Machine-Learning-Based Enterprise Risk Classification and Mitigation Using Predictive Analytics

Shourya Gupta

University of Bath, United Kingdom

[1]*Received: 28/03/2023; Accepted: 29/08/2023; Published: 09/10/2023*

ABSTRACT

Business risks are increasingly shaped by fast-changing markets, complex supply chains, digital operations, and evolving regulation. Traditional risk management approaches (workshops, qualitative scoring, periodic audits) remain essential, but they often struggle with early detection, real-time monitoring, and scaling across many business units. Machine learning (ML) can strengthen risk management by (1) identifying weak signals of emerging risks, (2) estimating likelihood and impact with data-driven models, (3) improving detection of anomalies and fraud, and (4) supporting better, faster mitigation decisions. This paper proposes an end-to-end ML risk management framework that connects risk identification, quantification, explainability, and control selection. We review common business risk categories (operational, supply chain, cyber, compliance/fraud, and financial/credit), map them to ML problem types, and outline model development choices (supervised, unsupervised, NLP, time series, causal and probabilistic models). We also present a comparative analysis of model families (logistic regression, random forest, gradient boosting, deep learning, Bayesian networks, and anomaly detection methods) across accuracy, interpretability, data needs, and deployment complexity. Practical issues including data quality, concept drift, fairness, governance, and integration into Enterprise Risk Management (ERM) processes are discussed. Finally, we provide implementation guidance and metrics aligned with risk outcomes, not only predictive performance.

1. Introduction

Risk management aims to protect value and enable confident decision-making. Most firms follow ERM-style cycles: identify risks, assess likelihood/impact, prioritize, mitigate, monitor, and report. However, modern risk environments generate continuous digital traces: transaction logs, operational events, supplier performance metrics, IT telemetry, customer interactions, and external signals such as news, weather, and public vulnerability databases. ML provides methods to detect patterns and predict outcomes from these data streams.

In practice, ML can support three “risk leverage points”:

1. **Early warning**: detecting abnormal patterns before losses occur (e.g., unusual process delays that correlate with incidents).
2. **Prioritization**: ranking risks by predicted probability and expected loss.
3. **Mitigation targeting**: revealing which drivers most influence risk, enabling focused controls and resource allocation.

Research and industry applications show ML’s role across risk classes. For example, operational risk modeling can use probabilistic approaches (e.g., Bayesian networks) to quantify how causal factors change incident likelihood [1]. Supply chain risk work demonstrates ML-based prediction from structured and unstructured signals, but also highlights interpretability needs for practitioner trust [2]. Cyber risk research emphasizes the data challenge and the need for better datasets, while proposing ML-based assessment approaches using public signals (e.g., CVE data) [3–4].

[1] *How to cite the article:* Gupta S (2023); Machine-Learning-Based Enterprise Risk Classification and Mitigation Using Predictive Analytics; *International Journal of Innovations in Applied Sciences and Engineering*; Vol 8, 94-102

Table 1— Traditional vs ML-Enhanced Risk Management

Dimension	Traditional approach	ML-enhanced approach
Signal detection	Periodic reviews, audits	Continuous monitoring, early warning models
Assessment	Qualitative scoring, expert judgment	Data-driven probability/impact estimation
Coverage	Limited by human bandwidth	Scales across processes, suppliers, systems
Adaptation	Slow to update	Retraining + drift monitoring
Explainability	Narrative rationale	Explainable models + driver analysis (e.g., SHAP for tree models) [5]

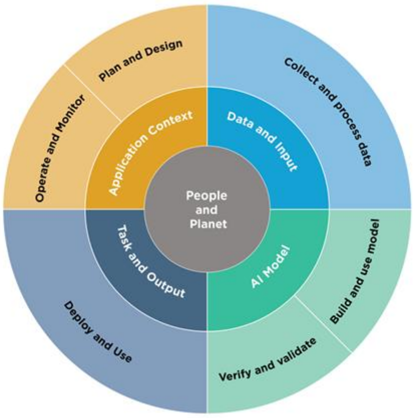

Fig. 1. Lifecycle and Key Dimensions of an AI System

2. Related Work and Theoretical Background

2.1 ML in Risk Management: Core Areas

A widely cited overview of ML and AI in risk management discusses applications in credit risk, market risk, operational risk, and compliance, while noting limitations around transparency and skills [6]. Recent work in operational risk pushes beyond periodic qualitative reviews toward data-driven, dynamic modelling of causal factors using Bayesian network approaches [1]. In supply chain risk, ML methods are used for early identification of production, transport, and supply risks, often using new data sources (including external data) [7], with dedicated work highlighting the performance–interpretability trade-off [2]. Cyber risk literature underscores the lack of open, high-quality data and the difficulty of measuring impacts, which constrains modelling and benchmarking [4]. Explainability is increasingly treated as a prerequisite for high-stakes risk decisions, with methods that connect local explanations to global understanding for tree-based models [5].

