




Neural Early Warning for Counterparty Risk and Compliance

1st Srinivas Kalyan , 2nd Carlos Fernnandez

Srinivas Kalyan Yellanki, Software Engineer 3 , yellankisrinivaskalyan@gmail.com

Carlos Fernnandez , Laboratory Research Assistant

Abstract

This work defines predictive risk and compliance automation for multi-counterparty financial data platforms and reviews the relevance of deep learning models in this context. Predictive risk models face three principal challenges: a scarcity of historical examples of adverse events, limited availability of temporal sequences capturing large negative macroeconomic shocks, and the need to predict low-probability, high-impact events in order to satisfy regulatory and operational resilience requirements. The choice of learning paradigm and model architecture strongly influences the coverage, generalisation, and predictive capability for extreme but consequential events. Supervised methods aim to produce predictive models that generate risk signal estimates, while unsupervised approaches focus on identifying and characterising salient patterns within the data. Multi-counterparty financial data platforms enable the ingestion, integration, harmonisation, storage, and exposure of data generated by the various counterparties participating in a financial ecosystem.

The architecture of such a platform is outlined, with particular attention to the components responsible for ingesting and harmonising data from different counterparties. The paper also discusses the definition of a risk and compliance ontology, and the semantic alignment of these two ontologies across the data ecosystem. This ontology may encompass both a taxonomy of risk signals and a mapping of those signals to fundamental regulatory concepts in risk and compliance, as defined by local and global financial regulators.

Keywords: Multi-Counterparty Data Platforms, Predictive Risk Modeling, Compliance Automation, Financial Risk Analytics, Deep Learning for Risk Prediction, Extreme Event Modeling, Low-Probability High-Impact Risk, Unsupervised Risk Detection, Supervised Learning for Risk Signals, Temporal Risk Modeling, Macroeconomic Shock Analysis, Risk Signal Taxonomy, Compliance Ontology Frameworks, Semantic Data Alignment, Financial Data Integration, Counterparty Data Harmonization, Regulatory Risk Mapping, Data-Driven Compliance Systems, Financial Ecosystem Analytics, AI-Driven Risk Intelligence.



1. Introduction

In multi-counterparty financial data platforms that support consolidated risk management and compliance, distinct risk signals, compliance decisions, or regulatory reports are automatically constructed from the information held about each counterparty. The automated functions are achieved using deep learning models and pipelines able to focus on specific counterparties and scenarios. A main pipeline consolidates data from multiple institutions in a governed manner and for risk analysis, econometric and deep learning models are built, evaluated, and refined. These contributions are typically combined with insights drawn from domain expertise or econometric analysis and formalized as predictive signals or scenario indications. The predictive signals can then be employed in realistic but still synthetic stress-test or scenario analyses. The resulting inputs can finally feed into counterparty-specific and automated regulatory-reporting pipelines.

Construction of predictive credit risk signals automatically leveraging market data across counterparties is a use case in the domain of supervised learning, whereas operations using only the considered counterparty's data fall within unsupervised domain adaptation. The development of compliance automation takes advantage of the broad set of information, the platform's ability to aggregate it in a semantically consistent manner, and the design of the compliance ontologies and pipelines. Deep learning offers several families of models that are applicable to the construction of a variety of credit, market, and operational risk signals. These models can also be employed for stress testing and for scenario-based analysis.

2. Background and Motivation

Deep learning is poised to profoundly affect risk management and compliance functions, even within multi-counterparty financial data-sharing platforms. The empirical evidence supporting this claim is substantial, as is the growing number of deep learning benchmarks for common risk signals, covering credit, market, operational, and liquidity risk, model validation, and KYC controls (including customer transaction screening). Despite these advances, a comprehensive survey detailing the theoretical foundations, empirical support, and practical applications of deep learning for risk and compliance is still lacking. Gaps also remain in the data ingestion and harmonization pipelines essential in a multi-counterparty context—especially in the specification of regulatory concepts.



Developing predictive models on the data shared by multiple counterparties is expected to influence risk and compliance functions greatly. This expectation is underpinned by the strong empirical evidence backing the application of deep learning techniques to risk and compliance topics, as well as a growing number of deep learning benchmarks—covering credit, market, operational, and liquidity risk, validation of risk models, and Know-Your-Customer controls (including transaction screening). However, a cohesive, consolidated, and objective framework for guiding the application of deep learning techniques to risk and compliance is still needed. Furthermore, many previous deep learning studies consider only data from a single entity, leaving the essential data ingestion and harmonization pipelines for multi-counterparty predictive modeling underexplored—especially the specification of regulatory concepts and their formal mapping to the data.

2.1. Significance and Rationale for Deep Learning in Risk and Compliance

Deep learning offers innovative and flexible solutions for predictive risk modeling problems and for automating regulatory compliance processes with respect to publicly available datasets. Results from the generalization of predictive risk signal generation, calibration, uncertainty estimation, and validation, and from showcasing the applicability of deep learning to an increasing number of risk domains provide empirical evidence of its potential value in the risk management practice.

Deep learning can change the process of regulatory compliance automation, allowing the massive utilization of data coming from third parties, enabling Know-Your-Counterparty (KYC) processes to be automated on data other than the obligors volume or helping screening and monitoring for sanctions lists by feeding on data available in the public domain. These aspects, in conjunction with their application to Regulatory Reporting, can open up the sector to the creation of new supervising companies that can be easily scaled in their offer and capacity by serving an increasing number of financial institutions. Along with the scalable and predictive nature of regulatory compliance in the Risk function and the Natural Language Processing (NLP) applied to unstructured risk and compliance controls, deep learning may enable External Auditors to be equipped to face the new supervisory demands from the Basel Committee while maintaining a reasonable price for customers.






Fig 1: Implement AI In Risk Management

3. Theoretical Foundations of Deep Learning for Risk and Compliance

The recent success of deep learning across many disciplines, in particular visual and temporal signal processing, has motivated its adoption for financial risk modeling. Deep learning architectures and learning methodologies are now available to address a full spectrum of financial risk modeling tasks, including credit risk horizon setting, market risk factor evolution, operational risk signal prediction, regulatory reporting generation and validation, sanction screening, and Know-Your-Counterparty objectives. A formal analysis of core models, commonly used loss objectives, and learning paradigms is presented. Where possible, theoretical guarantees are highlighted, and limitations for practical implementation are also noted.

