The American Journal of Interdisciplinary Innovations and Research
ISSN 2642-7478 | Open Access

TYPE Original Research
PAGE NO. 138-144

Cross-Domain AI-Driven Risk Assessment: Integrating Financial Forecasting And Hazardous-Material Transportation Risk Paradigms

Arjun R. Mehta

Department of Systems Engineering, Global Institute of Technology, India

OPEN ACCESS

SUBMITTED 13 August 2025
ACCEPTED 08 September 2025
PUBLISHED 30 September 2025
VOLUME Vol.07 Issue09 2025

CITATION
Arjun R. Mehta. (2025). Cross-Domain AI-Driven Risk Assessment: Integrating Financial Forecasting And Hazardous-Material Transportation Risk Paradigms. The American Journal of Interdisciplinary Innovations and Research, 7(09), 138–144. Retrieved from https://www.theamericanjournals.com/index.php/tajiir/article/view/6992

Abstract: The rising complexity of modern socio-economic systems demands sophisticated risk assessment frameworks capable of anticipating adverse outcomes across heterogeneous domains. This paper proposes a novel, cross-domain conceptual framework for AI-driven risk assessment that synthesizes insights from two apparently disparate fields: financial market risk forecasting and hazardous-material (hazmat) transportation risk analysis. Drawing on the latest developments in augmented analytics and predictive modeling within finance (Hassan, 2025; Pala, 2023; Oko-Odion, 2025; Kesarpu & Dasari, 2025) and on extensive literature on risk modeling for hazardous-material transport (Liu et al., 2021; Dong et al., 2020; Li et al., 2020; Ditta et al., 2019; Guo & Luo, 2022; Huang et al., 2018; Liu et al., 2018; Erkut et al., 2007; Ma et al., 2020), the framework aims to bridge methodological gaps and deliver a unified paradigm. The proposed framework foregrounds privacy-preserving federated learning techniques (Kalejaiye, 2025) and real-time event sourcing architectures (Kesarpu & Dasari, 2025) as foundational components. We argue that many of the structural and methodological challenges in financial risk forecasting—such as non-stationarity, data sparsity, and regulatory constraints—mirror those in hazmat transportation risk modeling, such as accident unpredictability, environmental variability, and inter-jurisdictional data silos. Through this convergence, the paper outlines theoretical implications, exposes limitations, and sets a research agenda for future empirical validation. The integration of cross-domain methods promises enhanced predictive accuracy, improved resilience, and holistic risk oversight in industries where both financial and physical risks are significant.

Keywords: AI-driven risk assessment; financial forecasting; hazardous-material transportation; federated learning; augmented analytics; real-time

event sourcing; predictive modeling.

Introduction: In recent years, the application of artificial intelligence (AI) and machine learning (ML) techniques for risk assessment has proliferated across a variety of domains. On one end of the spectrum lies financial markets, where volatility, fraud, and systemic risk continue to challenge traditional econometric models. On the other end is the domain of hazardous-material (hazmat) transportation, where accidents can precipitate catastrophic environmental, human, and economic losses. Despite the surface-level differences between these domains, both share deep structural similarities: they must operate under uncertainty, high stakes, regulatory oversight, and often incomplete or siloed data sources. This observation motivates an integrative examination of whether methodologies from one domain (e.g., financial risk forecasting) can be adapted to the other (e.g., hazmat transportation), and vice versa.

In financial contexts, scholars have recently advocated for "augmented analytics," whereby AI-driven techniques produce insights that go beyond traditional statistical models (Hassan, 2025). These methods aim to improve risk management and investment strategies by delivering timely, data-driven predictions and anomaly detection. Meanwhile, predictive analytics has long been recognized as critical in assessing market risk, enabling institutions to anticipate downturns and adjust portfolios accordingly (Pala, 2023; Oko-Odion, 2025). The need for real-time responsiveness has further spurred the adoption of event-streaming and event-sourcing architectures (Kesarpu & Dasari, 2025).

