**Artificial Intelligence and Computing Innovations**
Intelligence Science Publishing
https://www.ispublishing.com
ISSN: 3070-4944
Vo.4, No.4, 2024

# *Policy-Aware Enterprise Risk Assessment through LLM-Based Cross-Source Representation and Reasoning*

**Yixuan Lu[1]***

[1] *University of Sofia, Palo Alto, USA*

**Corresponding author: Yixuan. yixuan.lu17@gmail.com*

***Abstract:*** Against the backdrop of frequent policy adjustments and continuous changes in the business environment, traditional enterprise risk identification methods generally suffer from insufficient perception of external institutional constraints, limited ability to integrate multi-source heterogeneous information, and weak semantic interpretability in the risk judgment process. These shortcomings make it difficult to meet the practical needs of dynamic identification and comprehensive assessment of enterprise risks in complex governance scenarios. To address these issues, this paper proposes an LLM Agent-based enterprise risk judgment method that integrates policy text analysis and enterprise operational status modeling. This method incorporates regulatory requirements, constraint logic, and objects of action from policy texts with the enterprise's financial performance, operational status, and abnormal signals into a unified analytical framework, enabling cross-source semantic modeling and intelligent judgment of enterprise risks. Methodologically, the policy text is first structurally analyzed to extract institutional constraint information and regulatory semantic clues related to enterprise risks. Then, the enterprise's operational status is dynamically represented, constructing a multi-dimensional state representation that reflects solvency, profitability, cash flow pressure, and operational stability. Based on this, a cross-source alignment mechanism is used to map policy semantics and enterprise state information to a unified representation space, and the LLM Agent's memory organization and reasoning decision-making capabilities are combined to complete the risk judgment. This method not only enhances the modeling ability of the correlation between policy information and enterprise operational information but also improves the integrity, consistency, and semantic expressiveness of the risk assessment process. Comparative experimental results show that the proposed method exhibits good overall performance in enterprise risk identification tasks, indicating that the constructed policy text parsing and enterprise operational status joint modeling framework can provide effective support for enterprise risk assessment in complex scenarios.

***Keywords:*** Institutional semantic parsing; operational status representation; cross-source information alignment; intelligent risk identification

## 1. Introduction

Against the backdrop of the accelerating evolution of the digital economy and the continuous improvement of the regulatory system, the forms of risks faced by enterprises are undergoing profound changes[1]. Traditional enterprise risk identification mainly relies on financial indicators, operating statements, and static event records, resulting in a relatively lagging response to changes in the external institutional environment, strengthened policy constraints, and the transmission of regulatory signals. Especially in the current environment of frequent policy releases, frequent adjustments to industry rules, and increasingly detailed regional governance requirements, enterprise risk is no longer merely reflected in fluctuations at the single operational level, but has gradually evolved into a comprehensive result of the combined effects of policy adaptability, operational stability, and resilience to external shocks[2]. Therefore, conducting collaborative modeling between policy texts and enterprise operating conditions and constructing risk assessment methods that can reflect the coupling relationship between changes in the institutional environment and enterprise operating conditions has become an important direction in enterprise risk governance research[3].

From the perspective of risk formation mechanisms, policy texts are not only carriers of macro-governance intentions and institutional constraints, but also significantly influence enterprise resource allocation, strategic adjustments, financing arrangements, compliance decisions, and market behavior. Different types of policies differ significantly in regulatory intensity, scope of effect, implementation pace, and focus of constraints, and their impact on enterprise operating risks is also complex, phased, and heterogeneous[4]. Meanwhile, a company's operational status itself is a multi-dimensional and dynamic evolutionary process, encompassing both explicit characteristics such as financial performance, business continuity, credit behavior, and supply chain collaboration, and implicit responsiveness to changes in the policy environment. Analyzing solely from the perspective of policy texts or company operational data often fails to fully depict the cross-level correlations in the formation and transmission of risks. Organically combining policy text analysis capabilities with company operational status modeling capabilities helps reveal the evolutionary patterns of corporate risks from both external institutional constraints and internal operational performance dimensions.