The critical nature of risk and compliance automation in a multilateral trading environment, combined with their importance for financial stability, regulation, and supervision, justify the pursuit of deep learning approaches for these tasks. These applications, together with technical improvements, have been established as suitable for deep learning and are supported by a variety of empirical results. Risk signals for credit, market, and operational risk in multilateral environments share sufficient similarity to permit a common synthesis; within each class of risk, the desired predictive signal may be obtained through both supervised and unsupervised learning methods.



Equation 1. Supervised risk prediction

Start with counterparty features:

$$X = [x_1, x_2, \dots, x_n]$$

The deep learning model maps features to a hidden representation:

$$h = f_{\theta}(X)$$

The final layer produces a risk score:

$$z = w^T h + b$$

Convert the score into probability:

$$PD = \sigma(z) = \frac{1}{1 + e^{-z}}$$


So the complete equation is:

$$PD = \frac{1}{1 + e^{-(w^T f_{\theta}(X) + b)}}$$

3.1. Supervised and Unsupervised Learning in Financial Risk

Supervised and unsupervised data-driven learning approaches provide complementary sets of risk and compliance signals for multi-counterparty financial data platforms. They are both grounded in D. M. J. Tax et al.'s foundational concepts of label-space learning, and may use similar internal representations and architectures but differ in targets and thus underlying objectives. In supervised learning, the model learns to predict the internal label from input features, while unsupervised learning seeks to reconstruct the input using its internal representation.

Unsupervised representation learning attempts to extract the internal label invariant to counterparty. Probabilities of high-level events such as defaulting counterparties or detected money nationalism can be inferred from the joint density—limited to observed range—aided by density-estimation tasks (e.g., reconstruction of random noise from the latent space in VAEs).



High-density regions represent healthy combinations of the half-space inputs while outliers (measured by minimum autoencoder cost) signify unusual behavior.


Table 1: Multi-Counterparty Data Platform Architecture Components

Component Name	Input Data	Deep Learning Role	Output
Data Ingestion Layer	Market, counterparty, KYC data	Feature extraction & preprocessing	Raw unified dataset
Data Harmonization Layer	Heterogeneous schemas	Representation learning	Standardized schema
Risk Ontology Layer	Regulatory concepts	Semantic mapping	Risk taxonomy
Predictive Modeling Layer	Features + labels	Supervised / unsupervised DL	Risk signals
Stress Testing Layer	Historical + synthetic scenarios	Scenario simulation	Stress loss estimates
Compliance Engine	KYC + sanctions data	Classification & matching	Compliance decision
Monitoring Layer	Model predictions	Drift detection	Alerts / rollback

3.2. Sequence Modeling and Temporal Dependencies

While supervised learning with tabular training signals is well established, temporal dependencies and sequence modeling remain more of an open problem for directly predicting risk factors and signals. Given the complexity of risk creation and evolution—credit and counterparty defaults being merely the most extreme events—any predictive signal that could enable early action is of potential interest. Recursive neural networks (RNN) and their advanced versions (long short-term memory networks [LSTM] and Gated recurrent units [GRU]) are the natural supervised models of choice for sequence prediction problems with a fixed-length training dataset. However, the long crumbling times associated with risk data make this a significant challenge, and relatively few applications have appeared hitherto.

The temporal fusion Transformer introduced by Huang et al. could be of interest for risk or compliance tasks that draw on diverse signals across different time granularity levels.




Sequence modeling has further been explored in a non-supervised setting by researchers investigating the predictability of high-dimensional temporal processes. Unsupervised signal learning is slightly more mature in the context of risk, and it includes a variety of applications able to autonomously discover latent risk-generating factors. Other recent work has proposed low-complexity signal learning through Kalman filters, but the technology is limited to the hidden state being linear, although theoretically applicable to prediction with arbitrary temporal resolution.

4. Architecture of Multi-Counterparty Data Platforms

The architecture of multi-counterparty data platforms comprises pipelines for data ingestion from diverse market participants, followed by harmonization processes that utilitarianly align disparate data sources. Leading financial institutions, with their periodic reporting commitment to supervisory organizations governing bank capital and liquidity buffers, have established redundant pipelines for collecting numerous risk and compliance-related meta-data. These frictionless meta-data-sharing platforms aim to upload the ingested information into a common risk and compliance ontology before data quality checks and schema unification. Data required for computation of fully-fledged risk and compliance assessment models is generated, where diverse scenario-analytic pipelines allow continual risk assessment through feeding various risk models with diverse synthetic and historic market conditions.

Deep learning methods of sequence modeling and spurious sequence-analytic pipelines for credit risk, market risk, and operational risk are presented in subsequent sections. Dropping the latent non-policy-driven linkage among market participants enables sanction check, know-your-counterparty check, and regulatory reporting automatization by aligning the required ontologies with the deep-learning-driven data ingestion pipelines of a multi-counterparty coordinate system. An ontological mapping from the left-hand side of the know-your-counterparty check matrix to the relevant domains of the Basel Capital Accord and to various anti-money laundering laws is also provided. Transfer-based, artificial neural network-based, and doctor-check system concepts have been incorporated within the deep-learning paradigm and aim to utilize the framework of predict-without-repeat for supporting supervisory authorities of banks.



4.1. Data Ingestion and Harmonization

Data ingestion pipelines are required for acquiring and ingesting data from external sources (e.g., social media, news, etc.) and internal information systems of each counterparty (e.g., payment systems, settlement systems, trading platforms, etc.). These systems need to support high data quality in terms of accuracy, completeness, consistency, timeliness, and correctness while taking into consideration different quality requirements imposed by the downstream applications, and they must be supported by continuous monitoring and auditing. Ingestion systems also require semantic alignment and harmonization of the ingested data with risk and compliance ontologies to guarantee facilitation of internal audit processes.

The information ingested from external sources is also needed for subsequently assessing KYC and sanction screening decisions aligned with the applicable governmental regulations. Automated data extraction and unification also significantly reduce KYC costs by enabling small transactions to be processed without any human intervention. As the digital services used by both the counterparties and the sanctioned parties become more proliferated and sophisticated, automation of KYC decisions for the large number of small transactions is important for minimizing compliance costs without compromising on the respective filters.





Fig 2: Business Benefits Of AI In Risk Management


4.2. Risk and Compliance Ontologies

Formal ontologies are increasingly used for risk and compliance management since they enable the automated mapping of heterogeneous data. A formal ontology is a specification of a conceptualization together with a collection of information—axioms—that together determine an interpretation. The depth of their semantics makes it possible to precisely reason over the information contained in them. Consequently, they are often used to support intelligent systems.