2.2 Risk as a Prediction-and-Decision Problem

ML typically optimizes predictive metrics (accuracy, AUC, RMSE), but risk management needs **decision metrics**: expected loss reduction, control effectiveness, false-alarm cost, and regulatory defensibility. A practical framing is:

- **Risk event** (E) occurs with probability (P(E|X))
- **Loss** (L(E)) depends on severity/impact
- **Expected risk** (R = \mathbb{E}[L] = P(E|X)\cdot \mathbb{E}[L(E)|X])

ML can estimate components of (R), while mitigation policy selects actions (a) (controls) to minimize expected loss subject to cost and constraints.

Table 2— Representative Studies Used in This Paper

Area	Study (year)	Contribution	DOI
Operational risk	Cornwell et al. (2023) [1]	Bayesian network CFA for operational risk events	10.1016/j.pacfin.2022.101906
Supply chain risk	Baryannis et al. (2019) [2]	ML framework; performance vs interpretability	10.1016/j.future.2019.07.059
Supply chain risk review	Schroeder & Lodemann (2021) [7]	Systematic review of ML in SCRM	10.3390/logistics5030062
Cyber risk data	Cremer et al. (2022) [4]	Systematic review of cyber risk data availability	10.1057/s41288-022-00266-6
Cyber risk prediction	Kia et al. (2024) [3]	Cyber risk prediction from CVE signals	10.1016/j.eswa.2023.121599
Explainability	Lundberg et al. (2020) [5]	Global understanding from local explanations for trees	10.1038/s42256-019-0138-9
Enterprise risk assessment	Huang et al. (2021) [8]	Enterprise risk assessment with ML classifiers	10.1155/2021/6049195
Fairness in risk models	Kozodoi et al. (2022) [9]	Profit–fairness trade-offs in credit scoring	10.1016/j.ejor.2021.06.023
Anomaly detection	Agyemang (2024) [10]	Comparative evaluation of unsupervised anomaly detection	10.1016/j.sciaf.2024.e02386
General risk management	ASME Open Engineering (2025) [11]	Risk management based on ML methods (engineering focus)	10.1115/1.4069023

3. Proposed ML-Driven Risk Management Framework

We propose a framework that aligns ML work with ERM operations. The key idea: **models must plug into a decision loop**, not sit as isolated dashboards.

Step A: Risk taxonomy and use-case selection

Define risk classes and measurable outcomes:

- Operational incidents, process failures, losses
- Supply disruption events and lead-time spikes
- Cyber incidents, exploit likelihood, downtime
- Fraud/compliance violations, suspicious activity
- Credit default, churn, liquidity stress (depending on business)

Step B: Data and feature architecture

Unify signals across:

- Internal: ERP, CRM, ticketing, logs, audits, HR/attendance, finance
- External: supplier news, weather, macro indicators, vulnerability databases (cyber), shipping data

Step C: Modelling strategy

Match ML approach to the risk problem:

- **Supervised** (when labeled events exist): classification/regression

- **Unsupervised/semi-supervised** (rare events): anomaly detection, one-class models
- **NLP**: risk mining from text (policies, incidents, emails, news)
- **Probabilistic/causal**: Bayesian networks for driver analysis and scenario testing [1]

Step D: Explainability, controls, and actionability

Explain *why* the model flags a risk (feature attribution, counterfactuals, rule extraction). Explainability for tree models can be built using methods that aggregate local explanations into global insights [5]. Then map drivers to **controls**: monitoring thresholds, policy changes, supplier diversification, access restrictions, QA gates, etc.

Step E: Monitoring and governance

- Drift detection and periodic recalibration
- Model risk management (validation, documentation, audit trails)
- Fairness and compliance checks (especially for customer-impacting decisions) [9]

Table 3 — Mapping Risk Stages to ML Deliverables

ERM stage	ML deliverable	Example output	Owner
Identify	Signal detection, NLP risk mining	Emerging risk themes, anomaly clusters	Risk + Data team
Assess	Predictive scoring, severity models	(P(event)), expected loss	Risk analytics
Prioritize	Portfolio ranking	Top 20 risks by expected loss	CRO/ERM
Mitigate	Control recommendation	Which levers reduce risk most	Process owners
Monitor	Drift + KPI dashboards	Alert precision, loss reduction	Risk ops + IT

4. Data Sources, Labelling, and Feature Engineering

4.1 Data challenges by risk type

- **Operational risk**: event logs may be incomplete; “near-miss” data is valuable but often missing.
- **Supply chain risk**: disruptions are influenced by external shocks; integrating external data improves foresight [7].
- **Cyber risk**: strong modelling is constrained by limited open loss data and inconsistent reporting [4].
- **Compliance/fraud**: labels may be delayed (confirmed cases), creating leakage risks.
- **Credit/financial**: richer labels exist but fairness and regulatory constraints are strict [9].