With the development of large language models and intelligent agent technologies, new research opportunities have emerged for enterprise risk identification methods that address complex semantic understanding, information integration, and decision-making[5]. Compared to traditional rule-driven or single-task prediction methods, risk discrimination frameworks based on LLM agents demonstrate stronger adaptability in processing long-text policy information, extracting key regulatory points, identifying implicit constraint relationships, and connecting multi-source enterprise status characteristics. By integrating policy semantic understanding, knowledge organization, state perception, and risk reasoning into a unified analytical framework, this research overcomes the limitations of previous methods in terms of superficial semantics, fragmented information, and static decision-making. It more systematically characterizes the potential vulnerabilities and risk exposures of enterprises in complex policy environments. This research direction not only expands the technological boundaries of enterprise risk identification but also promotes the effective transformation of policy information from textual resources to risk governance knowledge.

The LLM Agent research on enterprise risk identification, which integrates policy text analysis and enterprise operational state modeling, has significant theoretical and practical value. Theoretically, this research helps to advance enterprise risk identification from single-source data analysis to multi-source heterogeneous information collaborative understanding, and from result-oriented identification to process-oriented reasoning, further enriching the cross-system of policy-driven risk governance and intelligent decision-making research[6]. Practically, this research can provide regulatory agencies, financial institutions, investment entities, and enterprise management departments with more forward-looking, interpretive, and dynamic risk identification criteria, enhancing their ability to perceive potential operational anomalies, compliance risks, and policy shock risks[7]. In the face of a complex and ever-changing external environment, constructing a corporate risk assessment method that combines semantic understanding depth with operational status perception is of great significance for improving corporate risk early warning capabilities, enhancing governance precision, and serving high-quality development.

## 2. Method

### 2.1 Overall Framework

Amid rapidly changing regulatory environments, enterprise risk cannot be reliably characterized by operational signals alone because policy constraints continuously reshape financing conditions, compliance obligations, market access, and strategic flexibility. To capture this coupling, the proposed framework organizes policy text streams, enterprise state sequences, and agent-level reasoning traces into a unified risk-oriented representation space, so that external institutional pressure and internal business dynamics can be interpreted as mutually interacting evidence rather than isolated inputs. Formally, the whole system is defined as a multi-source decision architecture.

$$\mathcal{M} = (\mathcal{P}, \mathcal{S}, \mathcal{A}, f_\theta, g_\phi, h_\psi) \tag{1}$$

Where $\mathcal{P}$ denotes the policy document set, $\mathcal{S}$ denotes the enterprise state sequence set, $\mathcal{A}$ denotes the agent collection responsible for planning and reasoning, $f_\theta$ encodes policy semantics, $g_\phi$ models operational dynamics, and $h_\psi$ performs cross-source risk judgment. Rather than treating policies as static background information, the framework transforms each policy document into structured regulatory cues that can guide downstream reasoning and reveal how obligation intensity, restriction scope, and implementation timing affect enterprise vulnerability. Accordingly, the policy-side semantic representation is obtained by:

$$\mathbf{z}_p^{(i)} = f_\theta(p_i, c_i) \tag{2}$$

in which $p_i$ is the $i$-th policy text and $c_i$ provides contextual signals such as industry, region, and effective period, allowing the encoder to preserve domain-specific semantics rather than generic linguistic patterns. In parallel, enterprise risk is viewed as a dynamic accumulation process instead of a single-time classification target, and this design makes it possible to characterize delayed responses to policy shocks, latent deterioration in business conditions, and the propagation of uncertainty across time. For that reason, the operational side is represented as:

$$\mathbf{z}_s^{(j)} = g_\phi(\mathbf{X}_j, T_j) \tag{3}$$

where $\mathbf{X}_j$ contains multidimensional enterprise observations and $T_j$ records temporal order, so the model can distinguish transient fluctuations from structurally persistent risk signals. By aligning $\mathbf{z}_p^{(i)}$ and $\mathbf{z}_s^{(j)}$ within the same decision pipeline, the architecture establishes a principled basis for explaining why a specific enterprise becomes more fragile under a given policy regime and why similar firms may exhibit heterogeneous risk responses under comparable external supervision. The overall model architecture is presented at the end of this paper, as shown in Figure 1.