In the context of regulatory compliance, a taxonomy of global financial regulation is built, together with an ontology mapping the taxonomy to regulatory obligations, banks' internal policies, and risk requirements leading to controls. The mapping enables automated checks of risk and compliance across jurisdictions.

The need for effective regulatory reporting is supported by investment banks at the same time that ontologies facilitate regulatory reporting automation. This process requires a well-defined mapping from the data available within the institution to the reporting requirements and is supported by a dedicated ontology. Ontologies thus also enable the semantic alignment of compliance-related Know-Your-Counterparty requirements and sanctions screening, facilitating automation through pipelines whose decisions can be captured in code.

5. Predictive Risk Modeling with Deep Learning



Given the architecture of multi-counterparty financial data platforms, a variety of predictive models for credit, market, and operational risk can be developed using deep learning based on the data and decision-making pipelines. The predictive signals generated by these models are formally treated as risk signals, as they are not architecturally integrated into a complete risk model but rather represent a level of risk that is being predicted using a data-driven approach. These risk signals can consist of well-defined risk measures, such as the default probability of a counterparty or the expected money lost in the case of a default by the counterparty, which can then be aggregated for a financial institution. More complex risk signals might not have a defined risk measure, yet they are still useful, such as the prediction of whether a financial institution will experience more than 20% losses or market risk VaR at a particular confidence threshold.

Deep Learning-Driven Predictive Risk and Compliance Automation in Multi-Counterparty Financial Data Platforms. For the predictive risk models employed, the input features, predicted target, and evaluation setup should be defined. Stress testing and scenario-level modeling enables the production of “what-if” analyses, while traditional modeling approaches focus on the description of the behavior of risk drivers over longer time horizons. Within a data platform connecting multiple counterparties, different modeling strategies can support these complementary applications. Synthetic scenarios for stress testing and scenario analyses are created by advancing the temporality of the data, while historical events act as natural stress test scenarios. These models can also be used for adapting the risk landscape for given stress test scenarios by identifying the possibility of strengthening, mitigating, or mutually mitigating the potential damages stemming from the events.

Equation 2. Expected loss equation


Credit-risk loss depends on three components:

PD = probability of default LGD = loss given default EAD = exposure at default

Expected loss is:

$$EL = PD \times LGD \times EAD$$

Substitute the deep-learning probability:



$$EL = \frac{1}{1 + e^{-(w^T f_{\theta}(X) + b)}} \times LGD \times EAD$$

5.1. Credit, Market, and Operational Risk Models


Credit, market, and operational risk signals can be generated either through supervised or unsupervised methods. In supervised learning, predictive models leverage labelled examples to learn risk signals. In contrast, unsupervised learning techniques focus on generating anomaly scores that highlight atypical patterns in the data, a property beneficial in itself and conducive to detection of change points, which can subsequently be classified or labelled.

A notable example of unsupervised modelling is the application of Generative Adversarial Networks (GANs) for replicating time series in non-image settings. Through independence and temporal coherence, anomalies can be identified in both directionality and magnitude by Real-valued GANs. Risk signals derived from K-means clustering could serve for credit risk, as an unsupervised approach is feasible when data is structured. In supervised perspective, Deep Neural Networks with Random Forest could be leveraged for market risk. Such architecture can produce the ability to capture the non-linear relationship between inputs and outputs, while also considering the ensemble and feature importance effects incorporated via Random Forest.

5.2. Stress Testing and Scenario Analysis

Stress testing and scenario analysis evaluate compliance with regulatory requirements and the impact of severe events on financial institutions. Regulatory bodies mandate banks and insurers to conduct rigorous internal testing and grant supervisory authorities the power to impose scenario analyses. The recently published Basel guidelines delineate a framework for implementing stressed testing and scenario analysis processes.

Deep learning supports both synthetic and historical scenario analysis. Synthetic scenarios are produced by stochastic processes resilient to market regimes, generating significant stress without injecting additional perfection. The Bank of England's reverse stress-testing mechanisms, coupled with edge intelligent synthesis, establish a minimal-event set with maximal adverse impact. For historical scenarios, major incidents are screened on liquidity, volatility dispersion, and recovery profile blends to prevent erroneous-shaped recoveries.



6. Compliance Automation Strategies


Effective SCF deployment requires integration with compliance functions. Automation processes for regulatory reporting necessitate data indicating financial and trading activities across multiple jurisdictions. Additionally, Know-Your-Customer (KYC) obligations and sanctions screening for parties involved in financing and related transactions impose heavy operational burdens, particularly for non-bank lenders acting as investors and issuers in the capital markets. Such models can incorporate features for both soft and hard capabilities, support validation through signal consistency cross-validation among counter-parties, and include simple decision rules for broad region-level screening.

Financing counterparties must share information regarding the jurisdictions of all financing parties and high-risk activities meant to be excluded from the SCF transactions, such as trading with parties identified within a detected sanction region. Thus, for each financing transaction, logical protection rules can be set to define the necessary information to perform due diligence. These rules can be based on sanctions issued by regulatory agencies or market practice, such as Public Information Notices (PIN) from the World Bank. Validating these screening rules requires mapping the list of countries from the parties involved to the sanction regions.

Table 2: Supervised vs Unsupervised Learning in Risk Modeling

Learning Type	Input	Objective	Output	Use Case
Supervised Learning	Labeled data (X, Y)	Predict risk labels	Probability / classification	Credit risk, reporting
Unsupervised Learning	Unlabeled data (X)	Learn hidden patterns	Anomaly score	Fraud detection, KYC
Semi-supervised	Partial labels	Combine both	Hybrid signals	Rare event modeling
Generative Models (GANs)	Time-series data	Data generation	Synthetic scenarios	Stress testing

6.1. Regulatory Reporting Automation



Deep learning models are being developed to automate the production of regulatory reports with a high level of fidelity across counterparties. Regulatory reporting obligations generally fall into one of two categories. The first requires organizations to submit structured reporting templates periodically, while the second involves continuous qualification against requirements established in regulations. The range of template-based reporting is vast, governed by multiple regulators globally, and the penalties for non-submission can be severe. These templates typically divide information into sections that must be consolidated and reconciled prior to a submission. High-quality reconciled data for any selected template can be assembled, and deep learning tools can be trained directly against prepared template tables for any regulator. The direct-training approach alleviates the need for the underlying components to be modeled explicitly, circumventing the risk of accumulating noise across many shallow models.

