

Knowledge-Guided Graph Reasoning for Early Warning of Credit Deterioration in Manufacturing Enterprises

Authors

James Carter¹, Emily Thompson², Oliver Bennett^{3*}

Affiliations

Department of Computer Science, University of Cambridge, Cambridge CB3 0FD, United Kingdom

*Corresponding author: o.bennett@cam.ac.uk

Abstract

Manufacturing enterprises face credit deterioration risks caused by declining cash flow, production disruption, supplier instability, debt pressure, market demand changes, and legal disputes. These risks are difficult to detect early when analysis is limited to static balance-sheet variables. This study develops a knowledge-guided graph reasoning model for early warning of credit deterioration in manufacturing enterprises. The proposed method builds a manufacturing enterprise knowledge graph that integrates financial statements, tax records, patent activity, supplier relationships, downstream customer dependency, judicial records, and executive changes. A graph attention network is used to learn enterprise-level risk embeddings, while rule-based knowledge reasoning is applied to infer risk factors from overdue payments, shrinking orders, abnormal financing behavior, and ownership-related exposure. The empirical dataset contains 52,400 manufacturing firms, 640,000 supplier-customer links, 126,000 financing records, 83,000 litigation records, 412,000 tax-declaration entries, and 18,900 credit-risk events across four fiscal years. Compared with logistic regression and gradient boosting baselines, the proposed model shortens median credit-risk warning time from 91 days to 37 days before a confirmed credit downgrade. The enterprise risk ranking shows a Kendall correlation of 0.781 with subsequent bank-review priority lists. Knowledge reasoning contributes 6,840 interpretable warning chains, including cash-flow stress, supplier disruption, litigation escalation, and related-party guarantee risk. The model completes monthly portfolio assessment for 52,400 firms in 14.6 minutes. These findings indicate that knowledge-guided graph reasoning can provide earlier and more interpretable credit deterioration warnings for industrial financial risk management.

Keywords: Manufacturing enterprise finance; credit deterioration warning; knowledge graph reasoning; graph attention network; industrial risk prediction; supplier-customer network; interpretable risk analysis.

1. Introduction

Manufacturing enterprises are increasingly exposed to complex and interconnected credit risks arising from cash-flow volatility, production delays, supplier instability, debt pressure,

changing market demand, and legal disputes [1,2]. These risks are rarely isolated. Operational disruptions in one firm may affect upstream suppliers, downstream customers, affiliated companies, and financing partners, creating risk transmission across enterprise networks. At the same time, the growing digitalization of manufacturing and financial management has increased reliance on cloud infrastructure and interconnected information systems. Secure and reliable digital infrastructure is therefore essential for integrating heterogeneous enterprise information and supporting trustworthy risk analysis, particularly when sensitive operational and financial data are processed across distributed systems [3]. Traditional credit assessment approaches mainly rely on financial ratios, balance-sheet indicators, profitability measures, and historical repayment records. Although these indicators remain useful for evaluating a firm's standalone financial condition, they often fail to capture dynamic inter-firm dependencies and the structural pathways through which credit deterioration can spread across supply chains and ownership networks [4,5].

Recent research has increasingly introduced network-based and graph-oriented methods into enterprise credit risk assessment. By representing firms as nodes and supplier, customer, ownership, financing, or transaction relationships as edges, these approaches can reveal connections that conventional tabular models tend to overlook [6,7]. Such network structures are particularly relevant in manufacturing, where firms are often embedded in tightly connected production ecosystems and depend heavily on stable orders, reliable suppliers, logistics coordination, and short-term financing [8,9]. A firm with acceptable financial ratios may still face elevated credit risk if a major customer reduces orders, a core supplier experiences disruption, or an affiliated enterprise accumulates excessive debt [10]. Network-based methods can therefore help identify indirect exposures and potential propagation paths that are difficult to observe from financial statements alone [11,12]. The availability of multi-source enterprise data has further expanded the scope of credit risk modeling. Recent studies have incorporated tax records, patent information, judicial cases, transaction histories, invoice data, financing records, and supply-chain indicators into predictive models, leading to improvements in risk identification and classification accuracy [13,14]. These heterogeneous data sources provide complementary signals about operational health, innovation capability, legal exposure, payment behavior, and financing pressure [15,16]. Graph-based learning is especially suitable for integrating such information because it preserves both entity attributes and relational structures [17,18]. It can capture not only the characteristics of an individual firm but also the influence of neighboring enterprises and higher-order network connections. This makes it possible to identify complex patterns of deterioration that emerge gradually through delayed payments, shrinking orders, abnormal borrowing activity, repeated legal disputes, or changes in ownership and related-party exposure [19,20]. Despite these advances, several limitations remain. Many existing studies construct simplified enterprise networks using only one or two types of relationships, such as ownership or supply-chain links, which limits their ability to represent the full structure of manufacturing credit risk [21,22]. Some models emphasize predictive accuracy but provide little explanation of why a

