Empowering Supply Chain Risk Monitoring with Ontology-Guided Knowledge Graph Extraction by LLMs

Shuhan Zheng[1,*], Keita Mizushina[1] and Ken Naono[1]

[1]*Research and Development Group, Hitachi, Ltd., Tokyo, Japan*

Abstract

With business globalization and increasing product complexity, companies often operate supply chains distributed around the world. Such globally distributed supply chains face various disruption risks, highlighting the need for procurement officers to effectively monitor these risks. A popular paradigm is to apply information extraction technologies to open data for risk extraction. Here, we introduce an ontology-guided method for supply chain risk extraction that leverages large language models. Our method iteratively extracts a supply chain risk knowledge graph from unstructured open data, guided by a user-specified ontology. We also developed knowledge graph verification and formatting modules. Our wholistic methods enable consistent and automated identification and extraction of risk knowledge, thereby empowering procurement officers to monitor supply chain risks.

Keywords

Knowledge graph, Ontology, Large language model, Supply chain,

1. Introduction

Business background: Modern enterprise supply chains are distributed around the world. With the broad geographical distribution, supply chains are facing various disruption risks, ranging from natural hazards to regional conflicts [1, 2]. To achieve efficient supply chain risk monitoring, extracting semantic information of risk events from a large volume of open data (e.g., news, government report) is a promising approach. By informing procurement officers (POs) about risk semantic information, POs are able to take targeted countermeasure to mitigate potential risk impacts on their supply chains.

Investigations on existing solutions: There has been a sustained interest in the supply chain research community in developing techniques to extract supply chain risks. Our literature investigation identifies two main categories of approaches. The first category comprises ontology-based semantic methods. These methods utilize static ontology metadata as data schema to organize extracted risk data [3, 4]. The second category reflects the emerging trend of leveraging artificial intelligence (AI) for risk extraction and assessment. This includes the application of deep learning models [5, 6], and state-of-the-art natural language processing (NLP) models [7, 8], as evidenced in a recent survey [9]. Although previous methods differ in detailed algorithms and input data, they share two notable drawbacks.

First, previous ontology-based semantic approaches did not effectively incorporate ontology within the risk extraction process. While these methods rely heavily on experts to build the ontology, the resulting ontology serves as a static metadata of data schema [3]. It does not play an active role in guiding or improving the extraction, which limits the integration of expert knowledge in the extraction process.

Second, recent AI-based models face challenges related to low explainability, limited interpretability, and poor adaptability to dynamic environments [9]. In supply chain risk assessment, stakeholders require clear insights into why a model makes a specific prediction or decision. However, many existing AI models function as black boxes. Furthermore, the supply chain environment is inherently dynamic and constantly evolving [10]. Adaptability is not just desirable—it is essential. Yet, many AI-based models are trained for a specific domain and remain fixed once trained, making them difficult to be adjusted. As a result, POs are unable to tailor these models easily to new risk challenges.

ISWC 2025 Companion Volume, November 2–6, 2025, Nara, Japan

*Corresponding author.
shuhan.zheng.zr@hitachi.com (S. Zheng); keita.mizushina.ez@hitachi.com (K. Mizushina); ken.naono.aw@hitachi.com (K. Naono)
0000-0002-2349-0984 (S. Zheng)

2. Methods

We propose an ontology-guided, iterative method that uses Large Language Models (LLMs) to extract a supply chain risk knowledge graph (KG), addressing the limitations of previous approaches. Compared with similar methodologies [11, 12], our method is distinguished by two key innovations: (**i**) the ontology is user-specified, allowing it to be tailored to specific risk-monitoring needs; and (**ii**) the ontology guides the LLM in a stepwise, chain-of-thought [13] iterative extraction process.

The process consists of three main stages:

1. **Ontology Specification.** An ontology should be specified to guide the LLM in the knowledge graph extraction process. While we provide templates (e.g., *NaturalHazard* ontology), users can create custom ontology like *MilitaryConflict* (e.g., involving defined labels like `ConflictEvent`, `ConflictSideA`, `ConflictSideB`) for monitoring conflict risks.
2. **Iterative Knowledge Extraction.** The extraction begins when source documents (e.g., news) are collected. Initially, the LLM is provided with a system prompt containing the complete ontology and instructions for KG extraction. Then, the extraction process iterates through each relation triple defined in the ontology (e.g., `<ConflictSideA>-involvesIn-<ConflictEvent>`). In each iteration, the LLM is prompted with the target relation triple, the source document, and the KG extracted so far. The LLM extracts new KG nodes and relationships in the text. An illustration is given in Figure 1. This iterative feedback loop creates an explicit chain-of-thought process, progressively assembling a comprehensive and context-aware KG.
3. **Verification and Formatting.** Here, we verify and format extracted nodes and relationships. First, a dedicated judge LLM verifies nodes and relationships against the source text. The judge LLM returns a confidence score reflecting how faithfully the relation is supported by the text. Nodes and relationships with scores below a threshold are discarded, while higher-scoring ones are kept. Second, another LLM formats key entities to a standard format (e.g., dates to ISO 8601) to enhance consistency across the final KG.