The compilation of supervised-reporting templates is challenging and error-prone owing to the richness, complexity, and scalability of the requirements. As organizations and jurisdictions mature their internal KYC processes and systems, the common checklist presented in the template will recede in importance as decision trees for screening counterparties against sanctions lists and high-risk jurisdictions, for determining risk categorization, and for establishing credit limits become the primary KYC support. Furthermore, similar new counterparties often still require additional manual testing for the early phases of onboarding. With sufficient training data, deep-learning-based automation now offers the potential to support the KYC processes.





Fig 3: Leveraging Advanced AI For Risk Management In Banking

6.2. Know-Your-Counterparty and Sanctions Screening

Automating the Know-Your-Counterparty (KYC) process with identification pipelines conforms to the first principle of ensuring adequate counterparty transparency and anti-money laundering provisions. Further benchmarking against global regulatory guidance on KYC reveals an explicit focus on primary objectives, supplemented by decision procedures. By employing these categories, combined with the risk ontology, semantic lineages across the multi-counterparty platform and linked databases create a visual association for supporting external regulatory compliance.

By contrast, due to data Protection by Design and by Default requirements, sanctions screening tends to occur in privately created single databases, which typically match counterparty names against known sanctioned parties. The screening decision is binary, hence rapid execution is paramount. Nevertheless, empirical studies indicate reconciliation checks enhance detection accuracy. When an exact match occurs, the sanction applies; however, for


near-matches (e.g. minor differences in spelling, additional identifiers, or name changes) research-numerated rules facilitate operators’ final decisions.



7. Data Governance, Privacy, and Security

Adequate data governance is a prerequisite for trust in any predictive modeling or compliance automation system that relies on proprietary multi-Counterparty Data Platforms. Policies for enforcing data lineage and provenance, limiting access to restricted data, ensuring data ownership, and preserving data privacy must therefore be defined. These policies must also be communicated to all internal and external Stakeholders who are consuming Insights generated by the systems deployed on the Data Platform.

In addition to adequate data governance, security architectures that limit access to sensitive information while still providing access to data that can help derive a reasonable Insight for a specific Target such as explicit labeling of a counterparty or embargo screening are also required. Moreover, different Counterparty operating in different geographical locations are bound by geographical legislation that mandates that some sensitive information regarding the Counterparty should not even leave a specified geographic location. These architectural and



governance considerations must also be taken into account as part of the data platform design effort.


7.1. Data Lineage and Provenance

Data lineage denotes the flow of data through its lifecycle, from its origin to its final destination. For multi-counterparty platforms, data lineage visualization is crucial to establishing compliance with industry regulations. In automated Know Your Customer (KYC) and Transaction Monitoring (TM) systems, provenance traces how and why learned models arrive at certain decisions, significantly determining model interpretability.

The data lineage of a multi-counterparty financial data platform encompasses every step of the data assimilation process, combining risk and compliance data in a harmonized manner. Provenance defines the sources from which risk and compliance data is acquired, the models and mechanisms used to cleanse and augment this data, and the alignment across multiple counterparties. As data enters the platform, lineage and provenance automations ensure that all requisite data quality checks are performed and all input sources and processing modules are recorded. These checks guarantee, among other things, that individual models are never applied using defective inputs.

7.2. Privacy-Preserving Learning

Data privacy and confidentiality remain critical concerns in many modern data ecosystems, especially in cross-institutional projects such as federated learning or multi-party computation. In these settings, the data is distributed across participants, and each party only has access to its local portion. Cloud-based architectures are very much in vogue but depend on the reliability and trustworthiness of the data host. Privacy-preservation allows participants to jointly learn a shared model without revealing their respective training data; it can also be leveraged across institutions participating in the data federation, where the data host is optimally avoided. Privacy-preserving learning can be defined in terms of gossip-learning principles, where each peer independently learns a model; guarantees such as differential privacy, where noise is injected into the model or gradients before broadcast; or secret-sharing principles, where model parameters are encoded and split across nodes to secure them from malicious eavesdroppers. Other techniques include homomorphic encryption, which introduces additional computational overhead, and secure multiparty computations, costly to execute and requiring the participants to be trusted not to cheat. For many monolithic learning tasks, pre-processing the sensitive training



records using k-anonymity allows anonymizing training and keeping those records off-limits to the model trainer.

Equation 3. Binary cross-entropy loss

For one observation:

$$y \in \{0,1\} \quad p = P(Y = 1 \mid X)$$

Loss is:

$$L = -[y \log(p) + (1 - y) \log(1 - p)]$$

For N counterparties:


$$L = -\frac{1}{N} \sum_{i=1}^N [y_i \log(p_i) + (1 - y_i) \log(1 - p_i)]$$

7.3. Security Architectures for Multi-Party Platforms

Model development and deployment introduce multiple points of exposure in a risk/benefit analysis for the assets involved. Attackers can try to steal the model either as a whole or by leakage of parts. For example, by feeding the model with the previously captured output, it may be possible to capture the architecture, parameters, and weights of the hidden layers. Another attack consists of generating adversarial samples that succeed in classifying certain types of input, making the model unusable in the context they were created. A natural concern for the model provider is the risk of its intellectual property being exploited by unauthorized users. A possible reply to this threat is model watermarking, a process that allows the identification of the source of origin, limiting the possibility of payment evasion as well as loss of value and credibility.

Access to the computation of heavy models may also be abused. Excessive requests from the same address, request from other addresses within a specified area, high requests on specific hours or periods of time, all must be detected. A Denial of Service (DoS) attempt may aim at making an application or service unreachable to users by overwhelming it with requests. Indeed, in the specific case of financial applications, the transaction volume may be temporally


concentrated in specific periods, making requests from the same area legitimate. For this type of risk analysis, mechanisms must be in place to correlate events. The detection of a high accessed API in a certain period from different addresses can be considered a risk. A mitigation may consist of static or dynamic throttling, prioritizing requests or introducing a toll per app/partner.



8. Risk Management and Operational Resilience

A model risk management framework is necessary to ensure the quality of predictive risk and compliance automation. The framework should specify the governance roles, planning, and organization required to minimize the risk of poor-performing models. Governance rules define the key responsibilities for the management, supervision, and validation of models and descriptions of the functional lines of responsibility. Audit comparable principles are also required to assess the functioning of the framework.