In the hazmat transportation domain, a rich body of work has accumulated over the past three decades, focusing on risk modeling, environmental impact assessments, accident causation, and mitigation strategies (Ditta et al., 2019; Guo & Luo, 2022; Erkut et al., 2007). More recent studies emphasize time-varying environmental conditions (Liu et al., 2021), network optimization for multi-vehicle transportation (Dong et al., 2020), and real-time fuzzy Bayesian networks to assess the risk of flammable liquid tankers (Li et Xu & Shuai, 2020). Additionally, spatial risk assessment of highway-adjacent areas (Huang et al., 2018) and cause-specific risk analysis in rail transport (Liu et al., 2018) highlight the complexity and multi-dimensionality of risk in transportation systems.

However, despite the methodological maturity in each domain independently, there is little to no published work that attempts to bridge financial risk forecasting models with hazmat transportation risk assessment frameworks. Such a bridge could yield mutual benefits. For financial institutions involved in commodity trading, supply-chain financing, or infrastructure development, a deeper understanding of transportation risks could materially inform asset valuations. Conversely, transportation stakeholders could benefit from advanced predictive risk analytics, anomaly detection, and real-time monitoring techniques developed in the financial sector.

The purpose of this paper is to propose a unified "Cross-Domain AI-Driven Risk Assessment Framework" that integrates core methodological advances from both financial forecasting and hazmat transportation risk modeling. Specifically, the framework leverages:

- Federated learning techniques to enable privacy-preserving, collaborative modeling across organizational boundaries while respecting data sovereignty (Kalejaiye, 2025).
- Real-time event sourcing architectures to capture streaming data and enable rapid detection of anomalies or risk events (Kesarpu & Dasari, 2025).
- Hybrid modeling approaches combining statistical, machine learning, fuzzy logic, and Bayesian networks, to handle non-linear, non-stationary, and stochastic risk factors observed in both domains (Li et Xu & Shuai, 2020; Liu et al., 2021).
- Augmented analytics pipelines for risk scoring, scenario simulation, and decision support (Hassan, 2025; Pala, 2023).

By elaborating this integrative framework, we aim to fill a crucial gap in the literature and pave the way for empirical implementation and validation. The rest of the paper proceeds as follows: we first delineate the methodology underlying the proposed framework; then we present a conceptual "results" section in the form of expected outcomes and potential performance metrics; followed by a detailed discussion addressing challenges, limitations, and future research directions; and conclude with summarizing arguments and recommendations for stakeholders in both domains.

Methodology

Given the conceptual nature of this paper, the "methodology" refers to the architecture, components, and procedures of the proposed cross-domain risk assessment framework. The design draws heavily on the methodological strengths and conventions articulated in the foundational literature from both fields.

1. Federated Learning Backbone

A foundational pillar of the framework is the use of federated learning (FL) to enable multiple stakeholders—be they financial firms, logistics companies, regulatory agencies, environmental monitoring bodies—to collaboratively train models without sharing raw data. As demonstrated by Kalejaiye (2025), FL facilitates privacy-preserving collaborative modeling across geopolitically sensitive organizational boundaries. In hazmat transportation, data silos are commonplace: companies may guard shipment records; regulators may hold incident logs; environmental agencies maintain pollution data; but rarely is this data aggregated for joint modeling. FL can bridge these silos: each entity trains a local model on its data, shares only model gradients or updates, and receives a federated aggregate model. This aggregate model benefits from diverse data sources without compromising privacy or competitive advantage.

In financial contexts, federated learning can integrate data from different market participants—hedge funds, exchanges, credit institutions—to build a more holistic picture of systemic risk. For example, a joint model could simultaneously capture liquidity risk, credit risk, and counterparty exposure without exposing proprietary trading data.

2. Real-Time Event Sourcing and Streaming Analytics

To handle the temporal dynamics and non-stationarity inherent in both financial markets and transportation systems, the framework adopts an event-sourcing architecture rooted in real-time streaming. This design is inspired by the architectural proposition of Kesarpu & Dasari (2025), who applied Kafka event sourcing to real-time risk analysis. In practice, streaming data might include: GPS and sensor data from vehicles in transit; environmental sensor readings (e.g., atmospheric conditions, pollution levels); external factors (e.g., weather, traffic conditions); in the financial domain, real-time trade data, order-book changes, credit-default swap spreads, macroeconomic indicators, and news sentiment streams.