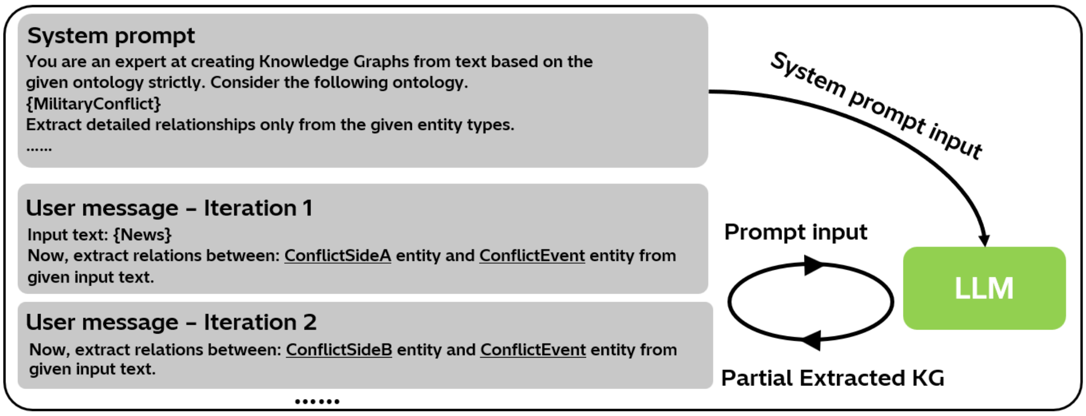

Figure 1: Example prompts for the iterative knowledge extraction process. {MilitaryConflict} is the ontology in text format. At each iteration, one target relation triple is included in the user message.

3. Applications and Case Study Validation

The adaptability of our ontology-guided method enables diverse applications. By allowing users to specify the ontology, the resulting KG can be tailored to unique monitoring needs. We demonstrate this with the following applications:

1. **Versatile Risk Monitoring:** Our method allows for supply chain risk monitoring adaptable to various risk domains by changing the guiding ontology.
   - **Natural Hazards:** Guided by a *NaturalHazard* ontology (see Figure 2), the LLM extracts natural hazard risk events from open data. The resulting KG, containing location information, can be linked to supplier locations, alerting POs to potential disruptions. For explainability, source URLs are included as metadata, enabling information traceability.

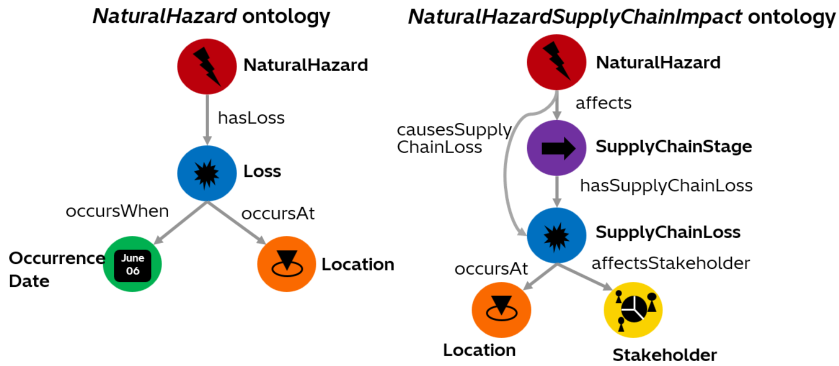

Figure 2: Example ontology. The left-hand-side shows the *NaturalHazard* ontology, which represent and organize information about natural hazards that pose threats. The right-hand-side shows the *NaturalHazardSupplyChain-Impact*, which focuses more on the impact of natural hazards on supply chains.

- **Conflicts:** To monitor conflict risks, users can design a *Conflict* ontology (e.g., with defined labels like `ConflictEvent`, `ConflictSideA`, `ConflictSideB`), guiding the LLM to extract relevant knowledge from news data.

2. **Impact Annotations:** Our method supports analyzing the potential impact of future risks using historical data. For instance, to help domestic procurement teams contextualize overseas events (e.g., a cyclone in India), a *NaturalHazardSupplyChainImpact* ontology (see Figure 2) can guide the LLM to extract the semantic details of how past cyclones have affected supply chains in that region. This structured historical knowledge enables the annotation of future risk alerts with meaningful impact predictions for better-informed decision-making.

To validate our method, we introduce a brief case study of **Impact Annotations**. We focused on cyclone hazards in the Indian Ocean region. We collected over 7,000 news articles via an open data API and applied the *NaturalHazardSupplyChainImpact* ontology to guide an LLM in extracting a KG. The process yielded a KG with 30,881 entities and 51,893 relationships. From this graph, we generated 9,592 specific impact annotations through Cypher queries [14], including 2,058 detailing power disruptions and 839 related to flight suspensions. This case study illustrates our framework's ability to produce rich, actionable impact annotations, enabling POs to better anticipate and manage future supply chain disruptions.