A comprehensive deployment pipeline for production systems should be set up to allow for convenient, supervised model deployment, and model performance should be monitored over time with suitable drift detection techniques and procedures for rollout rollback. Possible rollback situations include the failure of a drift detection system or a temporary unavailability of



the model that would prevent normal operations. Risk mitigation options might involve rolling back to a previous version or using a fallback strategy.

The deployment of machine learning models in critical applications must be accompanied by monitoring of their production behavior to ensure proper performance in the decision-making environment. Environmental conditions are subject to variation and models may not operate correctly over the entire range of situations to which they are exposed. While drift detection is the direct response to such a problem, change in the decision environment may equally lead to a non-performing model. The monitoring layer should allow for both types of detection.

Table 3: Compliance Automation Tasks and AI Mapping

Compliance Task	Input Data	Model Type	Output	Regulatory Purpose
KYC Verification	Identity, transactions	Classification	Risk category	AML compliance
Sanctions Screening	Names, entities	Similarity models	Match / no match	Legal compliance
Regulatory Reporting	Financial records	Deep learning	Structured reports	Basel compliance
Transaction Monitoring	Payment streams	Anomaly detection	Suspicious activity	Fraud detection
Counterparty Risk Scoring	Financial metrics	Predictive model	PD score	Credit assessment

8.1. Model Risk Management Frameworks

Risk management and governance framework for model development and use, including model risk appetite, model classification and categorization, approval and validation processes, and roles and responsibilities. It specifies the model lifecycle and the overarching model risk limits, which balance the use of ML techniques with the funding cost, complexity, custom-made solutions, and availability of training data. All models should be reviewed regularly and amended in response to new validation findings, changed business conditions, updated



regulations, or changes in the underlying theory or technology. For production deployment, the automated MLOps pipelines integrate code and non-code artifacts. Centralized monitoring detects performance drifts or quality issues, triggering alerting, review, and corrective actions. Additional checks on model predictions can be performed based on updated inputs (what-if scenarios). If adjustments are not sufficient, the models can be rolled back to previous versions.



Fig 4: AI Enhances Compliance In Finance Sector

8.2. Deployment, Monitoring, and Drift Detection

Management of risks arising from model-based decision making is an essential component in the deployment of machine learning and deep learning models. A model risk management framework should define the role of the risk management function and governance framework, specify responsibilities for validation and verification of new models, establish processes for ongoing monitoring of model performance, define how model drift detection will be performed, and allow for the timely rollback of models to previous stable versions. Model risk management also entails examining the adequacy of the data used for training model predictions, the interpretability and explainability of the model output, and the level of human oversight around the model output.

The success of the previous steps addressing these points should give model risk management the ability to go beyond standard performance monitoring and set rules on how the models should be used to discover findings that would have otherwise remained hidden. This knowledge should also facilitate the monitoring of drifting models and, consequently, define under what conditions they would require retraining or replacement or even a full rollback to the



previous version. Adding this capacity to the standard operational dashboard would close the model-risk confidence loop.

9. Evaluation and Validation Methodologies

Proper evaluation and validation are essential for assessing the quality and suitability of predictive models. Toward this end, distinct metrics, testing methods, and backtesting protocols can be envisaged.

Metrics for Predictive Performance and Calibration


When quantifying the performance of predictive models, the following widely used metrics can be calculated: area under the receiver operating characteristic curve (AUC), area under the precision-recall curve (AUPRC), classification accuracy, Brier score, and Matthews correlation coefficient (MCC). However, classification accuracy is generally not informative for imbalanced classification problems. The Brier score also becomes less effective in the presence of class imbalance. Therefore, if the model output requires calibration, a well-calibrated softmax classifier that returns probabilities can be preferred in these scenarios. The calibration curve and the expected calibration error (ECE) metric can thus be computed. In addition, if the task requires uncertainty quantification, the model can be evaluated against the expected value of information and external validation data collected through online experimentation.

9.1. Metrics for Predictive Performance

Metrics designed for KYC automation, sanctions screening, and credit and operational risk signal calibration differ from typical machine-learning objectives. Predictive risk modeling must encompass the risk of the unseen counterparty when modeling default probability. For a given time range of available features represented by X and a target Y , predictive performance can be quantified according to the modifying asymptotic bias:

where $Error$ quantifies the expected loss of a trained model over the test data and $denotes$ the asymptotic bias expected from moment corrections of a supervised learning approach to producing.

The penalty risk is related to the phenomenon of sample smoothing, which measures, in the context of probability density estimation, the amount of smoothing that occurs because the



training set contains probative (condition on the sample space, $P(Y|X)$) information about an unseen object (not in the training set). Such considerations of hidden markers help provide an intuition for KYC automation.

9.2. Calibration and Uncertainty Assessment

Two broad categories of decision decisions potentially affected by model uncertainty can be considered: (1) decisions based on predicted probabilities such as regulatory reporting, model validation, trigger mechanisms for risk appetite, and risk-based capital allocation; and (2) decisions based on predicted classes such as Know Your Counterparty decisions and expected loss provisions.


Calibrated predictions help assess model reliability and support meaningful interpretation of uncertainty in the predicted probabilities. The one-dimensional calibration plots provide a straightforward way to analyze how accurate predicted probabilities are, but the multi-dimensional counterpart does not easily lend itself to interpretation. Reliability diagrams are thus used as an alternative. To quantify the degree of calibration, the Brier score decomposition formula can be leveraged, as it explicitly separates the calibration and refinement components. An overall uncalibrated probability distribution is the starting point for assessing model uncertainty and can be combined with the above-mentioned subsequent decision-making rules and uncertainty quantification techniques.

Backtesting, including re-calibration and alignment with other sources/models such as market-implied signals, help take into account temporal evolution and assess the degree of model uncertainty. Validation scenarios can also be drawn from historical data, possibly also influenced by additional unobserved sources of uncertainty. The uncertainty in the predicted probabilities can be enriched through a variety of ways, such as ensembling or quantile prediction techniques, in addition to those previous. Finally, the results can be externally validated by an additional random electorate.