Event sourcing enables the system to log all changes in state (e.g., a truck's location changes, a tanker's status update, a sudden spike in oil price, or a regulatory news release). These temporal logs can then feed into models for anomaly detection, rolling-window risk scoring, or scenario simulation. The choice of a streaming architecture ensures that risk assessments remain current and responsive—a critical feature for high-stakes decision-making.

3. Hybrid Modeling Approaches

The core predictive engine of the framework employs a hybrid modeling strategy. This means combining statistical models (e.g., time-series forecasting), machine learning models (e.g., gradient boosting machines, neural networks), fuzzy logic systems, and Bayesian networks. There is precedent for such hybridization: in hazmat transportation, Li et al. (2020) utilized a fuzzy Bayesian network to perform real-time risk analysis for tankers carrying flammable liquid, capturing both aleatory uncertainty (e.g., environmental fluctuations) and epistemic uncertainty (e.g., incomplete information about driver behavior or maintenance status). Also, Liu et al. (2021) demonstrated dynamic risk assessments under time-varying conditions in road transportation. In financial markets, augmented analytics frameworks typically rely on a mix of classical econometrics and more flexible ML models (Hassan, 2025; Pala, 2023). The hybrid approach offers robustness: statistical models ensure interpretability and stability, while ML and fuzzy/Bayesian components provide adaptability to non-linearities, outliers, and uncertain or missing data.

The modeling pipeline can proceed as follows:

- Preprocessing and feature engineering: cleaning, normalization, encoding categorical variables, generating time-series features, computing domain-specific metrics (e.g., tanker load, atmospheric volatility, vehicle maintenance history, financial leverage, liquidity ratios).
- Model ensemble/training per stakeholder node (in federated setup): each node trains a local hybrid model on its subset of data.
- Federated aggregation: periodic or continuous aggregation of model updates into a global model, with differential privacy or secure aggregation to preserve confidentiality.
- Real-time scoring and anomaly detection: as new events stream in, the aggregated model computes risk scores, flags outliers, and updates scenario simulations.
- Feedback and adaptation: model retraining or fine-tuning triggered by significant events (e.g., accidents, near-misses, financial shocks), enabling continuous learning.

4. Augmented Analytics and Decision Support

The final layer of the framework is the augmented analytics interface, which transforms model outputs into actionable insights for decision-makers. This layer draws inspiration from the financial forecasting domain, where AI-driven insights guide risk management and investment strategies (Hassan, 2025; Oko-Odion, 2025). For hazmat transportation, the augmented analytics layer could deliver risk dashboards, route advisories, early warning alerts, "what-if" scenario simulations (e.g., "If tanker X traverses Route Y under weather condition Z, risk increases by 30%"), and incident debrief reports.

Decision support can also inform regulatory oversight:

for example, identifying high-risk corridors, scheduling enforcement resources, or recommending policy interventions. For financial institutions, this may involve adjusting asset allocations, hedging strategies, or counterparty exposures based on real-time risk assessments.

Together, these components—federated learning, event sourcing, hybrid modeling, and augmented analytics—constitute the proposed Cross-Domain AI-Driven Risk Assessment Framework (CAD-RAF).

RESULTS

Because the framework is conceptual, "results" here refers to projected benefits, performance metrics, and hypothetical case-study outcomes rather than empirical data. Nonetheless, by grounding expectations in findings from prior studies, it is possible to outline likely gains and challenges.

1. Enhanced Predictive Accuracy and Early Warning

Prior work in hazmat transportation has demonstrated that fuzzy Bayesian networks can effectively model the risk of tanker accidents under uncertain conditions (Li et al., 2020), and that time-varying environmental factors significantly influence risk (Liu et al., 2021). In financial contexts, augmented analytics has improved risk forecasting beyond static econometric models (Hassan, 2025). Therefore, a hybrid model trained on diverse, federated data is likely to outperform traditional single-source models in both domains. We anticipate:

- Lower false-negative rates (i.e., fewer missed risk events) in route-specific transportation contexts and asset-level financial risk assessments.
- Earlier detection of anomalous patterns (e.g., environmental conditions + route + vehicle status in hazmat transport; or sudden liquidity spikes + counterparty load in finance).
- Improved calibration of risk scores resulting in better alignment between predicted risk and observed outcomes.