4. Conclusions and Future Works

We proposed an LLM-based, ontology-guided method for supply chain risk monitoring. Our ontology-guided method provides semantically rich knowledge graph extraction and achieves adaptability to different risk scenarios.

Future works include modeling the severity of events to make the KG not only semantically rich but also quantitatively informative. Enhancing data orchestration with active metadata is another direction. To fully realize the potential of the proposed method, developing a data platform capable of supporting active ontology metadata is necessary.

Acknowledgment

We would like to thank Qi Xiu (Hitachi, Ltd.) for providing mentorship to Shuhan during his first year of research. We also thank Tsunehiko Baba (Hitachi, Ltd.) and Mika Takata (Hitachi, Ltd.) for their supports during this research.

Declaration on Generative AI

During the preparation of this work, the author(s) used ChatGPT in order to: Grammar and spelling check. The author(s) reviewed and edited the content as needed and take(s) full responsibility for the publication's content.

References

[1] M. Qin, C.-W. Su, M. Umar, O.-R. Lobonţ, A. G. Manta, Are climate and geopolitics the challenges to sustainable development? novel evidence from the global supply chain, Economic Analysis and Policy 77 (2023) 748–763.

[2] J. Woetzel, M. Rogers, B. Boland, H. Samandari, Could climate become the weak link in your supply chain?, McKinsey Sustainability (2020). URL: https://www.mckinsey.com/capabilities/sustainability/our-insights/could-climate-become-the-weak-link-in-your-supply-chain.

[3] S. Emmenegger, K. Hinkelmann, E. Laurenzi, B. Thönssen, Towards a procedure for assessing supply chain risks using semantic technologies, in: International Joint Conference on Knowledge Discovery, Knowledge Engineering, and Knowledge Management, Springer, 2012, pp. 393–409.

[4] C. Palmer, E. N. Urwin, A. Niknejad, D. Petrovic, K. Popplewell, R. I. Young, An ontology supported risk assessment approach for the intelligent configuration of supply networks, Journal of Intelligent Manufacturing 29 (2018) 1005–1030.

[5] M. M. Bassiouni, R. K. Chakrabortty, O. K. Hussain, H. F. Rahman, Advanced deep learning approaches to predict supply chain risks under covid-19 restrictions, Expert Systems with Applications 211 (2023) 118604.

[6] S. Chen, Y. Song, Y. Ding, M. Zhang, R. Nie, Using long short-term memory model to study risk assessment and prediction of china's oil import from the perspective of resilience theory, Energy 215 (2021) 119152.

[7] A. Deiva Ganesh, P. Kalpana, Supply chain risk identification: a real-time data-mining approach, Industrial Management & Data systems 122 (2022) 1333–1354.

[8] M. Shahsavari, O. K. Hussain, M. Saberi, P. Sharma, Event identification for supply chain risk management through news analysis by using large language models, The Review of Socionetwork Strategies 18 (2024) 255–278.

[9] M. A. Jahin, S. A. Naife, A. K. Saha, M. Mridha, Ai in supply chain risk assessment: A systematic literature review and bibliometric analysis, arXiv preprint arXiv:2401.10895 (2023).

[10] A.-M. Nitsche, C.-A. Schumann, B. Franczyk, K. Reuther, Artificial intelligence inspired supply chain collaboration: a design-science research and system dynamics approach, in: 2021 IEEE International Conference on Engineering, Technology and Innovation, IEEE, 2021, pp. 1–9.

[11] N. Mihindukulasooriya, S. Tiwari, C. F. Enguix, K. Lata, Text2kgbench: A benchmark for ontology-driven knowledge graph generation from text, in: International Semantic Web Conference, Springer, 2023, pp. 247–265.

[12] Z. Cauter, N. Yakovets, Ontology-guided knowledge graph construction from maintenance short texts, in: Proceedings of the 1st Workshop on Knowledge Graphs and Large Language Models, 2024, pp. 75–84.

[13] J. Wei, X. Wang, D. Schuurmans, M. Bosma, B. Ichter, F. Xia, E. H. Chi, Q. V. Le, D. Zhou, Chain-of-thought prompting elicits reasoning in large language models, in: Proceedings of the 36th International Conference on Neural Information Processing Systems, 2022, pp. 24824–24837.

[14] N. Francis, A. Green, P. Guagliardo, L. Libkin, T. Lindaaker, V. Marsault, S. Plantikow, M. Rydberg, P. Selmer, A. Taylor, Cypher: An evolving query language for property graphs, in: Proceedings of the 2018 International Conference on Management of Data, 2018, pp. 1433–1445.