Equation 4. Autoencoder anomaly score

Encoder compresses input:

$$z = g_{\phi}(X)$$



Decoder reconstructs input:

$$\hat{X} = d_{\psi}(z)$$

Reconstruction error:

$$A(X) = \|X - \hat{X}\|^2$$

Expanded:

$$A(X) = \sum_{j=1}^n (x_j - \hat{x}_j)^2$$


Decision rule:

$$A(X) > \tau \Rightarrow \text{anomaly / elevated risk}$$

10. Implementation Considerations and Case Studies

Deployment of deep learning algorithms for risk and compliance in a multi-counterparty setting may follow several paths. These instruments can serve as prototypes for user testing and model discovery for production-ready alternatives. Often, the initial phase serves the less data-rich counterparties or the smaller financial institutions in the ecosystem. The results may also support decision-making even if significant detection and prediction errors are present, as humans still play an important role in the final assessment.

A typical approach involves consolidation of data across parties for a specific aspect of risk. Such cross-institution collaboration can foster the development of data-related trust models and simplify auditing processes, thereby enabling more efficient data-sharing frameworks. The first joint study should incorporate a comprehensive assessment of data-sensitive topics across all participating entities, which would strengthen the privacy-preserving nature of the analyses and such privacy-related conditions as Fair Information Practice Principles. Regulation and legislation determine ethical and lawful data-sharing frameworks to support multiple-party data holistically.




10.1. Practical Deployment Scenarios

Deep-learning technology has opened new avenues for risk management solutions, such as credit risk, market risk, and operational risk modeling, across multiple counterparties. Recent developments in ML architecture and training approaches allow for accurate predictive and generative modeling of financial events, making it possible to build an ecosystem of predictive services on a deep-learning foundation.

Practical deployment scenarios span the architecture of services interfacing with users and counterparty stakeholders. First, a classification service is trained to provide predictive results for credit defaults, including early warning signals; losses given default; and facility- and counterparty-level probability-of-default estimates with temporal dependencies, stress-testing capabilities, and risk-scenario monitoring. Risk-scenario generation may benefit, for instance, from denoising diffusion probabilistic models. Alternatively, language models can also be useful by defining the prompt appropriately and conditioning on a suitable set of input features.

Second, the risk signals—structured or semi-structured—are fed into a model risk management core responsible for change detection and performance monitoring. Prior decisions



and evaluations may also be revisited, together with the introduction of additional automatic or semi-automatic testing techniques. Integration of backtesting guarantees enables financial compliance and model governance.

10.2. Cross-Counterparty Data Collaboration

Data collaboration across counterparties enhances predictive performance and supporting analytics. Counterparties operating within trusted networks can share data for better modeling and richer signal generation. For example, banks may calibrate models based on information pooled from within the group but applied to each entity's data. Such pooling may be operated in a secure environment, where the actual data does not leak during the transfer, and applies privacy-preserving models, such as federated learning. Another common use case in collaboration is the adaptation of a Beyond-Data-regime model, trained in a significantly different data distribution but in a similar business domain (telecommunication, biometrics, etc.) to a private domain with limited volume. This can be achieved by leveraging Transfer Learning to reduce Overfitting risk.





Fig 5: AI-Powered Market Research Tools For Financial Analysts

Throughout the collaboration process, trusted analytics will build the necessary level of partnership to allow progressively sharing more sensitive information to achieve more substantial benefits. In the Know-Your-Counterparty structuring process, some signals can be computed by accessing only public information, while others depend on specific counterparties' information. For determining sanctions lists, filters can screen known relationships or information shared privately, such as in the delivery network, but for the ultimate confirmation, such as in telecommunication, all possible relationships should be checked. However, these verification signals must never break trust with the involved actors, supplying again the importance of fair play, progressiveness, and transparent risk analysis along the partnership process.



Equation 5. Stress-test loss

Baseline loss:

$$EL_0 = PD_0 \times LGD_0 \times EAD_0$$

Stress scenario changes risk parameters:

$$PD_s = PD_0 + \Delta PD \quad LGD_s = LGD_0 + \Delta LGD \quad EAD_s = EAD_0 + \Delta EAD$$

Stress loss:

$$EL_s = PD_s \times LGD_s \times EAD_s$$


Substitute:

$$EL_s = (PD_0 + \Delta PD)(LGD_0 + \Delta LGD)(EAD_0 + \Delta EAD)$$

11. Ethical, Legal, and Regulatory Implications

Deep learning has become the de facto standard for a multitude of predictive modeling tasks across a variety of domains, leveraging large amounts of data originating from a broad ensemble of interacting agents. With many of the models deployed within a financial services context being wholly or partially supervised, or addressing regulatory or compliance-related tasks, it is essential to evaluate their ethicalness, fairness, accountability, and transparency. Different fairness definitions exist, no algorithm can guarantee fairness in its objectives, and fairness definitions may conflict. Furthermore, unethical models tend to operate in violation of fiduciary responsibilities, exposing those responsible for sacrificing those obligations for the sake of greater profits to civil legal liability.

Regulation of AI, Large Language Models (LLMs), and Deep Learning is a rapidly evolving space in the European Union, United States, and Singapore, including attempts at global collaboration through the Global Partnership on AI. The financial services industry is already subject to substantial regulation, much of which is indirectly applicable to AI, such as the European Union's General Data Protection Regulation, the Markets in Crypto-assets regulation, the Internal Market Act, the Payment Services Directive, and the EU Securitisation Regulation. AI and ML models are used to flag AML red flags for review by analysts, allowing banks to



invest in the other areas of KYC and Sanctions Screening, detecting ‘worrying patterns’ that don’t break laws yet appear odd. Model risk management provides a structured framework for controlling the risk of using these models and ensuring their deployment is appropriate and reviewed on a regular basis.


Table 4: Evaluation Metrics for Predictive Risk Models

Metric Name	Formula	Purpose	Interpretation
Probability of Default (PD)	$P(Y=1 X)$		Default prediction
Expected Loss (EL)	$PD \times LGD \times EAD$	Financial risk	Expected loss value
Brier Score	$\frac{1}{N} \sum (p_i - y_i)^2$	Calibration accuracy	Lower is better
AUC-ROC	Area under ROC curve	Classification quality	Closer to 1 is better
MCC	Correlation metric	Balanced evaluation	Robust to imbalance
ECE	Calibration error	Probability reliability	Lower is better

11.1. Fairness, Accountability, and Transparency

The increasing use of Machine Learning (ML) in decision-making processes has raised ethical considerations regarding model fairness, accountability, and transparency. Several factors contribute to potential discrimination when ML is applied, such as biased training data or miscalibrated distributions for sensitive groups. An ethical problem arises in cases where ML models impact individuals' lives but the decisions made are difficult for humans to grasp. Such black-box approaches hinder accountability and can perpetuate societal bias. Therefore, it is essential to assess the model's fairness, define sensitive attributes during training, and wrap the algorithm in a transparent solution that indicates the model's decisions with sufficient reliability and faithfulness.