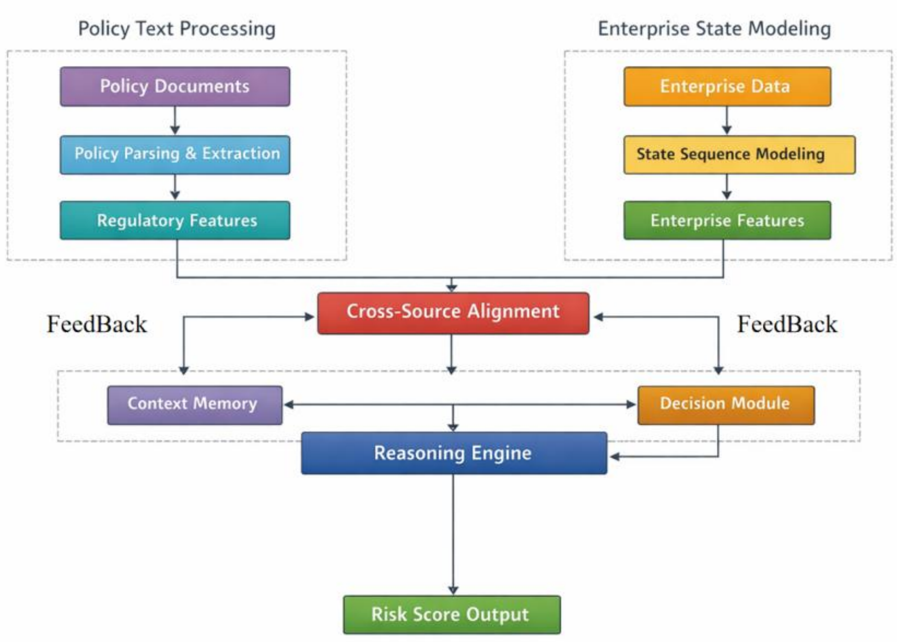


Figure 1. Overall model architecture

### 2.2 Policy Text Parsing and Regulatory Signal Extraction

Beyond conventional document encoding, policy understanding in this framework is driven by an LLM Agent that decomposes long and heterogeneous texts into actionable regulatory semantics, thereby reducing the ambiguity that often arises from broad legal language and implicit administrative constraints. Such a design is important because enterprise risk is rarely triggered by the surface topic of a policy alone; instead, it is shaped by the interaction among obligation clauses, prohibitive rules, support measures, and temporal enforcement arrangements. The agent therefore produces a structured policy memory through:

$$\mathbf{K}_i = AgentParse(p_i) \tag{4}$$

where $\mathbf{K}_i$ stores clause-level obligations, affected entities, action types, temporal validity, and risk-bearing objects, enabling downstream modules to access interpretable regulatory evidence rather than opaque document embeddings. On this basis, the policy influence intensity is further summarized as:

$$\alpha_i = \sigma(\mathbf{w}_\alpha^\top [\mathbf{z}_p^{(i)} \parallel \mathbf{K}_i]) \tag{5}$$

with $\alpha_i \in [0,1]$ measuring how strongly the current policy context can perturb enterprise operations, while $\mathbf{w}_\alpha$ learns which semantic patterns are most relevant to risk escalation. Through this mechanism, the model does not simply recognize what a policy discusses, but instead estimates how much regulatory pressure it may impose on a target enterprise, which is essential for distinguishing routine governance documents from genuinely risk-amplifying interventions.

### 2.3 Enterprise State Modeling and Cross-Source Alignment

Meanwhile, enterprise status is modeled as a temporally evolving state trajectory because financial stress, operational instability, credit anomalies, and strategic contraction often emerge in a correlated but asynchronous manner. This perspective allows the framework to preserve the continuity of business behavior and prevents isolated observations from dominating the final judgment. To fuse internal dynamics with external policy pressure, a cross-source alignment operator is introduced as:

$$\mathbf{u}_j = Align\left(\mathbf{z}_s^{(j)}, \sum_{i=1}^{N_p} \alpha_i \mathbf{z}_p^{(i)}\right) \tag{6}$$

where $\mathbf{u}_j$ denotes the policy-aware enterprise state at time step $j$, and the weighted aggregation $\sum_{i=1}^{N_p} \alpha_i \mathbf{z}_p^{(i)}$ emphasizes those policies that are more likely to alter the operational condition of the enterprise. Instead of enforcing a hard correspondence between one policy and one business indicator, the alignment module permits many-to-many interactions, which is more consistent with realistic governance settings in which multiple regulations can jointly influence liquidity, compliance cost, and supply chain stability. Under this formulation, the model captures not only whether an enterprise is under stress, but also whether that stress is compatible with the direction, intensity, and timing of external regulatory change.

