud. 2023. Decision Making Model for Identifying the Cyber Technology Implementation Benefits for Sustainable Residential Building: A Mathematical PLS-SEM Approach. *Sustainability* 15: 2458.
- [24]. Melnyk, Oleksiy, and Svitlana Onyshchenko. 2022. Ensuring Safety of Navigation in the Aspect of Reducing Environmental Impact. In *International Symposium on Engineering and Manufacturing*. Cham: Springer International Publishing, pp. 95–103.
- [25]. Melnyk, Oleksiy, Oleg Onishchenko, and Svitlana Onyshchenko. 2023. Renewable Energy Concept Development and Application in Shipping Industry. *Lex Portus* 9: 15.
- [26]. Memon, Mumtaz Ali, T. Ramayah, Jun-Hwa Cheah, Hiram Ting, Francis Chuah, and Tat Huei Cham. 2021. PLS-SEM statistical programs: A review. *Journal of Applied Structural Equation Modeling* 5: 1–14.
- [27]. Mohammed, Musa, Nasir Shafiq, Ali Elmansoury, Al-Baraa Abdulrahman Al-Mekhlafi, Ehab Farouk Rached, Noor Amila Zawawi, Abdulrahman Haruna, Aminu Darda’U. Rafindadi, and Muhammad Bello Ibrahim. 2021. Modeling of 3R (Reduce, Reuse and Recycle) for Sustainable Construction Waste Reduction: A Partial Least Squares Structural Equation Modeling (PLS-SEM). *Sustainability* 13: 10660.
- [28]. Mohd Dzin, Najah Hazirah, and Yoon Fah Lay. 2021. Validity and Reliability of Adapted Self-Efficacy Scales in Malaysian Context Using PLS-SEM Approach. *Education Sciences* 11: 676.