4.2 Labelling strategies

- Confirmed incidents (binary classification)
- Loss amount (regression / severity)
- Time-to-failure (survival analysis)
- Proxy labels: SLA breaches, exception counts, audit flags, customer complaints

4.3 Feature engineering patterns

- Aggregations over time windows (7/30/90 days)
- Ratios and trend deltas (week-over-week changes)
- Network features (supplier/customer graph)
- Text embeddings from incident descriptions/policies
- Interaction terms and monotonic constraints (where needed for policy)

Table 4 — Example Features for Business Risks

Risk type	Example raw data	Example engineered features
Operational	Tickets, process timestamps	Queue length trend, rework rate, exception frequency
Supply chain	Lead times, OTIF, vendor metrics	Lead-time volatility, supplier concentration index
Cyber	CVE feeds, patch logs, IDS alerts	Patch lag, exploitability-weighted exposure (time series)
Fraud/compliance	Transactions, user activity	Velocity rules, peer-group deviation, device mismatch score
Credit/finance	Payments, behaviour	Utilization trend, delinquency history, stability metrics

5. Modelling Approaches for Risk Identification and Prediction

5.1 Supervised learning (event prediction, severity estimation)

When labelled outcomes exist, strong baselines include logistic regression and tree ensembles (random forest, gradient boosting). In supply chain risk prediction, a key practical issue is interpretability vs performance for decision-making [2]. Enterprise risk assessment can also be framed as a supervised classification task using common ML models [8].

5.2 Unsupervised anomaly detection (rare events, unknown patterns)

When incident labels are sparse, anomaly detection is common. A comparative evaluation of unsupervised methods shows meaningful differences across One-Class SVM, Isolation Forest, LOF, and robust covariance approaches, with Isolation Forest often offering a good precision–recall balance under certain conditions [10]. In risk operations, anomaly detection is valuable for early warnings but must be tuned to manage false positives.

5.3 Probabilistic and causal models for actionable insights

Operational risk work illustrates Bayesian network-based modeling to quantify how causal factors influence incident likelihood, improving targeting of mitigations [1]. Such models can support scenario testing ("if control X improves, how does risk change?").

5.4 NLP for risk sensing

NLP can extract risk signals from incident narratives, audit notes, vendor communications, and external text. This often supports:

- Topic detection of emerging risks
- Classification of incident types
- Entity linking (suppliers, systems, products)

5.5 Explainability for high-stakes risk decisions

Explainable AI is essential in risk contexts. Tree-based explanation methods can combine local explanations into global structure, supporting both analyst validation and stakeholder trust [5]. For credit-related models, fairness and governance are central because decisions affect individuals and can trigger regulatory scrutiny [9].

Table 5 — When to Use Which Model Type

Scenario	Recommended model family	Why
Labeled incidents; structured data	Gradient boosting / RF	High accuracy, handles nonlinearity
Need simple, auditable baseline	Logistic regression	Transparent, stable, easy to govern

Scenario	**Recommended model family**	**Why**
Rare events; weak labels	Isolation Forest / One-Class SVM	Works without dense labels [10]
Need scenario reasoning	Bayesian networks	Driver quantification + what-if analysis [1]
Text-heavy risk signals	NLP classifiers / embeddings	Converts narratives/news into measurable signals
Strict fairness constraints	Constrained models + fairness processors	Manage bias and profit trade-offs [9]

6. Evaluation Metrics and Validation in Risk Contexts

6.1 Why classic ML metrics are not enough

Accuracy alone can be misleading when incidents are rare. Risk teams care about:

- **Recall at top-k** (catch the riskiest cases)
- **Precision** (control false alarms)
- **Cost-weighted loss** (false negatives may be far more expensive)
- **Expected loss reduction** after mitigation

6.2 Backtesting and stress testing

Backtesting compares predicted risk vs realized incidents/losses over time. Stress testing evaluates model behavior under plausible extreme conditions (supplier shock, cyber vulnerability surge, demand spikes).