2. In practice, these improvements could translate into reductions in accident rates, lower insurance premiums, fewer environmental incidents, decreased financial losses, and more resilient operational strategies.

3. Data Democratization While Preserving Privacy

The federated learning backbone ensures that organizations with sensitive or proprietary data can still contribute to and benefit from a shared model. This helps circumvent one of the major barriers in both domains: data silos. For example, hazard data, driver logs, environmental sensor data, and maintenance records often reside in separate institutions (logistics companies, regulators, environmental agencies). Similarly in finance, trading firms may be reluctant to share proprietary transactional data. By enabling model updates without raw-data sharing, CAD-RAF helps create a collaborative risk intelligence ecosystem without compromising confidentiality. This can foster cross-industry or cross-jurisdictional risk assessments and lead to more comprehensive risk maps—especially vital for global supply chains, multinational logistics companies, and international financial institutions.

4. Operational Responsiveness Through Real-Time Analytics

With event sourcing and streaming analytics, stakeholders can respond rapidly to emerging risk signals. For hazmat transportation, real-time monitoring can trigger alerts—such as rerouting vehicles during sudden weather deterioration, issuing warnings when a tanker passes through high-risk zones, or suspending shipments if maintenance data indicate elevated risk. For finance, real-time analytics could signal liquidity crunches, margin call clusters, or cascading defaults—enabling institutions to adjust positions preemptively.

Over time, such responsiveness could reduce the duration and severity of risk events, mitigate cascading failures, and improve system resilience.

5. Cross-Domain Risk Insights

Perhaps the most novel potential outcome of CAD-RAF is the ability to correlate risk across domains. For example, a logistics company financing tanker transportation might benefit from a combined risk score that accounts for both market volatility (e.g., commodity price swings) and transportation risk (e.g., accident probability). This could inform underwriting decisions, commodity hedging strategies, supply-chain financing, or infrastructure investment. Regulators and policymakers could similarly benefit from a unified risk intelligence platform that spans environmental, transportation, financial, and commodity-market risk.

Discussion

The proposed CAD-RAF framework offers a bold, integrative vision. However, as with any conceptual proposal, it is essential to examine potential limitations, challenges, and assumptions, and to outline a rigorous research agenda for future empirical validation.

1. Data Quality, Heterogeneity, and Labeling Challenges

While federated learning addresses privacy and data silos, it does not automatically solve issues of data quality, heterogeneity, or labeling. In hazmat transportation, data might come from different countries, using different formats, units, and reporting standards. Environmental sensors may vary in calibration; maintenance logs may be incomplete or unstructured; incident logs may have missing information. Similarly, financial data may vary across markets, jurisdictions, and institutions—especially when combining data from credit, derivatives, and real economy exposure.

Inconsistent or noisy data can degrade model performance. Unlabeled data (e.g., near-misses, unreported incidents) complicate supervised learning. While unsupervised anomaly detection and semi-supervised learning might help, the framework's effectiveness hinges on the availability of sufficiently clean and representative data across stakeholders. Without this, federated aggregation could incorporate systemic biases or propagate inaccurate risk estimations.

2. The Challenge of Concept Drift and Non-Stationarity

Both financial markets and hazmat transportation environments are subject to concept drift—a phenomenon where the statistical properties of the data generating processes evolve over time. In finance, macroeconomic shifts, regulatory changes, or structural market transformations can alter risk dynamics. In transportation, new routes, emerging regulations, technological upgrades (e.g., safer tankers), and climate change can shift risk patterns.

The proposed streaming architecture and periodic retraining help address drift; however, determining the optimal retraining frequency, detecting when models have become stale, and ensuring convergence across federated nodes remain non-trivial. Without careful management, models might oscillate, overfit to recent anomalies, or underperform in previously unseen conditions.