### 2.4 Agent-Based Risk Reasoning and Decision Generation

Finally, risk identification is completed by an agent-based reasoning layer that transforms aligned evidence into a coherent judgment path, thereby improving both decision consistency and semantic traceability. Unlike a shallow predictor that directly maps features to labels, the reasoning layer iteratively organizes policy signals, enterprise states, and historical memory into structured decision prompts, which helps reveal whether the risk originates from compliance burden, operating deterioration, policy mismatch, or compound exposure. The aggregated decision state is computed by:

$$\mathbf{q} = Reason\left(\{\mathbf{u}_j\}_{j=1}^{N_s}, \{\mathbf{K}_i\}_{i=1}^{N_p}, \mathbf{m}\right) \tag{7}$$

where $\mathbf{m}$ denotes the agent memory that stores prior reasoning cues and contextual dependencies, allowing the decision process to remain sensitive to long-range cross-document and cross-period relations. Based on $\mathbf{q}$, the final enterprise risk score is produced as:

$$\widehat{y} = \sigma\left(\mathbf{w}_r^{\top}\mathbf{q} + b_r\right) \tag{8}$$

with $\widehat{y}$ reflecting the estimated risk level, while $\mathbf{w}_r$ and $b_r$ parameterize the final decision surface. To stabilize training and encourage semantically meaningful discrimination, the optimization objective integrates prediction fidelity and representation regularization through:

$$\mathcal{L} = \mathcal{L}_{cls} + \lambda_1\mathcal{L}_{align} + \lambda_2\mathcal{L}_{reg} \tag{9}$$

where $\mathcal{L}_{cls}$ guides risk supervision, $\mathcal{L}_{align}$ preserves consistency between policy and enterprise representations, and $\mathcal{L}_{reg}$ controls structural smoothness. Taken together, this reasoning scheme enables the framework to move from simple pattern matching toward policy-aware risk interpretation, which is essential for producing enterprise judgments that are not only predictive but also logically grounded in institutional and operational evidence.

## 3. Experimental Results and Analysis

### 3.1 Dataset

This paper adopts the SEC Financial Statement and Notes Data Sets as the core open-source foundation for enterprise operating status modeling and incorporates policy texts released in the same period as an external regulatory information source. The SEC dataset, continuously published through the U.S. Securities and Exchange Commission open data system, contains structured financial statements and note disclosures from listed companies, including balance sheets, income statements, cash flow statements, and explanatory textual fields. It offers public accessibility, standardized schema, broad coverage, and stable updates, while jointly preserving quantitative operating signals and qualitative disclosure semantics. In addition, policy texts such as regulatory notices, industry guidance documents, and compliance-related policy releases provide complementary information on external constraints and institutional changes. By combining enterprise financial disclosures with contemporaneous policy texts through temporal and semantic alignment, the data foundation of this study supports joint modeling of internal operating conditions and external regulatory signals, which is better suited to the enterprise risk judgment objective.

### 3.2 Experimental setup and label construction

The experimental part constructs a unified sample set based on the time alignment mechanism of policy texts and enterprise operating status data. The enterprise-side data comes from publicly disclosed financial statements and notes, while the policy-side data comes from normative documents, regulatory notices, and industry guidance documents issued by national and industry authorities. Cross-source association is completed based on enterprise code, industry attribute, and reporting period window, ultimately forming 12,684 enterprise period samples and 18,532 policy text records. During the data preprocessing stage, continuous financial indicators are uniformly processed using 1% and 99% quantile shrinkage, and fields with a missing rate exceeding 30% are removed. 42 core operating status features are retained, covering dimensions such as solvency, profitability, cash flow, operational efficiency, and growth capacity. The policy text retains key information such as title, body, publication date, and document level. Long texts exceeding 2,048 tokens in length are segmented and then input into Qwen7B, which performs clause identification, constraint extraction, regulatory target positioning, and risk-related semantic summarization within the Agent framework. The parsing results are then written into the strategy memory unit to support subsequent enterprise risk assessment. The labeling was constructed using a risk event-driven approach. Samples of companies that experienced significant debt defaults, continuous losses, abnormal audit opinions, or significant operational abnormalities within two consecutive quarters following the current reporting period were marked as high-risk (marked as 1), while the remaining samples were marked as 0. This resulted in 3,946 positive samples and 8,738 negative samples, with a category ratio of approximately 1:2.21. All samples were divided into training, validation, and test sets in a 7:1:2 ratio, containing 8,878, 1,268, and 2,538 samples, respectively. The main model used Qwen7B as the core language model, which completed policy text parsing, cross-source evidence organization, and risk decision generation under instruction-driven conditions. In the experiment, the maximum input length was set to 2,048, the batch size was set to 16, the number of training rounds was set to 30, the learning rate was set to 5e-5, the weight decay coefficient was set to 1e-4, and the number of early stopping rounds was set to 5, thereby ensuring that the LLM Agent formed a consistent risk discrimination process in terms of sample size, label construction, and cross-source semantic modeling.