6.3 Drift, calibration, and reliability

Risk environments drift. Operational processes change, suppliers change, attackers adapt. A governance plan should include drift monitoring, recalibration, and performance reporting by segment.

Table 6— Risk-Aligned Metrics

Metric	**Best for**	**Notes**
AUC / PR-AUC	Ranking cases	PR-AUC is better for rare events
Recall@k	Triage workflows	Measures capture rate among limited investigation capacity
Expected cost	Business value	Incorporates false positive/negative cost
Calibration (Brier, reliability)	Probability-based decisions	Needed when thresholds tie to policy
Drift metrics (PSI, KS, error drift)	Monitoring	Triggers retraining or review

7. Comparative Analysis of ML Approaches for Business Risk Management

This section compares methods across key deployment concerns: interpretability, data requirements, robustness, and operational fit.

7.1 Cross-model comparison

Table 7 — Model Trade-offs for Risk Management

Model family	Strengths	Weaknesses	Best-fit risks
Logistic regression	Highly interpretable; easy governance	Limited nonlinear capture	Credit baselines, compliance scoring
Random forest	Robust; handles mixed features	Harder to explain than linear	Operational risk, fraud triage
Gradient boosting (e.g., XGBoost-like)	Strong accuracy; flexible	Needs careful tuning; explainability needed	Supply chain prediction [2], enterprise risk scoring [8]
Deep learning	Strong for unstructured data	Data-hungry; harder governance	NLP risk mining, complex sensor/telemetry
Bayesian networks	Scenario reasoning; causal factor analysis	Requires structure assumptions; setup effort	Operational risk CFA [1]
Anomaly detection (iForest, OCSVM, LOF)	Works with limited labels	False positives; tuning sensitive	Cyber/ops early warning [10]

7.2 Domain comparison

Supply chain research shows ML improves early identification of disruptions and can integrate external signals, but adoption barriers include data standards and systems integration [7]. Cyber risk work highlights that limited open loss datasets restrict validation, pushing many models to rely on proxies such as vulnerabilities and telemetry [4]. A cyber risk prediction approach using CVE-based signals demonstrates one path to reduce expert bias and automate forecasting [3]. For operational risk, Bayesian network approaches can link operational conditions to incident likelihood, helping prioritize mitigations [1].

Table 8 — Domain Constraints vs Modeling Choices

Domain	Data reality	Practical modeling choice
Supply chain	Multi-source, external shocks	Boosted trees + interpretable features [2,7]
Cyber	Sparse impact labels, proxy-heavy	Time series + supervised proxies; anomaly detection [3,4]
Operational	Rich internal logs; causal ambiguity	Bayesian networks + supervised triage [1]
Credit/finance	Strong labels; strict regulation	Interpretable models + fairness controls [9]

8. Implementation and Mitigation: Turning Predictions into Controls

A useful ML risk system must connect predictions to mitigation actions.

8.1 Control mapping

Once top drivers are identified, mitigation can be framed as:

- **Prevent**: reduce probability (patching, training, process redesign)
- **Detect**: increase detection speed (monitoring thresholds, alerts)
- **Respond**: reduce impact (playbooks, redundancy, insurance transfer)

Explainability helps translate model outputs into control levers. For tree models, explanation tooling can support both local case investigation and global control strategy design [5].

8.2 Human-in-the-loop workflows

Risk teams often need analyst review before action. A strong workflow:

- Model produces risk score + top drivers
- Analyst validates and annotates outcomes
- Feedback loop improves labels and retraining
- Policy defines when automation is allowed vs review required

8.3 Governance and model risk management

ML introduces “model risk”: errors, drift, hidden bias, and operational failure. This is why many risk frameworks emphasize documentation, validation, and monitoring, especially in regulated domains [6,9].

Table 9 — Practical Checklist for Deployment

Category	Checklist items
Data	Lineage, quality tests, leakage checks
Model	Benchmark baselines, calibration, stress tests
Explainability	Driver stability, case-level explanations [5]
Monitoring	Drift, alert volumes, incident capture rate
Governance	Approval gates, audit trails, retraining policy
Mitigation	Control playbooks tied to risk drivers

9. Challenges, Ethics, and Future Directions

9.1 Key challenges

- **Data limitations**: cyber risk in particular suffers from limited open data and inconsistent reporting [4].
- **Interpretability vs performance**: especially visible in supply chain risk prediction where practitioner trust matters [2].
- **Concept drift**: attackers adapt, suppliers change, processes evolve.
- **Fairness and accountability**: credit and customer-impacting risk models must manage bias and profit–fairness trade-offs [9].
- **Integration**: risk tools must fit existing ERM governance and reporting.