Fairness in algorithmic decisions can be defined in multiple ways: statistical sameness among sensitive groups, Pareto efficiency, difference invariance, and counterfactual fairness. During model development, statistical fairness metrics can indicate sensitivity for a particular decision-making problem. Conversely, establishing accountability requires model explanations that reveal the factors influencing the final decision. Shapley additive explanations provide a unified framework to evaluate model compliance with fairness metrics. In sensitive decisions, such as credit applications, the bank's final decision should consider the individual probability of default, the discriminatory factor of the logistic regression in default predictions, and Shapley values of the decision-maker variables. Transparent systems should also allow conditioning explanations to different counterfactual scenarios, enabling a fairer and more ethical final decision. Initiatives like the Local Interpretable Model-agnostic Explanations framework can also offer model agnosticism at a cost.



11.2. Compliance with Global Standards


The successful automation of compliance requires the execution of multiple data-intensive processes. Regulatory mandates from institutions such as the Financial Action Task Force (FATF), Bank für Internationalen Zahlungsausgleich (BIS), and European Banking Authority (EBA), among others, drive the implementation of periodic Know Your Customer (KYC) and Sanctions Screening procedures. In the case of KYC, in-scope transactions often

require verification of the business relationship and are classified according to rationale for the association. Required checks include ownership structure, identification of politically exposed persons, description of business and economic purpose of the relationship, location and products, and expected account activity. The use of a structured questionnaire helps to provide a clear and well-documented rationale to support decision-making, enable easy assessment of information quality, and ensure peer review.

For Sanctions Screening, institutions are tasked with Globally Active Adult Sanctioned (GAAS) entities and must reject business with these parties. Linkedin and similar platforms are sources of information to capture sanctioned entities with aspirational business relationships. For all other parties, table checks against available legal summaries, news articles, and sanctions lists are required. Explicit do's and don'ts typically support the final decision and create an immutable trail for audit purposes. While Challenge Question Answers provide a means to trouble-shoot the response in case of investigation by regulators, the ever-increasing complexity of crypto-currency and the platforms enabling it is straining these compliance processes, demanding more heuristic than analytic scrutiny.



Fig 6: Deep Learning Impact On Fraud Detection In Large Retail Bank




12. Conclusion

Machine learning, and deep learning in particular, have demonstrated state-of-the-art predictive performance in a variety of application domains. Nevertheless, concerns related to the validation, governance, and acceptable use of learnt models remain. These concerns are even more pronounced in risk and compliance applications, where insufficient, incorrect, or biased predictions can have severe consequences. Poor predictive performance may also hinder industry acceptance and negates the key advantages of modern predictive techniques. The systematic assessment of these issues is an essential step toward building deep learning-driven predictive risk and compliance automation systems on multi-counterparty data platforms. Consequently, metrics have been proposed for predictive performance, calibration, and uncertainty quantification of risk and compliance models and protocols for monitoring, validation, backtesting, and stress-test evaluation have been discussed.


Machine learning-based risk and compliance automation systems can facilitate more efficient regulatory reporting and allow greater focus on high-risk counterparties, scenarios, and decisions. The scalable nature of these predictive pipelines also supports cross-counterparty data collaboration without disclosing sensitive information. Multi-party platforms can leverage a large volume of data within established data governance models to train and validate more accurate and reliable models. Deployed models are thus better aligned with the expectations of regulatory authorities, and the consolidation of monitoring and validation efforts can improve model risk governance across all participating data owners and users.

References


1. Bussmann, N., Giudici, P., Marinelli, D., & Papenbrock, J. (2021). Explainable machine learning in credit risk management. *Computational Economics*, 57, 203–216.
2. Wang, L., & Zhang, W. (2023). A qualitatively analyzable two-stage ensemble model based on machine learning for credit risk early warning: Evidence from Chinese manufacturing companies. *Information Processing & Management*, 60(3), Article 103267.
3. Motie, S., & Raahemi, B. (2024). Financial fraud detection using graph neural networks: A systematic review. *Expert Systems with Applications*, 240, Article 122156.
4. Kumar, S. S., Gadi, A. L., Sheelam, G. K., Kummari, D. N., Reddy Koppolu, H. K., & Pamisetty, A. (2026). AI-Driven Compliance and Audit Framework for Manufacturing Infrastructure in Automotive Connected Services and Financial Ecosystems. In 2026 IEEE



- International Conference for Convergence in Computing Technology (I3CTCON) (pp. 1–6). IEEE. 2026 IEEE International Conference for Convergence in Computing Technology (I3CTCON). <https://doi.org/10.1109/i3ctcon68242.2026.11507231>
5. Deprez, B., Wei, W., Verbeke, W., Baesens, B., Mets, K., & Verdonck, T. (2025). Advances in continual graph learning for anti-money laundering systems: A comprehensive review. *WIREs Computational Statistics*, 17(3), Article e70040.
 6. Cheng, D., Zou, Y., Xiang, S., & Jiang, C. (2025). Graph neural networks for financial fraud detection: A review. *Frontiers of Computer Science*, 19, Article 199609.
 7. Nandan, B. P., Kumar, M. V. K., Garapati, R. S., Bandi, V. D. V. K., Davuluri, P. S. L. N., & Mangalampalli, B. M. (2026). AI-Enhanced Semiconductor Yield Optimization Using Hybrid Deep Learning and Edge Data Analytics. In 2026 IEEE International Conference on AI Engineering and Innovations (AIEI) (pp. 1–6). IEEE. 2026 IEEE International Conference on AI Engineering and Innovations (AIEI). <https://doi.org/10.1109/aiei69164.2026.11497190>
 8. Bussmann, N., Giudici, P., Marinelli, D., & Papenbrock, J. (2021). Explainable machine learning in credit risk management. *Computational Economics*, 57, 203–216.
 9. Li, H., & Wu, W. (2024). Loan default predictability with explainable machine learning. *Finance Research Letters*, 60, Article 104867.
 10. Raut, B., Garg, A., Mangalampalli, B. M., Jindal, N., Mahajan, V., & Saraswati, K. (2026). Multidisciplinary Strategies for Improving Maternal and Reproductive Health Outcomes and Mental Outcomes. *International Journal of Drug Delivery Technology*, 16(34s). <https://doi.org/10.25258/ijddt.16.34s.36>
 11. Wang, L., & Zhang, W. (2023). A qualitatively analyzable two-stage ensemble model based on machine learning for credit risk early warning: Evidence from Chinese manufacturing companies. *Information Processing & Management*, 60(3), Article 103267.
 12. Motie, S., & Raahemi, B. (2024). Financial fraud detection using graph neural networks: A systematic review. *Expert Systems with Applications*, 240, Article 122156.
 13. Meda, R., Motamary, S., Dodda, A., Kummari, D. N., Kaulwar, P. K., & Gadi, A. L. (2026). Edge-AI Integration for Accelerated Decision-Making in Industrial Robotics and Manufacturing Plants. In *Lecture Notes in Networks and Systems* (pp. 332–344). Springer Nature Switzerland. https://doi.org/10.1007/978-3-032-19185-4_27
 14. Deprez, B., Wei, W., Verbeke, W., Baesens, B., Mets, K., & Verdonck, T. (2025). Advances in continual graph learning for anti-money laundering systems: A comprehensive review. *WIREs Computational Statistics*, 17(3), Article e70040.
 15. Bussmann, N., Giudici, P., Marinelli, D., & Papenbrock, J. (2021). Explainable machine learning in credit risk management. *Computational Economics*, 57, 203–216.