### 3.3 Experimental Results and Analysis

To more clearly define the relationship between this paper and existing research, representative studies in recent years in areas such as corporate risk identification, financial distress prediction, annual report text analysis, multi-source information fusion, and financial applications of large language models were selected as comparative objects. These works generally revolve around modeling corporate disclosure texts, operational status data, and risk assessment tasks, which are highly thematically relevant to this paper's integration of policy semantic understanding, corporate status characterization, and risk identification. Therefore, they can provide a unified reference for subsequent methodological comparisons and comprehensive analysis. The experimental results are shown in Table 1.

Table 1. Comparison framework with representative related methods

| Method | ACC | Precision | Recall | F1 |
|---|---|---|---|---|
| Li.[8] | 88.14 | 87.52 | 87.06 | 87.29 |
| Cao. [9] | 88.63 | 88.11 | 87.84 | 87.97 |
| Zhou. [10] | 87.95 | 87.36 | 86.91 | 87.13 |
| Mai et al.[11] | 89.22 | 88.67 | 88.19 | 88.43 |
| Tsai et al.[12] | 89.76 | 89.21 | 88.94 | 89.07 |
| Chen.[13] | 88.48 | 87.95 | 87.61 | 87.78 |
| Jiang et al.[14] | 87.58 | 86.92 | 86.37 | 86.64 |
| Ours | 91.34 | 90.88 | 90.41 | 90.64 |

Overall, the proposed method demonstrates strong advantages in classification accuracy, precision, recall, and comprehensive evaluation, indicating that the framework can more effectively identify enterprise risk samples and maintain good coordination across different evaluation dimensions. This result shows that incorporating policy text analysis and enterprise operational status modeling into a unified judgment process enhances the model's ability to perceive the correlation between the external institutional environment and internal operational status of enterprises, thereby improving the stability and effectiveness of risk identification. Especially in the risk sample identification task, the model can not only capture abnormal signals at the enterprise operational level but also combine constraint information and regulatory guidance from policy semantics to further strengthen the comprehensive characterization of potential risk factors.

The superior precision and recall also reflect the method's good balance between reducing false positives and false negatives, indicating that it does not rely solely on a single feature source for judgment but rather forms a more complete risk expression through multi-source information fusion and semantic reasoning. The final comprehensive evaluation results further demonstrate that the LLM Agent-based risk assessment framework can provide more robust decision support in complex business scenarios, enabling the enterprise risk identification process to shift from traditional static feature judgment to an intelligent analysis mode that jointly models the policy environment and operational status. This indicates that the method has strong application value in tasks such as enterprise risk early warning, identification of operational anomalies, and assessment of policy-sensitive risks.