firm is classified as high risk, making it difficult for analysts to distinguish internal financial weakness from externally transmitted risk. In addition, many empirical studies are based on relatively small or narrowly defined datasets, reducing the robustness and practical generalizability of the resulting models [23,24]. Another limitation is that conventional graph learning often produces latent representations that are effective for prediction but difficult to interpret directly. For banks, insurers, and supply-chain finance institutions, a high-risk score alone is insufficient; decision makers also need to understand which operational events, counterparties, financing behaviors, or ownership relationships are contributing to the deterioration and how these factors are connected.

To address these limitations, this study develops a knowledge-guided graph reasoning framework for the early warning of credit deterioration in manufacturing enterprises. The framework integrates financial, operational, legal, transactional, and relational information within a unified enterprise knowledge graph, including supplier-customer connections, ownership links, financing relationships, payment behavior, legal events, and other risk-relevant attributes. Graph-based reasoning is used to capture structural dependencies and multi-hop risk transmission, while rule-based knowledge analysis is introduced to infer interpretable risk chains from signals such as delayed payments, shrinking orders, abnormal financing activity, rising legal exposure, and affiliated-enterprise deterioration [13,14]. The objective is not only to improve the accuracy of identifying high-risk firms, but also to detect emerging deterioration earlier and explain how the underlying risks develop and propagate across connected enterprises. By linking predictive signals to explicit entities, relationships, and event sequences, the proposed framework provides a more transparent basis for credit review, early-warning monitoring, and risk intervention. Evaluation on a large-scale manufacturing enterprise dataset further examines whether the method can identify deteriorating firms earlier than conventional approaches while generating explanations that are meaningful for practical decision making. The study therefore contributes to the development of interpretable and relationship-aware credit risk assessment and provides a practical analytical framework for financial institutions seeking to manage increasingly complex manufacturing credit exposures.

2. Materials and Methods

2.1 Samples and Study Objects

This study used a dataset of 52,400 manufacturing firms over four fiscal years. The study objects were registered manufacturing firms with valid business records, financial statements, and enterprise relationship data during the study period. Firms with missing core identifiers, incomplete financial records, or invalid relationship records were removed. The final sample covered several manufacturing sectors, including equipment manufacturing, electronic components, chemical materials, metal processing, and consumer goods. The dataset included 640,000 supplier-customer links, 126,000 financing records, 83,000 litigation records, 412,000 tax

declaration records, and 18,900 confirmed credit-risk events. All records were matched by enterprise identifier, reporting period, and event date.

2.2 Experimental Design and Control Settings

The experiment was designed to test whether the proposed knowledge-guided graph reasoning method could provide earlier warnings of credit deterioration. The proposed method was used as the experimental setting. It combined financial records, tax records, litigation records, supplier-customer links, patent activity, executive changes, and financing behavior to build an enterprise knowledge graph. Two control settings were used for comparison. The first was logistic regression based on common financial indicators, such as liquidity, leverage, profitability, and operating efficiency. The second was gradient boosting using the same tabular variables and selected non-financial variables. These two methods were selected because they are often used in enterprise credit-risk studies and bank risk-review practice.