9.2 Emerging directions

- **Hybrid systems**: combine rules + ML + causal models (better governance and robustness).
- **Scenario generation**: probabilistic models for what-if planning (building on CFA approaches) [1].
- **Better datasets and reporting standards**: especially for cyber loss data [4].
- **Operationalizing explainability**: using global explanation methods to shape policies and controls [5].

Table 10 — Risks Introduced by ML and Mitigations

ML risk	Example	Mitigation
Drift	Supplier behavior shifts post-contract	Drift detection + retraining cadence
Bias	Disparate impact in scoring [9]	Fairness evaluation + processors
Over-alerting	Too many anomalies	Threshold tuning + cost-based optimization
Leakage	Using post-incident info	Strict feature timing rules

ML risk	Example	Mitigation
Governance gap	No audit trail	Model documentation + approvals

10. Conclusion

Machine learning can materially strengthen business risk management by detecting early signals, quantifying risk more consistently, and supporting targeted mitigation. However, success depends less on “the best algorithm” and more on building an end-to-end system: risk taxonomy, data pipelines, model selection aligned to risk economics, explainability, and governance. Comparative evidence across operational, supply chain, and cyber risk shows consistent themes: interpretability and actionability are essential, data constraints shape feasible methods, and continuous monitoring is mandatory in dynamic environments. Implemented well, ML shifts risk management from periodic assessment toward continuous, decision-centric risk intelligence.

References

Agyemang, E. F. (2024). Anomaly detection using unsupervised machine learning algorithms: A simulation study. *Scientific African, 26*, e02386. https://doi.org/10.1016/j.sciaf.2024.e02386

Aziz, S., & Dowling, M. (2018). Machine learning and AI for risk management. In T. Lynn, J. Mooney, P. Rosati, & M. Cummins (Eds.), *Disrupting finance: FinTech and strategy in the 21st century* (pp. 33–50). Palgrave Macmillan. https://doi.org/10.1007/978-3-030-02330-0_3

Baryannis, G., Dani, S., & Antoniou, G. (2019). Predicting supply chain risks using machine learning: The trade-off between performance and interpretability. *Future Generation Computer Systems, 101*, 993–1004. https://doi.org/10.1016/j.future.2019.07.059

Cornwell, N., Bilson, C., Gepp, A., Stern, S., & Vanstone, B. J. (2023). Modernising operational risk management in financial institutions via data-driven causal factors analysis: A pre-registered report. *Pacific-Basin Finance Journal, 77*, 101906. https://doi.org/10.1016/j.pacfin.2022.101906

Cremer, F., Sheehan, B., Fortmann, M., Kia, A. N., Mullins, M., Murphy, F., & Materne, S. (2022). Cyber risk and cybersecurity: A systematic review of data availability. *The Geneva Papers on Risk and Insurance - Issues and Practice, 47*(3), 698–736. https://doi.org/10.1057/s41288-022-00266-6

Huang, B., Li, Y., & Liu, Y. (2021). Enterprise risk assessment based on machine learning. *Computational Intelligence and Neuroscience, 2021*, Article 6049195. https://doi.org/10.1155/2021/6049195

Kia, A. N., Fortmann, M., Mullins, M., Murphy, F., Cremer, F., Sheehan, B., & Materne, S. (2024). A cyber risk prediction model using common vulnerabilities and exposures. *Expert Systems with Applications, 238*(Part B), 121599. https://doi.org/10.1016/j.eswa.2023.121599

Kozodoi, N., Jacob, J., & Lessmann, S. (2022). Fairness in credit scoring: Assessment, implementation and profit implications. *European Journal of Operational Research, 297*(3), 1083–1094. https://doi.org/10.1016/j.ejor.2021.06.023

Lundberg, S. M., Erion, G., Chen, H., DeGrave, A., Prutkin, J. M., Nair, B., Katz, R., Himmelfarb, J., Bansal, N., & Lee, S.-I. (2020). From local explanations to global understanding with explainable AI for trees. *Nature Machine Intelligence, 2*(1), 56–67. https://doi.org/10.1038/s42256-019-0138-9

Schroeder, M., & Lodemann, S. (2021). A systematic investigation of the integration of machine learning into supply chain risk management. *Logistics, 5*(3), 62. https://doi.org/10.3390/logistics5030062

Wang, L., Zhang, Y., & Zhang, J. (2025). Risk management based on machine learning: A review of applications and frameworks. *ASME Open Engineering, 3*(1), 020001. https://doi.org/10.1115/1.4069023