3. Governance, Incentives, and Adoption Barriers

A federated, cross-domain platform presumes collaboration among actors who may be competitors—or even adversaries. Financial institutions might fear that participating in a federated risk model reveals vulnerabilities; logistics firms may worry about regulatory scrutiny; governments may hesitate to allow cross-border data sharing due to sovereignty and national security concerns.

Incentive alignment is crucial. Stakeholders must perceive sufficient benefit—such as lower insurance premiums, regulatory relief, improved operational efficiency, or competitive advantage—to engage. Designing governance mechanisms, data-sharing agreements, access controls, and privacy guarantees will be essential. Without clear institutional governance and trust, the theoretical advantages of CAD-RAF may remain unrealized.

4. Interpretability vs. Complexity Tradeoffs

The use of hybrid models (statistical + ML + fuzzy logic + Bayesian networks) improves flexibility and predictive power, but may impair interpretability. In high-stakes domains like hazmat transport or financial regulation, decision-makers often demand explanations for risk predictions. A "black box" model that flags a route or tanker as high-risk without explaining the contributing factors may not meet regulatory or organizational requirements.

Balancing predictive power with interpretability may require sacrificing some performance for transparency, or investing in Explainable AI (XAI) techniques. Yet, the added complexity may reduce model efficiency or complicate governance.

5. Scalability, Infrastructure, and Computational Costs

Implementing a federated, streaming, hybrid analytics platform across multiple organizations and jurisdictions demands robust infrastructure: distributed compute nodes, secure communication channels, fault-tolerant streaming pipelines, and storage for large volumes of event data. Particularly for smaller logistics firms or regional regulators, the cost of building or integrating into such a system may be prohibitive. There may also be regulatory and compliance costs associated with cross-border data sharing or system audits.

For financial institutions, computationally intensive real-time analytics and scenario simulations may require high-performance computing resources, possibly increasing operational costs.

6. Validation, Benchmarking, and Ground-Truth Scarcity

To move beyond conceptual design, empirical validation is indispensable. However, establishing ground truth for risk events is challenging. In hazmat transport, many near-misses go unreported; environmental impacts may manifest long after an incident; or incidents may be suppressed for liability reasons. In finance, systemic risk events (e.g., market crashes) occur infrequently, making them hard to predict or use for training.

Benchmarking the federated global model against baseline local models—or against traditional risk assessment frameworks—will require careful design, long-term data collection, and possibly synthetic data generation. Without rigorous validation, stakeholders may remain unconvinced of the framework's efficacy.

Future Research Agenda

To realize the promise of CAD-RAF, we propose the following research steps:

- Pilot Studies: Begin with small-scale pilots in controlled environments. For instance, select a consortium of logistics firms operating in a single country or region; integrate their hazmat transportation, maintenance, and environmental data into a federated model; monitor risk scores and incident rates over 6–12 months.
- Benchmarking Against Baselines: Compare the federated hybrid model's predictions to those produced by classic risk assessment methods—such as string-based safety audits, historical-frequency models, or expert judgment. Evaluate metrics like precision, recall, false positives/negatives, lead time for warnings, and stability over time.
- Explanability and Stakeholder Feedback: Incorporate interpretability modules (e.g., SHAP values, rule extraction) and gather feedback from domain experts—logistics managers, regulatory officials, environmental analysts—to assess whether risk outputs are credible, actionable, and understandable.
- Governance and Incentive Modeling: Research organizational and economic structures that encourage participation—such as consortium models, data trusts, shared savings mechanisms, regulatory incentives, or insurance premium reductions for participating firms.
- Scalability and Infrastructure Development: Explore cost-effective architectures (e.g., edge computing, lightweight streaming platforms, container-based deployment) that can democratize access for smaller firms, and test the architecture under real-world load and failure conditions.
- Cross-Domain Case Studies: Develop integrated risk assessments for complex systems—e.g., supply chains for dangerous goods financed by commodity traders—to test the viability of cross-domain risk scoring in decision-making (e.g., loan underwriting, pricing, regulatory compliance).
- Regulatory and Ethical Frameworks: Collaborate with policymakers, legal scholars, and ethics experts to draft guidelines for data sharing, privacy, accountability, and transparency—particularly given the cross-border, multi-stakeholder nature of federated risk platforms.