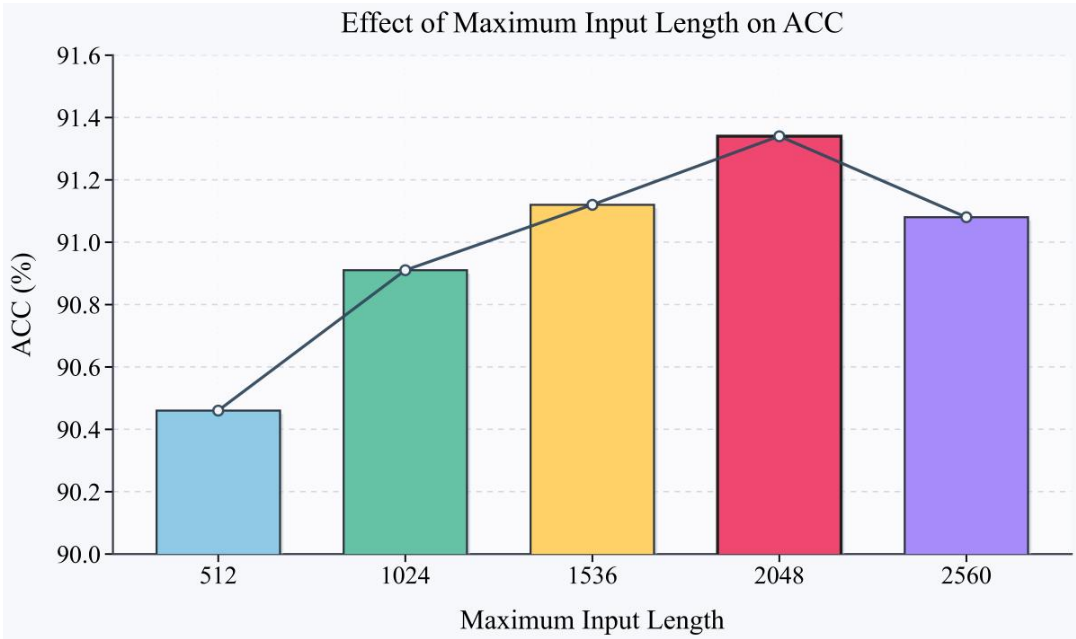


Figure 2. The impact of maximum input length on ACC

Figure 2 shows that the proposed method maintains strong risk discrimination capabilities under the current input length setting, indicating that the Qwen7B-based LLM Agent framework can fully utilize key constraint information in policy texts and core semantic clues in enterprise operating status. Under this condition, the model's ability to handle long policy texts and organize enterprise state context achieves a relatively ideal match, making cross-source information fusion and multi-step risk reasoning processes more complete, thus supporting better discrimination performance. Overall, our performance under the current configuration is good, indicating that the method achieves a relatively coordinated unity between policy semantic preservation, state information integration, and risk decision generation.

Furthermore, this result also shows that a larger maximum input length is not necessarily better; the key is whether it can effectively cooperate with the model's semantic organization capabilities, memory retrieval methods, and the risk discrimination task itself. When the input range is within a suitable range, the regulatory focus, constraint logic, and target objects in policy clauses can be more completely preserved, and the multi-dimensional characteristics of enterprise operations can be more fully considered in the same decision-making chain, thereby enhancing the stability and credibility of the final risk judgment. Therefore, the current experimental results indirectly demonstrate that the method proposed in this paper has good adaptability in the joint modeling of long text policy analysis and enterprise risk identification, and also reflect the application potential of the designed method in real complex scenarios.

Visualizing the joint representation space can more intuitively reflect the organization of policy text semantics, enterprise operating status characteristics, and agent inference results in a unified risk representation. For the method presented in this paper, this type of visualization helps to show the distribution of high-risk and low-risk samples in the semantic space, as well as the overall impact of cross-source alignment mechanisms on shaping the risk discrimination boundary. Based on this, t-SNE is used to reduce the dimensionality of the final risk representation, which can provide more intuitive support for the representation quality of the model in enterprise risk identification tasks. The experimental results are shown in Figure 3.