16. Li, H., & Wu, W. (2024). Loan default predictability with explainable machine learning. *Finance Research Letters*, 60, Article 104867.
17. Kolla, S. K., Bandi, V. D. V. K., & Meda, R. (2026). Comment on “From laboratory promise to decision-grade practice: strengthening reproducibility and casework relevance in AI-assisted bloodstain ageing.” *Forensic Science, Medicine and Pathology*.
<https://doi.org/10.1007/s12024-026-01258-x>
18. Advancing credit risk modelling with machine learning: A comprehensive review of the state-of-the-art. (2024). *Engineering Applications of Artificial Intelligence*, 137, Part A, Article 109082.
19. Explainable AI for credit assessment in banks. (2022). *Journal of Risk and Financial Management*, 15(12), Article 556.
20. Krishnan, M., Yandamuri, U. S., Recharla, M., Challa, K., Kolla, S. H., & Valiki, S. (2026). Adaptive Machine Learning Ecosystem for Personalized Medicine and Digital Finance: Leveraging Azure Cloud, Data Engineering, and Workflow Automation. In *2026 5th OPJU International Technology Conference (OTCON) on Smart Computing for Innovation and Advancement in Industry 5.0* (pp. 1–6). IEEE. 2026 5th OPJU International Technology Conference (OTCON) on Smart Computing for Innovation and Advancement in Industry 5.0.
<https://doi.org/10.1109/otcon69332.2026.11629728>
21. An explainable AI framework for credit evaluation and analysis. (2024). *Applied Soft Computing*, 153, Article 111307.
22. Graph learning-empowered financial fraud detection: Progress and future directions. (2025). *Intelligent Computing*, 4.
23. Sudha Rani, P. R., Amistapuram, K., Pamisetty, V., Singireddy, S., Kummari, D. N., & Sheelam, G. K. (2025). Hybrid Knowledge Graph–Deep Learning Framework for Automated Exception Handling and Investigation in Complex Insurance Claims. In *2025 IEEE 3rd Global Conference on Wireless Computing and Networking (GCWCN)* (pp. 1–6). IEEE. 2025 IEEE 3rd Global Conference on Wireless Computing and Networking (GCWCN).
<https://doi.org/10.1109/gcwcnc66157.2025.11448301>
24. Network analytics for anti-money laundering—A systematic literature review and experimental evaluation. (2025). *INFORMS Journal on Data Science*.
25. Detecting anomalies in FinTech: A graph neural network and feature selection perspective. (2025). *Computers, Materials & Continua*, 86(1), 1–40.
26. Graph neural networks for financial fraud detection: A review. (2025). *Frontiers of Computer Science*, 19, Article 199609.



27. Mahadevan, S. (2026). Governed Serverless Automation Ecosystem for AI-Driven Enterprise Integration and Sustainable Cloud Operations. *International Journal of Computer Information Systems and Industrial Management Applications*, 18(16s), 1561-1570.
28. Financial fraud detection using graph neural networks: A systematic review. (2024). *Expert Systems with Applications*, 240, Article 122156.
29. Advancing credit risk modelling with machine learning: A comprehensive review of the state-of-the-art. (2024). *Engineering Applications of Artificial Intelligence*, 137, Part A, Article 109082.
30. Explainable AI for credit assessment in banks. (2022). *Journal of Risk and Financial Management*, 15(12), Article 556.
31. An explainable AI framework for credit evaluation and analysis. (2024). *Applied Soft Computing*, 153, Article 111307.
32. Loan default predictability with explainable machine learning. (2024). *Finance Research Letters*, 60, Article 104867.
33. Graph learning-empowered financial fraud detection: Progress and future directions. (2025). *Intelligent Computing*, 4.
34. Yandamuri, U. S. Next-Generation Industry Platforms Integrated Systems For Predictive & Intelligent Operations. *Solaris Vantage*.
35. Network analytics for anti-money laundering—A systematic literature review and experimental evaluation. (2025). *INFORMS Journal on Data Science*.
36. Advances in continual graph learning for anti-money laundering systems: A comprehensive review. (2025). *WIREs Computational Statistics*, 17(3), Article e70040.
37. Detecting anomalies in FinTech: A graph neural network and feature selection perspective. (2025). *Computers, Materials & Continua*, 86(1), 1–40.
38. Explainable machine learning in credit risk management. (2021). *Computational Economics*, 57, 203–216.
39. Engineering Intelligent Cloud-Native Data Ecosystems for Predictive Decision-Making in Industry. (2026). *Journal of European Economic History*, 7(2), 68–88.
<https://doi.org/10.61336/jeeh/26-2-7> (Original work published as Engineering Intelligent Cloud-Native Data Ecosystems for Predictive Decision-Making in Industry)
40. A qualitatively analyzable two-stage ensemble model based on machine learning for credit risk early warning: Evidence from Chinese manufacturing companies. (2023). *Information Processing & Management*, 60(3), Article 103267.