Conclusion

This paper has proposed a conceptual Cross-Domain AI-Driven Risk Assessment Framework (CAD-RAF) that unites methodological advances from two historically distinct domains: financial market risk forecasting and hazardous-material transportation risk modeling. By leveraging federated learning, real-time event sourcing, hybrid modeling, and augmented analytics, CAD-RAF aims to deliver enhanced predictive accuracy, privacy-preserving collaboration, real-time responsiveness, and cross-domain risk insights.

Despite its promise, the framework faces significant challenges related to data quality, governance, interpretability, scalability, and validation. Overcoming these will require coordinated efforts from researchers, industry stakeholders, regulators, and policymakers.

If successfully implemented and empirically validated, CAD-RAF could transform how organizations assess, monitor, and respond to risk—breaking down silos between financial, environmental, and operational domains. In a world where supply chains, commodity markets, and logistics are increasingly interconnected and volatile, such holistic risk intelligence may no longer be optional—but essential.

References

1. Hassan T. Augmented Analytics in Financial Forecasting: Enhancing Risk Management and Investment Strategies through AI-Driven Insights. 2025 Mar 14.
2. Kalejaiye AN. Federated learning in cybersecurity: privacy-preserving collaborative models for threat intelligence across geopolitically sensitive organizational boundaries. Int J Adv Res Publ Rev. 2025;2(07):227–50.
3. Pala SK. Role and importance of predictive analytics in financial market risk assessment. International Journal of Enhanced Research in Management & Computer Applications. 2023;1:2319-7463.
4. Kesarpu S., Dasari H. P. Kafka Event Sourcing for Real-Time Risk Analysis. International Journal of Computational and Experimental Science and Engineering. 2025;11(3).
5. Oluwagbade E. Bridging the healthcare gap: the role of AI-driven telemedicine in emerging economies. Int J Res Publ Rev. 2025 Jan;6(1):3732–43.
6. Oko-Odion C. AI-Driven Risk Assessment Models for Financial Markets: Enhancing Predictive Accuracy and Fraud Detection. International Journal of Computer Applications Technology and Research. 2025;14(04):80–96.
7. Anietom Ifechukwu Chelsea. Evaluating machine learning-based polypharmacy risk prediction in multigenerational households under family medicine, internal and pediatric care interfaces. Int J Sci Res Arch. 2025;16(01):1288–1306.

8. Liu L., Wu Q., Li S., Li Y., Fan T. Risk assessment of hazmat road transportation considering environmental risk under time-varying conditions. Int. J. Environ. Res. Public Health. 2021;18:9780.
9. Dong S., Zhou J., Ma C. Design of a network optimization platform for the multivehicle transportation of hazardous materials. Int. J. Environ. Res. Public Health. 2020;17:1104.
10. Li Y., Xu D., Shuai J. Real-time risk analysis of road tanker containing flammable liquid based on fuzzy Bayesian network. Process Saf. Environ. Prot. 2020;134:36–46.
11. Ditta A., Figueroa O., Galindo G., Yie-Pinedo R. A review on research in transportation of hazardous materials. Socio-Econ. Plan. Sci. 2019;68:100665.
12. Guo J., Luo C. Risk assessment of hazardous materials transportation: A review of research progress in the last thirty years. J. Traffic Transp. Eng. (Engl. Ed.) 2022;9:571–590.
13. Erkut E., Tjandra S., Verter V. Hazardous materials transportation. Handbook of Operations Research and Management Science 2007;14:539–621.
14. Huang X., Wang X., Pei J., Xu M., Huang X., Luo Y. Risk assessment of the areas along the highway due to hazardous material transportation accidents. Nat. Hazards. 2018;93:1181–1202.
15. Liu X., Turla T., Zhang Z. Accident-cause-specific risk analysis of rail transport of hazardous materials. Transportation Research Record: Journal of the Transportation Research Board. 2018;2672:176–187.
16. Ma C., Zhou J., Xu X., Pan F., Xu J. Fleet scheduling optimization of hazardous materials transportation: A literature review. J. Adv. Transp. 2020;2020:4079617.