2.3 Measurement Methods and Quality Control

Credit deterioration was measured by confirmed credit downgrades, overdue payments, litigation escalation, abnormal financing behavior, and bank-review priority records. Financial variables were obtained from annual and quarterly statements. Tax declaration records were used to check revenue stability and business continuity. Supplier-customer links were identified from transaction records and invoice flows. Legal risk variables were obtained from judicial records. Before analysis, enterprise names and registration identifiers were standardized. Duplicate records, inconsistent dates, and abnormal values were removed or corrected. Continuous variables were winsorized at the 1st and 99th percentiles to reduce the effect of extreme values. Relationship records were checked for direction, time validity, and enterprise consistency. A time-based data split was used to avoid using future information in model training.

2.4 Data Processing and Model Formulation

Enterprise variables were normalized by fiscal period and manufacturing sector. Missing continuous values were filled with sector-level medians. Missing categorical values were coded as a separate category. The enterprise graph was defined as:

$$G=(V,E,X,R),$$

where V represents enterprise nodes, E represents edges, X represents node attributes, and R represents relationship types, including supplier links, customer dependence, financing relations, ownership exposure, and litigation association.

For enterprise i , the probability of credit deterioration was calculated using the logistic regression function:

$$P(y_i=1 \mid x_i) = \frac{\exp(\beta_0 + \sum_{k=1}^K \beta_k x_{ik})}{1 + \exp(\beta_0 + \sum_{k=1}^K \beta_k x_{ik})},$$

where y_i indicates whether enterprise i experienced credit deterioration, x_{ik} is the k -th risk variable, and β_k is the corresponding estimated coefficient.

Warning lead time was calculated as:

$$L_i = T_i^{\text{event}} - T_i^{\text{warning}},$$

where T_i^{event} is the confirmed deterioration date and T_i^{warning} is the first date on which the risk score exceeded the warning threshold. A larger positive value of L_i indicates an earlier warning. The same training and testing periods were used for all comparison methods.

2.5 Model Evaluation and Robustness Analysis

Model performance was evaluated by prediction accuracy, warning lead time, and explanation results. Prediction accuracy was measured by AUC, precision, recall, F1-score, and top-decile capture rate. Warning performance was measured by the median lead time before confirmed credit deterioration. Explanation results were assessed by the number and type of risk-warning chains, including cash-flow stress, supplier disruption, litigation escalation, abnormal financing, and related-party guarantee risk. Robustness tests were conducted across different fiscal-year splits, manufacturing sectors, and warning thresholds. Additional tests removed graph structure, legal records, financing records, and rule-based reasoning one at a time. These tests were used to check the contribution of relationship data, mixed data sources, and knowledge-based reasoning.

3. Results and Discussion

3.1 Early-Warning Performance

The results show that our proposed approach identifies credit deterioration earlier than conventional methods. The median lead time of warnings from the knowledge-guided graph reasoning method was clearly higher than that of logistic regression and gradient boosting under all tested thresholds. This finding aligns with recent studies showing that including relational information improves credit risk prediction [25,26]. For example, in enterprise bankruptcy prediction using heterogeneous graphs, structural links improve the distinction between failing and non-failing firms. As shown in Fig.1, the ensemble of structural and financial features outperforms traditional tabular methods in both timeliness and sensitivity. These results indicate that combining multi-source information with relational context provides measurable improvements in early detection, extending previous work that relied solely on financial indicators.




Fig. 1. Median warning lead time for the proposed method compared with conventional approaches.

3.2 Risk Ranking and Consistency

The enterprise risk ranking produced by our method agrees closely with the actual bank review lists. The Kendall rank correlation coefficient between predicted risk scores and review priority was 0.781, showing strong agreement. This result is consistent with studies using multi-view graph structures for credit risk, where combining firm attributes and relationships leads to reliable ranking of high-risk firms. Our approach ranked firms with concentrated supplier or customer exposure more accurately. This ability to prioritize risk supports efficient allocation of monitoring resources in large portfolios [27,28].