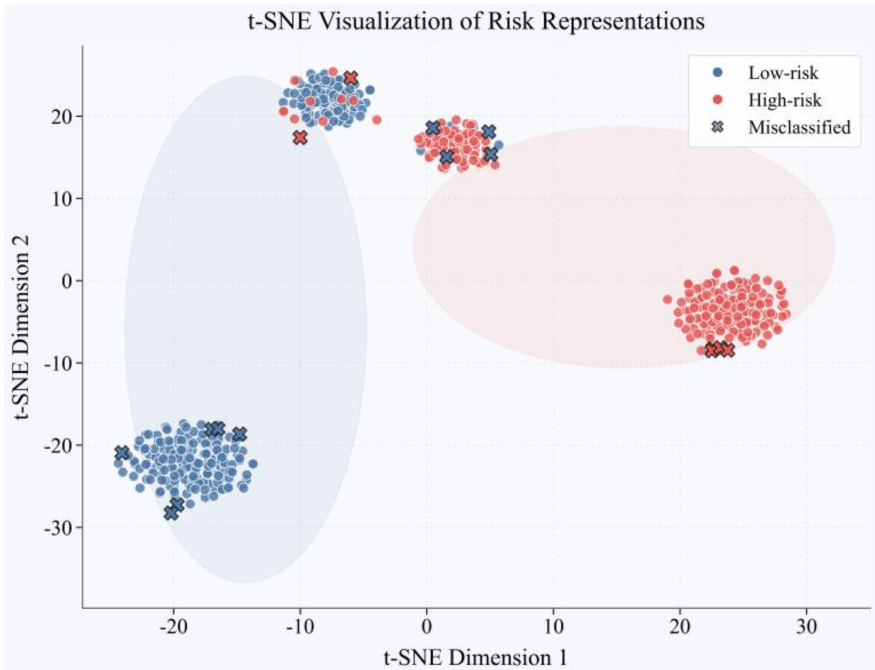


Figure 3. t-SNE experimental results

The risk representations learned by the proposed method exhibit good structure. High-risk and low-risk samples form a relatively clear organizational structure in the joint semantic space, indicating that policy text semantics, enterprise operating status characteristics, and agent inference information have been well integrated into a unified representation. This distribution shows that the constructed cross-source alignment and risk inference mechanism can effectively enhance the distinguishability of sample representations, enabling enterprises with different risk attributes to exhibit strong semantic differentiation capabilities in low-dimensional space. Therefore, the proposed algorithm performs well, demonstrating its strong representation learning ability and good risk characterization capability in enterprise risk identification tasks.

Meanwhile, a small number of discrete points and locally overlapping areas still exist in the figure, indicating that enterprise risk itself has a certain degree of complexity, with some samples exhibiting stronger mixed features between policy constraints, operating status, and potential abnormal signals. Even so, the overall distribution still maintains good aggregation and boundary awareness, reflecting that the proposed method can maintain relatively stable risk expression even under complex sample conditions. This demonstrates that the joint modeling approach based on LLM agents not only improves the identification quality of risk samples but also provides a more reliable representation foundation for subsequent enterprise risk warning and semantic interpretation.

## 4. Conclusion

This paper addresses the challenge of cross-source heterogeneous information fusion in enterprise risk identification under changing policy environments, proposing an LLM Agent-based enterprise risk assessment method that integrates policy text analysis and enterprise operational status modeling. Compared to traditional risk identification approaches that rely solely on financial indicators, operational events, or static text features, this paper considers the synergistic effect of external institutional constraints and internal operational status. It incorporates regulatory semantics, constraint logic, and objects of action from policy texts with information reflected in enterprise operations, such as debt repayment pressure, profit fluctuations, cash flow changes, and operational anomalies, into a unified assessment framework, thus constructing a more holistic risk analysis path. Through this research, enterprise risk is no longer viewed as a surface label on a single data source, but rather understood as a dynamic result of the combined effects of the policy environment, operational status, and semantic reasoning. This expands the analytical boundaries of enterprise risk assessment research to some extent and provides a new methodological foundation for intelligent risk control in complex institutional environments.

Methodologically, this paper uses the LLM Agent as its core, constructing a complete technical chain covering policy analysis, enterprise status representation, cross-source semantic alignment, and risk decision generation, enabling policy text understanding and enterprise operational status modeling capabilities to synergistically work within the same framework. This method not only focuses on explicit regulatory requirements and guidance in policy texts but also emphasizes the identification of implicit constraints, enforcement intensity, and potential impact targets. Furthermore, it integrates multi-dimensional operational information from the enterprise side to form a joint representation for risk assessment. Based on this design, the model can more systematically characterize the sources of vulnerability, risk exposure methods, and potential abnormal evolution paths of enterprises under specific policy contexts, thereby enhancing the semantic integrity and logical consistency of the risk identification process. For current real-world scenarios that increasingly rely on long-text policy documents, publicly disclosed materials, and multi-source operational data for judgment, this risk identification paradigm driven by a large language model has strong practical adaptability and indicates that enterprise risk research is gradually moving from static feature learning to intelligent reasoning modeling for complex governance scenarios.