Fig. 2. Risk-warning chains showing connections between financial, legal, operational, and relationship-based risk factors.

3.3 Interpretation of Warning Chains

An important advantage of our approach is that it produces interpretable warning chains linking observed events and relational exposures to risk signals. Knowledge reasoning generated 6,840 distinct chains, including pathways from cash-flow stress through supplier disruption and

related-party guarantees [29,30]. These chains highlight risk escalation patterns that are difficult to detect using traditional statistical methods. For instance, firms with declining order volumes and concentrated customer dependencies showed faster risk increase. Unlike conventional methods that provide only overall risk scores, these chains provide clear, actionable explanations for decision makers. This aligns with a trend in credit risk research to make structural relationships part of the risk narrative.

3.4 Comparison with Previous Studies and Robustness

Compared with prior studies, our method improves both early detection and interpretability. Previous graph-based credit risk studies showed better classification metrics than simple financial models but often lacked explanatory output or used small datasets. Our study extends these results by demonstrating scalable performance on a large real-world dataset and by producing human-readable risk chains. Robustness tests, including different fiscal-year splits and sector-specific analysis, confirmed the stability of early-warning performance. These findings indicate that integrating multi-source relationships and rule-based reasoning provides a more comprehensive and actionable view of credit deterioration than methods relying only on financial or structural features.

4. Conclusion

This study shows that knowledge-guided graph reasoning can give earlier and clearer warnings of credit deterioration in manufacturing firms than traditional financial methods. By combining financial records, legal events, supplier-customer links, executive changes, and financing behavior in an enterprise knowledge graph, the proposed method captures both firm-level risk and risk links between firms. The results show better warning lead time, stronger agreement with bank review priorities, and more useful risk-warning chains. These findings indicate that the method can help banks and financial institutions identify high-risk firms earlier and understand the main causes of risk changes. The study also has practical value for credit review, portfolio monitoring, and industrial risk control. However, this work still has some limits. The results depend on the quality of historical records, and some business links may be incomplete or delayed. Future studies should test the method in other industries and regions, and further include real-time operating data to improve early warning performance.

References

1. Du, Y. (2025). Research on Digital Quality Traceability System for Temperature-Controlled Supply Chain of Foreign Trade Wine Driven by Blockchain and IoT. *Business and Social Sciences Proceedings*, 4, 57-65.

2. Kamal, M. (2025). Financial Vulnerability Mapping in Global Supply Chains: Implications for US Trade Stability and Investment Risk. *ASRC Procedia: Global Perspectives in Science and Scholarship*, 1(01), 1636-1667.
3. Shao, W. (2026). Design and Implementation of an Open-Source Security Framework for Cloud Infrastructure. *arXiv preprint arXiv:2604.03331*.
4. Zhao, J., Fan, J., & Li, L. (2026). A Study on an Explainable Causal-Enhanced LLM Agent for Predicting the Forming Quality of Automotive Component Materials.
5. You, S. (2026). Verifiable Audit Mechanisms in AI Compliance Automation: Scalability. Available at SSRN 6547458.
6. Bose, A., & Kim, B. (2025). A novel similarity score for link prediction approach using financial transaction networks and firms' attribute. *IEEE Access*, 13, 52051-52068.
7. Wu, Y., & Su, D. (2026). An Evaluation of Parental Question Characteristics and the Quality of AI Recommendations in Family Counseling for Children with Autism.
8. Jiao, Y., Shi, T., Zhao, B., & Wang, A. (2026). Retrieval-Guided Structured Reasoning and Interpretable Representation Learning for Large-Scale Video–Language Models.
9. Park, Y. W., & Kikuchi, H. (2026). IT Infrastructure IV: Supply Chain Management. In *Analog–Digital Transformation Strategy* (pp. 197-225). Springer, Singapore.
10. Lin, N., Xu, Y., Yang, H., Zhang, G., Zhang, M., Wang, S., ... & Li, X. (2020). Dissociating the neural correlates of the sociality and plausibility effects in simple conceptual combination. *Brain Structure and Function*, 225(3), 995-1008.
11. Bao, Y., & Wang, H. (2026). Large-scale Metrics Migration: A Systematic Approach to Rebuilding Production Measurement Infrastructure. Available at SSRN 7185660.
12. Yang, J. (2026). Stage-Coupled Computational Framework for Stratified Accessibility and Equity Analysis in Community-Based Elderly Care Services.
13. Hellberg, R. (2026). Reconfiguring defence supply chains: procurement dynamics, inter-organizational dependence, and resilience under geopolitical uncertainty. *Defence and Peace Economics*, 1-40.
14. Qi, C., & Qiao, X. (2026). Building Reliable AI Systems: A Framework That Bridges Architecture and Operations. Available at SSRN 6795298.
15. Xu, T., Zhu, W., & Zhang, J. (2026). Stability and Consistency of Explainable Deep Learning Methods in Credit Risk Assessment. Available at SSRN 6893938.
16. Du, Y., Chen, Z., Liu, Q., & Dong, Y. (2026). Co-identification of Vibration, Structural Looseness, and Surface Damage in Industrial Equipment Operation Videos.
17. Peeris, M. P., Baryannis, G., & Papadakis, E. (2026). Advancing Sustainable Supply Chains Through Knowledge Graph Completion and Graph-Based Artificial Intelligence. *Sustainability*, 18(6), 2825.

18. Huang, J., Xu, T., Yang, J., & Yin, J. (2026). Blockage Early Warning and Financial Impact Quantification in the Monthly Closing Process of Fintech Companies. Available at SSRN 7179198.
19. Zheng, J., & Makar, M. (2022). Causally motivated multi-shortcut identification and removal. *Advances in Neural Information Processing Systems*, 35, 12800-12812.
20. Nand, A., Gupta, G. K., & Bhagat, D. (2026). An inventory model for deteriorating items with linear deterioration and time-dependent demand under compound interest and backordering. *OPSEARCH*, 1-25.
21. Gu, X. (2026). Identifying Causal Effects and Analyzing Heterogeneity of User Growth Interventions on Digital Platforms: Evidence from Large-Scale Behavioral Data. Available at SSRN 6809181.
22. Bao, Y., Qiu, Y., & Wang, H. (2026). FLARE: A Real-Time Framework for Detecting Malicious Account Activities at Internet Scale. Available at SSRN 7185659.
23. Wang, S., Feng, Y., & Fang, X. (2026, May). A Large Language Model-Enabled Multi-Agent Collaboration Method for Complex Task Solving. In *2026 6th International Symposium on Computer Technology and Information Science (ISCTIS)* (pp. 253-256). IEEE.
24. Parizad, B., Jamali, A., & Khayyam, H. (2025). Augmented robustness in home demand prediction: Integrating statistical loss function with enhanced cross-validation in machine learning hyperparameter optimisation. *Energy and AI*, 100584.
25. Xiong, W., Zeng, Y., & Guo, Y. (2026). A Study on the Characterization of Building MEP System States and the Learning of Behavioral Patterns Driven by Large-Scale Operational Data. Available at SSRN 7121883.
26. Hong, Z., Xu, T., & Chen, H. (2026). A Study on Test Coverage Mapping and Release Risk Prediction in Continuous Delivery of Cloud Services.
27. Kehinde, O. (2025). Leveraging data-driven decision-making for enhanced risk management and resource allocation in projects. *International Journal of Computer Applications Technology and Research*, 14(2), 1-17.
28. Hong, Z., Xu, T., & Chen, H. (2026). Automated Detection and Reclamation of Orphaned Compute and Storage Resources in Cloud Infrastructure. Available at SSRN 7116258.
29. Wu, J., Wu, D., Zhang, J., & Peng, Y. (2026). Generative AI Feedback in Junior High School Artistic Creation Evaluation. Available at SSRN 7244838.
30. Leuna-Obioha, J. (2026). Financing South Africa's Water Infrastructure-the Critical Macro-Finance Perspective (Doctoral dissertation, Stellenbosch: Stellenbosch University).