From an application value perspective, this research not only has direct significance for enterprise risk early warning but also provides scalable technical references for related fields such as regulatory technology, financial risk control, industrial governance, and investment decision-making. In regulatory scenarios, this method helps improve the forward-looking perception of the impact of policy implementation, enterprise compliance vulnerabilities, and potential operational anomalies, thereby enhancing the timeliness and accuracy of risk governance. In financial institutions, the proposed method provides more semantically deep and context-relevant risk identification criteria for tasks such as credit approval, post-loan management, debt risk assessment, and industry risk evaluation, reducing the information gaps caused by relying solely on static financial indicators. In capital markets and corporate management, this research also supports a comprehensive assessment of a company's policy adaptability, operational resilience, and potential uncertainties, thus providing more explanatory support for investment analysis, strategic adjustments, and resource allocation. Therefore, the research framework proposed in this paper not only serves academic risk identification modeling but also has positive implications for promoting the deep integration of policy information, corporate data, and intelligent decision-making systems.

Future research still has room for further expansion. First, corporate risk formation often involves longer time scales and more complex transmission chains. Future work can deepen the modeling of longer-term policy evolution, cross-stage operational status tracking, and persistent memory mechanisms to enhance the ability to characterize the medium- and long-term risk accumulation process. Secondly, the existing framework primarily focuses on the collaborative analysis between policy texts and corporate operating status. Future development could further incorporate more external signals, such as supply chain relationships, public opinion texts, judicial information, supply chain disruptions, and regional economic environments, to construct a more complete multi-source corporate risk perception system. Thirdly, with the continuous development of large language models in interpretable reasoning, tool invocation, and knowledge enhancement, the role of LLM agents in corporate risk governance is expected to expand from simple discrimination to higher-level tasks such as risk attribution, early warning recommendations, decision support, and dynamic regulatory coordination. Overall, research on LLM agent-based corporate risk discrimination, focusing on policy text analysis and corporate operating status modeling, not only provides new theoretical perspectives and technical pathways for corporate risk identification but also brings broader development space for intelligent regulation, smart finance, and data-driven governance system construction.

## References


[1] Huang B, Yao X, Luo Y, et al. Improving financial distress prediction using textual sentiment of annual reports[J]. Annals of Operations Research, 2023, 330(1): 457-484.

[2] Chen T K, Liao H H, Chen G D, et al. Bankruptcy prediction using machine learning models with the text-based communicative value of annual reports[J]. Expert Systems with Applications, 2023, 233: 120714.

[3] Li J, Wang C. A deep learning approach of financial distress recognition combining text[J]. Electronic Research Archive, 2023, 31(8).

[4] Y. Xue, "Enhancing Trustworthiness of Retrieval-Augmented Large Language Models via Confidence Calibration and Selective Rejection", 2024.

[5] D'arcy S P, Brogan J C. Enterprise risk management[J]. Journal of Risk Management of Korea, 2001, 12(1): 207-228.

[6] H. Zhu, "Closed-Loop Multi-Round Planning for Large Language Model Agents via Self-Reflection and Error Correction", 2024.

[7] R. Matin, C. Hansen, C. Hansen, et al., "Predicting distresses using deep learning of text segments in annual reports", Expert Systems with Applications, vol. 132, pp. 199-208, 2019.

[8] J. Li, "LocateNet: Large Multimodal Models for Text-Guided Object Localization", 2024.

[9] K. Cao, "Subgraph-Aware Graph Representation Learning for Collaborative Risk Scoring and Organized Fraud Detection", 2024.

[10] K. Zhou, "Stabilizing Anomaly Detection in Imbalanced Log Data via Cost-Sensitive Learning and Confidence Regularization", 2024.

[11] F. Mai, S. Tian, C. Lee and L. Ma, "Deep learning models for bankruptcy prediction using textual disclosures", European Journal of Operational Research, vol. 274, no. 2, pp. 743-758, 2019.

[12] C. F. Tsai, Y. F. Hsu and D. C. Yen, "A comparative study of classifier ensembles for bankruptcy prediction", Applied Soft Computing, vol. 24, pp. 977-984, 2014.

[13] C. Chen, "Interpretable Graph-Based Anomaly Detection for Related-Party Transactions in Auditing Systems", 2024.

[14] Jiang C, Ma L, Wang Z, et al. Financial distress prediction using the Q&A text of online interactive platforms[J]. Electronic Commerce Research and Applications, 2023, 61: 101292.